Exploring IBM @server iSeries and AS/400 Computers

TENTH EDITION 10

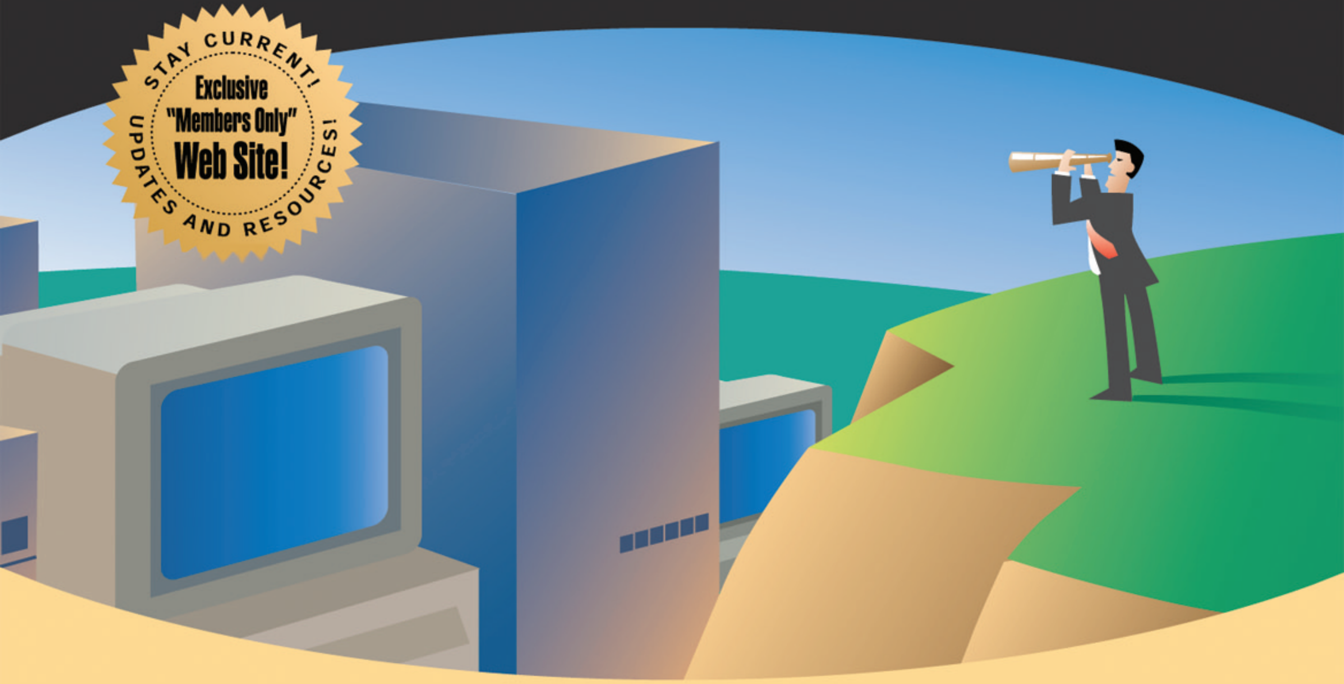

"...Anyone with the resposibility of staying close to the newest technological changes and protecting your employer's investment must have a copy."

—System/3X World

The Instant Insider's Guide to IBM's Popular Mid-range Computers

Jim Hoskins and Roger Dimmick

Table of Contents

Chapter 3: iSeries and AS/400 Software 161

Chapter 5: AS/400 and Communications—An Introduction 342

Chapter 6: iSeries, AS/400 and Your Business 397

1

What Is an iSeries or AS/400 System?

This chapter provides an overview of the IBM AS/400 family of computers, covering the highlights of these systems and then moving in for a closer look at the details. The characteristics of the AS/400 computers are compared with those of the IBM System/3X family.

The IBM Application System/400 (AS/400) family of products represents IBM's newest generation of midsize business computing systems. Like their predecessors, the System/3X family, they are **multiuser** computer systems, meaning that a single computer can interact with more than one user at a time. In developing the AS/400 systems, designers drew from the ease-of-use features of the System/36, combined these with the advanced architecture and productivity of the System/38, and then added new functions. In addition to the many application programs developed directly for execution on the AS/400, many of the application programs developed for the System/36 and System/38 computers can be migrated to and used on AS/400 systems by applying the migration tools available.

Many users have no concept of what equipment makes up the computer system they use daily. Fortunately, it is not necessary for them to, just as it is not necessary to understand the inner workings of a carburetor to drive a car. However, it is helpful to have a fundamental view of what general elements make up an AS/400. Figure 1.1 shows the components of a very simple AS/400 system configuration. The heart of the system is the system unit, which contains the "brain" that runs the computer programs and controls all activities. People interact with the computer system through terminals—or personal computers (PCs) acting as terminals—that display computer information and allow for keyboard entry.

The terminal shown on the left side of the figure is the **system console.** The system console is a specially designated terminal used by the system operator to manage the day-to-day operations of the computer system. The other terminals are for general-purpose use. The printers shown in the figure are used to generate reports, documents, graphs, and the like. A printer can be a workstation used to fill the needs of specific user(s), or all users can share it. Both terminals and printers were initially attached to the system unit via twinaxial cable (or twinax), typically laid in the building's walls or ceiling. In today's environment, the terminals (or PCs) and printers are attached through many other media including radio communications and telephone wiring.

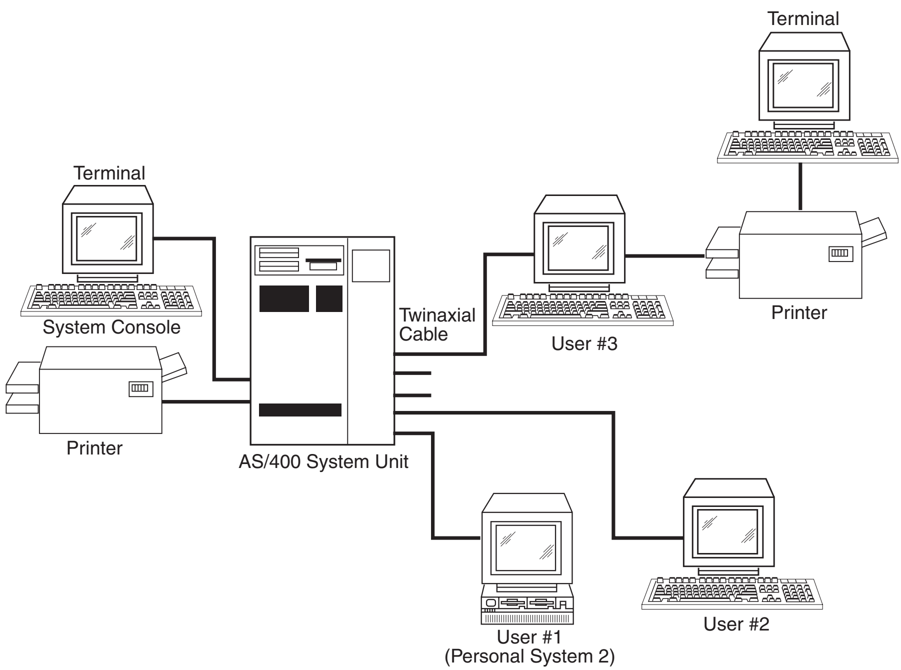

Figure 1.1. Components of a simple AS/400 system.

Figure 1.2 shows the packaging used for the AS/400e server 250. This packaging is typical of the AS/400 system line in that it encompasses the primary components of the computer system including everything necessary to compose an entry system except the terminal functions identified above. Among the elements that distinguish the system from other systems available in the midrange marketplace are the rounded rear cover which allows it to be placed against walls, and the air intake scoops at the bottom and mid-area of the unit's front. The rounded rear cover allows air exhaust and cabling to exit from the rear. And the fact that all air intake is through the front cover allows the unit to be placed against walls and desks at either side.

Application-Centric Computing

As a result of recent releases, the AS/400 system has been transformed from a host-centric computing system to an **application-centric** com-

Figure 1.2. AS/400e server 250.

puting system. This transformation has been accomplished while supporting and expanding the host-centric applications that execute on the AS/400 system. Those host-centric applications constitute the core of the 350,000 installed systems. The system configuration shown in Figure 1.1 is an example of what constitutes a **host-centric** computing system. All the computational capability resides in a centralized processor called the host, and that host is surrounded by nonintelligent terminals that provide data input to the application programs, which both reside and execute on the host.

An application-centric computing system is one that will execute and support an application's data and programs on the operating system, regardless of the operating system of origin (Windows, Macintosh, UNIX, Lotus, Java, etc.). The support may include network-centric computing systems—in which case the data and/or programs may reside on separate computing systems, some of which may be Internet servers—and the application, which is the composite of the data and program, and may execute on a different system. In an application-centric computing system, the application itself may have been defined on one vendor's computing system and be executing on a different vendor's computing system.

Consider the system configuration shown in Figure 1.1. If the terminals were replaced randomly by both AS/400 systems and PCs, with the vendors of the PCs selected at random, and the terminal connection interface was replaced by a **Local Area Network (LAN),** then the hardware portion of the application-centric paradigm would have been met. If in addition the applications to be executed within the configuration of LAN-interconnected systems could be developed on any of the systems, regardless of the vendor source, and would still execute on any other of the systems (regardless of the vendor source), then the software part of the application-centric paradigm would have been met.

Within the context of this description of application-centric computing, not only is host-centric computing included, but *open systems computing, client/server computing,* and *distributed computing* are also included. The AS/400 Advanced Series achieved the transformation to an application-centric computing system by supporting all of these forms of computing system structures. Host-centric computing was defined earlier; open systems computing, client/server computing, and distributed computing are defined in the following paragraphs and are described in greater detail in later chapters.

The AS/400 Advanced Series achieved the transformation to an application-centric computing system by supporting all of these forms of computing system structures. Figure 1.3 illustrates the architecture implemented to achieve the open and client/server portions from what was previously a host-centric base with distributed computing.

Open Systems Computing

Open systems computing has, in the minds of some people, come to mean UNIX. In reality, open systems computing means that applications developed on an open system will provide the two characteristics of **interoperability** and **portability.** *Interoperability* means that both pro-

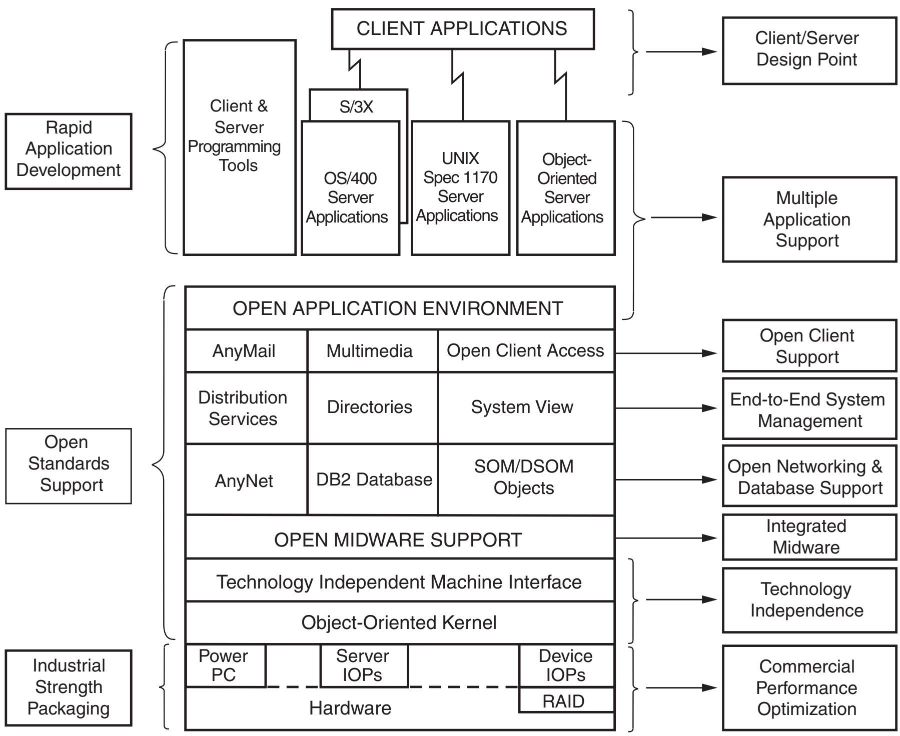

Figure 1.3. Advanced applications architecture.

grams and people can exchange information in a meaningful way. *Portability* means that it is possible to move applications, data, and users from one vendor or computer architecture to another. Together, interoperability and portability mean that an application will achieve the same results from the same data and present the users with the same interfaces at the system level—no matter what vendor's system is executing the application.

Businesses should perceive four benefits from the achievement of the open systems computing capability: freedom of choice, flexibility and change management over time, lasting value, and investment protection:

- *Freedom of choice* means that the business can select from among many vendors the hardware and software needed.

- *Flexibility and change management over time* means that businesses can recombine and redeploy their open systems applications and information technology infrastructure as business needs dictate without requiring perfect foresight. Here, they know that the application can be moved to several different platforms over its useful life.

- *Lasting value* means the business is not locked into a single vendor's hardware or software.

- *Investment protection* is provided to the business because new software and retraining are not required if the hardware base is changed.

The AS/400 has achieved open system computing by providing an open application environment, which supports interoperability and portability not only for nearly all of the other vendors' de facto standard application interfaces, but also for most UNIX applications.

Client/Server Computing Systems

Client/server computing systems are those in which applications or resources (programs, data, and sometimes both) are spread across more than one intelligent system. Usually, the client is a workstation or a PC, and the server(s) is a larger system that controls resources such as data or hardware. The outboard intelligence is called the client, and the centralized intelligence is called the server. Most of today's LAN-connected client/server computing environments include PCs or network stations and any or all of the following: file servers, fax servers, database servers, application servers, print servers, communications servers, and Internet servers.

In many of the PC-based client/server computer systems, each of the server functions is performed by a separate PC, and the user's problem is one of management to keep all of this separate activity in synchronization with the application. In the AS/400, all of the server functions are integrated and come under the control of a single operating system. As a result, performance is optimized, connectivity issues that arise from the presence of multiple servers are avoided, and when more performance is needed, users can easily add memory or other features to the AS/400e systems.

The AS/400e series supports over 1,600 different client/server applications under the many languages that are available. A list of abstracts from those applications is available from the Internet at this address: *http://www.software.ibm.com/solutions/isv*

The client/server environment for the AS/400 has been extended with the support for Novell NetWare 4.10 on the **Integrated PC Server** (formerly called the File Server Input/Output Processor, or FSIOP). This includes file serving, print serving, and data sharing with NetWare Loadable Modules (NLMs) while preserving NetWare commands to install, configure, and control the environment.

In addition, Novell's Internet Packet Exchange (IPX) protocol is now available as a native protocol in the AS/400 suite of communications protocols. This enables AS/400 to be plugged into existing Novell networks with fewer problems and administrative changes.

Distributed Computing Systems

Distributed computing systems at first glance look like client/server computing systems, but instead of a full copy of all data and programs residing on the server, the data may be segmented and may reside across many system boundaries in the network of systems, and each application program may reside on the computer at which it is normally executed.

Although this is the normal situation, the programs may be called and executed at any computer in the network. Ideally, which computer in the distributed system network of computers is actually executing the program is unknown to the user. Except in the case of some very short programs with limited data sets, the ideal has not yet been achieved in most distributed computing system networks.

Advanced Application Architecture

Figure 1.3 illustrates the **advanced application architecture,** which was implemented on the AS/400 system to achieve the open client/server computing system capabilities described in the preceding paragraphs. The following paragraphs introduce the layers of this architecture. Later chapters of this book provide additional detail for the reader.

There are seven layers in the advanced application architecture for client/server and distributed computing. The *top layer* is the client application layer. This primarily consists of Client Access Family for Windows for AS/400 and Client Access Family for AS/400, the PC-resident offering that allows the widest variety of clients to take advantage of the AS/400 server resources. This includes support for both the Microsoft ODBC interface and the Java DBC interface at both the client and at the server, as well as **Apple's** DAL. Client operating systems supported are:

- Extended DOS
- 32-bit Windows
- Macintosh SNA*ps
- Novell Windows 3.1
- IBM Connection
- 16-bit OS/2
- Windows NT
- Windows 95/98
- Windows 3.1

- Program/400 for RS/6000
- Base DOS
- UNIX (SUN or HP)
- 32-bit OS/2

The *second layer* is the server layer, which supports client and server programming tools, OS/400 server applications, PC server applications, **A-OPEN** (UNIX Spec 1170) server applications, and **object-oriented** server applications. In general, this layer enables the distributed computing function, whereby the client application might also execute at the server in whole or in part. This layer in combination with the next layer down (the open application environment) enables the AS/400 to support more than 80 percent of the most commonly used commercial APIs in UNIX applications.

The fact that these server functions are equivalent to each other allows them to share data through the integrated file system, including the DB2 for AS/400 relational database. This, along with Data Propagator/400, allows flat files and relational files to coexist and PC applications to readily access data, change it, and put it back again, as well as replicate it in whole or in part to a different application. Thus, legacy applications, client/server applications, and object-oriented applications can coexist in one AS/400.

Implemented within the *third layer* as part of the integrated file system are triggers, stored procedures, declarative referential integrity, two-phase commit, and long file names. These functions are available not only for the client/server and open systems interconnect file structures but for all of the file structures supported, improving the total function of the AS/400 including legacy applications.

The *fourth layer* is the integrated midware layer, which allows the AS/400 to reduce the system management headaches associated in general both with client/server computing and with distributed computing. The integrated midware layer reduces complexity in the following areas:

- Network protocols
- Database management

- Security
- Access to coded and noncoded data (open file system)
- Enablers for advanced applications such as multimedia
- Enablers for mail and directory services

This midware is integrated and tested before it is delivered on the system as part of OS/400.

Among the tools available in the midware are Systems Manager, Managed Systems Services, OS/2 Warp Server for AS/400, and ADSM/400. Systems Manager and Managed Systems Services allow the AS/400 to deliver new releases of PC applications to all authorized users. OS/2 Warp Server for AS/400 allows the AS/400 and Novell networks to share resources such as printers and storage devices and allows central administration of the networks. And ADSM/400 allows the backup and recovery services of the AS/400 to be extended to PC users.

The *fifth layer* is the technology-independent machine interface, which allows the AS/400 to change major hardware and software components of the system without affecting business applications. This allows the AS/400 to change the functions and hardware below this interface without causing customers to rewrite or recompile their applications.

The *sixth layer* is the System Licensed Internal Code (SLIC), now designed in C++, which runs on 64-bit microprocessors. This layer allows new hardware to be introduced without affecting the applications above the machine interface layer until those applications are ready to exploit the new functions provided.

The final layer, the *seventh layer,* is the hardware layer. In the August 1997 release of this layer, integrated DASDs were continued and a RAID (Redundant Array of Independent DASDs) capability for the Multi-Function I/O Processor (MFIOP)-attached integrated DASDs were introduced. The Integrated NetFinity Server, and a family of IOP, and controllers, are included under the new systems **PCI** (Portable Computer Interface) packaging.

Because RISC is the new evolving technology for the AS/400 family main processor, before discussing the remainder of the hardware, we will spend a few paragraphs briefly explaining what RISC technology is and the PowerPC AS implementation of that technology. Then we will point out the differences from other RISC implementations necessitated by the commercial processing nature of the AS/400 family.

Why RISC?

Although additional improvements could have been achieved with the **CISC** (Complex Instruction Set Computer)-based microprocessor family, which constituted the base for all of the pre-1995 AS/400 systems, the AS/400 was moved to a **RISC** (Reduced Instruction Set Computer) base because RISC provides extended future growth, is mainstream and strategic, can be optimized for commercial usage, and offers several advantages at a complete system level. This involves more than chips—RISC also better enables an optimizing compiler and simplifies the instruction decode function.

Before discussing those driving elements, let's define CISC and RISC somewhat more fully. As implemented in the IMPI (Integrated Micro Programmed Interface), the CISC microprocessor had 392 instructions that were 2, 4, or 6 bytes in length, including a group of eight different complex operations. The RISC PowerPC AS (PowerPC AS is the specific implementation of the PowerPC RISC architecture implemented for the AS/400) has 250 instructions, all of which are 4 bytes in length and require only two move-assist complex operations.

The net result is that the RISC PowerPC AS allows for growth to a larger address space, improved performance from 64-bit data and instructions, more than 20 GB (gigabytes, or billion bytes) of main storage, a larger page size (4 KB), and significant I/O growth (greater than 4 GB addressability, more I/O buses, and three times faster bus performance). What this all means is that you get more processing power with RISC. That power will be needed as the computing paradigm moves toward emerging technologies such as object-oriented programming, multimedia computing, network computing, and distributed computing.

The RISC PowerPC AS technology for the AS/400 provides synergy with the remainder of IBM and is strategic to the long-term direction of the AS/400. The question was never whether the AS/400 would move to RISC, but *when* the move would be made. The RISC PowerPC AS was determined to be compatible with commercial applications if some things could be added, including a Tags Active mode for single-level storage, decimal-assist functions, move-assist operations for fast memory management, and vectored supervisor calls.

The PowerPC AS with a wider I/O bandwidth allowed the implementation of a larger number of I/O buses, each of which could have a greater functional bandwidth. This allows systems to grow in I/O capacity commensurate with the performance capabilities of the processor and memory components.

RISC PowerPC AS Microprocessors

A microprocessor is a computer chip containing millions of microscopic transistors that work together to form the "brain" of a computer system. The internal structure of a microprocessor is called its architecture, and many different architectures are in use today. Current AS/400 systems (and other computers, such as the IBM RS/6000 family) use one or more PowerPC microprocessors as the basis for their computing engines. These microprocessors are special implementation of the PowerPC architecture, which is based on the Reduced Instruction Set Computer (RISC) concept. The idea behind RISC (pronounced "risk") is to gain a performance advantage by utilizing a simple set of instructions executed very quickly to do all work.

The PowerPC architecture enables high levels of performance through its superscalar design, using pipelining, hardwired operations, new op codes, and formats for optimal decoding and branch prediction to improve superscalar scheduling and cycle times.

The *op code* is the portion of the instruction that tells the microprocessor what operations to perform on the data to be manipulated by the instruction. *Superscalar* means that more than one instruction can be executed in a single cycle of the processor. *Pipelining* allows instructions to be processed as if they were on an assembly line. *Optimal decoding* adds a structure to the instruction layout that makes interpretation (within the microprocessor) faster. *Branch prediction* means that (during the compiling activity) potential jumps from one section of a program to another have been inspected to predict if they will be taken, in order to minimize the delay caused by the need to empty the pipeline if the branch requirements are met. Many compiler optimization techniques have been implemented to minimize design bottlenecks while maximizing parallelism.

AS/400 PowerPC microprocessors have on-chip cache memory. These very high speed memory areas serve as temporary storage and reduce the amount of time the microprocessor spends waiting for information to be pulled in from main storage, where access speed may be as much as thirty-five times slower. There are separate instruction and data cache areas to help streamline information flow through the microprocessor.

The AS/400 versions of PowerPC microprocessors (called RISC PowerPC AS microprocessors) have several unique features not found in the base PowerPC architecture (decimal support, move assist, vectored supervisory calls, tagged operations, etc.). These extensions help RISC-based AS/400 systems maintain and enhance commercial transaction performance.

Commercial computing workloads have different characteristics than engineering/scientific computing workloads (a traditional strength of RISC). The commercial environment typically has an increased number of concurrent users, longer instruction path lengths in both application code and the operating system, decreased predictability in branches, mostly fixed-point arithmetic functions, and randomly organized I/O activity. These differences are illustrated in Figure 1.4.

The RISC PowerPC AS systems also have a wider I/O bandwidth, allowing for AS/400 systems with more I/O buses. This makes for AS/400 systems with greater I/O capacity, which is necessary to maintain balance with the higher levels of performance the RISC PowerPC AS microprocessors enable. The five PowerPC AS microprocessors used in current AS/400 systems are the A10 Microprocessor, the A35 Microprocessor, the A50 Microprocessor, the Pulsar Microprocessor and the Istar Microprocessor.

Commercial Workloads	Scientific/Engineering Workloads
Many concurrent users	Few concurrent users
Longer path length over larger set of instructions	Smaller instructions working sets
More execution time in operating system code	More execution time spent in application work
Fewer loop iterations—more branches	Tight loops
Extensive manipulation of data structures through integer arithmetic and strong operations	Extensive use of floating point arithmetic
Random I/O activity	Sequential I/O activity

Figure 1.4. Commercial vs. scientific/engineering workloads.

The A10 Microprocessor

The **A10 Microprocessor** is a single-chip **CMOS** (Complementary Metal-Oxide Semiconductor) technology implementation that performs at a 77 MHZ cycle rate as a pipelined superscalar design with separate fixed-point, floating-point, load/store, condition register, and branch units. In entry and midsize AS/400 systems, the A10 Microprocessor has about 4.7 million transistors. It can execute up to three instructions per cycle at a peak rate of 231 **million instructions per second (MIPS).** There are two on-chip caches, a 4 KB (1 **KB,** or kilobyte, is about 1,000 bytes) instruction cache and an 8 KB data cache, and the A10 can support an optional 1 MB (1 **MB** is about 1 **million bytes**) off-chip cache. Figures 1.5 and 1.6 illustrate the interconnection paths for the memory and I/O for the various processors using the A10 Microprocessor chip as their base.

In Figure 1.5, one additional chip is used to generate the copper I/O bus used in the system unit that houses the processor, one additional chip is used to generate each pair of external optical buses used to house I/O in the external expansion towers, and the memory is driven directly by the A10 Microprocessor chip. In Figure 1.6, the I/O is driven in the same way as in Figure 1.5, but the memory bus is partitioned into two buses, each of which is supported by a separate storage control chip that manages the main storage accesses, and up to 0.5 MB of intermediate-level cache can be supported.

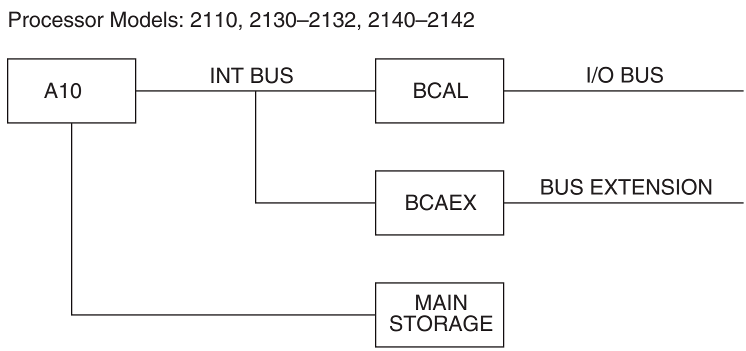

Figure 1.5. A10 Microprocessor I/O and main storage interconnect for entry and midrange processors.

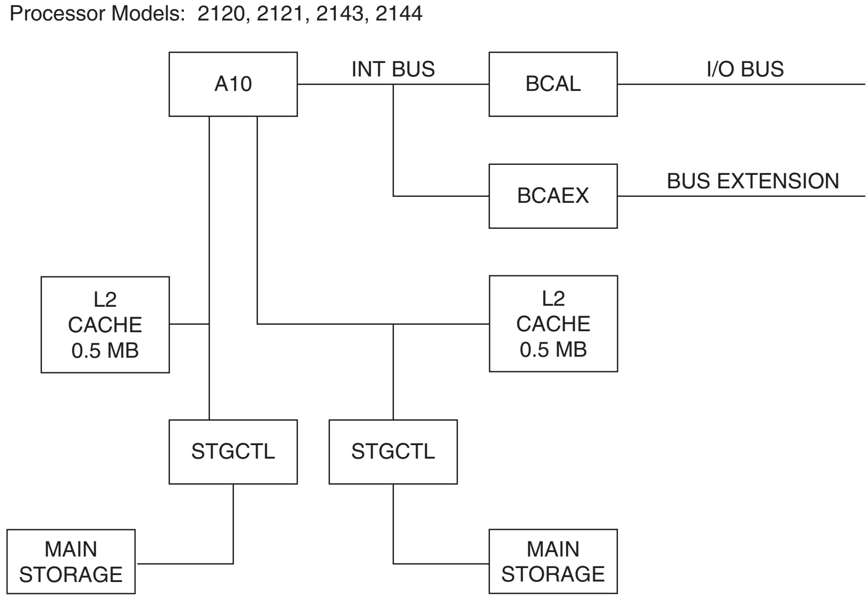

Figure 1.6. A10 Microprocessor I/O and main storage interconnect for midrange and high-end processors.

The A50 Microprocessor

The A50 Microprocessor is a single-chip CMOS technology implementation of the RISC PowerPC AS that performs at a 262 MHz cycle rate as a four-way superscalar pipelined design with separate fixed-point, floating-point, load/store, condition register, and branch units in a single chip carrier for midsize to high-end AS/400 systems.

The A50 Microprocessor also comes in two different multi-chip carriers that include the A50 Microprocessor and a single chip-storage controller and I/O hub for the low-end AS/400 systems. The first package does not have support for an external cache and the processor operates at 200 MHz. The second package supports a 4 MB external cache and operates at about 255 MHz.

The A50 Microprocessor has about 12.7 million transistors. It can execute up to four instructions per cycle with a cycle time of 3.81 nanoseconds. There are two on-chip caches, a 64 KB instruction cache and a 64 KB data cache. There is an optional external four-way set-associative cache with a 32-byte high-speed data-pumped interface that can support a 4 MB or an 8 MB off-chip cache. This cache-to-processor performance allows accessing 32 bytes at up to 262 MHz with a system bus bandwidth of 16 bytes at up to 87.3 MHz.

In low-end and midsize AS/400 systems (Figure 1.7), the A50 microprocessor is connected to main storage through a storage control chip, and to the I/O interfaces through a bus control adapter logic (BCAL) chip. The BCAL chip creates an intermediate interface from which other chips create a PCI bus interface. The base system unit in these systems always contains a PCI bus interface.

In the high-end symmetrical multiprocessor system units, the main storage is again managed through storage control chips, but the storage cards are arranged in either two or four banks and the I/O interface is driven across a high-speed serial/parallel interface bus referred to as an HSL (High-Speed Link) bus. The implementation of the HSL bus allows the I/O devices to be packaged in a separate unit from the System Processor/Memory unit. The E-Server Model 820 uses the separation to package the I/O in a separate tower.

The Pulsar Microprocessor

The Pulsar Microprocessor is the A50 Microprocessor implemented with copper interconnections technology. The major difference as a result is reduced resistance in the interconnections between transistors. This allows the Pulsar Microprocessor to run at clock speeds of 400 and 450

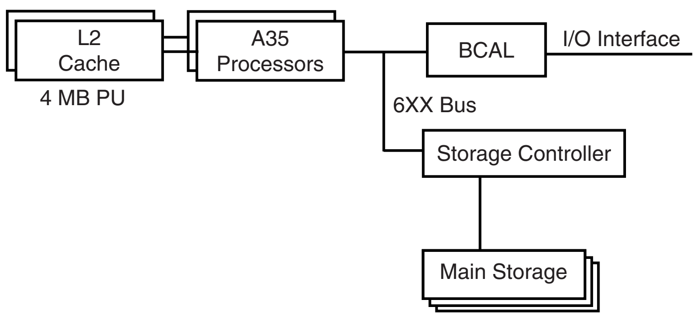

Figure 1.7. A35 Microprocessor/main storage and I/O interface for 2136 and 2119 Processors.

MHz compared to the A50 clock speed of 262 MHz. The Pulsar Microprocessor is used in the Model 270 and lower-performance Model 820 systems, and incorporates a high-speed link for connection to memory and I/O interfaces. The high-speed link is discussed in greater detail under the section on hardware architecture later in this chapter.

The IStar Microprocessor

The IStar Microprocessor is the Pulsar Microprocessor with clock speeds of 500 and 540 MHz. The IStar Microprocessors use the silicon-on-insulator and copper-interconnection technologies. The Silicon-on-insulator technology reduces the junction capacitance in the transistor switches, increasing performance while decreasing the power dissipation on the chips. The IStar Microprocessor is used in the higher-performance Model 820 systems and the Model 830 and 840 systems. All versions of the IStar Microprocessor support the high-speed link, which is described under hardware architecture later in this chapter.

Meet the Family

Twelve basic computers form the core of the IBM AS/400e series family: the entry Model 150, available in both general-purpose and server configurations; the Model 250, which expands the growth possibilities for the entry-level customer; the Model 170 server versions with seven processor options; Model 270 server versions with four processor options; the six server Models 720, 730, 740, 820, 830, and 840; the dedicated Domino Servers Models 170D, 270D, and 820D; and the custom application servers SB1, SB2, and SB3. The Domino Servers and the custom application servers are special cases either through hardware constraints or software constraints of previously identified models.

IBM has chosen to separately classify the 270, 820, 830, 840, the dedicated Domino servers 270D and 820D and the third tier servers SB2 and SB3 as iSeries 400 EServers. The “i” in iSeries 400 stands for intelligent integration, featuring open systems functions which are built-in to the operating system and pretested for reliable, turnkey function. The integrated functions include web servers (IBM HTTP server and Apache server), Web application server (WebSphere Standard Edition), Java Virtual Machine (JVM), database (IBM DB2

Universal database for AS/400), communications (TCP/IP), OS/400 PASE, e-business security, and Linux in the future. This intelligent integration enables businesses to deploy solutions faster, with greater reliability, and lower cost of ownership while providing superior performance across a wide range of e-business applications based on industry standards. Eservers as a group offer: new tools for managing e-business, application flexibility, and innovative technology which offers leading edge server performance. The iSeries 400 systems microprocessors are based on copper interconnect and silicon on insulator technologies, which are innovations in the semiconductor industry. These technologies result in denser packages, lower power consumption, and higher performance. Also implemented in these systems is the a memory switching technology with switch speeds up to 36 GB per second. High Speed Links on these systems allow data transfer between I/O devices and other systems at rates up to 1 GB per second. The iSeries intelligent integration of these technologies allows businesses to benefit from enterprise class computing without the "enterprise size" IT support staff.

The iSeries architecture features a flexible operating environment that concurrently runs any combination of AS/400, ported UNIX applications, Linux (in the future), Windows, Domino, or Java applications. When that flexibility is combined with the inherent iSeries 400 workload management capabilities, Logical Partitioning (LPAR), and the integration of **IBM @server** xSeries, the iSeries 400 enables businesses to run several diverse application environments on one physical machine. Choices are provided with regard to whether a business should expand an existing server farm, maintain the existing server farm, or consolidate server footprints on the iSeries 400 for simplified management, increased availability, or lower total cost of ownership. The iSeries 400 also offers businesses the freedom to grow without disruption through Capacity Upgrade on Demand features on the Model 840. These features enable a customer to activate immediate incremental processing power non-disruptively, when it is needed, and pay for it only when initiated.

The AS/400e server 150 continues to use the same packaging and has received some additional capabilities in the software that is available. The AS/400e server 250 expands the performance and the I/O capabilities available in the entry-level price range. The Model 170 has a new lower-performance entry-point processor, and also a two-way increased performance high-end processor. The Model 270 extends the midrange offering of the Model 170 in performance, memory, and I/O capacity.

The six server Models 720, 730, 740, 820, 830, and 840 provide increased server performance and a base interactive performance level of 35, 70, and 120 CPW, respectively, and can be featured to support a significantly larger interactive workload. The three Domino dedicated servers Models 170D, 270D, and 820D provide specially tuned Domino environments.

The number of steps that may be taken from the base interactive performance to the maximum offered by a particular processor varies with the processor. In no case can the maximum client/server processor CPW performance value be exceeded when using the featured processor in interactive mode. (See the paragraphs on benchmarking later in this chapter for additional detail.)

The capability to feature the amount of interactive performance available on each of the 720, 730, and 740 servers allows them to support the ISV preload environment previously provided by the AS/400e custom mixed-mode server Models S20-ISV, S30-ISV, and S40-ISV. Upgrades from the previous custom mixed-mode server models to the V4R5 8XX servers are supported. The custom mixed-mode software packages continue to be provided by J. D. Edwards, Software Systems Associates, Intentia International's Movex V10.5, and International Business Systems (IBS), as well as others. The preload of the above business partner software packages continue to be offered on the AS/400e server 170.

The custom application servers SB1, SB2, and SB3 are a main storage, DASD, and other-I/O-device-constrained version of the eight-way, twelve-way, and twenty-four-way processor versions of the 740, 830, and 840 servers specifically tuned to provide high performance in the compute-intensive environment required by many multitiered computing environments. The SB1, SB2, and SB3 are intended to be used as **SAP R3** application servers to a separate second-tier server.

For those unfamiliar with the AS/400e server Models 150, 170, 720, 730, and 740, the authors refer the reader to the nineth edition of this book, *Exploring IBM AS/400 Computers*.

Figure 1.8 contains a photograph of the AS/400e series family. Let us briefly look at each of these. The "e" in the AS/400e series names for these systems signifies that the systems are capable of supporting the

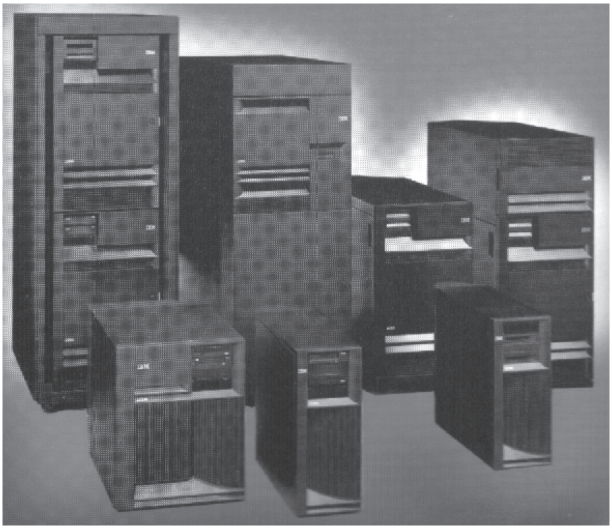

Figure 1.8. AS/400 V4R5 family showing (from left-rear) 1.8 I/O Tower, Models 840 and 830; (from left front) models 820, 270, and 250.

electronic business environment, which means that users may establish a storefront on the World Wide Web and conduct business with the assurance for both the user and the customer that all transactions are secure and private.

System Unit Members

This introduction to the members of the AS/400 family describes in greater detail each of the system unit members in the following paragraphs.

AS/400e Server Model 250 Specifics

The AS/400e Server Model 250 is the smallest, lowest-performing version of the AS/400 system. It is available in a server version with capability to perform a limited amount of interactive processing. There are two processors available: the #2295 with server workload capability of 50 CPW and interactive capacity of 15 CPW, and the #2296 with server workload capability of 75 CPW and interactive capacity of 20 CPW.

Both processors start with 256 MB of main storage and 8.58 GB of disk storage, and can have up to thirty communications lines. The base main storage can be grown to 1,024 MB and the disk storage can be grown to 175.4 GB. The software integrated with the system includes support for improved performance, new Internet functions, new versions of Client Access Family for Windows, support for Windows NT Server, new encryption and security capabilities, and wireless capabilities for a small business.

The AS/400e server Model 250 has been enhanced over the previously available AS/400 advanced entry Model 150 with faster processors, a larger main storage capacity, larger disk storage capability, higher-capacity tape drives, a higher performance CD-ROM, an improved-performance 333 MHz Integrated Network Server (INS), and support for an increased number of LANs. The Model 250 is targeted for the 2- to 240-user environment, including distributed locations and small businesses. The package, the smallest of the AS/400 systems and described in the following paragraphs, is shown in Figure 1.2 earlier in this chapter.

Restraints on positioning the system unit against walls or desks do not apply unless the user removes the rounded rear cover. Air intake is entirely from the front, and exhaust is entirely through the rear cover. The rounded rear cover allows for both cable exit and air exhaust independent of system positioning. The control panel on the front of the system unit is used by the system operator and service personnel to control the Model 250. From this panel, power can be turned on or off, the system can be initialized, and system problems can be analyzed. Many of these functions, including turning on the system power, can be performed remotely through communications lines.

The base system unit can be expanded to achieve its maximum capacities through the addition of the 7102 System Expansion Unit, which increases the base system capabilities by six disk units and supports up to nine PCI adapter cards driven by two PCI controllers and one integrated PC server. One PCI controller card is included with the base expansion unit.

The standard tape unit and CD-ROM can also be seen from the front of the system unit. The Model 250 can be plugged into a standard electrical outlet (90–140 VAC) or a high-voltage outlet (180–260 VAC) and meets “quiet office” guidelines (under 5.5 dB) for operational noise. Utility power failure protection can be added via a 9910 UPS, which will hold the system up in the presence of a utility power failure for five minutes while data from main storage are being written to disk. This, along with its attractive appearance, allows the Model 250 to fit easily beside a desk or in a corner.

The base package includes a 12/20x CD-ROM or an 8 GB ¼-inch cartridge tape unit, one 8.58 GB DASD unit, a 64-bit PowerPC Advanced System Microprocessor, and two industry-standard 168-pin **Dual In-line Memory Modules** (**DIMM**) with **Error Correction Code** (**ECC**) for a base main storage capacity of 256 MB, and slots for six Portable Computer Interface (PCI) I/O adapters, and up to three additional DASD units. The main storage of the Model 250 may be increased to 1,024 MB in three steps of 256 MB each. The system can accommodate both commercial processing, in which case some of the workstations must be twinax based, and server-based processing, in which case some of the workstations must be PCs and LAN connected, and a Communications Console must be present. A port is provided for an external UPS.

The system capacities and performance capabilities of the Model 250 are summarized in Figure 1.9. As shown in the figure, the Model 250 system has both a base unit and a #7102 expansion sidecar. The sidecar can support six additional DASD units, two external tape connections, and nine PCI card slots. The two external tape units may be either tape units or tape libraries. The available PCI slots can support a maximum of six twinax controllers, thirty communications lines, two high-speed ATM connections, two DES/L Crypto IOPs, two Integrated Network Servers, and a maximum of six LAN cards.

iSeries 400 Model 270 Specifics

The iSeries 400 Model 270 extends the price performance design point of the Model 170. Four processor options are allowed, ranging in performance from about three times the performance of the Model 170 processor to almost double the performance offered in the AS/400e server 170. A single interactive feature card is offered on each of the four processors, and if at any time the option is taken to use the interactive feature, any upgrade must include an interactive feature. It is not possible to upgrade to the Model 270 from any other model or from the Model 270 to any other model. Interactive feature cards and main storage feature cards are not customer installed; if a customer orders an interactive feature card or additional main storage, IBM will install the

Description / Function			
Processor Number	2295	2296	
Performance			
CPW-Batch	50	75	
CPW-Interactive	15	20	
Main Storage			
Min/Max-MB	256/1024	256/1024	
DIMMs (Min/Max)	2/8	2/8	
Description / Function	**Base CEC**	**Sidecar**	**Total**
DASD Storage (GB)			
Min	8.58	0	8.58
Max	70.1	105.2	175.4
DASD Arms (Max)	4	6	10
CD-ROM	1	0	1
External Tape (Max/System)	0	2	2
Tape Libraries (3)	0	2	2
Optical Libraries (Max)	0	2	2
System I/O Card Slots (PCI)	6	9	15
Workstation Attach (Max)			
Twinax Controllers	2	5	6
Twinax Devices	40/80	200	240
Communications Lines-Max	12 (1)	18	30
ATM-Max	0	2 (2)	2
Crypto IOP (DES/L)-Max	0	2	2
INS	1 (2)	1 (2)	2
LAN-Max	2	4	6
Non INS LAN-Max	1	4	5
Low Speed TR/Ethernet	1	4	5
10/100 Ethernet-Max	1 (2)	2 (2)	3
INS LAN-Max	2		4
Low Speed TR/Ethernet	2	2	4
10/100-Max	1	2	2

Notes:
1. One line used by operations console if installed
2. The INS is mutually exclusive with the high speed slot for LAN and ATM. Each INS can support two LANs.
3. The total number of external tape drives does not increase.

Figure 1.9. Advanced Server Model 250 capacities.

main storage and any other features also ordered. The iSeries 400 Model 270 has been tuned for Web, Java and data mart applications, which differentiate it from the previous Model 170 which while sufficient for running your e-business web applications had no specific characteristics for Java and data mart applications.

Like the Model 170, the base Model 270 is packaged into a small deskside unit, which is then cabled to the necessary AS/400 workstations and printers. A sidecar is supported to increase the DASD capacity of the base system. The #2252 and #2253 processor versions of Model 270 support a maximum of one high-speed link to which can be attached a small tower that can be used to increase the DASD and I/O capacity of the system. Like the predecessor models, the processor performance within the AS/400e series models can be upgraded without having to change any other element of the system. The primary differences in the Model 270 relative to the 170 server models are the much improved performance offered, the total I/O capacity of the systems, and the capability to hot plug (it is necessary to quiesce the system) the DASD and PCI cards while the system continues to operate. The Model 270 uses the same OS/400 operating system as the rest of the AS/400 family. This means that application programs and data on a 9406 Model 270 system can be used unchanged on larger AS/400e systems as a business grows.

The Model 270 processors are differentiated from each other by clock rate, by the server performance provided, and by the interactive performance capacity options. The two lowest-performance processors (#2248 and #2250) operate at a clock rate of 400 MHz, whereas the two highest-performance processors (#2252 and #2253) operate at a clock rate of 450 MHz. The #2248 processor can support a single processing thread; the three other processors each can support two threads, which with thread switching improves the performance capabilities. The #2248 processor has a CPW rating of 150 with a base interactive capacity of 25 CPW. The #2250, #2252, and #2253 processors have CPW ratings of 370, 950, and 2000, respectively, and can be purchased without an interactive card. In that case, the interactive capacity, although rated at zero, is adequate to support setting up and managing the system. One interactive feature is available for each of these processors with interactive ratings of 30, 50, and 70 CPW, respectively. The highest-performing two-way processor significantly overlaps the Model 820 processor performance in client/server computing, although it cannot be extended as far as the lowest-performing Model 820 processor in interactive performance.

The Model 270 is shown in Figure 1.2 at the beginning of this chapter and in appearance is identical to the Model 250. The Model 270 comes in a base unit and can be expanded by attaching a bolt-on expansion sidecar. In appearance, the Model 270 looks like a reduced-height Model 720 but is slightly wider when the expansion sidecar is attached. The Model 270 comes standard with a PowerPC Pulsar Microprocessor, 256 MB of main storage, a **Multi-Function I/O Processor** (**MFIOP**), an 8.58 GB 3½-inch disk drive, ¼-inch cartridge tape, a CD-ROM drive, either a console or a workstation IOA, and a communications adapter.

All processors contain a base main storage of 256 MB. The base main storage can be grown to a maximum of 4 GB with processors #2248 and #2250 by installing either pairs of 128 MB DIMMS or 256 MB DIMMS. The main storage with processors #2252 and #2253 can be grown to 8 GB, again by installing those DIMMS in pairs. The main storage expansion options used on these models attach within the processor cage so that they don't consume expansion slots.

The Model 270 base structure can house up to six disk drives in a disk cage assembly for a maximum disk capacity of 105.2 GB. The base disk cage assembly can also house one tape drive, one CD-ROM, and an operator panel. The internal tape, disk, and CD-ROM drives are all supported via an Ultra-SCSI interface. The disk capacity can be increased by the addition of a sidecar that can hold twelve additional disk drives for an additional capacity of 210.5 GB. The base also supports up to 160 twinaxial workstations spread across four controllers, twenty communications lines, and three LAN ports attached to a single IPCS. A maximum of three Crypto IOPs, either DES or PKA, can be inserted in the base. The base can also attach three external tape devices, three tape libraries, and three optical libraries.

In addition to the preceding capacities which fit in the base tower or sidecar attachment to the base tower, the #2253 and #2252 processors can attach a small tower using the High-Speed Loop (HSL). The small tower can contain six additional DASD units for an additional disk capacity of 105.2 GB, and up to eight PCI cards (one of which must be an IOP). This can increase the attached workstations to a maximum of 280, the communications lines to a maximum of fifty, and the LAN ports to a maximum of eight spread across two IPCS controllers. System capacities are summarized in Figure 1.10.

The Model 270 system unit and the system unit expansion option can be individually plugged into standard electrical outlets (110–260

Description / Function				
Processor Number	**2248**	**2250**	**2252**	**2253**
Performance				
CPW-Batch	150	370	950	2000
CPW-Interactive	25	0/30 (5)	0/50 (5)	0/70 (5)
Technology				
Clock Rate	400 MHz	400 MHz	450 MHz	450 MHz
L2 Cache	0	0	2 × 4	4 × 4
Threads	1	2	2	4
Processor	Uni	Uni	Uni	2-Way
Main Storage				
DIMMs or Cards Min/Max	2/8	2/8	2/16	2/16
Minimum (MB)	256	256	256	256
Maximum (GB)	4	4	8	8
DASD Storage (GB)				
Integrated MIN	8.58	0	0	8.58
Integrated Max	105.2	210.5	105.2	421.1
Maximum External				0
Total Maximum	105.2	210.5	105.2	421.1
DASD Arms (Max)	6	12	6	24
CD-ROM/Internal Tape (3)	2	—	0	2
External Tape (Max/Sys)	3	—	3	3
Tape Libraries (Max) (1)	3	—	3	3
Optical Libraries (Max)	3	—	4	4

Figure 1.10. Advanced Server Model 270 capacities. *(continued on next page)*

VAC). Base utility power failure protection is provided with a #9910 external CPM/UPS.

iSeries 400 Model 820 Specifics

The iSeries 400 Model 820 reuses the Model 720 package. The Model 820 can be ordered from the factory in high-availability versions (referring to the ability to continue operation in the event of component failures). The high-availability features on these models include redundant

Description / Function	Base CEC Tower	Sidecar Maximum	Small	System
Physical Packaging				
External HSL ports	2	—	—	2
External HSL loops	1	—	1	2
Small Towers supported	1	—	—	1
Large Towers supported	0	—	—	0
Embedded IOP	1	—	1	2
PCI Card Slots	7	—	8	8
Maximum				
PCI IOA Cards	6	—	7	13
Workstation Attachments (Max)				
Twinax Controllers	4	—	6	6
Twinax Devices	160	—	240	240
Communication Lines (Max) (2)	20	—	34	50
Crypto IOP (DES/PKA) (Max)	3	—	5	8
LAN ports (Max)	3	—	5	8
IPCS (Max)	1	—	2	2

Notes: 1. Total number of tape drives does not increase.
2. One line used for Operations console, if installed
3. Must be 1 CD-ROM per system.

Figure 1.10. Advanced Server Model 270 capacities. *(continued from previous page)*

power supplies, concurrent maintenance of disk/tape devices, hardware RAID-5/mirroring data protection, and hot plugging of the PCI I/O cards.

The Model 820 is a deskside system (Figure 1.11). It uses the previous Model 720 mechanical package and comes standard with a PowerPC Pulsar RISC Microprocessor #2395 (CPW performance rating 370 client/server, 35 interactive), 256 MB of main storage, a MultiFunction I/O Processor (MFIOP), four 8.58-GB 3½-inch disk drives to support the RAID default, a CD-ROM drive, a two-line PCI WAN communications adapter, and either a twinax workstation controller or a systems console on Op Console and an Alternate IPL device. The interactive performance can be enhanced by swapping the interactive base feature card #1521 for either a #1522 with 70 CPW interactive performance, a #1523 with 120 CPW interactive performance, or a #1524 with 240 CPW interactive performance.

Figure 1.11. AS/400e server 820 with expansion.

The Model 820's performance can be increased by selecting one of three processor upgrades (feature codes #2396, #2397, #2398) with CPW performance ratings of 960.0, 1,950.0, and 3,160.0 client/server, respectively, and a base interactive performance of 35 CPW. The #2396 Pulsar-based microprocessor interactive performance can be enhanced by swapping the base interactive feature card for increased interactive performance features cards #1522, #1523, or #1524 with the aforementioned interactive performance values, or #1525 with 560 CPW interactive performance. The #2397 Istar-based processor is a two-way Symmetrical Multi-Processor (SMP) implementation. The #2397 processor interactive performance can be enhanced by swapping the base interactive feature card for increased interactive performance features cards #1522, #1523, #1524, or #1525 with the aforementioned interactive performance values, or #1526 with 1050 CPW interactive performance. The #2398 Istar-based processor is a four-way SMP implementation. The #2398 processor interactive performance can be enhanced by swapping the base interactive feature card for increased interactive performance features cards #1522, #1523, #1524, #1525, or #1526 with the aforementioned interactive performance values, or #1527 with 2000 CPW interactive performance.

The Model 820 base structure can house up to twelve disk drives in two disk cage assemblies (one base and one featured). It can also house one tape drive, one CD-ROM, an operator panel, and a nine-adapter Portable Computer Interface (PCI) I/O card cage. Up to 16,384 MB of main storage can be packaged with the #2397 and #2398 processors. Up to 4,096 MB of main storage can be packaged with the #2395 and up to 8,192 MB of main storage can be packaged with the #2396 processor. The main storage expansion options used on these models attach within the processor cage so that they don't consume expansion slots. Main storage must be added in matched pairs using 128 MB, 256 MB, or 512 MB DIMM storage cards. Disk drive capacity can be increased to 4,159.1 GB.

It is also possible to attach up to 2,480 twinaxial workstations, 160 communications lines, 3 Crypto IOPs, and 30 local area networks. (Some of the LANs may be displaced by Integrated PC Servers, with a maximum of twelve allowed on a system.) In addition, eight tape libraries (either ½-inch cartridge or 8 mm tape), twelve CD-ROMS or internal tape units, and fourteen optical libraries can be attached. Figure 1.12 illustrates the maximum system capacities and performance capabilities for the Model 820.

To achieve the maximum system capacities, it is necessary to attach the #5075 PCI Expansion small tower, which can increase the total number of disk drives by six with seven additional PCI adapter card slots and/or the #5074 Disk Expansion Tower, which can hold forty-five disk units plus two internal tape or CD-ROM devices. The #5074 disk expansion tower includes internal battery backup to allow the completion of data writes in case of power failure. The expansion towers are attached using the High-Speed Link technology and may include migrations towers with attached expansion units for customers upgrading from earlier AS/400 systems. The maximum capacities of the Model 820 system can be attained only by using these expansion towers.

The Model 820 system unit can be plugged into standard high-voltage electrical outlets (180–260 VAC). Utility power failure protection is via a separately purchased external Universal Power Supply (UPS).

Description / Function					
Processor Number		**2395**	**2396**	**2397**	**2398**
Performance					
CPW-Batch		370	960	1950	3160
CPW-Interactive					
Feature #	1521	35	35	35	35
	1522	70	70	70	70
	1523	120	120	120	120
	1524	240	240	240	240
	1525		560	560	560
	1526			1050	1050
	1527				2000
Technology					
Clock Rate		400 MHz	450 MHz	500 MHz	500 MHz
L2 Cache		0	2 × 4	4 × 4	4 × 4
Threads		1	2	4	8
Processor		Uni	Uni	2-Way	4-Way
Main Storage					
DIMMs or Cards Min/Max		2/8	2/16	2/32	2/32
Minimum (MB)		256	256	256	256
Maximum (GB)		4	8	16	16

Description / Function	Base CEC	Small Tower #5075	Large Tower #5074	Migrated Total w/ *503 × (4)	System Maximum
DASD Storage (GB)					
Integrated MIN	8.58	0	0	0	8.58
Integrated Max	210.5	105.2	789.7	1625.9	4159.1
Maximum External	0	0	0	1595.3	1595.3
Total Maximum	210.5	105.2	789.7	1625.9	4159.1
DASD Arms (Max)	12	6	45	210	237
Diskette	0	0	0	2	2
CDROM/Internal Tape (3)	2	0	2	18	12
External Tape (Max/Sys)	6	7	8	8	8

Figure 1.12. Advanced server model 820 capacities. *(continued on next page)*

Description / Function	Base CEC	Small Tower #5075	Large Tower #5074	Migrated Total w/ *503× (4)	System Maximum (5)
Tape Libraries (Max) (1)	6	7	8	8	8
Optical Libraries (Max)	6	7	14	14	14
Physical Packaging					
External HSL ports	2	—	—	—	2
External HSL loops	1	—	—	—	1
Small/Large Towers supported	5	—	—	—	5
SPD Towers supported		—	—	4	5 (6)
Embedded IOP	1	1	8	1	6
PCI Adapter Card Slots	12	8	14	86	82
Maximum PCI IOA Cards	9	7	11	70	63
Workstation Attachments (Max)					
Twinax Controllers	7	7	11	66	62
Twinax Devices	280	280	440	2628	2480
Communication Lines (Max) (2)	44	34	52	128	160
Crypto IOP (DES/PKA) (Max)	3	3	3	3	3
Lan Ports Max	6	5	8	24	30
IPCS (Max) 2	2	2	2	16	12

Notes:
1. Total number of tape drives does not increase.
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system
4. Includes the migration tower and all SPD towers attached to the migration tower.
5. New build-only, does not apply to migrated system.
6. Includes migration tower.

Figure 1.12. Advanced server model 820 capacities. *(continued from previous page)*

iSeries 400 Model 830 Specifics

The Model 830 makes up the high end of the middle range of the AS/400e series family. The Model 830 processors include the two-way, four-way, and eight-way processor versions of AS/400 computers using the most powerful of the IStar Microprocessors. This provides for higher performance, more users, more communications, more storage, and greater growth capability than the Model 730.

The Model 830 uses a package common with the previous Model 640 and S/30, approximately the size of the 1.6-meter rack in the previous S/30 through F 9406 models. The Model 830 uses a High-Speed Link (HSL) for connecting to the I/O function, the base portion of which is totally contained within the system unit tower.

The Model 830 (Figure 1.13) is a tall-tower system. The Model 830 comes standard with a PowerPC IStar RISC Microprocessor #2400 (CPW performance rating 1850 client/server, 70 CPW interactive), 1,024

Figure 1.13. AS/400e server 830.

MB of main storage, a Multi-Function I/O Processor (MFIOP), four 8.58-GB 3½-inch disk drives attached to a RAID controller, a CD-ROM drive, eight HSL ports, a PCI two-line WAN communications adapter with V.90 modem, either a twinax workstation controller or a system console on Op Console option, and an alternate IPL device.

The Model 830's performance can be increased by selecting one of two processor upgrades (feature codes #2402 and #2403) with CPW performance ratings of 4,200 and 7,350, respectively. The #2400, #2402, and #2403 are, respectively, two-way, four-way, and eight-way Symmetrical Multi-Processors (SMPs). Each of these microprocessors has a base interactive performance of 70 CPW. The base interactive performance can be increased by the swapping of the base interactive feature card for one of the optional interactive feature cards.

In no case can the interactive performance exceed the client/server CPW performance value supplied by the processor. The interactive feature cards supported by the Model 830 are #1532, #1533, #1534, #1535, #1536, and #1537, offering interactive performance values of 120, 240, 560, 1,050, 2,000, and 4,550, respectively. With the #2400, #2402, and #2403 processors, the interactive performance can become as large as two-thirds of the client/server performance. The section on performance later in this chapter explains what happens when an attempt is made to get greater interactive performance than is available from the processor.

The Model 830 base structure can house up to forty-five disk drives, one tape drive, one CD-ROM, an operator panel, up to 32,384 MB of main storage , and an eleven-adapter card PCI I/O card cage. Up to four High-Speed Links can be generated for the attachment of expansion towers. The main storage expansion options used on these models attach within the processor cage so that they don't consume expansion slots. Main storage cards must be installed in matched groups of eight DIMMS. Available main storage card sizes are 128, 256, and 512 MB. It is also possible to attach up to 6,080 twinaxial workstations, 300 communications lines, 3 Crypto devices, and 72 LANs (16 Integrated PC Servers). Twenty-two optical libraries can be attached. The system capacities and performance capabilities are summarized in Figure 1.14. Maximum total disk storage capacity for the system is 11 Terabytes (TB).

To achieve the maximum system capacities, it is necessary to attach expansion towers to the High-Speed Links. The system maximum can contain a total of thirteen large towers.

The Model 830 system unit can be plugged into standard high-voltage electrical outlets (180–260 VAC). Internal battery backup and Continu-

Description / Function			
Processor Number	**2400**	**2402**	**2403**
Performance			
CPW-Batch	1850	4200	7350
CPW-Interactive			
Feature # 1531	70	70	70
1532	120	120	120
1533	240	240	240
1534	560	560	560
1535	1050	1050	1050
1536		2000	2000
1537			4550
Technology			
Clock Rate	400 MHz	540 MHz	500 MHz
L2 Cache	2 × 4	4 × 4	4 × 4
Threads	4	8	16
Processor	2-Way	4-Way	8-Way
Main Storage			
DIMMs or Cards Min/Max	8/64	8/64	8/64
Minimum (GB)	1	1	1
Maximum (GB)	32	32	32

Description / Function	Base Small Tower CEC	Large Tower #5075	Migrated Total w/ #503 × (4)	System Maximum w/#5077	(5)
DASD Storage (GB)					
Integrated MIN	8.58	0	0	0	8.58
Integrated Max	789.7	789.7	1625.9	2499.6	11055.8
Maximum External	0	0	1595.3	2473.9	2473.9
Total Maximum	789.7	789.7	1625.9	2499.6	11055.8
DASD Arms (Max)				596	630
Diskette	0	0	2	2	2
CDROM/Internal Tape (3)	2/1	2/2	18/17	18/17	18/17
External Tape (Max/Sys)	8	10	8	10	10

Figure 1.14. Advanced server model 830. *(continued on next page)*

Description / Function	Base CEC	Small Tower #5075	Large Tower #503 × (4)	Migrated Total w/ w/#5077	System Maximum (5)
Tape Libraries (Max) (1)	8	10	8	10	10
Optical Libraries (Max)	8	11	14	22	22
Physical Packaging					
External HSL ports	8	—	—	—	8
External HSL loops	4	—	—	—	4
Small Towers supported	0	—	—	—	0
Large Towers supported	13	—	—	—	13
Embedded IOP	0	0	2	0	2
PCI Adapter Card Slots	14	14	86	270	452
Maximum PCI IOA Cards	11	11	70	216	359
Workstation Attachments (Max)					
Twinax Controllers	9	11	66	175	152
Twinax Devices	360	440	2628	7000	6080
Communication Lines (Max) (2)	40	52	128	250	300
Crypto IOP (DES/PKA) (Max)	3	3	3	3	3
LAN Ports (Max)	6	8	24	48	72
IPCS (Max)	2	2	16	16	16

Notes:
1. Total number of tape drives does not increase.
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system.
4. Includes the migration tower and all SPD towers attached to the migration tower.
5. Includes second migration tower and all SPD towers attached to that tower.
6. New build only, does not apply to migrated systems.

Figure 1.14. Advanced server model 830. *(continued from previous page)*

ously Powered Main Storage (CPM) provide base utility power failure protection.

iSeries 400 Model 840 Specifics

The Model 840 makes up the high end of the AS/400 family. The Model 840 includes twelve-way and twenty-four-way versions of Istar microprocessors. Using the Istar microprocessor provides higher performance, more users, more communications, more storage, and greater growth capability than either the Model 830 or the Model 820. The Model 840 uses a common package approximately the size of the 1.6-meter rack in the previous 740 Model plus a separate I/O function tower just as in the 740 Model. Both the processor unit and the I/O function tower operate from separate 240-V line sources. The I/O function tower connects to the processor unit tower by means of a new technology interface cable called High-Speed Link (HSL), which allows a separation between any two units on a loop of up to fifteen meters.

The Model 840 reuses the same package as the Model 740 (Figure 1.15). It comes standard with a twelve-way Symmetrical Multi-Processor (SMP), PowerPC Istar RISC Microprocessor #2418 (CPW performance rating 10,000 client/server, 120 CPW interactive), 4,096 MB of main stor-

Figure 1.15. AS/400e server 840 and e-server 840-SBI.

age, a Multi-Function I/O Processor (MFIOP), four 8.58 GB 3½-inch disk drives, a CD-ROM drive, a sixteen-port HSL card, a PCI two-line **Wide Area Network (WAN)** communications adapter, either a twinax workstation controller or a system console on Op Console, and a logic cage capable of containing eleven PCI adapter cards and battery backup units.

The Model 840 can be upgraded to a twenty-four-way PowerPC Istar RISC Microprocessor #2420 (CPW performance rating 16,500 client/server, 120 interactive). The base interactive performance can be increased by swapping the base interactive processor function for one of the interactive feature processors #1541, #1542, #1543, #1544, #1545, #1546, or #1547, with interactive CPW value ratings of 240, 560, 1,050, 2,000, 4,550, 10,000, or 16,500, respectively. In no case can the interactive performance value exceed the client/server performance value. The section on performance later in this chapter explains what happens when an attempt is made to get greater interactive performance than is available from the processor.

Three new processor features offer "On Demand" capacity for ASPs, ISPs, and dynamic growth customers. If one of these processor options are ordered, activation can be temporary, at the time of server installation and after each new activation, for a period of fourteen days, and requires no IPL activity, or permanent, in which case the activation code is both posted to the Web and mailed, and requires the entry of the activation code at the server. The processor feature codes are: #2416 with eight active processors and 4 on demand standby processors, #2417 with twelve active processors and six on demand standby processors, and #2419 with eighteen active processors and six on demand standby processors. On Demand standby processors are not paid for until activated.

The Model 840 base structure can house up to forty-five 3½-inch disk drives, one tape drive, one CD-ROM, an operator panel, up to 97,152 MB of main storage in increments of four main storage cards each, and an eleven-adapter PCI card cage. The main storage expansion options used on these models attach within the processor tower so that they don't consume I/O expansion slots.

Main storage must be plugged in in groups of four cards. Main storage card sizes are 1,024, 2,048, 4,096, and 8,192 MB. It is also possible to attach up to 7,000 twinaxial workstations, 400 communications lines, 3 Crypto devices, and 96 LANs. Up to sixteen Integrated PC Servers may be used instead of the LANs, where an IPCS can contain two LANs. Twenty-six optical libraries can be attached. The #5074 Storage Expansion Units may be added to the base I/O tower to increase the total number of 3½-inch disk drives for each unit. A maximum disk capacity of 18,952 TB is supported. The total system capacities and performance capabilities of the Model 840 are summarized in Figure 1.16. To achieve the maximum system capacities, it is necessary to attach expansion towers to each of the eight possible High-Speed Loops.

The Model 840 system unit processor tower and the base I/O tower must be plugged individually into standard high-voltage electrical outlets (180–260 VAC). Internal battery backup and Continuously Powered Main Storage (CPM) provide base utility power failure protection. To extend the retention interval, the external battery backup feature must be installed.

To accommodate users with the previous nineteen-inch rack-mounted equipment, the new 8X0-based models can interface compatibly with that equipment. It is also possible to perform processor performance upgrades within the new models. These models use the same OS/400 operating system as the rest of the AS/400 family. In fact, the operating system is now bundled with the hardware, which means that when one of these-systems, whether it is used as a server or a general-purpose system, is ordered, the operating system is included with that order. For previous hardware models, an operating system license was ordered separately from the hardware. For clarity, we refer to all AS/400e servers as Application Systems. When we discuss a specific product, the specific name is used.

Server Preload Packages

The server's software preload capability continues to be supported on the Models 720, 730, 740, 820, 830, and 840. The customer can define the interactive workload desired to be supported by selecting the interactive workload feature. Applications software can be selected from a set of tested business partners. The applications software in general fits into an Enterprise Resource Planning environment. The following list identifies the companies whose software can be preloaded, the program name, a summary of application areas addressed by the program, and links to the program provider's Internet site. This software preload capability is also available on the AS/400e servers mod-

Description / Function				
Processor Number		**2418**	**2420**	
Performance				
CPW-Batch		10000	16500	
CPW-Interactive				
Feature # 1541		120	120	
1542		240	240	
1543		560	560	
1544		1050	1050	
1545		2000	2000	
1546		4550	4550	
1547		10000	10000	
1548			16500	
Technology				
Clock Rate		500 MHz	500 MHz	
L2 Cache		8 × 4	8 × 4	
Threads		24	48	
Processor		12-Way	24-Way	
Main Storage				
DIMMs or Cards Min/Max		4/16	4/16	
Minimum (GB)		4	4	
Maximum (GB)		96	96	
Description / Function	**Base CEC**	**Large Tower #5074**	**Migrated Total w/#5077**	**System Maximum (5)**
DASD Storage (GB)				
Integrated Min	8.58	0	0	8.58
Integrated Max	789.7	789.7	4292.9	18952.9
Maximum External	0	0	4260.6	4260.6
Total Maximum	789.7	789.7	4294.6	18952.9
DASD Arms (Max)			596	1080
Diskette	0	0	2	2
CD-ROM	2/1	2/2	18/17	24/26
External Tape (Max/Sys)	8	11	14	26
Tape Libraries (Max) (1)	8	11	14	26
Optical Libraries (Max)	8	11	22	26

Figure 1.16. Advanced server model 840 capacities. *(continued on next page)*

Description / Function	Base CEC	Large Tower #5074	Migrated Total w/#5077	System Maximum (5)
Physical Packaging				
External HSL ports	16	—	—	16
External HSL loops	8	—	—	8
Small Towers supported	0	—	—	0
Large Towers supported	23	—	—	23
SPD towers supported	—	—	18	19
Embedded IOP	0	0	0	0
PCI Adapter Card Slots	14	14	270	336
Maximum PCI IOA Cards	11	11	216	264
Workstation Attachments (Max)				
Twinax Controllers	9	11	175	175
Twinax Devices	360	440	7000	7000
Communication Lines (Max) (2)	40	52	300	400
Crypto IOP (DES/PKA) (Max)	3	3	3	3
LAN port (Max)	6	8	72	96
IPCS (Max)	2	2	16	16

Notes:
1. Total number of tape drives does not increase
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system.
4. Includes the migration tower and all SPD towers attached to the migration tower.
5. Includes second migration tower and all SPD towers attached to that tower.
6. New build only, does not apply to migrated system.

Figure 1.16. Advanced server model 840 capacities. *(continued from previous page)*

els 170 and 270. IBM Solutions Packaging website = *http://www.ibm.com/as400/developer/packaging*.

SSA—BPCS Client/Server or Mixed Mode

BPCS Client/Server is a comprehensive set of integrated client/server applications that addresses the core system needs of industrial-sector enterprises on a global scale. BPCS Client/Server covers Configurable Enterprise Financials applications, Supply Chain Management applications, Multi-Mode Manufacturing applications, and CIM applications, as well as Electronic Commerce applications such as EDI. BPCS Client/Server is based on SSA's proven object technology foundation, which ensures that it operates identically from an end-user's perspective across any supported AS/400 server.

SSA website = *http://www.ssax.com*

MAPICS—MAPICS XA

MAPICS is the most widely used manufacturing software system in the world. Its modular design offers more than forty applications to custom-configure the right solution for a specific manufacturing situation.

MAPICS website = *http://www.mapics.com*

JBA International—JBA System 21

This software contains an integrated enterprise-wide suite of supply chain applications that provide manufacturing, financial, service management, customer service, and logistics solutions with expanded functional fit for the automotive, beverage, food, and apparel and footwear sectors.

JBA website = *http://www.jbaworld.com*

International Business Systems—IBS–ASW

ASW covers an extensive range of functions in modules for:

- Management information
- Data warehousing

- Sales and distribution
- Financial management
- Manufacturing

IBS website = *http://www.ibs.se* and *http://www.ibsus.com*.

San Francisco Project (discussed in Chapter 4) = *http://www.softmall.com/sf*

Intentia—The Movex System

The Movex System is a fully integrated ERP system with a high degree of functionality for manufacturers and distributors.

Intentia website = *http://www.intentia.com*

Baan—BaanERP

BaanERP is a suite of components including sales, warehousing, purchasing, manufacturing, projects and service, finance, technology, and architecture. Baan is a world leader in innovative application solutions for the extended industry with over 4,000 customer sites in over sixty countries. The BaanERP solution can be preloaded on the AS/400e servers (except Model 150).

Baan/400 newsletter (subscribe by emailing) = *ibmbaan@us.ibm.com*

Baan website = *http://www.baan.com*

Infinium—Infinium Suite for IBM AS/400

The Infinium Series for the IBM AS/400 has been designed, and optimized, for the AS/400. The Infinium Suite for AS/400 runs native to the AS/400. It includes financial management, materials management, human resources, and process manufacturing solutions. The Infinium Suite for AS/400 includes five packages, which may be ordered separately or as a suite:

- Infinium Human Resources
- Infinium Financial Management
- Infinium Materials Management

- Infinium Process Manufacturing
- Infinium AS/400e Extensions

Infinium website = *http://www.infinium.com/html/as'400'products.html*

Acacia Technologies—Acacia Suite for AS/400

Acacia Technologies offers a full suite of enterprise-wide applications and tools designed to run manufacturing and distribution enterprises. These products include Enterprise Resource Planning (ERP), Warehouse Management System (WMS), Advanced Planning and Scheduling (ASP), Data Warehousing, and a variety of powerful tools. All Acacia Technologies products are designed exclusively for the AS/400 client/server platform. Also, all current releases of Acacia products are year 2000 compliant.

Acacia website = *http://www.acaciatech.com*

Lilly Software Associates—Visual Manufacturing

Visual Manufacturing is Lilly Software Associates' comprehensive ERP solution for manufacturers. Designed to accommodate job-shop, engineer-to-order, make-to-order, make-to-stock, and mixed-mode manufacturing environments, VISUAL Manufacturing provides the technology needed to improve productivity, control costs, and increase profits.

Lilly Software Associates website = *http://www.lillysoftware.com*

SAP-Base R/3 Package

For over two years, IBM has been delivering preloaded R/3 software to AS400 customers all over the world. This offering provides a quick, simple start in working with SAP R/3. It provides the option to have the base R/3 application installed in a two-tier or three-tier configuration, offering a savings of at least three fewer days of on-site install time. Upon arrival, the system is ready for consultants to begin R/3 customization and configuration.

SAP—Ready to Run R/3

Customers will find SAP's Ready to Run R/3 solution on AS/400e to be among the easiest to use and fastest to deploy, utilizing a unique selection of hardware and software management tools and support materials as part of the package. Ready to Run R/3 offerings on AS/400e in easy-to-order packages for 25, 50, 75, 100, 150, 250, and 500 R/3 users are available as a single footprint solution, eliminating the hidden multi-tier costs of competitive offerings. Predefined R/3 Users and Profiles, Operating Modes, System Maintenance and Backup Schedules, and Alert Monitor threshold values are just a few of the software customization options available as part of the preload package.

SAP website = *http://www.sap.com.menu*

Custom Application Servers

The custom application servers are the SB1, SB2, and SB3—second-tier servers for compute-intensive environments. The SB1 was described in *Exploring IBM AS/400 Computers,* 9th edition (IBM document G325-0400, Maximum Press), which for the sake of brevity will not be repeated in this edition. As a result of their intended use in a compute-intensive environment, the custom application servers have a large main storage with tight restrictions on the other I/O functions available.

iSeries 400 Custom Application Server Model SB2

The Model SB2 is a constrained version of the eight-way processor Model 830 and comes standard with an eight-way PowerPC Istar RISC Microprocessor (125,888 normalized FI Dialog Steps per hour at 65 percent CPU utilization), 12,288 MB of main storage, a Multi-Function I/O Processor (MFIOP), four 8.58-GB 3½-inch disk drives, a CD-ROM drive, two High-Speed Loops, a PCI two-line WAN communications adapter with V.90 modem, either a Twinax workstation controller or a system console on Op Console, and a PCI card cage capable of eleven PCI adapters, one of which must be a RAID controller.

The Model 830-SB2 base structure can house one tape drive, one CD-ROM, and an operator panel. Main storage may not be increased beyond the base 12,288 MB. It is also possible to attach up to twenty-eight twinaxial workstations, thirty-two communications lines, three Crypto devices, and six LANs. Two optical libraries or four external tape libraries can be attached. The Model SB2 capacities are summarized in Figure 1.17.

Description / Function	
Processor Number	**2315**
Performance	
CPW-Batch	7350
CPW-Interactive	0
Technology	
Clock Rate	540 MHz
L2 Cache	8 × 4
Threads	16
Processor	8-Way
Main Storage	
DIMMs or Cards Min/Max	48/48
Minimum (GB)	12
Maximum (GB)	12

Description / Function	**Base CEC**	**Migrated Total w/#5077**	**System Maximum**
DASD Storage (GB)			
Physical Base	34.3		34.3
Physical Max	70.1		70.1
Logical Max (Raid 17.54 GB)	52.6		52.6
DASD Arms (Max)	4		4
Diskette	0	2	2
CDROM	2/1	3	3
External Tape (Max/Sys)	5	4	5
Tape Libraries (Max) (1)	4	4	4
Optical Libraries (Max)	2	2	2

Figure 1.17. Third tier server Model SB2 capacity. *(continued on next page)*

The Model SB2 system unit processor tower and the base I/O tower must be plugged individually into standard high-voltage electrical outlets (180–260 VAC). Internal battery backup and Continuously Powered Main Storage (CPM) provide base utility power failure protection.

Physical Packaging			
External HSL ports	4	—	4
External HSL loops	2	—	2
Small Towers supported	0	—	0
Large Towers supported	0	—	0
5077 Towers supported	1	—	1
Embedded IOP	14	—	14
PCI Adapter Card Slots	11	—	11
Maximum PCI IOA Cards	3	2	3
Workstation Attachments (Max)			
Twinax Controllers	1	1	1
Twinax Devices	28	28	28
Communication Lines (Max) (2)	32	12	32
Crypto IOP (DES/PKA) (Max)	3	1	3
LAN port (Max)	6	6	6
IPCS (Max)	2	2	2

Notes: 1. Total number of tape drives does not increase.
2. One line used for Operations console, if installed.
3. Must be 1 CDROM per system.

Figure 1.17. Third tier server Model SB2 capacity. *(continued from previous page)*

To extend the retention interval to forty-eight hours, it is necessary to install the external battery backup feature.

The SB2 application server runs SAP-BaseR/3 Package and BaanERP software meeting the full range of business requirements, including financial accounting and controlling, sales and distribution, materials management, production planning, and human resources management. Information and early-warning systems are also available. Business Information Warehouse can edit external and internal data to support decision making at necessary corporate levels.

iSeries 400 Custom Application Server Model SB3

The Model SB3 is a constrained configuration of the Model 840 and comes standard with a twelve-way PowerPC Istar RISC Microprocessor (185,533 normalized FI Dialog Steps per hour at 65 percent CPU utilization), 16,384 MB of main storage, a MultiFunction I/O Processor (MFIOP), four 8.58-GB 3½-inch disk drives, a CD-ROM drive, four High-Speed Loops, a PCI two-line WAN communications adapter with V.90 modem, either a twinax workstation controller or a system console on Op Console, and a PCI card cage that can hold eleven PCI I/O adapters.

The Model SB3's performance may be increased by swapping the twelve-way processor with a twenty-four-way processor (feature code #2318) for 325,878 normalized FI Dialog Steps per hour at 65 percent CPU utilization. Both of the processor configurations supported on the Model SB3 are Symmetrical Multi-Processors (SMPs).

The Model SB3 base structure can achieve a maximum disk capacity of 122.7 GB to enable mirroring of the disk data, one tape drive, one CD-ROM, and an operator panel. Main storage on the #2318 processors accommodates 24,576 MB of main storage. It is also possible to attach up to twenty-eight twinaxial workstations, thirty-two communications lines, three Crypto devices, and four LANs with two IPCS controllers. Two optical libraries can be attached, as can four tape libraries. The Model SB3 capacities are summarized in Figure 1.18.

The Model SB3 system unit processor tower and the base I/O tower must be plugged individually into standard high-voltage electrical outlets (180–260 VAC). Internal battery backup and Continuously Powered Main Storage (CPM) provide base utility power failure protection. To extend the retention interval, it is necessary to install the external battery backup feature.

The SB3 application server runs SAP-BaseR/3 Package and BaanERP software, meeting the full range of business requirements, including financial accounting and controlling, sales and distribution, materials management, production planning, and human resources management. Information and early-warning systems are also available. Business Information Warehouse can edit external and internal data to support decision making at necessary corporate levels.

Description / Function			
Processor Number		**2316**	**2318**
Performance			
CPW-Batch		10000	16500
CPW-Interactive		0	0
Technology			
Clock Rate		500 MHz	500 MHz
L2 Cache		8 × 4	8 × 4
Threads		24	48
Processor		12-Way	24-Way
Main Storage			
DIMMs or Cards Min/Max		8/8	12/12
Minimum (GB)		16	24
Maximum (GB)		16	24
DASD Storage (GB)			
Physical Base		34.3	34.3
Physical Max		105.2	140.3
Logical Max (Raid 17.54 GB)		87.7	122.7
Arms		6	8
Description / Function	Base CEC	Migrated Total w/#5077	System Maximum
Diskette	0	2	2
CD-ROM	2/1	3	3
External Tape (Max/Sys)	7	4	7
Tape Libraries (Max) (1)	4	4	4
Optical Libraries (Max)	2	2	2
Diskette	0	2	2
CD-ROM/Internal Tape (3)	2/1	3	3
External Tape (Max/Sys)	7	4	7
Tape Libraries (Max) (1)	4	4	4

Figure 1.18. Third tier server Model SB3 capacity. *(continued on next page)*

Description / Function	Base CEC	Migrated Total w/#5077	System Maximum
Optical Libraries (Max)	2	2	2
Physical Packaging			
External HSL ports	8	—	8
External HSL loops	4	—	4
Small Towers supported	0	—	0
Large Towers supported	0	—	0
5077 Towers supported	1	—	1
Embedded IOP	14	—	14
PCI Adapter Card Slots	11	—	11
Maximum PCI IOA Cards	3	2	3
Workstation Attachments (Max)			
Twinax Controllers	1	1	1
Twinax Devices	28	28	28
Communication Lines (Max) (2)	32	12	32
Crypto IOP (DES/PKA) (Max)	3	1	3
LAN ports (Max)	6	6	6
IPCS (Max)	2	2	2

Notes:
1. Total number of tape drives does not increase.
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system.
4. As only 3 controllers and 2 towers are allowed on the SB3, max of 4 HSL ports and 2 HSL loops are unable to attach towers. Other ports and loops are usable for clustering only.

Figure 1.18. Third tier server Model SB3 capacity. *(continued from previous page)*

iSeries 400 Dedicated Servers for Domino

There are three dedicated servers for Domino offered by the AS/400 product line.

Those servers include the 170D, the 270D, and the 820D. These servers have been tuned to match the requirements for Lotus Domino R5 workloads, including Domino web serving, applications, and mail. Within the design point, non-Domino workloads should consume less than 10–15 percent of the available CPU CPW capability, whereas Domino workloads can consume 100 percent of that capability. No interactive processing capability is available for application use beyond the impact of the activities of a single systems administrator.

When deciding whether to select a Dedicated Server for Domino model or a traditional AS/400e server model, you should consider the following:

Dedicated Server for Domino is best for . . .

- Domino Mail
- Domino applications
- Mixed Domino workload
- Domino web serving
- Domino web front-end to AS/400 Line-of-Business (LOB) server

Traditional AS/400e server is best for . . .

- Domino and other AS/400 Line-of-Business (LOB) applications on the same server
- Large-enterprise user populations
- DB2 database integration on the same server

Dedicated Server for Domino Model 170D

Three processor options are allowed, ranging in performance from 1,300 Simple Mail Users (SMU) to 4,300 Simple Mail Users. It is not possible to upgrade to the Model 170D from any other model or from the Model 170D to any other model. Main storage feature cards are not customer installed; if the customer orders additional main storage, IBM will install the main storage and any other features also ordered.

The #2407 Processor supports 1,300 SMU and can be upgraded to the #2408 or #2409 processor, which supports 2,300 and 4,300 SMU, respectively. The Main Storage maximum capacity for each processor is 4,096 MB with a maximum disk capacity of 175.1 GB. System capacities in all other respects are the same as for the Model 170e server.

iSeries 400 Dedicated Server for Domino Model 270D

Three processor options are allowed, ranging in performance from 2,400 Simple Mail Users (SMU) to 7,580 Simple Mail Users. The Model 270D also supports the Mail and Calendar User environment ranging from 1,600 MCU on the lowest-performing processor to 5,050 MCU on the highest-performing processor offered. It is not possible to upgrade to the Model 270D from any other model or from the Model 270D to any other model. Main storage feature cards are not customer installed; if the customer orders additional main storage, IBM will install the main storage and any other features also ordered.

The #2422 Processor supports 1,300 SMU and 1,600 MCU and can be upgraded to the #2423 or #2424 processor, which supports 3,860 and 7,580 SMU, or 2,570 and 5,050 MCU, respectively. The Main Storage maximum capacity for the #2422 processor is 4,096 MB with a maximum disk capacity of 421.1 GB. The Main Storage capacity for the #2423 and #2424 processors is 8,192 MB with the same maximum DASD capacity of 421.1 GB. The relative processor ratings are 50 CPW, 100 CPW, and 200 CPW for the #2422, #2423, and #2424 processors, respectively. These processor ratings presume a Domino workload with no interactive workload other than that needed by the system administrator. System capacities in all other respects are the same as for the Model 270e server. LPAR is not available on the Model 270D. Figure 1.19 shows the capacities for the Dedicated Server for Domino Model 270D.

iSeries 400 Dedicated Server for Domino Model 820D

Three processor options are allowed, ranging in performance from 3,860 Simple Mail Users (SMU) to 14,840 Simple Mail Users. The Model 270D also supports the Mail and Calendar User environment, ranging from

Description / Function				
Processor Number	**2422**	**2423**	**2424**	
Performance				
CPW-Batch	370	950	2000	
CPW-Interactive	0	0	0	
Technology				
Clock Rate	400 MHz	450 MHz	450 MHz	
L2 Cache	0	2 × 4	4 × 4	
Threads	2	2	4	
Processor	Uni	Uni	2-Way	
Main Storage DIMMs or Cards Min/Max	2/8	2/16	2/16	
Minimum (GB)	256	256	256	
Maximum (GB)	4	8	8	
Description / Function	**Base CEC**	**Sidecar**	**Migrated Total w/#5077**	**System Maximum**
DASD Storage (GB)				
Integrated Min	8.58	0	0	8.58
Integrated Max	105.2	210.5	105.2	421.1
Maximum External				0
Total Maximum	105.2	210.5	105.2	421.1
DASD Arms (Max)	6	12	6	24
CD-ROM	2	—	0	2
External Tape (Max/Sys)	3	—	3	3

Figure 1.19. Dedicated server for Domino Model 270D capacities. *(continued on next page)*

2,570 MCU on the lowest-performing processor to 9,890 MCU on the highest-performing processor offered. It is not possible to upgrade to the Model 820D from any other model or from the Model 820D to any other model. Main storage feature cards are not customer installed; if the customer orders additional main storage, IBM will install the main storage and any other features also ordered.

The #2425 Processor supports 3,860 SMU and 2,570 MCU and can be upgraded to the #2426 or #2427 processor, which supports 8,420 and 14,840 SMU, or 5,610 and 9,890 MCU, respectively. The Main Storage maximum capacity for the #2425 processor is 8,192 MB, with a maximum disk capacity of 4 TB. The Main Storage capacity for the

Description / Function	Base CEC	Sidecar	Migrated Total w/#5077	System Maximum
Tape Libraries (Max) (1)	3	—	3	3
Optical Libraries (Max)	3	—	4	4
Physical Packaging				
External HSL ports	2	—	—	2
External HSL loops	1	—	—	1
Small Towers supported	1	—	—	1
Large Towers supported	0	—	—	0
Embedded IOP	1	—	1	2
PCI Card Slots	7	—	8	15
Maximum PCI IOA Cards	6	—	7	13
Workstation Attachments (Max)				
Twinax Controllers	4	—	6	6
Twinax Devices	160	—	240	240
Communications Lines (Max)	20	—	34	50
Crypto IOP (DES/PKA) (Max)	3	—	3	3
LAN ports (Max)	3	—	5	8
IPCS (Max)	1	—	2	3

Notes:
1. Total number of tape drives does not increase.
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system.

Figure 1.19. Dedicated server for Domino Model 270D capacities. *(continued from previous page)*

#2426 and #2427 processors is 16,384 MB, with the same maximum DASD capacity of 4 TB. The relative processor ratings are 100 CPW, 200 CPW, and 300 CPW for the #2422, #2423, and #2424 processors, respectively. The #2426 processor is a two-way, and the #2427 is a four-way processor. These processor ratings presume a Domino workload with no interactive workload other than that needed by the system administrator. System capacities in all other respects are the same as for the Model 820e server. Figure 1.20 shows the capacities for the Dedicated Server for Domino Model 820D.

Description / Function			
Processor Number	**2425**	**2426**	**2427**
Performance			
CPW-Batch	950	2000	3200
CPW-Interactive Feature #	0	0	0
Technology			
Clock Rate	450 MHz	500 MHz	500 MHz
L2 Cache	2 x 4	4 x 4	4 x 4
Threads	2	4	8
Processor	Uni	2-Way	4-Way
Main Storage			
DIMMs or Cards Min/Max	2/16	2/32	2/32
Minimum (MB)	256	256	256
Maximum (GB)	8	16	16

Description / Function	**Base CEC #5075**	**Small Tower #5074**	**Large Tower (5)**	**System Maximum**
DASD Storage (GB)				
Integrated MIN	8.58	0	0	8.58
Integrated Max	210.5	105.2	789.7	4159.1
Maximum External	0	0	0	0
Total Maximum	210.5	105.2	789.7	4159.1
DASD Arms (Max)	12	6	45	237
Diskette	0	0	0	0
CD-ROM/Internal Tape (3)	2	0	2	12

Figure 1.20. Dedicated server for Domino Model 820D capacities. *(continued on next page)*

System Availability

Since their inception, the AS/400 systems have had availability as a target. The average single AS/400 system in 1997 had an availability of greater than 99.9 percent, with average unplanned outages of less than nine hours per year. Daily reports are accumulated on greater than 87,000 AS/400 systems in the United States. These reports show that on AS400

Description / Function	Base CEC #5075	Small Tower #5074	Large Tower (5)	System Maximum
External Tape (Max/Sys)	6	7	8	8
Tape Libraries (Max) (1)	6	7	8	8
Optical Libraries (Max)	6	7	14	14
Physical Packaging				
External HSL ports	2	—	—	2
External HSL loops	1	—	—	1
Small/Large Towers\ supported	5	—	—	5
Embedded IOP	1	1	8	6
PCI Adapter Card Slots	12	8	14	82
Maximum PCI IOA Cards	9	7	11	63
Workstation Attachments (Max)				
Twinax Controllers	7	7	11	62
Twinax Devices	280	280	440	2480
Communication Lines (Max) (2)	44	34	52	160
Crypto IOP (DES/PKA) (Max)	3	3	3	3
Lan Ports Max	6	5	8	30
IPCS (Max)	2	2	2	12

Notes: 1. Total number of tape drives does not increase.
2. One line used for operations console, if installed.
3. Must be 1 CD-ROM per system.

Figure 1.20. Dedicated server for Domino Model 820D capacities. *(continued from previous page)*

systems shipped since January 1996, the average interval between hardware problems is sixty-one months. These results were accomplished through a combination of hardware and software actions. The hardware and software items described next identify the actions taken by the AS/400 to achieve that availability. The five items described after the hardware and software actions are what take the AS/400 to an availability of 99.99 percent.

On the *hardware* side, things like battery backup and continuously powered main storage protect against utility failures and the loss of data that could result from those failures. Redundant bulk power supplies and independent redundant local voltage regulators protect against power system failures. RAID-5, mirroring, and redundant write cache protect against data loss from disk unit failures, whereas concurrent maintenance allows the repair of those failures without loss of the system to your users. System Power Control Network (SPCN) turns the system and its remote components on and off at the same time and monitors the remote units to enable maintenance actions before they cause a catastrophic system failure.

From the *software* side, items like commitment control and journaling, auxiliary storage pools, system-managed access paths, save-while-active, S/R parallelism, hierarchical storage management, BRMS, and ADSM (combined with the pretesting of the integrated operating system and applications code) reduce the duration of the recovery actions needed if an outage does occur. Figure 1.21 describes the hierarchy of actions needed to achieve 100 percent availability.

The following paragraphs describe the options offered as well as the choices a customer must make to increase the AS/400 availability. Included in the choices are capabilities like logical partitioning, multisystem coupling, data and application resiliency, clustering of systems, and disaster tolerance.

OS/400 Subsystems and Logical Partitioning

Logical partitioning is the creation of several independently functioning systems within the covers of a single system. *N*-way **Symmetrical Multi-Processors (SMP)** (where *N* is replaced by the number of processors) and sufficient main storage are both required. The primary partition requires one processor and 256 MB of main storage. Creating and managing secondary partitions is done from the primary partition. Each secondary partition requires a minimum of one processor, 64 MB of memory, an I/O processor with disk resources, and its own console device. The number of partitions has been extended from twelve to twenty-four with the announcement of the twenty-four-way Model 840. It is possible to have OS/400 V4R4 running in one partition while OS/400 V4R5 is running in another partition. Support has been added

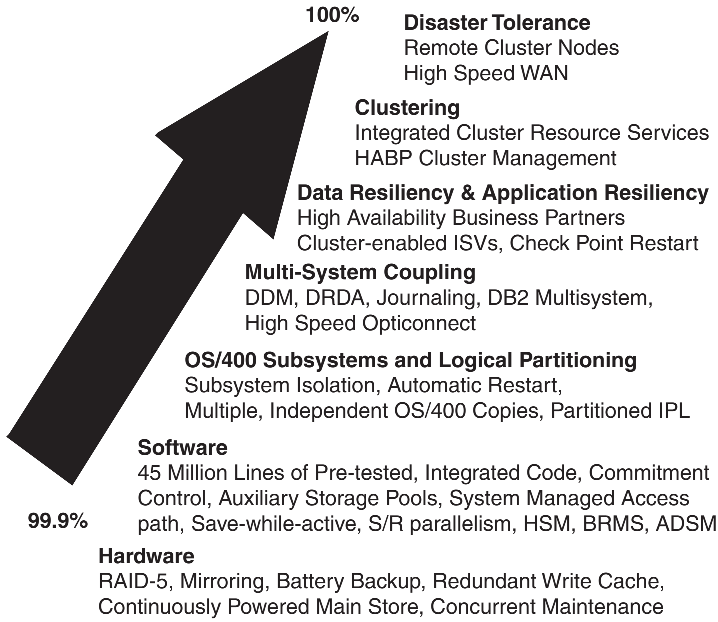

Figure 1.21. AS/400e availability solutions.

for an external stand-alone CD-ROM that can be switched between partitions to simplify the loading of software.

There can be multiple secondary partitions, each with its own I/O resources assigned from the list of total system resources. No resources can be shared between partitions. An installation of a secondary partition does not impact other partitions, but an IPL of the primary partition requires an IPL of all secondary partitions. Figure 1.22 illustrates how logical partitioning might appear on a single system.

New processor and memory resources can be allocated to a secondary partition by an IPL of that partition, not the entire system. Available I/O resources (e.g., disk, tape, and CD-ROM) can then be allocated to that secondary partition without an IPL. Tape, I/O Processors, and CD-ROM devices can be switched between partitions, and can also be dedicated to individual partitions. Adequate communications resources such as LAN adapters should be dedicated to each partition.

OS/400 is licensed once for the entire system by its normal processor group, independent of the number of partitions. V4R4 of OS/400 is

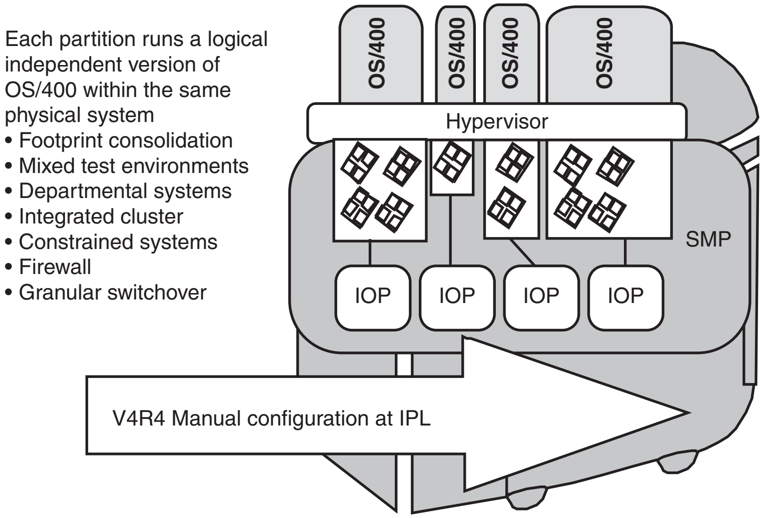

Figure 1.22. AS/400e logical partitioning.

the initial enabling release for logical partitioning, and can be installed on each partition. Pre-V4R4 releases are not supported for logical partitioning. License management has been improved to manage system licenses across the entire system rather than at the partition level.

An internal virtual bus provides the option to use OptiConnect for high-performance communications between partitions. OptiConnect allows applications on one partition to access the database of another partition using DDM remote SQL, although disk resources are not physically shared. So although logical partitions can be operated independently, the high-speed internal bus connection between partitions enables applications having a need for multiple AS/400 application servers to access a single database server.

System values for each partition are set independently. Each partition has a different system name and can use a different primary/secondary national language, or can even be operated in different time zones. Multinational companies can centralize multicountry operations in a single location while retaining the national characteristics of each system. For Logical Partitioning, a minimum of V4R4 must be installed on each partition.

The increased scalability of the AS/400 provides the opportunity for companies to consolidate multiple AS/400s into fewer, more manage-able systems. **Logical Partitioning (LPAR)** in OS/400 enhances the role of the AS/400 as a consolidated server. Logical partitioning provides both the power and flexibility to address multiple systems requirements in a single machine. Logical partitioning needs to be considered for environments requiring server consolidation, business unit consolidation, mixed production and test environments, and integrated clusters.

Product Preview—LPAR Enhancements

It will be possible in the future to run multiple partitions on a single processor. The current requirement that at least one processor must be dedicated to each partition will no longer be true. Resources will be capable of reallocation dynamically between partitions with no IPL required after the reallocation. Partitioning management will be managed using Operations Navigator to simplify the management process.

Multisystem Coupling

Multiple AS/400 systems can be coupled using the OS/400 resources associated with Distributed Data Management (DDM), **Distributed Relational Database Architecture (DRDA),** Journaling, DB2 Multisystem support for Symmetrical Multi-Processors, and high-speed OptiConnect.

Data Resiliency and Application Resiliency

Data resiliency and application resiliency are provided by the IBM high-availability business partners DataMirror, Vision Solutions, and Lakeview Technology. Checkpoint restart is a capability included within the software provided by these business partners. Also included within the business partner provided software are cluster management, data replication, and application availability tools.

Two types of cluster resource groups are supported, one for **data resiliency** and one for **application resiliency.** Data resiliency provides the capability to switch the point of access for a set of data over to a backup node that is maintaining an exact replica of that data. Resilient resources are object replicated to one or more nodes of a cluster. Application resilience provides the capability to switch an IP address (representing the application server) to a backup node and to restart the application in the event of a primary node failure.

Continuous Availability Clustering

Continuous availability cluster solutions are designed for those customers who require access to their data and systems 365 days a year, 24 hours each and every day. The continuous availability cluster solutions can use a pair of AS/400 systems consisting of any AS/400 system or server. The systems are connected together using a high-performance data path, either optical or ATM. Also, the systems are connected to either a LAN or a WAN containing either terminals or personal computers. Figure 1.23 illustrates the system interconnection for continuous availability. Figure 1.24 illustrates the supported cluster topologies including a remote cluster node connected by a high-speed WAN through a switch.

Cluster resource services within the OS/400 operating system reduce the complexity of managing AS/400 clustered systems. OS/400 keeps track of data and applications for those systems. Cluster-management and data-resilience applications (provided by high-availability business partners) protect customers' business from unplanned and planned outages as well as site loss disasters.

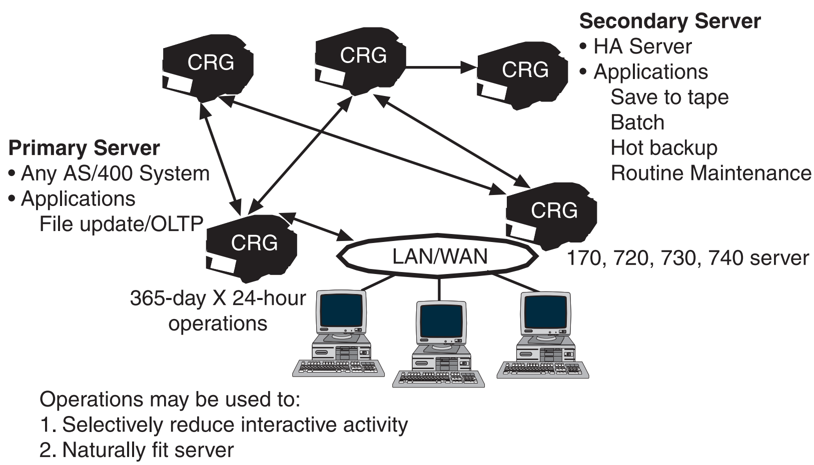

Figure 1.23. System interconnect for high availability.

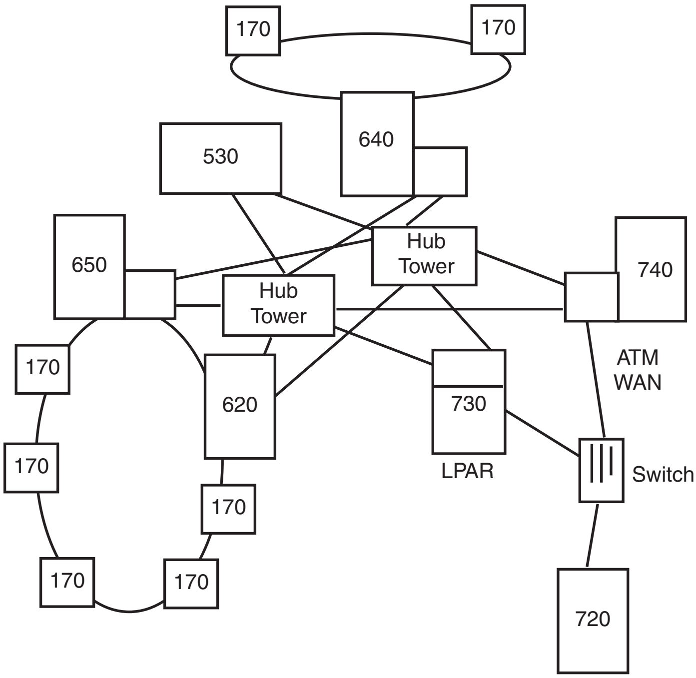

Figure 1.24. Cluster topologies supported for AS/400.

Cluster resource services contain an open set of APIs providing cluster facilities available to AS/400 application providers and customers to enhance application availability. These APIs allow the creation, configuration, and administration of the cluster. Systems are defined into the cluster as **cluster nodes.** Communication interface addresses are defined to form the cluster node interconnection links.

Cluster resource services also provide low-level facilities such as heartbeat monitoring, reliable message delivery, switchover administration, and distributed activities. The services are built on cluster topology and messaging functions, which keep track of each node in the cluster and ensure that all nodes have consistent information about the state of cluster resources. **Heartbeat monitoring** ensures the active state of each node. If the heartbeat for a node fails, the condition is reported and the cluster can automatically fail-over to the resilient resources on the backup node.

AS/400 clusters support up to 128 nodes (instances of OS/400), using any combination of the existing OptiConnect, WAN, and LAN connectivity options to build the cluster. A single workstation (containing the continuous-availability business partner cluster management application) manages all systems in the cluster. OptiConnect hardware is a connectivity method for high- and midrange models. ATM provides a high-performance connection to remote systems in a cluster. Ethernet and Token-Ring LANs can be used to connect low-end AS/400 models into a cluster.

Data resilience support for modern applications has been enhanced with the introduction of **Byte Stream File (BSF)** replication. Business partner solutions provide real-time recording of changes to BSF data replication to remote systems. Transfer of changes using remote journaling protects in-flight data at the point of failure.

Disaster Tolerance

Disaster tolerance is provided by the support of remote cluster nodes as replication points for resilient data. The remote cluster nodes can be off the customer site, and may even use facilities of other businesses like IBM to provide recovery mechanisms. High-Speed WAN facilities like ATM are used to support the remote cluster nodes. Disaster environments include hurricanes, tornadoes, and floods, wherein the business becomes unable to utilize its local facilities for a period of time.

A 1998 Gartner Group study compared the availability of various systems. The only systems that ranked higher than the AS/400 were clustered systems. The average NT server outages were forty times worse than the AS/400.

Upgrades

There are two methods available for upgrading within the AS/400e series product line: upgrading within a model and upgrading through a system unit swap.

Upgrading within a Model

This method is used if the need is for more performance or input/output (I/O) device functionality and the current system is not using the maximum performance features available within the specific AS/400 model.

These could be the processor, interactive processor feature, main storage size, attachable I/O device capability, or disk storage size. In this situation, performance can be upgraded within a model by swapping processors, adding additional main storage, or adding disk storage or additional IOPs without having to swap system units. In general, the system should be upgraded within the model until those maximums are reached. This method of growth is referred to as *horizontal growth*.

Upgrading through a System Unit Swap

Another way to upgrade within the AS/400 family is to replace the current AS/400 system unit with a more powerful model (e.g., you upgrade a Model 820 to a Model 830 or 840 by changing the system unit). This method is used if the current system is at its processor and storage maximums, or if the number of users has increased beyond the capabilities of the current model. This method would also be used if there is a need for application programs that run on newer RISC-based AS/400 models only, or if you are upgrading from an older AS/400 model. This method of growth is referred to as *vertical growth*.

Upgrading through Migration Towers

Upgrades to the AS/400 Models 820, 830, or 840 from 6XX/7XX/SXX models is performed by means of converting the system unit for those models to a migration tower. The type of migration tower involved in the conversion depends on the original model being upgraded. Figures 1.25 and 1.26 illustrate this migration activity. In all cases, the existing I/O previously included in the system unit as well as the I/O attached to the system unit through expansion towers remains in place and continues to be used with the new model. New I/O devices are placed in the new system unit, and any expansion towers required by the customer are supported by that model's system unit.

Upgrades between Models

No upgrades between Models 6XX/SXX/7XX and 8XX servers may be accomplished without exchanging the system unit. You can continue to use all the same workstations, modems, and so forth. Upgrades of CPUs from a Model 150, Model 170, Dedicated Server for Domino 170D, SB1, 4XX/5XX, or CISC to 8XX models are not supported. Upgrades

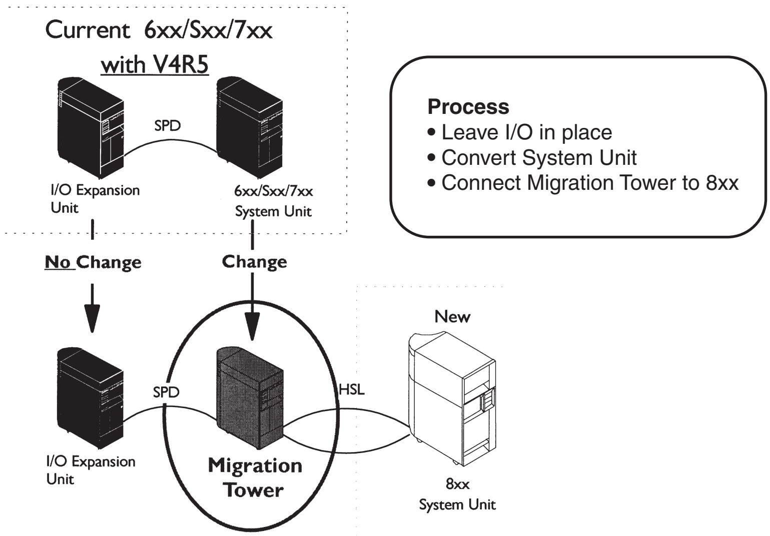

Figure 1.25. The migration process for 6XX/SXX7/XX servers to 8XX servers.

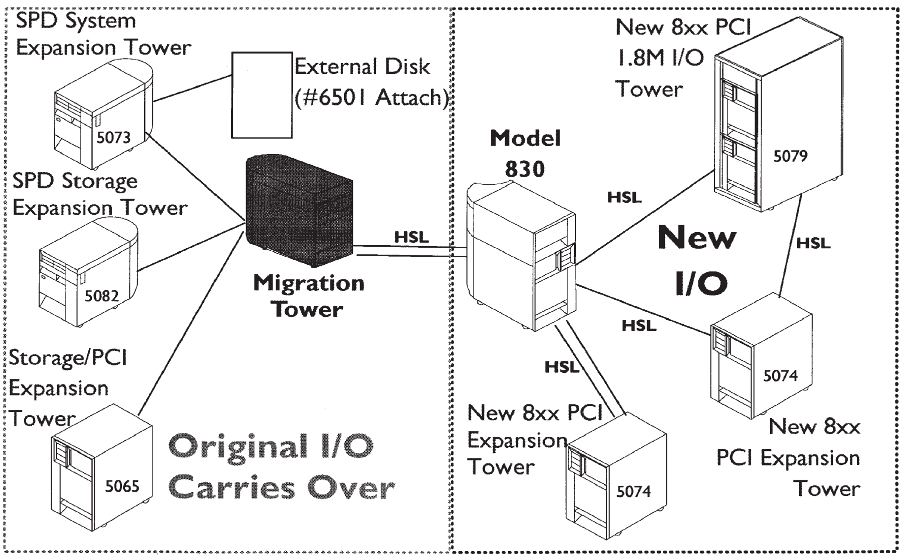

Figure 1.26. I/O tower coexistence with previous I/O on 8XX servers.

from these models will get a new serial number. Upgrades between the 6XX/SXX/7XX servers and 8XX servers, including migration tower identification, are illustrated in Appendix G. These upgrades will retain the same serial number for the new models.

Appendix F identifies the recommended migration paths between older AS/400 system/server models and the new system/server models.

Performance Overview

One important aspect of a computer system is its performance—that is, the speed at which the computer can perform work. The higher the performance, the more work the computer can do. Many things—such as the processor, main storage, disk storage, I/O bandwidth, and program design—affect the performance of a computer system. It is difficult and often misleading to predict the overall performance of a computer system by looking at selected specifications of the individual components that make up the system. Although things such as disk unit seek times and raw processor speed (usually measured in millions of instructions per second, or MIPS) are important, they do not provide the whole picture in terms of overall system performance. In general, the higher the MIPS value, the better. The group of things that fit into the category "more is better" include: number of processors, super scalar—number of execution units per processor, the presence or size of L1 cache, the presence or size of L2 cache, the size of memory, and the number of buses. Some things fit into the category "faster is better," including transfer rate (how fast information can be moved from cache to the processor), transfer rate (how fast information can be moved from the memory to the processor), transfer rate (how fast information can be moved from the I/O subsystem to memory), link rate (the speed of the link or bus to I/O devices—disk), and switch rate (total rate of moving information between processors, memory, and I/O subsystems). The one area where "smaller is better" is pipeline depth—the number of stages a machine instruction takes to execute. This becomes more important when related to the frequency at which the processor performs task switching or logic switching as opposed to sequential processing. With a task switch, the pipeline is flushed and needs to be refilled. The bigger the pipeline, the longer it takes to feed instructions through the pipeline to feed the processor, and the longer the processor waits. Commercial applications tend to have a lot of task and logic switching.

In addition to these considerations, in the case of the AS/400e Series computers, performance must be considered in relation to the environment in which the computer system is executing. The AS/400 may operate in a general-purpose (batch and interactive processing) environment or in a client/server environment. What makes a computer function well in one of these environments may be an inhibitor in the other environment. (A computer that has been tuned for the client/server environment will perform weakly in an interactive environment, and vice versa.) As a result, different methodologies are used to measure computer performance for those different environments.

When choosing a multiuser computer system, it is important to understand that the maximum number of active users is different from the maximum number of workstations that can be attached to a system. History has shown that when the active number of users is greater than 40 percent of the maximum number of workstations allowed, there is a strong probability that some of the users will start to experience a decrease in response. The maximum number of users is not a system configuration limitation, but rather a system performance limitation for the benchmark used. Therefore, when some of the user population includes only occasional users, it is normally desirable to put more workstations on a system than can simultaneously supported.

The response time of a computer system is at least as important as throughput. Users' satisfaction and productivity can be drastically affected by a computer's sluggish response after the user hits the Enter key. As was mentioned earlier, response time is related to the workload on the computer system. This is not intended to imply that every workload will see subsecond response time.

Benchmark Testing Overview

Benchmark testing has evolved as the best way to compare the overall performance of different computers. Benchmark testing involves loading the computer system with various programs designed to simulate a workload and then measuring how the system behaves under the load. Through this benchmark testing, all elements of the computer system come into play, and the overall performance of selected computer systems can be meaningfully compared.

To perform benchmark testing, the test group must make assumptions about the kind of work being done and the behavior of the users at each workstation. For this reason, the performance measurements derived from a benchmark may vary significantly from what users will get if their business environment does not match the set of test assumptions. For example, word-processing users typically load a system down more heavily than users performing business transactions. Therefore, a system with heavy word-processing activity will usually not be able to support as many users as a system performing order entry.

However, because all assumptions are the same for all computers included in the benchmark testing, it is possible to get a good idea of the relative performance of tested computer systems. IBM has conducted benchmark testing on AS/400 systems using the CPW benchmark, which is designed to simulate a typical business workload (e.g., order entry, accounts payable, and accounts receivable).

When IBM's AS/400 group benchmarks a system, the benchmark is against criteria that measure how that system is intended to be used. The general-purpose systems, like the previous Models 600 and 650, are benchmarked against their ability to support interactive users. In this context, whatever system capacity is not being used for interactive users is available for performing other work such as client/server–based file serving or print serving. Other systems, like the server-based systems, are intended to provide the majority of their capabilities for the batch mode server-based functions such as printing, file serving, communications serving, and so forth. They are benchmarked against criteria that measure how well the systems perform in that environment.

The AS/400e solutions packaging preload environments are measured against criteria that combine the need to get a larger number of interactive users applying a workload while still allowing a significant client/server workload. The AS/400e solutions packaging preload environments expect to perform their function with this increased interactive workload while supporting a limited set of business partner applications.

The chart shown in Figure 1.27 illustrates the expected performance for the described benchmark environments for the Models 8XX and 270 CPUs. If this chart were compared to the previous Models 170 and 7XX CPU charts, the difference would occur after the point where the interactive performance reached its maximum capability. With the new CPUs, the available client/server performance stays fixed after the inter-

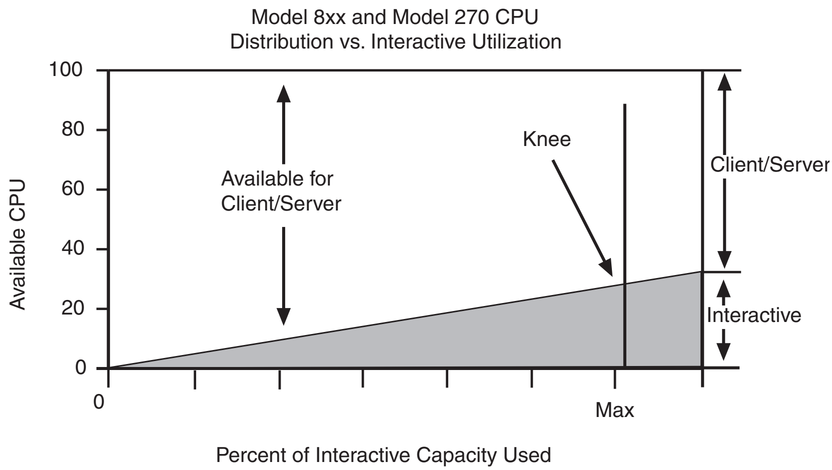

Figure 1.27. Interactive Performance Model—What's This?

active performance has peaked, and continuing demands on interactive performance cause the interactive performance to decrease. There is no interference from system overhead caused by excess interactive workload for any or all client/server workloads.

As an example, for the AS/400e 840-#2418 processor, total processor CPW rating is 10,000. (If the interactive feature #1544 is present, the maximum interactive performance is 2,000 CPW, and a minimum client/server performance capability of 8,000 CPW is available.) If, at any point in time, only 50 percent of the interactive performance is being used, then 9,000 CPW is available for client/server usage. Depending on how the job priorities are set, a big difference exists in the results.

8XX Model Priorities

The default for the Model 8XX processors when interactive cards are present is that when the interactive workload exceeds 115 percent of the design point, the interactive activity begins to slow down, and the server workload is unaffected.

CPW Benchmark Test Results

In 1996, IBM introduced a new benchmark testing parametric called Commercial Processing Workload (CPW), which replaced the RAMP-C benchmark previously used. IBM chose to make the change to CPW because RAMP-C was driven by the increased performance levels of current AS/400 models and the recognition that users were utilizing more OS/400 software function. (Contention, queuing, and journaling are examples of software functions that RAMP-C did not test.) In today's environment, IBM feels that CPW benchmarks provide a more representative result than RAMP-C.

Figure 1.28 shows the interactive CPW performance ratings of the V4R5 RISC-based models of AS/400 systems. Additional information on other benchmark results on the AS/400 and other vendors' general-purpose commercial processing systems' performance may be obtained from the Transaction Performance Council (TPC, *http:llwwwtpc.org*), which publishes test results on a wide variety of systems.

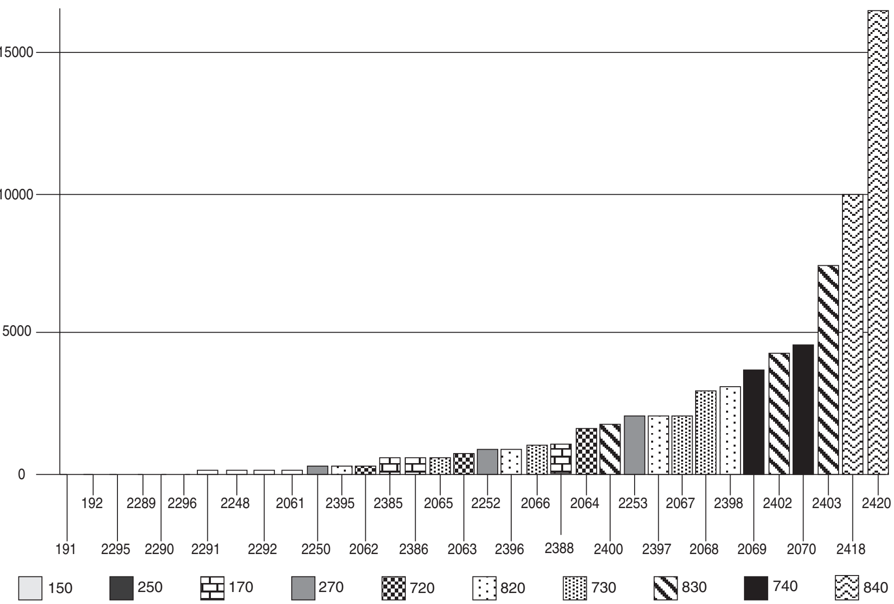

Figure 1.28. V4R5 processor CPW by model.

Symmetrical Multi-Processing (SMP)

Some of the AS/400e models use Symmetrical Multi-Processing (SMP), meaning that they use more than one processor in a single system. (The discussion "Hardware Architecture Overview" later in this chapter contains more on how multiple processors achieve this performance increase.) With SMP, the performance rating increases with each new processor added. For example, if you compare a single processor system to an SMP system with two processors of identical power, you get a performance increase of about 70 percent. You don't get a 100 percent improvement because the two system processors must spend some of their processing power coordinating activities with one another. They must also share the same main storage, I/O bus (or buses), disk storage, and so on. For these reasons, adding a second system processor does not double the performance of the system; however, SMP configurations do allow significant performance improvements. Figure 1.28 shows the performance advantage afforded by the use of two-, four-, and eight-way processors in the Model 730, along with eight-way (#2240) and twelve-way (#2243) processors in the Model 740 SMP configurations.

Server Performance

Client/Server Labs measures server performance on systems for many companies, using a benchmark suite specially tailored for servers. The server environment evolved as a result of two nearly coincidental events: (1) the personal computer (PC) penetration into the generalized business world and (2) the natural desire to share the information contained on those PCs. The PC put the power of computers into the hands of each critical individual within a business.

The need to share the results on one PC with the users on several PCs resulted in the definition of the local area network (LAN). An example of an application across a LAN would be this: If one secretary at a business location reserved a conference room for 1:00 p.m. on Tuesday, other secretaries would see that the conference room was reserved (including when, and by whom). As in this example, a LAN should be used to connect together groups of users with some common interest. The first LAN to be defined was Ethernet (which eventually became IEEE 803.2), followed later by the token ring (which became IEEE 803.5). LAN definitions, speeds, and protocols that could not talk to each other proliferated until finally the distributed computing environment emerged, settling upon TCP/ IP and APPC as the LAN access method and management structure.

Another problem of networking only PCs is that if one user wants a specific piece of information and the PC that has that information is turned off at the time the information is needed, the information cannot be accessed. Also, since PCs have limited disk storage space, individual users were asking why they should spend money and time creating a repository for data that, once created, might be of no further interest. Finally, the question of controlling access to the information on the network was begging for an answer because of privacy (and need-to-know) issues.

An answer that evolved was the client/server relationship, in which the PCs were the clients and the server was a large centralized repository of information. In this environment, the client requested either programs or data from the server, which checked the client's authorization to receive that information and sent the data to the client. The client then processed the data and sent the results back to the server repository for long-term storage. In order to check the authorization, find the programs and data, and provide enough storage repository to be useful, the server required the processing characteristics of the client. In fact, in the case of certain long-running applications, it would be preferable to execute the application at the server and send only the results of the application to the client, who could then decide whether the data generated should reside on the server or at the client. Since the long-running application executed on the server, this freed up the client resources for execution of short-running applications during the long application's execution time.

It follows, then, that minimizing paths in the server for APPC and TCP/IP, and minimizing database search structures (via good search algorithms, improved base structures, some capability for parallel I/O processing, efficient file serving, and efficient batch-processing capability) are the characteristics of a good server. The performance of the server models is shown in Figure 1.28, with the tallest column showing the server batch-processing performance at Version 4 Release 4 and Version 4 Release 5. It is no longer necessary to limit the configuration and the available application set to provide improved interactive performance as was done for the custom mixed-mode S20-ISV, S30-ISV, and S40-ISV models.

Interactive Performance

The interactive performance capabilities of e-servers 720, 730, 740, 820, 830, and 840 systems is established by which of the interactive feature cards is plugged into the system processors. The table in Figure 1.29 identifies the interactive features that can compatibly be plugged into each of the processors for the 8XX systems and the interactive performance associated with each feature. The interactive feature cards for the 7XX models is identified in *Exploring IBM AS/400 Computers*, 9th edition.

Dedicated Server for Domino Benchmarks

A new performance metric has evolved for Domino. In the past, Simple Mail Users (SMU) was used as a benchmark. To arrive at a measure of Typical Mail Users for Domino, a measurement was made of Simple

AS/400e server 820 Interactive Cards

Feature	Client/Server CPW	1521	1522	1523	1524	1525	12526	1527
2395	370	35	70	120	240			
2396	950	35	70	120	240	560		
2397	2000	35	70	120	240	560	1050	
2398	3200	35	70	120	240	560	1050	2000

AS/400e server 830 Interactive Cards

Feature	Client/Server CPW	1531	1532	1533	1534	1535	1536	1537
2400	1850	70	120	240	560	1050		
2402	4200	70	120	240	560	1050	2000	
2403	7350	70	120	240	560	1050	2000	4550

AS/400e server 840 Interactive Cards

Feature	Client/Server CPW	1540	1541	1542	1543	1544	1545	1546	1547
2418	10000	120	240	560	1050	2000	4550	10000	
2420	16500	120	240	560	1050	2000	4550	10000	16500

Figure 1.29. Interactive feature cards by processor and model.

Mail Users, and that result was divided by three. The new performance metric, Mail and Calendar Users (MCU), evolved because the previous Notesbench Mail benchmark is no longer representative of the more complex environment for Domino R5 users. Domino R5 includes calendaring and scheduling, a 60 percent heavier workload with 30 KB more traffic per hour. To arrive at Typical Mail Users using the MCU benchmark, which is closer to the Typical Mail Users environment, the rule of thumb is to divide by two. Typical Mail Users values allow the comparison of older AS/400 servers with new servers, as Typical Mail Users represents a constant metric. Figure 1.30 illustrates the benchmark values for the Dedicated Server for Domino models.

OptiConnect Performance

The AS/400e server 7XX and 8XX models support a high-performance optical bus connection called OptiConnect. Through this connection, AS/400 systems can be connected together and cooperate (exchange programs and data) without suffering the performance degradation typi-

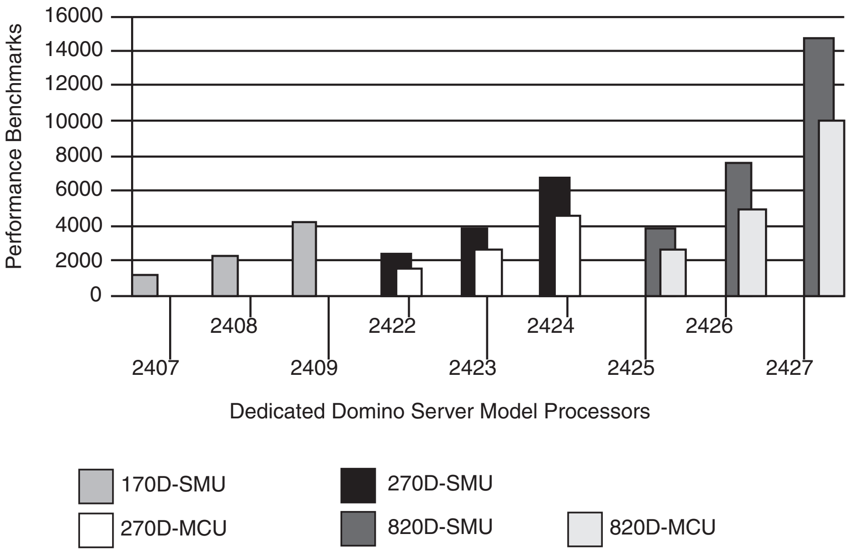

Figure 1.30. Dedicated server for Domino performance capacities.

cally associated with system interconnection via other methods (e.g., TCP/IP, Ethernet, or token-ring networks).

Measured performance capacity increases (1.5 to 1.7 times the performance of a single system, depending upon application efficiency) have resulted from connecting two AS/400 systems together via OptiConnect. Currently, up to thirty-two AS/400 systems can be connected via OptiConnect to provide loosely coupled distributed computing.

Loosely coupled computing is an environment in which data and programs are exchanged between computers over a serial communications line, which from a hardware aspect of the systems is the only thing shared. The coupling is referred to as “loose” because the potential processing occurring in the system (from which the information is being moved) is significant during the period of the transfer, and because of the resulting problems involved in keeping the information in the sharing systems synchronized.

A Closer Look

Many elements provide the functions and performance of IBM AS/400 computers. This is perhaps the appropriate place to mention that many of the concepts used in the AS/400 system were built on those of IBM’s System/38 computer. This fact is a testimonial to the rich function and growth capability built into the original System/38!

The remainder of this chapter provides a closer look at the following aspects of the IBM AS/400 systems: hardware architecture, clustering technology, OptiConnect technology, main and auxiliary storage, packaging, and the fiber optic bus.

Hardware Architecture Overview

The underlying arrangement and interconnection of a computer system’s electrical components is called its hardware architecture. This architecture is the fundamental structure upon which all system functions are built and has the largest effect on how the computer system will behave. A basic understanding of the AS/400 system architecture (Figure 1.31) makes it possible to compare AS/400 computers with other systems and to understand important aspects of system performance and capacity.

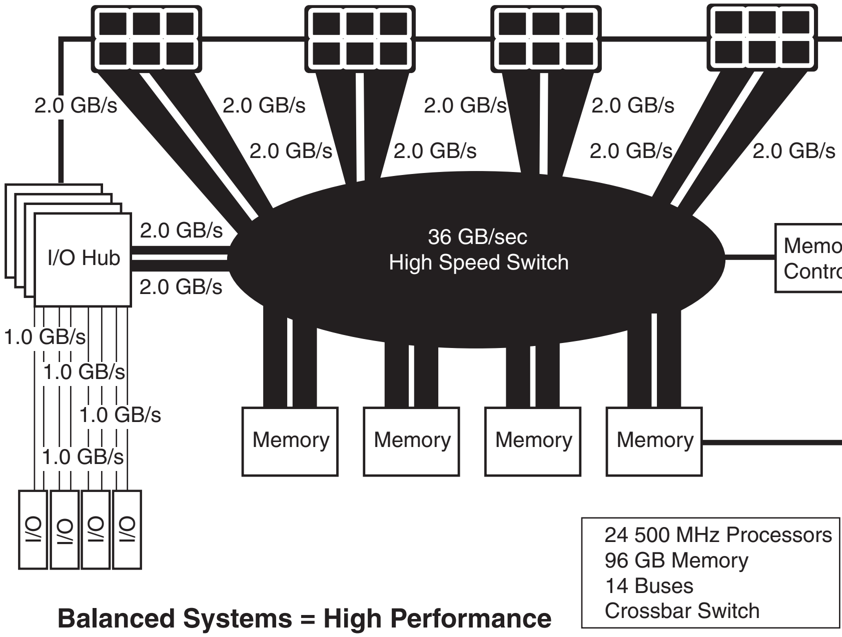

Figure 1.31. AS/400 model 840 architecture.

The core of the AS/400 computer (as in all computers) is the system processor (shown across the top of the figure). The system processor is the circuitry that actually executes a computer program's instructions and does all the mathematical calculations. The smallest piece of information (data) in the computer is called a bit. Bits are grouped into bytes (8 bits), half words (16 bits), full words (32 bits), and double words (64 bits) inside the computer. These groupings form the computer's representation of numbers, letters of the alphabet, and instructions in a program. AS/400 system processors move information around one double word (64 bits) at a time. Since much of a computer's time is spent moving information around, the double-word organization helps improve overall system performance.

Other bits inside the system processor are used to uniquely identify or address storage and I/O devices (e.g., a disk unit) within the computer system. AS/400 system processors group 64 bits together to form a unique address. This 64-bit addressing provides 18,446,744 trillion (2^{64}) unique addresses, which is more than any other IBM computer system from PCs to the largest System/390 computers. This is more than enough addresses for today's midsize computer environment (and even for the foreseeable future). In fact, the largest AS/400 systems and servers today have just identified the capability to use 18+ trillion of those addresses. This shows the kind of growth potential inherent in the AS/400 architecture.

The "memory," or main storage (shown at the bottom of the figure), provides a work space for the system processor. Since much of a computer's time is spent moving information to and from main storage, the speed of main storage can be a limiting factor to the overall performance of any computer system. The speed of storage is measured by the time it takes to respond to a request to store or recall information (the cycle time). The main storage cycle time for AS/400 computers varies depending on the model.

The shorter the cycle time, the better the system performance. The largest AS/400 computers can have up to 96,000 MB of main storage. The main storage in all AS/400 systems provides error detection and error correction. This main storage error detection and correction works to protect the all-important integrity of user information in the computer system.

All AS/400 system processors also use cache memory to help increase the effective cycle time of main storage. A cache is a small and very high speed memory area that sits between the processor and main storage. The idea is to keep the information most likely to be needed next in cache to avoid the time delay associated with main storage. AS/400 systems have data and instruction caches on the processor to accelerate performance when accessing information/program instructions (respectively).

Another important part of the AS/400 architecture is the System Licensed Internal Code (SLIC). SLIC is a set of extremely simple instructions (never seen by the computer programmer or user) that are directly performed by the electronic circuits within the system processor. All user program instructions are automatically converted into a series of these SLIC instructions, which are then executed by the system processor.

In the high-end AS/400, the separation of main storage and I/O is possible because of the high-speed 36 GB switch. The 36 GB switch makes possible the four independent 128-bit data buses between the processors, I/O controllers, and the switch network. As shown, the twenty-four processors are distributed uniformly across the three data buses, with the fourth data bus dedicated to the I/O controllers. A single 64-bit address bus connects all of the processors plus the I/O controllers with the switch. Although the address bus is heavily loaded and physically traverses the system board to connect the processor cards, the bus utilization and latency were kept sufficiently low by configuring the system bus arbiter located within the switch to allow only one address for every two data cycles. The I/O is driven from an HSL bus (the HSL bus is an ultra-high-speed serial interface), a copper bus limited to fifteen feet in length for a single segment, having a bandwidth of 10 GB/sec. The copper bus connects to at least one I/O tower, which can be populated with disk units, tape units, and I/O interface cards. The main storage bus is partitioned into multiple buses, each of which is supported by a separate storage control chip (STG/CTL) that manages the main storage accesses and up to 0.5 MB of intermediate-level cache. This multiple-memory bus architecture helps improve the performance of these systems.

The I/O processors (shown at the bottom of the figure) are responsible for managing any devices attached to the AS/400 system. Each of these specialized processors has independent responsibilities and performs tasks in coordination with the system processor. A computer that has multiple processors working together with the system processor like this has a multiprocessor architecture. The advantage of having multiple processors performing work simultaneously is simply that more work can be done in a given period of time. For example, the workstation I/O processor manages the detailed processing associated with the multiple terminals and printers attached to the system, allowing the system processor to concentrate on doing more productive work for the user. The same is true of the other specialized I/O processors such as the storage I/O processor, which manages disk, diskette, and tape devices attached to the AS/400 system.

The I/O processors communicate with the system processor over an I/O bus (called the PCI bus), a group of wires that carry information very quickly from one area to another inside the computer system. As indicated in the figure, some AS/400 systems have a single I/O bus, whereas others have multiple I/O buses. Because only one information transfer can occur on any one bus at any one time, systems with multiple buses have the advantage of allowing overlapping transfers between I/O processors and the system processor or main storage. Therefore, multiple buses contribute to the overall system performance advantages of larger AS/400 systems. The PCI I/O buses in the 8XX models provide an I/O bandwidth of greater than 10 GB for the movement of data and instructions between the I/O devices and the processor/memory complex.

Various controllers and adapters plug into physical slots in each of the packages used to provide electrical connections to the bus. In addition to I/O processors, a service processor (shown in the upper right of the figure) is built into every AS/400. It is responsible for starting the system and constantly monitoring the health of the entire computer. It interacts with the system operator through the control panel and helps with such things as system fault isolation, error detection, and error reporting. It is the equivalent of having a built-in service person who watches over things with relentless consistency.

All AS/400 systems employ a multiprocessor architecture in that they have a system processor and multiple specialized processors (e.g., workstation and I/O processors) to handle specific tasks. However, larger AS/400 models (e.g., Models 170, 720, 730, 740, 820, 830, and 840) employ multiple system processors to cooperatively execute a single copy of the operating system (OS/400), thus appearing to be a single large processor. This multiple system processor architecture is called the **N-way multiprocessor architecture** (where *N* is replaced by the number of processors), also referred to as the Symmetric Multi- Processor (SMP) architecture.

Figure 1.32 shows how N-way models are organized. (*Note:* All system processors share the same I/O buses, I/O processors, and main storage.) Symmetric Multi-Processors process in parallel, sharing a task list. Each processor in a symmetric multiprocessor set has its own data and instruction cache, and its own virtual storage view of the system. In the case of a query, each processor in the SMP group will process in parallel against a segment of disk storage to resolve the query.

Clustering Technology

The AS/400 leverages the clustering capability both to provide high-availability systems and to allow for horizontal growth of the AS/400 systems. The following paragraphs discuss the performance and connection characteristics provided by the OptiConnect capability for horizontal growth. Of course, these characteristics apply also in the high-availability connection for resilient data and resilient applications.

What Is OptiConnect/400?

OptiConnect/400 is the system-to-system interconnection fabric that allows distributed functions and applications to operate at high speed. It can be used to form a high-performance, multisystem, loosely coupled cluster.

Two things differentiate OptiConnect from traditional communications-based distributed operations. The first is a system bus

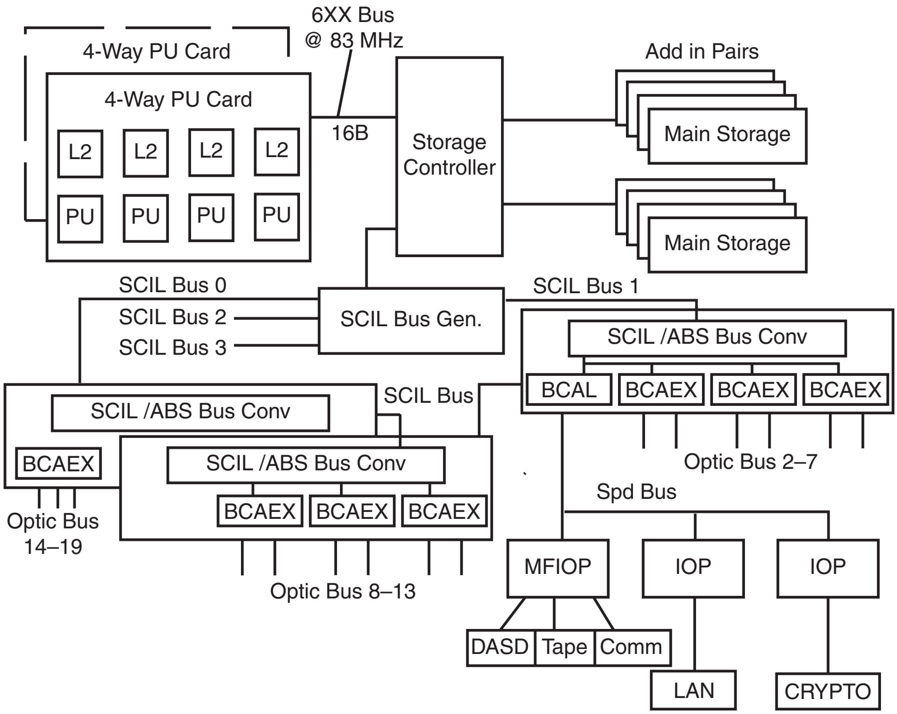

Figure 1.31. N-Way hardware architecture block diagram.

connection between multiple systems using high-speed fiber optic technology. The second is an I/O driver embedded in the operating system that streamlines the application access to data on a remote system.

To accomplish this, OptiConnect provides a shared bus on which systems communicate using a peer-to-peer protocol. Once OptiConnect establishes system connections on the shared bus, much of the APPC (Advanced Peer-to-Peer Communication) protocol stack is bypassed. The OptiConnect fast-path connection for database transactions provides DDM (Distributed Data Management) access to databases anywhere in the OptiConnect cluster at a fraction of the standard communications code path. Data warehouse, DRDA, and data propagation functions can use this technology.

Traditional communication protocol overhead is too impractical for heavy-workload-distributed applications. To minimize latency, or the time it takes to send a message and receive a response, the protocol chosen must be efficient in the number of code steps required. To maximize bandwidth, or the amount of data that can be transferred in a unit of time, high-performance hardware must be used. The OptiConnect channel is very efficient and is the best solution for both latency and bandwidth. The length of the cable affects how low latency can get. As distance increases, the speed of light becomes a limiting factor for a given bandwidth. OptiConnect is limited to 500 meters over the 1,063 Mbps (megabits per second) used on RISC models and 2 kilometers over the 220 Mbps link supported on CISC models.

OptiConnect/400–Enabled Functions

An OptiConnect cluster can consist of up to fourteen systems with full system-to-system connectivity and up to thirty-two systems in complex structures where all satellite systems must communicate with one or two hub systems. Figure 1.33 illustrates the clustering of three systems through an expansion unit hub.

The OptiConnect for OS/400 software provides a streamlined communications path across the dedicated system buses. Together with the hardware, the following functions can utilize the high-speed system-bus level connections between AS/400 systems:

- All **Distributed Data Management** (DDM) operations for supported object types can run across OptiConnect, including data files, data areas, and data queues.

- All **Structured Query Language** (SQL) support provided by DRDA will run across OptiConnect, including distributed unit of work and remote unit of work.

- **DB2 Multisystem** with its DB2/400 support for multinode files will run across OptiConnect, providing data warehouse functions of Query/400 support and two-phase commit.

- ObjectConnect/400 will operate over OptiConnect to provide high-speed system-to-system save/restore functions.

- Standard APPC conversations are available over OptiConnect with an OptiConnect communication controller. This allows for **System**

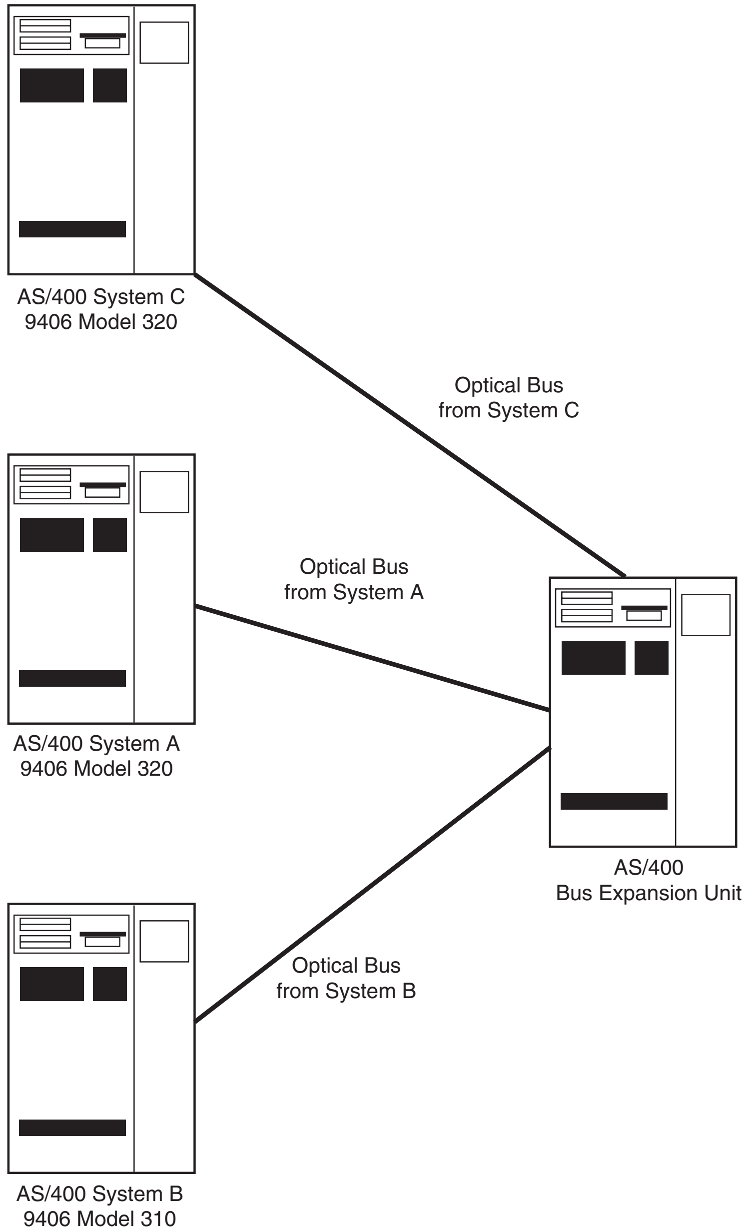

Figure 1.33. Three-system cluster with one bus given up in each system to connect to the expansion tower.

Network Architecture Distribution Services (SNADS), display station passthrough, network printer passthrough, and other functions.

- OptiConnect has an **Application Program Interface (API)** to which business partner software packages can code.

OptiMover for OS/400 is a special, low-cost PRPQ version of OptiConnect for OS/400 software. It enables non-DDM functions to utilize the OptiConnect high-speed link, allowing system-bus-level connections between AS/400 systems. "PRPQ" stands for **Programming Request per Quote**. Business-partner packages written to these APIs allow customers to take advantage of this lower-priced option.

OptiConnect/400 Environment

There are two primary uses for the OptiConnect technology. Horizontal growth is the traditional and most popular use. Separating database operations from applications workload allows multiple systems to operate as a cluster to grow computing power beyond what a single system can provide. It is important to understand that not all applications are conducive to this type of workload distribution. The horizontal growth/scalability is dependent on the database I/O intensity. The best implementation is the separation of the interactive application from the corresponding data while maintaining the batch application on the same system as the batch data. Techniques are available that transparently manage the batch job submission to the database server system.

The second use for OptiConnect technology is *high availability*. When used in conjunction with business-partner applications such as those offered by Lakeview Technology and Vision Solutions, efficient high-availability solutions can be achieved. The OptiConnect technology provides the most efficient use of CPU resources to allow fast replication of data between systems.

Main Storage

The "memory," or main storage, is the set of electronic circuits within an AS/400 system that provides a "workspace" for the system processor. Data and programming instructions are moved from disk storage to main storage for the programs and the processing of data. After execution or processing, modified data are moved back to main storage and kept there until no longer needed. At that point, the data are typically written back to disk storage for safekeeping.

AS/400 systems spend much time moving information between the system processor and main storage. There are two major reasons that this is so. First, all programs currently being executed by the system processor reside in main storage. Therefore, the system processor must (at some point) retrieve every instruction from main storage. Second, main storage holds the data to be acted on by the system processor. Since the information traffic between the system processor and main storage is heavy, the speed of main storage is important to overall system performance. That is why high-speed cache memory, which increases the effective speed of main storage, is included in the system design.

As we have seen, every main storage location in AS/400 main storage consists of sixty-four pieces of information, or bits. All information in main storage is encoded using these 64-bit groupings, called double words. In addition to the 64 bits of information, each word in main storage has several additional bits (called check bits) that are generated based on the value of that particular 64-bit word. In the event that one or two bits of the 64-bit words is somehow corrupted, the check bits notify the AS/ 400 computer that the error exists. If only one bit of the 64 bits is corrupted, as is usually the case, the check bits actually restore the corrupted bit and correct the error. This main storage error detection and correction works to protect the all-important integrity of user information in the computer system.

Storage Management

The methods used within a computer system to manage main storage and disk storage, called the computer's storage management, are fundamental to the capabilities of the computer. Understanding the basics of storage management provides insight into one of the unique features of AS/400 computers compared with traditional computer systems.

Figure 1.34 shows conceptually what the storage in AS/400 computers looks like. All programs and information currently being used by the computer system must be contained in main storage, which resides inside the computer's system unit. Main storage is relatively expensive and responds at very high speeds (compared to disk storage) when called on to provide or store information. Because main storage loses all information when the computer system is turned off, it is called volatile storage.

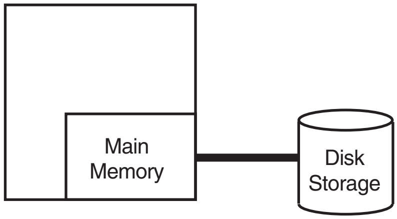

Figure 1.34. AS/400 main storage and fixed disk storage.

Disk storage is less expensive but cannot provide or store information as quickly as main storage. Disk storage is said to be nonvolatile because it does not lose its information when the power is turned off (or lost due to a power failure). As a result of this nonvolatility and relatively low cost, disk storage is commonly used to hold all information that must be readily available to the computer. Disk storage may reside either inside the system unit or inside a separate box cabled to the system unit (as depicted in the figure).

When the AS/400 computer is first turned on, information vital to an orderly startup and smooth operation is automatically copied from the disk to main storage. Once normal system operation is established, users can begin to do their work. During the course of this work, users will start various computer programs. As each program is started, it is copied from the disk to main storage and then executed. Depending on the work being done, the computer programs manipulate various sets of data that are also loaded from the disk as needed. It does not take long to realize that the main storage in a computer can quickly become filled up with programs and data as the system is called upon to do more and more work.

In the earlier days of computing, the main storage size limited the amount of work a computer could manage at any one time. This limitation capped the size of programs, the number of programs that could be run concurrently, the number of users that could share the system, and so on. In modern computers a technique called *virtual storage* (discussed more fully in the next section) alleviates the need to squeeze all active programs and data into main storage. In computers that support virtual storage, the computer basically "fakes out" the computer programs and appears to have much more main storage than it actually has. (The AS/400 allows a virtual storage size of 18 million terabytes). Virtual storage therefore allows more programs, data, and users to be simultaneously active on the system than could be supported in real main storage.

Although virtual storage is a powerful system feature, the "swapping" between disk and main storage is processing overhead that can reduce the overall system performance. A little swapping does not appreciably hurt performance, but increased swapping does. When the swapping performed by a virtual storage becomes excessive, the system is said to be "thrashing," or spending too much time swapping information between disk and main storage. Thrashing can be reduced by increasing the amount of main storage in the AS/400 system through the installation of main storage expansion options. Increasing the main storage in the system provides more room for programs and data, reducing the amount of virtual storage swapping. Thrashing can also be reduced through system management means such as rescheduling work for off-peak periods.

The virtual storage concept is implemented in many of today's computer systems to one degree or another. AS/400 systems implement their virtual storage scheme through a concept called *single-level storage*. This simply means that in AS/400 systems, no distinction is made between disk storage and main storage. All storage appears to be one homogeneous sea of main storage accessed in exactly the same way. This consistency provides for a simple and efficient virtual storage implementation that is the same for programs, data, temporary holding areas, and so forth. The simplicity of single-level storage results in a consistent and more complete virtual storage system than most other implementations.

Another difference between AS/400 storage management and that of conventional computer systems is object-oriented access. With this concept, all programs, databases, documents, and so on, stored in AS/400 computers are stored as independent entities called objects. The AS/400's object-oriented access again provides the user and the programmer with a simple and consistent way of managing all programs and information in the system. Users can access an object by simply referring to its name. The AS/400 security system will check to make sure that the user has authorization to use the object and that it is being used properly. This is called *capability-based addressing*. The AS/400 system manages the complexities associated with the physical location and addressing of the information.

AS/400's implementation of single-level storage and capability-based addressing spreads information through various disk units in a way that optimizes storage efficiency. Objects provide consistency in the areas of security, usage, and systems management for all items stored on AS/400 systems. Objects can be organized into groups called *libraries*. A library (which is also an object) is analogous to a drawer in a file cabinet (or a subdirectory, for those with PC disk management). A library might contain all programs related to the accounting function of a business to keep things organized. Because access to libraries can be restricted by the AS/400 security system, a payroll database, for example, might be kept in a library separate from other business information for security reasons.

Virtual Storage

Virtual storage is a technique for using main storage that alleviates the need to squeeze all active programs and data into main storage. As mentioned earlier, with virtual storage, the computer basically "fakes out" the system, appearing to have much more main storage than it actually has. The virtual storage supported by AS/400 systems is more than a whopping 18 million terabytes in size (over 18 quintillion [18,000,000,000,000,000] unique addresses). This 18 million TB of addressing capability is enough to keep track of the information contained on over 9 quintillion pages of single-spaced computer output—a stack of paper over 800 billion miles high that could stretch between the earth and the moon 3 million times. Virtual storage therefore allows more programs, data, and users to be simultaneously active on the system than could be supported in real main storage without virtual storage.

Virtual storage works as follows: A user tells the computer to start a word-processing program. The computer first attempts to load the needed portion of the program into main storage. If there is no space left in main storage, some space is made available by overwriting an inactive portion of some program, or by "swapping" out some inactive data to a temporary space on the disk. The necessary portion of the word-processing program can then be loaded in the available space, and the user can begin typing the memo. If the program that was overwritten or the data that were "swapped" out are again needed, they are reloaded from a disk unit to some other available main storage area. Therefore, a virtual storage computer system is constantly swapping programs and information between main storage and disk storage, "robbing Peter to pay Paul" and vice versa.

Virtual storage allows the maximum size program or combination of all programs and data to be limited only by the combined amount of main storage and disk storage space rather than by the amount of main storage alone. The advantage of virtual storage is that neither the programmer nor the user of any AS/400 system needs to be concerned with main storage size. The system seems to have as much main storage as needed, and users are never aware that information is constantly being swapped from main storage to disk storage and back again. This swapping is handled automatically.

Auxiliary Storage

Auxiliary storage, commonly used to keep data and program information in all computers, is an inexpensive way to retain and later access information. Information kept on auxiliary storage can be easily modified or kept unchanged over long periods of time as an archive. Because all auxiliary storage is nonvolatile, the information remains intact whether the computer is turned on or off. The AS/400 systems use four types of auxiliary storage: diskette, disk, optical libraries, and tape.

Diskette Storage

Diskettes are a portable magnetic storage medium that can be used to record and later retrieve computer information via a diskette unit. The diskettes consist of a flexible disk with a magnetic surface permanently enclosed in a square, protective outer jacket, as shown in Figure 1.35.

One of the primary functions of diskettes is to provide portable storage, allowing for the transfer of programs and data between computers. To this end, all similarly configured AS/400 computer systems can freely exchange programs and data via diskettes. Also, information

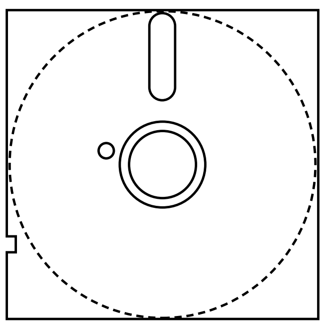

Figure 1.35. Diskette used with application systems.

on System/3X diskettes can be freely exchanged with a properly configured AS/400 computer.

Disk Storage

Earlier in the chapter we introduced another kind of auxiliary storage used with AS/400 systems called disk storage units or **Direct Access Storage Devices (DASDs)**. These are high-capacity magnetic storage devices commonly used in all types of computers from PS/2s to large mainframe computer systems. The basic anatomy of a disk unit is shown in Figure 1.36. Disks consist of a drive mechanism with permanently installed metallic disks, often called platters (because they are shaped like a dinner plate). These platters have a magnetic surface that can store information. Disk units are described in greater detail later in this book.

Disk unit performance is important to the overall performance of a computer system in most applications. This is particularly true in virtual storage and multiuser environments, in which there is heavy transfer of information between disk storage and main storage.

The *performance* of a disk unit refers to the rate at which information can be located and transferred between the disk unit and the main storage. The speed at which a disk unit can position the read/write head over the proper region of the platter is the average seek time, usually expressed in milliseconds (1/1,000 second). After the read/write head is properly positioned, the system must wait as the platter spins until the needed data begins to pass under the read/write head. The average time it takes for the platter to rotate to the proper position is called the average **latency** (also expressed in milliseconds).

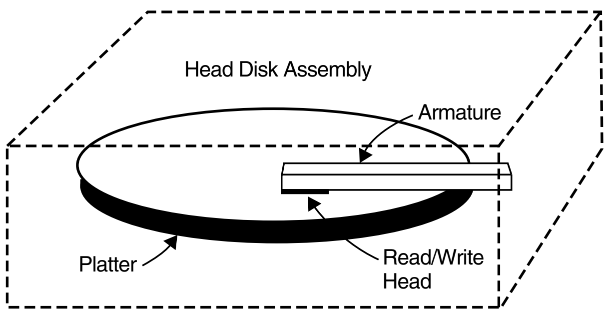

Figure 1.36. The anatomy of a disk unit.

Finally, once the read/write head is positioned and the data begin to pass by on the spinning platter, the information is transferred from the disk unit to the computer system. The speed at which this is done is called the data-transfer rate and is usually expressed in **millions of bytes per second (MBps).** The shorter the average seek time and the average latency, and the higher the transfer rate, the better the performance of the disk storage subsystem and the overall computer system. Figure 1.37 shows the average seek times, average latency, and data transfer rates for some internal disk units that are provided as standard equipment on AS/400 systems.

In addition to the disk unit specifications, there are other performance considerations when configuring the disk storage subsystem of an AS/400 system. Some disk units have multiple actuators, whereas others only have one. (The **actuator** is the accessing mechanism of the disk unit, which carries the read/write head and flies over the surface of the disk detecting or imprinting the magnetic bits.) Often, each actuator

Computer System	Standard Configuration	Disk Size	Averate Seek Time	Data Transfer Rate	Rotation Speed (rpm)	Number of Actuators/ Drives
150	4194 MB	3.5 inch	8.6	20	7200	1
720, 730, 740, 170	4194 MB	3.5 inch	8.6	40	7200	1
150, 170, 720, 730, 740, 250, 270, 820, 830, 840	8580 MB	3.5 inch	8.6	40	7200	1
150, 170, 720, 730, 740, 250, 270, 820, 830, 840	17500 MB	3.5 inch	8.6	40	7200	1
270, 820 830, 840	8580 MB	3.5 inch	8.6	40	10000	1
270, 820	17500 MB	3.5 inch	8.6	40	10000	1

Figure 1.37. Performance characteristics of internal fixed disks provided as a standard in AS/400 systems.

in a disk unit can perform independently, so the more actuators you have for a given amount of disk storage, the better the performance. In fact, choosing disk storage configurations that have more actuators for a given amount of storage can result in higher performance than selecting disk units with faster specifications but fewer actuators. This is particularly true with AS/400 systems because single-level storage tends to spread information over many areas of the disk units. This spreading of information produces more efficient operation because the actuators can all share the load, but it does emphasize the need to follow proper backup procedures.

Optical Libraries

Optical libraries consist of arrays of optical disks associated with one or more optical disk read/write units. In some cases, the optical storage read/write units also have one or more conventional magnetic disk storage units associated with them to improve the write performance from a system perspective. The optical disks may be CD-ROM, WORM, or WMRM technology, each of which imposes different requirements upon the read/write unit and upon the controller within the system. CD-ROM is an abbreviation for **Compact Disk–Read Only Memory,** and the technology presents digital data in a continuous serpentine path across the surface of the optical disk. WORM is an abbreviation for Write Once, Read Many. This technology presents data in circumferential paths across the surface of the optical disk. Because the data will only be written once, this technology generally has the header embedded on the raw media, and a sector corresponds to the data content that can fit in the shortest circumferential track. WMRM, sometimes referred to as erasable optical disk technology, is an abbreviation for **Write Many, Read Many** and also presents data in circumferential paths across the surface of the optical disk. But because the data written at one time must be erased before new data may be written to replace it, WMRM must follow the sectoring, header, trailer, and error-correction rules of magnetic disk technology, including bad track recovery and directory management.

Tape Storage

The last type of auxiliary storage to be covered is magnetic tape, or simply "tape." One primary purpose of tape is to provide a backup storage medium for information on the computer's disk storage. The low cost and high recording densities inherent in tape make it ideal for archiving information. Tape is also very useful in distributing programs and transferring information from one computer system to another. Diskettes can be used for these same functions, but the higher storage capacity of tapes is preferred if you are dealing with a large amount of information. Tape storage consists of a long flexible strip coated with magnetic material and rolled on a reel or into a cartridge.

Packaging Technology

Most of the circuitry in AS/400 systems was built using IBM's version of **Very Large Scale Integration (VLSI)** technology called **Complementary Metal-Oxide Semiconductor (CMOS).** This packaging technology builds circuits with 0.2 micron sized elements, allowing millions of high-speed transistors in a single chip. Main storage is implemented using IBM's 16 Mb and 64 Mb (16 and 64 million bit) memory chip technology. Sixteen-megabit chips are used on all except the largest-capacity memory cards, which use 64 Mb chips to achieve an improved packaging density.

Five basic mechanical designs are used in AS/400 computer system units. The sizes range from the deskside tower of the Model 150 to the dual tall-tower racks that house the Models 740 and 840. The system unit package contains the system processor, main storage, tape, disk units, a CD-ROM, and some number of I/O controllers.

The main storage for the A10 microprocessor-based processors is contained on a card that measures about 7 × 3 inches and plugs directly to the processor card. Main storage Models 170/720 is packaged on standard DIMM form factor cards. Main storage for the A50 microprocessor-based processors for Models 730/70 is contained on separate 9 × 11-inch cards that are plugged into an internal card chassis (called a cage), allowing for electrical connections to the bus. The design is modular, and the various elements can easily be added or replaced without the need for IBM service personnel.

All AS/400 systems use the cable-thru wiring scheme. Cable-thru allows multiple displays or printers (i.e., workstations) to be attached together in a daisy-chain fashion, as shown in Figure 1.38. With cable-thru, it is not necessary to run a separate cable from the computer to each workstation. Instead, a single cable from the computer can be used to attach up to seven workstations over a distance of up to 5,000 feet.

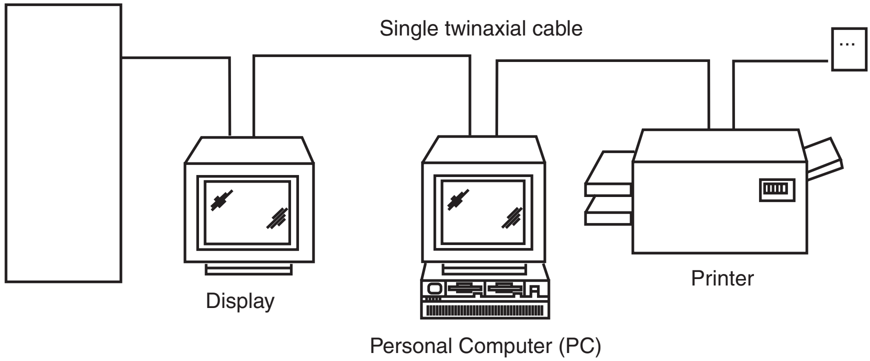

Figure 1.38. All applications use the cable-thru approach to attach to local workstations.

Fiber Optic Bus

The fiber optic bus design continued from the D, E, and F models permits the larger Advanced Series systems to add additional I/O buses to an AS/400 system. This allows the system to accommodate more I/O devices without loading down existing I/O buses. The fiber optic cable that attaches these additional SPD I/O buses (i.e., system unit expansion features) uses laser light rather than electrical signals to exchange information with the rest of the AS/400 system. Using light allows the additional I/O buses to operate at full speed over greater distances and eliminates the electrical interference inherent in electrical cables carrying high-speed signals. The fiber optic cables can carry light signals at a rate of 1,063 Mbps. Except for the rack-mounted 5042 units, V4R2 upgrades that contain the older 266 Mbps fiber optic cables must be upgraded to the 1,063 Mbps type. With these optical cables, the various towers of an AS/400 system can be located up to 100 meters apart rather than within the twelve-foot limit imposed by earlier electrical I/O bus cables. Furthermore, optical cable configurations are available that will allow intertower distances of up to two kilometers in some situations.