

Relevant Case Studies

Chapter 1 A

Case Study 1: The Xerox PeopleNet Story- Use of Information Technology in Business organisation

1. **Background**

With revenues of \$ 18.7 billion in the year 2000, Xerox is a global leader in the document management business. The company offers a vast range of document products, services and solutions in association with its joint-venture partner, Fuji Xerox Co. Ltd. of Japan. The products include digital printing and publishing systems; digital multifunction devices and copiers; laser and solid ink printers; fax machines; document-management software; supplies (toner, paper, ink cartridges, etc.); and comprehensive document-management services such as running in-house production centers and developing online document repositories.

2. **The Challenges**

In the early 1990s, Xerox saw media reports commenting that the functioning of the company's HR department was 'hardly a model of efficiency.' The company realized that the fully centralized functions of hiring, awarding bonuses and granting promotions were resulting in the central HR department being over-burdened. More often than not, HR personnel were tied up in paperwork that left them little time to use their skills elsewhere. As a result, tasks such as counseling managers on ways to handle employee problems were frequently postponed by HR managers. Xerox was not happy with its internal job posting process as well. In the existing setup, hiring managers filled in job description forms and sent them to the HR department, where the information was reentered on posting forms. It often took a week before employees could view the job listings, thus delaying the recruitment process considerably.

3. **The Solution**

In 1992, Xerox began to explore the possibilities of using a computer system to release HR professionals from the monotony of administrative tasks. The company decided to install an intranet application to act as a communication and productivity tool for its 50,000 managers and employees. The idea was to deliver traditional human resource information such as benefits, compensation, policy manuals, phone directory and training as well as the ability to change personal information like name and address, to minimize the HR administrative support needed by employees. It was thus decided to develop the required HRMS in-house. This marked the beginning of the Xerox PeopleNet initiative.

4. **Developing Xerox PeopleNet**

In June 1993 Xerox began with an internal HR survey to identify the desired objectives of the proposed system. The company's existing infrastructure comprised workstations and the basic networking hardware under a mainframe environment. However, mainframe technologies were not user friendly enough for the typical manager/employee skill set. Also, Xerox was shifting from the existing hardware setup to a personal computer (PC) setup. Therefore, client/server architecture was decided upon as the best choice. The development team along with the company's 'Global Process and Information Management Group' (GP&IM) began establishing the basic application/system standards for a technical framework. Prototyping was used to test different development approaches and design the overall graphical user interface (GUI).

The infrastructure support was outsourced from a leading IT services company, Electronic Data Systems. All the software developed was put through strict compatibility testing to ensure that it complied well with other Xerox applications and standard configurations.

By the end of 1994, a pilot was developed, which was approved by the company and funds were released for the production version. In January 1995, a read-only version was launched with simple applications. Xerox opted for a phased launch because it was shifting from mainframes to a PC based setup and the developers needed to understand the new infrastructure well before going in for a full-fledged implementation.

The project followed a phased development approach of prototyping, testing, re-testing and then rolling out on a continual basis. Visual Basic was selected as the client development tool because of its ability to facilitate prototyping. The Microsoft Access engine acted as the invisible middleware component, which tracked configuration and routing information and stored cached data.

To get the software to the users, Xerox adopted various delivery strategies, viz, shipping to local system administrators, using CDs and floppy disks, the company intranet. To access Xerox PeopleNet, employees had to enter both their social security number and a Xerox PIN number.

The software used application-level encryption for passwords as well as transmission-level encryption for data traveling over the network. As the company was shifting to a PC-based setup, implementation of the system was rather slow and had to be in a phased manner. As the hardware setup stabilized and more features were added, Xerox PeopleNet became increasingly popular within the organization.

#### 5. The Benefits

Xerox PeopleNet supported applications covering such areas as training, retirement fund performance and a corporate phone directory. In addition, employees could check Xerox's stock price as well as those of its competitors. The unique feature of Xerox PeopleNet was that unlike typical HRMSs, it did not restrict the availability of information HR staff alone. All employees could access information through any PC on the company's network. Xerox PeopleNet cost Xerox around \$2 million. However, its benefits far outweighed the investment. The solution helped Xerox accomplish its objectives of empowering its people, increasing satisfaction and boosting productivity. In addition, online publication of the human resources manual and other publications saved approximately \$1.5 million annually in printing costs.

Online transaction processing and electronic signature approval capabilities added later saved another \$1.1 million annually by eliminating manual forms and paper-based processing. On the hiring front, managers could open the Xerox PeopleNet application on the desktop, create a posting on an online form and post it immediately on a central electronic bulletin board. Any interested Xerox employee could then print an application form and submit it to the hiring manager in paper form. The paper element was to be completely eliminated over a period of time and internal job applications were to be processed entirely online.

The system included a feedback feature that let employees suggest new ideas and improvements. As a result, employees were able to monitor their profit sharing and retirement plans and change their contributions from their desktops itself. Interestingly, Xerox PeopleNet seemed to have had certain undesirable results as well. Commenting on the massive layoffs by Xerox during the 1990s, analysts said that so long as software such as Xerox PeopleNet continued to render personnel redundant, the trend of manpower trimming was likely to continue.

#### Case Study 2: NAS Case Study: The International Image Processing Company Overview

The International Image Processing Company, or IIP, provides a proprietary digital imaging service to academic, educational, and government institutions throughout the United States. Its computer imaging services provide clients with the ability to transfer physical documents, such as historical books, documents, and newspapers, as well as photographic images, such as historical photos, video images, and medical images, into digital media. The company's clients range from major universities to some of the largest archival government institutions.

Despite the scope of its clients, IIP is a small business with less than \$5 million in revenues, with plans to grow threefold in five years if it can increase its capacities to bring on more clients. IIP has a small but distributed business organization with headquarters in Texas and major client work sites in New York City, Washington, D.C., and San Francisco. Due to the nature of IIP's work, its IT staff is integrated into the specialized work force that the company employs—given the digital scanning software and processes that are integral to their work. The company also employs two dedicated IT technicians—one for hardware and another responsible for infrastructure software. Additional IT related activities are handled through the imaging software support and configuration personnel.

The company was challenged by its dependencies on storage, created through its software scanning processes that were increasingly exacerbated by new clients. Daily operations could quickly utilize a terabyte of data storage. Because the imaging scanning product was archived to clients using both optical and tape storage devices, if operations ran short on available storage, the imaging processes would slow and eventually stop until sufficient space was freed up for operations. This slowed billing and eventually cash flow.

The storage situation reached critical levels that started to impact the company's ability to handle new clients and subsequently impacted its planned business growth. IIP was in the classic predicament of needing a high-performance solution, while being restricted by a minimum budget and a limited IT staff. The evaluation of potential storage solutions ranged from storage area networks (SANs) to IDE RAID solutions. While a network attached storage (NAS) solution would satisfy the size requirement, NAS remained a problematic solution because it also had to be managed remotely at the New York, Washington, D.C., and San Francisco sites.

IIP chose a general-purpose solution, even though its processes were largely proprietary, which was a matter of thoughtful strategy on IIP's part to utilize standard hardware and OS platforms. The NAS solution was chosen for its ability to integrate easily into a small IP environment, while being able to deploy remotely with some level of remote management. This case study will discuss the IIP storage and related processing problems and challenges in detail. These drove the activities involved in identifying the company's workload and analyzing potential storage solutions. Finally, we discuss how IIP concluded with the decision to implement NAS and the subsequent transition to the new storage model.

**1. The Situation Analysis**

IIP developed its storage solution from years of experience in scanning and duplicating photographic images, ranging from military applications to historical documents. The challenges had always been the time required to scan an image, versus the quality required for large clients. This had given way to the amount of space required for each image, given that the IIP client base dealt with millions of images. The type of work IIP performed was not directed toward the market of flatbed scanners nor even larger scanners in terms of quality and production process requirements. Instead, IIP developed a process and proprietary software that used specialized digital cameras to scan images of various dimensions, physical states, and types-various papers, photographic techniques and types, and so on. The process, coupled with the scanning stations, provided a production-oriented environment where imaging could take place 24/7, if required. The process and software included a fully complete life cycle of imaging capture, correction, and quality assurance before it was placed on a CD or tape for shipment to the client.

Clients of IIP had requirements to digitize documents and photographs to provide a wider distribution and availability of these items through the Internet. Consequently, these clients had become part of the growing movement within both the academic community and public sector to save historical documents. For the most part, these markets are just emerging, given the tremendous amount of material that remains to be scanned and digitized. IIP's client requirements could run into the 500,000-plus number of images for a single project. That type of work drove the production-oriented environment introduced by IIP in the late 1990s.

**2. IIP Systems Infrastructure**

IIP maintained a distributed infrastructure with its headquarters in Texas and field locations in New York, Washington, D.C., and San Francisco. Figure A-1 illustrates the configurations at the headquarters' locations. Here you see the integration of capture stations, process servers, correction/quality assurance workstations, database servers, and archive servers. Each field location is set up in an identical fashion. Each is linked to the headquarters' network and web server through a virtual private network (VPN). E-mail and FTP services are handled in this manner. Architecturally, this setup was designed for future development of remote scan processing and diagnostic imaging services. The imaging process will describe the storage utilization scenarios and why the process is so data-centric.

**3. Problem Details**

What IIP had not foreseen were the systems infrastructure requirements for this type of production work. This required the calculation of both processing cycles and, most of all, the amount of storage space that would be needed on a daily operational basis. Because IIP is a small business, it had resisted a formal capacity plan and had relied on its ability to respond quickly when additional capacities were needed. That meant the additional servers were purchased on an 'as-needed' basis, with most of the hardware being 'do-it-yourself' built chassis and motherboard configurations.

With this orientation to developing and maintaining the hardware portion of the systems infrastructure, the storage challenges were met with larger and higher speed internal IDE disks. This gave rise to additional server installations that were needed to handle the post scan image processing. This then prompted the acquisition of dedicated media servers to write out the client images using CD or tape media. This is the archival system on the backside of the process. This meant that a faster network was necessary to speed the transmission of scanned raw files to the process servers, and ultimately it placed the problem back at the storage infrastructure once again as the image scans overtook the capacities on the servers.

A stopgap effort was a move to IDE RAID to provide adequate storage for the process servers. This was largely driven by the 'do-it-yourself' mode of the hardware and severe limitations of budgets constraints. Although the IDE RAID facilitated a quick fix, IIP's flexibility in providing reliability and backup protection was problematic. In many cases, the volatility of the data movement over the period of one week could easily surpass more than five terabytes running through a single process server. As the tremendous write activities continued, the IDE drives generally failed twice a month, with minimal success running data protection with RAID level 1. However, the space it provided offered a brief interlude to the space problems that shut down the process entirely. Given that additional business was coming in with more restrictive time constraints for completion, IIP concluded that a longer term solution had to be found.

**4. The Search and Evaluation**

The IIP hardware and software IT personnel researched the solution along with assistance and input from the imaging software specialist. They found that a SAN was a valid consideration, since it appeared to be the choice of others working with unstructured data such as video and audio projects. However, they found that the imaging system, although proprietary by its methodology, used open and commodity levels of hardware and operating environments and was further open to additional solutions that integrated well into the small business environment. Another alternative was to move to larger process servers with external

SCSI drive arrays, to scale up in both process and storage power. Yet another alternative was the consideration of a NAS solution, which would integrate easily with the existing network, would use file systems, and would have the capacity they needed.

5. **Estimating the Workload**

The IT personnel, working with an outside consultant, identify the company workloads. They further moved to understand the types of configuration needed for all three possible alternatives. First looking at the SAN configuration, followed by the larger server with external RAID, and finally the potential NAS configuration. The results are summarized in the following sections.

6. **Workload Identification:**

Looking at a year's history of scanning images, the IIP team concluded that the workload was complex and data-centric, and it fit somewhere between online transaction processing (OLTP) and data warehousing. The workload encompassed OLTP characteristics when scanning the image and then transmitting the write transaction to the process server. Although developed as a synchronous process, it was recently changed to an asynchronous process to facilitate greater throughput at the capture station. However, this still required a sequential write process at the process server as each image was scanned.

On average, the image scans were 300MB in size. The calculation of 300 images per shift × three capture stations working two shifts provided the necessary throughput. It was determined that at least 540GB of free space was needed to accommodate the daily scanning process. This required that the overall storage infrastructure be able to accommodate a 566MB per second throughput rate.

7. **Workload Estimates for SAN:**

we can quickly calculate that the required components for a SAN which could be handled by one, 16-port switch, given that a single point of failure is acceptable for the installation; or it could be handled by two, 8-port switches for some level of redundancy. Three HBA adapters with 2 ports each for redundancy and performance would be required for the process servers. However, not to be overlooked, this configuration will require additional Fibre Channel storage arrays to accommodate and be compatible with the new Fibre Channel storage network. Given that the total capacity of 540GB needs to be available every 24 hours, we can estimate that two storage arrays of 500GB each would provide the necessary capacity with sufficient free space to handle peak utilization as images are processed through the system.

8. **Workload Estimates for Direct Attached Storage:**

Aligning the requirements to new servers, we find that all the process servers would have to be upgraded. This would also require that the storage capacities be carefully aligned with each process server. Even with this alignment, specific workload affinity would have to be observed to utilize the storage effectively. On the other hand, the process server could more easily share storage across the network but would have to reflect some level of duplication for storage requirements to accommodate the total capacity, essentially doubling the entire storage requirement.

This would require, in addition to new servers installed, OS software upgrades, with appropriate maintenance and all the necessary activities of a major system installation. It would result in a normal disruption of service and reliability characterized by new system installations. However, the new servers would have to be configured to handle the I/O throughput of an aggregate of 566MB per second. This would require each server to handle 188MB per second if the workload is evenly distributed, which in most cases will not be the case; however, we will use this for estimating purposes. That relates to a minimum of six

Ultra-wide SCSI-3 adapters necessary to handle the sustained rate of 188MB per second. This requires the total storage to be divided among the servers, and subsequently the adapters, and places a limitation of approximately 120GB per Logical Unit Number. Thus, a more complex management problem in terms of flexibility of reconfiguration based on required storage would be necessary, given that one capture station could generate 180GB of images every 24 hours.

9. **Workload Estimates for NAS:**

We can calculate that our workload requirements are definitely within the mid-range NAS device configuration and probably just under the enterprise NAS solutions.

We can select mid-range even though our sizing factor is within the enterprise range. This is based on the special application circumstances and because the aggregate data is below a terabyte and would be physically segmented within the aggregate data capacity estimate. In addition, we considered the workload being further characterized by limited users working with an almost dedicated Gigabit Ethernet network.

The NAS solutions also offer the flexibility of storage incremental selection-for example, installing two large NAS servers and one small server, or one large and two medium-sized servers. These solutions also provide the flexibility of RAID processing, network compatibility, and non-disruption to the existing server configurations. In addition, these solutions can be easily configurable to support the scanning projects and mapped as network drives with the same flexibility. They will also provide a closed, yet remotely accessible, solution for the remote network configurations.

One last word on our estimating process: We recognize the characteristics of the small integrated IT staff and the company's lack of any formal capacity planning activities. The process of workload identification and estimates provides this company a level of direction and planning. The result of this exercise has identified that the mid-range NAS devices can meet the company's workload now and within a limited planning period. However, it also provides an insight into future challenges IIP will encounter, as its staff has become aware that it borders on moving into enterprise solutions of either the NAS type or probably a SAN if the budget for infrastructure can support either.

#### 10. The IIP NAS Solution

The final configuration provides increased storage capacity and enhanced performance. The IIP scanning process is now being up-scaled with new customers and imaging projects. In addition, the space is being utilized on an automated basis from new project initiation to deletion after delivery. Figure A-3 illustrates the NAS configurations that support the revised IIP systems infrastructure. Figure A-3's NAS configuration illustrates the increased capacity for storage of images, but it also provides the foundation for image access through the Internet. This was designed into the system to allow the remote work sites that have identical installations to upgrade to NAS devices. Over and above the increased capacity at all sites will be the potential ability to access images across the IIP storage network infrastructure. In other words, images scanned in New York could be evaluated in Texas, Washington, D.C., San Francisco, or other remote offices as they begin operation. This is an added business efficiency that takes advantage of the imaging expertise across the company without regard to location. It also allows clients eventually to be provided test and quality images through the Internet to further facilitate the delivery and client acceptance process.

In summary, the NAS solution turned out to be a good fit for IIP. Even though its application appeared to have many proprietary processes and software processes, NAS operated within the bounds of leveraging commodity infrastructures for networks, systems, and storage. With the exception of the database server, all the servers within the IIP infrastructure now share the available storage in the NAS devices. The acknowledgment of this value further expands the use of the NAS solution as the company began to configure its web and imaging software code development on the NAS devices.

The NAS solution for IIP provided a cost-effective solution for a small company, but it also provided the company with the necessary storage resources for expanding its business, which had grown dependent on a storage-centric product. The company's storage infrastructure is now poised to grow into the future with either more sophisticated products such as a SAN, or to continue to scale economically with further uses of NAS devices. Either way, IIP now has the business flexibility to meet the growing dynamics of the imaging business.

### CASE BASED QUESTIONS

1. Dr. Kartik is the director of an autonomous institute of Weather and Economic forecasting. He is planning to setup a new state of art research centre for weather forecasting that would predict the economic growth of the country based on agriculture production. He is also planning to setup an educational institute for students of weather and economic research, providing bachelors, masters and PhD degree. Dr. Kartik has the responsibility to setup the centre. He is involved in planning, coordinating, controlling and administrating all the activities including building and infrastructural facilities, purchasing hardware and other equipments and so on. He has all the budgetary approval for these facilities.

#### Questions:

- (a) What type of computers should Dr. Kartik procure for the research centre for weather and economic forecasting and why?
  - (b) What type of computers should Dr. Kartik procure for the education institute for weather and economic forecasting and why?
  - (c) What will be the budgetary implication of decisions to procure computer systems by Dr. Kartik? Explain.
- 
2. Mr. Rahul is working as a Research Officer in Internet based multinational consulting firm in Delhi. His main job is to provide the support to the existing clients and gather the information about future clients so that company would be able to give necessary clientele support in near future. For the purpose of making MIS reports on daily basis, it is very much required to connect to the people providing blogs over internet all the time, the system must be very fast and efficient enough to respond quickly and also equip with latest technology to work upon various application in integrated environment.  
Earlier he was working with PIII computer with 80GB HDD, 256 RAM, CD\_ROM, 1.44 inch Floppy drive, Color Monitor and Internet Dialup connection. The performance was not so good over the years. Mr. Rahul has requested to provide new computer system with latest technology. The company has given the approval to purchase a new system with latest configuration with minimum price. Rahul has started searching various computers in the market ranging from Rs. 20000 to Rs. 50000 based on whether they are branded or configured.

**Questions:**

- (a) What type of computer system Mr. Rahul would like to purchase? Is there any specific type of computer is available in the market? Explain.
- (b) What type of configuration Mr. Rahul would like to see and take the decision to purchase the system that satisfies the requirements of his job responsibility? Explain.  
What will be the budgetary implication of decisions to procure computer systems by Mr. Rahul? Explain.

**Chapter 1B**

**Case Study: Automated Scoring of General Assessment Test for Students**

**1. Introduction**

In response to education reform In India and accountability requirements, most recently, the Right to Education Act, national agencies, states, and large districts across the country has implemented far-reaching systems that reform educational standards and include large-scale assessments to monitor student achievement of those standards. Basic reform strategies emphasize the need for high academic standards describing what all students should know and be able to do and high-quality assessments aligned with those standards. States have developed and implemented assessment programs that correspond to curriculum standards and assess students in designated grade levels and subjects.

**2. Challenges**

The authorities has combined both traditional and innovative measures in the design of largescale assessment programs, most of which include multiple-choice, short answer, openresponse questions, and writing tasks for elementary, middle, high school, and college students. The teachers and local educators are usually involved in the processes of test development, standard setting, scoring, and interpretation of results. The objective of this case study is to address the challenges faced by the teacher in scoring of the assessment test papers that are returned after test administration.

The key challenges faced in the scoring process:

- (a) To process vast number of answer sheets.
- (b) To provide results in timely manner to meet Department of Education (DOE) deadlines.
- (c) To assure 99.9% accuracy of all student results.
- (d) To provide highly automated processing with minimum human intervention.
- (e) To provide scalable solution to meet peak processing.
- (f) To get cost effective solution.

**3. Kind of document to be processed**

Once the test development is completed, the end result is a student test booklet containing various content areas like Computer, Math, Science, etc. Each content area contains multiple test questions. The question could be in form of multiple-choice, short-answer, open-response, or writing. The first page of the booklet and the multiple choice questions has "bubble" areas to fill in the answers. The "bubble" area is a set of blank ovals or boxes that correspond to each question, usually on separate answer sheets of paper or underneath the question. Students use 2 HB pencils to mark their answers, or other information by darkening ovals areas. The student booklet also contains a bar code describing the serial number of the booklet. The serial number binds the student to the booklet. Each sheet of the booklet contains a page number and it is related to the configuration of bubble areas on each page.

**4. Solution**

The customer has implemented automated scoring solution based on Forms Processing application to recognize machine printer character using Optical Character Recognition (OCR), penciled bubbles using Optical Mark Recognition (OMR), handwriting using Intelligent Character Recognition (ICR) & Barcode technologies off the student test booklet. Students' booklets are cut into sheets and scanned by high speed scanning systems that use Kodak 840 scanners to create JPEG image files. These image files are then processed by Forms Recognition Processor that is calibrated to recognize various booklets types and the contents of each booklet type including bubble, bar-coded, and machine printed areas.

**Case Study 1: An Expert System - Mycin****1. Overview:**

Mycin was one of the earliest expert systems, and its design has strongly influenced the design of commercial expert systems and expert system shells. Mycin was an expert system developed at Stanford in the 1970s. This expert system was designed to identify bacteria causing severe infections, such as bacteremia and meningitis, and to recommend antibiotics, with the dosage adjusted for patient's body weight — the name derived from the antibiotics themselves, as many antibiotics have the suffix "-mycin". The Mycin system was also used for the diagnosis of blood clotting diseases.

**2. Challenges:**

MYCIN jobs was to diagnose and recommend treatment for certain blood infections. To do the diagnosis "properly" involves growing cultures of the infecting organism. Unfortunately this takes around 48 hours, and if doctors waited until this was complete their patient might be dead! So, doctors have to come up with quick guesses about likely problems from the available data, and use these guesses to provide a "covering" treatment where drugs are given which should deal with any possible problem.

**3. Solution:**

Mycin was developed partly in order to explore how human experts make these rough (but important) guesses based on partial information. However, the problem is also a potentially important one in practical terms - there are lots of junior or non-specialised doctors who sometimes have to make such a rough diagnosis, and if there is an expert tool available to help them, then this might allow more effective treatment to be given. Mycin represented its knowledge as a set of IF-THEN rules with certainty factors. The following is an English version of one of Mycin's rules –

IF the infection is primary-bacteremia AND the site of the culture is one of the sterile sites AND the suspected portal of entry is the gastrointestinal tract THEN there is suggestive evidence (0.7) that infection is bacteroid.

The 0.7 is roughly the certainty that the conclusion will be true given the evidence. If the evidence is uncertain the certainties of the bits of evidence will be combined with the certainty of the rule to give the certainty of the conclusion. Mycin was written in Lisp, and its rules are formally represented as Lisp expressions. The action part of the rule could just be a conclusion about the problem being solved, or it could be an arbitrary lisp expression.

However, Mycin used various heuristics to control the search for a solution (or proof of some hypothesis). These were needed both to make the reasoning efficient and to prevent the user being asked too many unnecessary questions.

**4. Mycin Strategy:**

One strategy is to first ask the user a number of more or less preset questions that are always required and which allow the system to rule out totally unlikely diagnoses. Once these questions have been asked the system can then focus on particular, more specific possible blood disorders, and go into full backward chaining mode to try and prove each one. This rules out a lot of unnecessary search, and also follows the pattern of human patient-doctor interviews. The other strategies relate to the way in which rules are invoked. The first one is simple: given a possible rule to use, Mycin first checks all the premises of the rule to see if any are known to be false. If so there's not much point using the rule. The other strategies relate more to the certainty factors. Mycin will first look at rules that have more certain conclusions, and will abandon a search once the certainties involved get below 0.2.

**5. Mycin Stages:**

A dialogue with Mycin is somewhat like the mini dialogue, but of course longer and somewhat more complex. There are three main stages to the dialogue. In the first stage, initial data about the case is gathered so the system can come up with a very broad diagnosis. In the second more directed questions are asked to test specific hypotheses. At the end of this section a diagnosis is proposed. In the third section questions are asked to determine an appropriate treatment, given the diagnosis and facts about the patient. This obviously concludes with a treatment recommendation. At any stage the user can ask why a question was asked or how a conclusion was reached, and when treatment is recommended the user can ask for alternative treatments if the first is not viewed as satisfactory.

**6. Results:**

Research conducted at the Stanford Medical School found MYCIN to propose an acceptable therapy in about 69% of cases, which was better than the performance of infectious disease experts who were judged using the same criteria. However, MYCIN was never actually used in practice. This wasn't because of any weakness in its performance. The greatest problem, and the reason that MYCIN was not used in routine practice, was the state of technologies for system integration, especially at the time it was developed. MYCIN was a stand-alone system that required a user to enter all relevant information about a patient by typing in response to questions that MYCIN would pose. The program ran on a large time-shared system, available over the early Internet (ARPANet), before personal computers were developed. In the modern era, such a system would be integrated with medical record systems, would extract answers to questions from patient databases, and would be much less dependent on physician entry of information. In the 1970s, a session with MYCIN could easily consume 30 minutes or more— an unrealistic time commitment for a busy clinician.

MYCIN's greatest influence was accordingly its demonstration of the power of its representation and reasoning approach. Rule-based systems in many non-medical domains were developed in the years that followed MYCIN's introduction of the approach.

7. **Future Scenario:**

In the 1980s, expert system "shells" were introduced (including one based on MYCIN, known as E-MYCIN and supported the development of expert systems in a wide variety of application areas.

Similarly, CADUCEUS was a medical expert system finished in the mid-1980s, intent to improve on MYCIN - which focused on blood-borne infectious bacteria - to focus on more comprehensive issues than a narrow field like blood poisoning ; instead embracing all internal medicine. CADUCEUS eventually could diagnose up to 1000 different diseases. While CADUCEUS worked using an inference engine similar to MYCIN's, and has been described as the "most knowledge-intensive expert system in existence". Also a new expert system called PUFF was developed using EMYCIN in the new domain of heart disorders. And system called NEOMYCIN was developed for training doctors, which would take them through various example cases, checking their conclusions and explaining where they went wrong. A difficulty that rose to prominence during the development of MYCIN and subsequent complex expert systems has been the extraction of the necessary knowledge for the inference engine to use from the human expert in the relevant fields into the rule base (the so-called knowledge engineering).

**Case Study 2: With SAAS ERP, Oxford Bookstore Increased its Revenue By Rs 3.4 Crore**

1. **Summary:**

The problems of excess purchases, over-stocked SKUs (stock keeping unit) and slow-moving inventory plagued the iconic, 90-year-old Oxford Bookstore. A SaaS-based ( Software As A Service), on-demand ERP came to its rescue.

2. **Organization:**

Standing regally in the heart of Calcutta's Park Street, the Oxford Bookstore has been a haven for book lovers and intellectuals since the 1920's. A part of Rs.1,500-crore Apeejay Surrendra Group, the iconic bookstore celebrates 90 years in 2010 (it also marks the centenary year of the group). And IT brought a big gift to the party: a Rs.340 lakh boost in revenue.

3. **Case Study Highlights:**

- The ERP solution has helped increase footfalls into Oxford's bookstores by 25 percent and swell bill sizes by 20 percent.
- The SaaS ERP solution has also made other aspects of running a retail business easier, including managing a CRM, a loyalty program and discount schemes.

But things weren't always this good. When CTO, Apeejay Group, joined in April 2008, the bookstore's IT infrastructure suffered from serious challenges. "There were significant inefficiencies with the point of sale function. And since the inventory management system was not integrated with the POS system, we lacked control over our inventory," he says.

4. **Business case:**

That's something the Oxford Bookstore's head of finance and accounts, has seen the bookstore battle for a long time. "Traditionally, the book retailing business has been prone to excess purchases, an over-stocking of SKUs, and a pile-up of non-moving or slow-moving inventory."

It's a problem that all Indian book retailers without IT systems face, given that India produces over 80,000 new titles a year. Oxford, for example, adds about 50,000 new titles a year to the 2 lakh it already maintains. What made it harder for Oxford was its decentralized approach. This allowed a single title to have multiple codes across each of Oxford's 25-plus stores, making it almost impossible to get a fix on the number of copies it had of a popular title like, say, Shantaram. Inevitably, this led to high inventories. "With a dispersed retail operation and increased volume of transactions, the challenge was to find a way to regulate processes and manage transactions," says head of finance and accounts. And with over 150 additional SKUs a day, Oxford's problem wasn't getting smaller. Despite these challenges, CTO's bosses in Kolkata gave him a tall order. "We wanted to encash on an early-mover advantage by rapidly expanding our retail operation, both in India and abroad," says head of finance and accounts. "The vision was to operate over 100 bookstores by 2010, the eve of centenary of the Apeejay Surrendra Group."

**5. Project:**

If the iconic bookstore was going to meet that target, it needed visibility along its entire supply chain, from its supplier to down to its stores — the sort of visibility that an ERP system could provide. The problem was that an SME like the Rs 25-crore Oxford Bookstore didn't have the deep pockets an ERP implementation called for. It's a growing problem for India with SMEs upping their contribution to the country's GDP from 17 percent to 22 percent by 2012, according to ASSOCHAM. "Management was looking for a model that would allow Oxford Bookstore to leverage its full IT potential without large one-time IT capital investments or issues of scalability," says CTO of Oxford Bookstore. Theoretically, that's exactly what a SaaS model offers. And although some solution providers have tested it with industries like apparel, admittedly, the idea remains unconventional in India — and it had almost never been tried in Indian retail before.

**6. Challenges:**

Also, few providers offered the complete set of services CTO wanted and there was no guarantee that once a deal had been inked CTO's small account wouldn't get jostled aside by larger ones. At the same time, Oxford's management was apprehensive about using a SaaS model because it had heard about security challenges, and had fears about vendor lock-in.

Most of these issues could be dealt with, if CTO could find the right vendor and more research and relationship later, he did. "We decided on a vendor only after we were sure they would not just be an implementer but a fully-involved business partner," says head of finance and accounts. "Now, we have one of the very few full-cycle implementations of SAP IS retail delivered via a SaaS model and we're a first-of-its-kind in the book retail industry."

**7. Benefits:**

Today, the ERP-SaaS project has centralized all Oxford titles across all branches into a single directory controlled by an ERP. Data from here flows to all modules including a store inventory module and the bookstore's points of sale. It also covers the company's payroll and accounting needs, and Web tools, says CTO. The results of the Rs 1.08-crore ERP-SaaS project (which is to be paid in monthly installments over five years) is the best gift Oxford could have asked for. Revenues have catapulted a full 20 percent adding a whopping Rs 3.4 crore to the bookstore's coffers.

The ERP-SaaS project also helped increase footfalls into Oxford's bookstores by 25 percent and swell bill sizes by 20 percent. Profits have been spectacular on the e-platform too. Online sales have jumped a full 100 percent and revenues have escalated by 50 percent. It's also made other aspects of running a retail business easier, including managing a CRM, a loyalty program and discount schemes. "We have optimized our merchandising mix, reduced the carrying cost of inventory, and increased our monitoring on non-moving and slow-moving inventory. Also we have managed to regulate discounts and credit limits," says head of finance and accounts. CTO estimates that Oxford will recover the cost of the ERP-SaaS implementation in four years.

---

**Case Study 3: Sheela Foam Boosts Bottom-line With An SMS Solution Executive Summary**

---

The Rs 900-crore Sheela Foam was struggling to focus on its most important business objective: improving profit margins. Using their in-house ERP, the IT team created a system that sent every employee daily and weekly reports via SMS.

**1. Organization**

A few years ago, the manufacturer of one of India's most popular mattresses Sleepwell, (and the leading producer of other industrial foam products such as Feather Foam), was having sleeping problems. The Rs 900-crore Sheela Foam was struggling to get the organization to focus on its most important business objective: improving profit margins.

**2. Case Study Highlights**

- Instead of using an expensive collaboration system, head-IT, Sheela Foam, employed SMS technology to bring focus to the company, raising turnover by 45 percent.
- Using their in-house ERP, the IT team created a system that sent every employee daily and weekly reports via SMS. It included the throughput generated during that week.

That's not to say Sheela Foam wasn't doing well. According to industry reports, the manufacturer of polyurethane foam products owns 40 percent of the Indian PU foam market and is well ahead of its competitors. But until recently its very size hobbled it.

### 3. Business case

Sitting in their headquarters in Ghaziabad, Sheela Foam's management was having the hardest time pointing their 2,000-plus strong organization (with 13 manufacturing plants) in one direction. "Over time, each department developed their own set of performance parameters," says head-IT, Sheela Foam. "We realized that we needed a system that could help us align our performance parameters and achieve an overall business objective." Another part of the problem was that because the company's reporting and reviewing mechanism worked on a monthly basis, it made course corrections possible only at a month's interval, says head-IT. Which meant that whatever solution head-IT came up with needed to allow the company to review its performance at a higher frequency. It is crucial for Sheela Foam to monitor its operating margin because of its high dependency on its primary raw material: polyol. According to a CRISIL (a premier credit and rating organization) analysis of its business, Sheela Foam's operating margins were under pressure in part because of fluctuations in polyol prices, which are linked to volatile international crude oil prices and foreign exchange rates. To deal with the problem, Sheela Foam's management adopted an approach called the Theory of Constraints, made popular by Dr. Eliyahu Goldratt in his book *The Goal*. "The Theory of Constraints provides focus in a world overloaded with information," says head-IT. "It helps organizations by focusing on a few issues: the constraints to ongoing profitability. The Theory of Constraints focuses on the ability of an organization to achieve its goal by improving its single, but most important, constraint."

### 4. Challenges

For Sheela Foam, those constraints boiled down to four parameters. One of these was throughput, the difference between unit sale price and unit variable cost. "We realized we were not able to keep up with the competition because we could not assess the margin at which our products were sold," says head-IT. "This acted as a major constraint because we could not decide which product should be given priority or if pricing needed to be changed." Aligning the entire organization to focus on margins was the solution.

### 5. Project

Faced with this challenge, head-IT could have decided to implement a complex collaboration project, linking various departments and installing a monitoring system. But he settled on a much simpler and more elegant SMS solution. Using their in-house ERP, the IT team created a system that sent every employee daily and weekly reports via SMS. It included the throughput generated during that week, the company's growth over last year, material dispatched the previous day, and outstanding orders, among others. Today, the dissemination of that data allows Sheela Foam to make better calls and stay on top of its margins. "The throughput for each transaction is calculated at the order level, so that orders that generate throughput below a defined threshold can be put on hold," says Managing Director, Sheela Foam. "This helps us take corrective and preventive steps faster." It also allows management to highlight cases that generated low throughput. And it motivates employees to realign their focus to controlling margins, and thus, the bottom line of the business. "The SMS brings everyone on the same page by providing a clear picture of everything that happened up to yesterday and where we are today," says head-IT. "We understood that 'measures drive behavior' and thus we aligned all our measures accordingly. These SMSes keep everyone on their toes," he says. "Today our entire sales staff — and the senior and top management — are updated everyday at 6 AM," says Managing Director. "Some people call it their morning alarm," adds head-IT smiling.

### 6. Benefits

Introducing a higher level of awareness is bringing in better returns for the business. After the Rs 5-lakh SMS project was implemented, Sheela Foam's average weekly throughput increased from 101.5 lakh to 272.9 lakh; a 168 percent growth. That drove the company's turnover up by an impressive 45 percent. The constant monitoring at multiple levels also reduced the company's operating expenses by 5 percent and upped its delivery rate from 76 percent to 89 percent. This healthy operating efficiency is one of the reasons credit rating agency CRISIL recently upgraded Sheela Foam's rating. "Monitoring throughput has enabled us to increase our profits and market share. It has helped us to identify weak areas both in market and product terms, which are now the main focus areas for improvement. We have seen huge growth because of focused weekly monitoring and a strong daily-update mechanism," says head-IT.

---

## Chapter 2

### Case Study 1: PNB deploys an Enterprise Wide Data Warehouse

#### 1. Executive Summary

Punjab National Bank is enabling over 55,000 employees in over 4,500 locations to serve customers and conduct regular banking duties having IT budget of Rs.285 crore and a 500 strong IT team members. Like other banks, PNB too had to adhere to risk compliance and have a 360 degree view of its customers and MIS. Not to forget, with the downturn pushing the BFSI sector into rough waters, cost optimization was a necessity.

## 2. Case Study Highlights

The EDW project not only brought dispersed data under one umbrella, it also enabled the bank to meet regulatory requirement and implement an anti-money laundering solution (AML). The CRM helps the bank now to better retain existing customers by cross-selling and upselling.

## 3. Challenges and Solution

The EDW was meant to leverage the bank's operational data available in multiple sources systems and give access to data required for regulatory, statutory reporting and analytical purposes. The existing heterogeneous data sources included a core banking solution and a total bank automation solution. On top of that, there were various other systems in different head office (HO) divisions across the country and data available was often incomplete and inaccurate. There was a need to cut manual intervention as much as possible.

"For providing a 360 degree customer view and aiding the bank in strategic decision making, the EDW needed to integrate data from these various source systems," said Chief General Manager, PNB. The EDW project not only brought dispersed data under one umbrella, it also enabled the bank to meet regulatory requirement and implement an anti-money laundering solution (AML).

With the AML, alerts regarding suspicious transactions are now generated automatically based on the set rules and are resolved at head offices. Prior to this, such information was sought from branches which were collated at head offices for identifying suspicious transactions, requiring much higher levels of manual intervention. "An additional benefit was the resulting CRM solution deployed across all the branches of 11 identified circles. This enabled the marketing teams to effectively implement their strategy, making good use of higher visibility," said Chief General Manager, PNB. The CRM helps the bank now to better retain existing customers by cross-selling and up-selling, to attract new customers by offering various value added products and services, and even convert loss making customers into profitable ones.

## Case Study 2: Godrej's Data Warehousing Initiative: A Warehouse of Value

### 1. Overview

In 1995 Godrej Consumer Products Limited (GCPL) decided to implement MFG/PRO, an ERP package from QAD Inc. Since its complete roll-out the company experienced numerous benefits. Inventory levels reduced, workflow processes standardized, business processes were more disciplined, manpower was better planned, and product cost information could be viewed instantly.

The use of the ERP also ensured that data was generated in a consistent and structured format which could be easily archived. This invaluable data was lying idle unless required for an ERP query or statistic. "We devised a way to extract the most out of this critical data. We decided to implement data warehousing applications and functions," said General Manager, Information Systems, Godrej Industries Limited.

After the creation of the data warehouse the company has been able to reap a number of benefits. It performs jobs like contribution analyses, profit and loss analyses, and sales breakup analyses. These help enhance product lines, build greater customization, and favorably impact the bottom line. This combination of ERP and data warehousing applications has also paved a future path for GCPL's e-commerce initiative.

### 2. The warehouse is born

The entire ERP roll-out in 120 Godrej group companies and locations was complete in 25 months. And GCPL's employees now depend heavily on the ERP for their daily operations. Vital data is collected and collated to the company's benefit. Since there's a lot of formatted standardized data in the system, it makes good business sense to use it to the company's benefit.

"The large amount of data was virtually jailed, and could have many potential productive uses," said General Manager, Information Systems, Godrej Industries Limited. GCPL decided to use data warehousing tools from Oracle to accomplish the task. The range of data warehousing products includes Oracle Express Server and the OLAP (Online Analytical Processing) client. The system serves as a decision support platform based on historical sales and cost pattern analysis. The data warehousing suite of products cost around Rs.10 lakh.

An RDBMS from Progress Software was implemented at the backend to support the ERP functionality. The company evaluated Oracle as well as Progress and settled on Progress because it was a more robust solution at that time. The ERP and warehousing applications are not linked online. Data is extracted from the ERP system with the help of extract routines and uploaded into an Oracle warehouse with the help of upload routines once at the end of every month. This creates a separate undisturbed database in the warehouse. This data is now routinely processed for a few days and results derived from it.

The OLAP draws from the collected database and performs analysis, calculation, and recalculation to support what-if scenarios and other strategy-setting aides. The Express Server uses a caching scheme to store, manage, and analyze relational data.

**3. Case Study Highlights**

- (a) **The Company:** Godrej Consumer Products Limited is the flagship company of the Mumbai-based Godrej Group. The company manufactures consumer products like soaps, detergents, and hair care solutions. The group has 18 factories and 120 locations all over India.
- (b) **The Need:** A lot of ERP data in a standardized format was lying idle in the company's servers. It would be a good idea to use the data for other intelligent applications that can enhance business productivity.
- (c) **The Solution:** The Company decided to implement data warehousing applications and functions to help extract maximum usefulness from the data. A range of tools from Oracle were used to make this possible.
- (d) **The Benefits:** The Company now can perform a number of contribution analyses, profit and loss analyses, and sales breakup analyses from the data in its warehouse. This has given it competitive advantage and the ability to manage resources better. The applications have also paved way for future e-commerce initiatives.

**4. Warehouse benefits**

Data warehousing tools are mostly used for analysis and trends that allow the company to create short- and long-term strategies and business problem solutions. "You can view interesting trends in sales and costs. And if you like, you can integrate the findings from standard reports and the warehouse analyses to create strategic advances. The limit of the capabilities of a warehouse is bound by the creativity of the end-user," said General Manager, Information Systems, Godrej Industries Limited.

The company can load data like sales figures of goods from the factory to the distributors (primary sales), from distributors to retailers (secondary sales), and from retailers to end-users into a common database in the warehouse. The data warehouse can now allow us to derive sales and inventory trends.

Suppose primary sales have increased over a period of time and secondary sales have decreased. This can lead to 'dumping' or excess inventory. The trend analysis can thus help avoid a situation like this and maintain a consistent flow of inventory.

**5. E-commerce initiative**

In July 2001 the management at Godrej decided to implement an e-commerce solution to manage the supply chain better and an e-CRM package to foster better ties with its distributors. After going through a long list of vendors and consultants, Godrej selected BroadVision as the provider and Satyam Infoway as the implementer of the solutions. "We will set up inventory norms for all the products at the distributor level and our systems will interact with systems at the distributor's end. This will allow us to extract figures like the closing stock and factory dispatch volumes. The system can now make suggestions like the amount of inventory to order next week and the minimum inventory to maintain," said General Manager, Information Systems, Godrej Industries Limited.

**Case Study 3: A Disaster Recovery Plan reduced outages and helped maintain uptime for Ericsson**

**1. Organization**

Ericsson has been associated with the India telecom industry for over 100 years. The association began back in 1903, when they supplied manual switchboard to the government. Ericsson today has a pan India presence and provides mobile networks to all major Government and private operators.

**2. Case Study Highlights**

The business' contention was simple: reduce the number of infrastructure and service outages the enterprise was facing. Based on the amount of bandwidth, infrastructure and manpower that could be tapped, CIO reckoned they could trim service outages to three days

**3. Business Case**

For almost a year, CIO at Ericsson, fought tooth-and-nail with unit heads and business users to arrive at a consensus, and failed. The business' contention was simple: reduce the number of infrastructure and service outages the enterprise was facing. CIO agreed to bring it down from double-digit outage days a year to three days, but the business wanted the number to go down to a couple of hours, at no extra investment. One day CIO at Ericsson announced that his team was going ahead with a DR implementation. It was a unilateral decision that would have been hard to get away with in good times and the slowdown made it harder.

The project would enable Ericsson to reduce the number of days of outage to three a year. "If we had continued to strive for consensus, the DR project, in whatever form, would not have commenced. I would have bled all my resources to chase zero-day downtime and would have still been apologizing to the business every time an outage occurred," says CIO.

4. **Project**

Nevertheless, IT decided to work around the problem and bring in a business SLA. That's where the problem began," recalls CIO at Ericsson. The business returned with a DR plan that blueprinted zero-day outage. Based on the amount of bandwidth at various locations and the infrastructure and manpower that could be tapped with the current investments, CIO reckoned they could trim service outages for Ericsson India and Sri Lanka to three days. "On paper, we still have the double-digit figure in the contract. But with our soft handling of the issue and vendor management practices, we managed to get our outsourcing vendors to offer us more value-add for our investment. However, we still can't penalize our vendors for not maintaining blackout under three days," he says.

5. **First steps**

It's been three months since IT has put in place all the tools, technologies, processes and practices to keep its promise. And as much as the business didn't like the idea of not having zero-day downtime, CIO's railroading approach opened its eyes to the efforts to curtail blackouts. "Users are beginning to realize the amount of effort required by IT to maintain uptime of critical systems. In fact, the business has slowly begun to release some more funds to invigorate the DR backbone," he says. The business now understands better, says CIO, what they can and cannot get for every investment they make in IT. With business' increasing openness to funding the DR project, CIO plans to plug more loopholes. "If investment comes in, we can gradually take the three-day outage further south," he says.

6. **Benefits**

"The business now realizes that maintaining uptime is not as simple as maintaining a home computer. It requires planning, investment and effort. The DR project that we forced on management has now started a chain reaction. Not only do they realize the extent to which they can make commitments to their clients, but it also opening up funds for more DR and business continuity projects," says CIO of Ericsson.

---

Case Based Questions

---

1. In groups, discuss different situations and or businesses various file organizations which could be used in different business scenarios. Typical examples could be:-
- (a) Student Enrolment System
  - (b) Clothing Manufacturer
  - (c) Automotive Manufacturer
  - (d) Hospital Systems
  - (e) Taxation Systems

Students are expected to prepare a detailed study note regarding the types of file used in different business scenario as discussed above. Also state the reason why the used file organization would be the best one in each case.

2. A University is running various technical courses for students with the help of faculty members working in various departments. To coordinate various programs, University officials would like to implement some kind of methodology where data can be shared and communication may be clear among various entities involved in this process. After detailed study by the officials of the university, they decided to implement E/R diagram which comprises of set of entities and their relationships among these entities. The following types of entity and their types of relationships exist between these entities.
- (a) Student - Program
  - (b) Student - Department
  - (c) Faculty - Department
  - (d) Faculty - Course

Design an E/R diagram that shows the attributes and the types of relationship among these entities.

3. IBM has more than 100000 employees working in different places across the world. The personal information of each employee is to be recorded into company's database called IMS. The database contains following information.
- (a) The company has set of departments
  - (b) Each department has a set of employees, a set of projects, and a set of offices.
  - (c) Each employee has a job history

The database contains following entity and data item as follows.

- (a) Department : Deptno, Deptname, Budget
- (b) Employee : Eno, Projno, Offno, Salary
- (c) Project : Projno, Projname, Budget

Design an appropriate Database and E/R diagram which represents information of the employees working in the company.

4. Hindustan Lever Ltd. is pioneer in producing various consumer durable products ranges from daily used items like soap, toothpaste etc. HLL has various distributors across the country distributing products to the consumers at their doorsteps. The distributors provide the selected range of items to the retailers so that they can sale these items to their customers. The distributor takes the help of sales persons to distribute the items in efficient manner. Each sales person is given a target of Rs 100000 to sale the products to the retailers on monthly basis. HLL maintains the databases of the sales persons who sale the goods to the retailers and accordingly, on completion of the target achieved by sales person, company gives them incentives. To list out eligible sales person who should be given incentives. For this, company use DBMS to filter the queries using structured query languages. The following tables have been designed with data item to find out the values.
- (a) Product : Prodid, Prodname, Prodprice
  - (b) Customer : Custid, Custname, City
  - (c) Salesperson : Spid, Spname, City, Custname, Prodname, Totalsales

The students are expected to perform following exercise in respect to above given case.

- (a) Design and join the tables using SQL statement
- (b) Generate list of product produces by the company.
- (c) Generate the output with Custname residing in Delhi.
- (d) Generate the output with Spname, Prodname, City where Totalsales is more than 100000 in ascending order.
- (e) Generate the output with Spname, Prodname using groups(group by) based on City.

