Unlocking the Corporate Data Resource

BUTLER COX FOUNDATION

March 1986



THE BUTLER COX FOUNDATION REPORT SERIES

UNLOCKING THE CORPORATE DATA RESOURCE

ISSUED MARCH 1986

Research Method

The research for this report was carried out during the second half of 1985 and early 1986, and was led by Paul Thorley, a consultant with Butler Cox in London who specialises in systems design and enduser computing. He was assisted by Rob Moreton, a consultant with Butler Cox who specialises in enduser computing and data management, and David Flint, a senior consultant with Butler Cox.

The response by Foundation members to the original research plan emphasised the interest and importance of this research topic. Many Foundation members suggested that the research should attempt to explain how organisations could improve access to corporate data. As a result of these initial comments and further discussions with members, the direction of the project was set.

Our research had a dual focus. First, we reviewed the organisational and technical infrastructures typically used by Foundation members, and the dataaccess problems that these are causing. We then investigated what needs to be done, both technically and organisationally, to unlock the corporate data resource so it can be more readily available to business users of microcomputers and other desktop computing devices.

The research began with a literature search, and any material that significantly influenced the content of this report is listed in the bibliography at the end of the report.

The views of Foundation members (and other organisations) were gathered during a series of 25 individual interviews conducted throughout Europe and the United States. Some of these discussions are reported in the case histories presented in the appendix of this report. To supplement the experiences of these commercial users, we also interviewed several researchers. These included people working at MIT in the United States and at the Alvey Directorate, a British research organisation sponsored by government and industry.

During September and October 1985, we held a series of focus group discussions in France, the Netherlands, Scandinavia and the United Kingdom. Six meetings were held, attended by 53 Foundation members. Each meeting followed the same format and posed 17 questions related to the research theme. These questions were grouped under the following headings:

- -The growing demand for access to data.
- -Choosing the best approach.
- -Technical infrastructure.
- -Organisational infrastructure.
- Management issues.
- -Future trends.

The mix of user organisations, suppliers and consultants at the focus groups led to some fascinating discussions that influenced the development of the report. Summaries of significant points to emerge from these discussions are included within the body of the report.

Additional report copies

Member organisations usually receive three copies of each report as it is published. Additional copies of this report (or previous reports) may be purchased from Butler Cox. THE BUTLER COX FOUNDATION REPORT SERIES

4

UNLOCKING THE CORPORATE DATA RESOURCE

CONTENTS

SU	MMARY OF RESEARCH FINDINGS	i
1.	ADOPT A NEW APPROACH TO DATA MANAGEMENT	1 2 2 3 4
	Summary	5
2.	ESTABLISH THE TECHNICAL INFRASTRUCTURE. The problems with today's infrastructure Migrate to a two-level system architecture Integrating the systems environments.	6 6 9 10
3.	ESTABLISH A USER-SUPPORT CENTRE. Roles and responsibilities of the user-support centre Establishing a successful user-support centre.	12 13 15
4.	SELECT THE DATA-ACCESS TOOLS AND TECHNIQUES	17 18 19 20 21 24 24
RE	PORT CONCLUSION	25
AP	PENDIX CASE HISTORIES. Avon Products Inc. The Chase Manhattan Corporation ISTAT Philips Rhône-Poulenc Surrey County Council	26 26 28 30 32 33 34
BII	BLIOGRAPHY	36

ang I ag

THE BUTLER COX FOUNDATION REPORT SERIES

UNLOCKING THE CORPORATE DATA RESOURCE

SUMMARY OF RESEARCH FINDINGS

In many large organisations, most of the routine transaction-based systems have now been installed. Most of the new computer applications being developed today fall into the category of enquiry and analytical systems, rather than the traditional bookkeeping and transaction-recording systems. From now on, the trend will be to install decisionsupport systems that will be used as an integral part of the corporate decision-making process. Many of these systems will be run on business microcomputers or workstations linked to an information centre. The main requirement for many of these new systems will be for access to corporate data held in the organisation's conventional mainstream information systems.

Another reason for the increasing demand for access to corporate data is the growing awareness of the role that data and information systems can play in providing the organisation with a competitive advantage. We know of several companies that are taking positive steps to exploit the data collected automatically as part of their mainstream operations. One of the best-known examples is the McKesson Corporation in the United States. McKesson is a drug-supply company, and in the process of supplying drugs to pharmacists, it compiled a valuable database about the usage of drugs. Using this database, McKesson has built a new business that supplies market research information to drug manufacturers.

Corporate data held in the central databases has been accumulated over many years by the mainstream operational systems (order processing, purchasing, management accounts, etc.), and is usually not structured in a way that makes it easy for business users with a microcomputer or workstation to get at what they need. Managers and professionals require facilities that allow them to access only the data that is specific to a particular problem. At present, this often cannot be achieved, because the required data is 'hidden' (or implicit) in a mass of irrelevant data. Moreover, once the data has been retrieved, these users require tools to manipulate the data until it is of direct use by a specific person for a particular problem. The requirements for making corporate data available to users can be summarised as follows:

- Users need to know what data resides where, and what its meaning is.
- They require easy access to the relevant data, no matter where it is stored.
- They require powerful (but easy-to-use) tools with which to manipulate the data.

The techniques for providing users with access to corporate data are therefore very much concerned with end-user (or personal) computing. Improvements in the capabilities of desktop computing devices, and pressure by the suppliers, have raised the general level of awareness of the potential benefits of enduser computing. As a consequence, there is a rapidly growing population of users who understand the benefits that personal computing can provide for decision making and productivity improvement.

The four main technological developments that have led to a greater awareness by business users of the potential benefits of personal computing are:

- Improved telecommunications facilities, particularly for linking microcomputers to mainframe systems for access to corporate data, and for interlinking facilities within a local site.
- -Easier access to external information services.
- -Wider availability of 'user-friendly' software.
- Decreasing hardware costs that have made more feasible the use of facilities that require considerable computing power (query languages and relational database management systems in particular).

Each of these developments enhances the benefits that personal computing can bring to a manager or professional. As users become familiar with the basic technology and begin to treat the facilities as an ordinary part of their everyday business lives, their attention will turn increasingly to the data they can access with the devices on their desks.

Unfortunately, there is no universal technical solution to the problem of providing these users with easy

SUMMARY OF RESEARCH FINDINGS

access to the data locked up in corporate systems. Many companies have attempted to solve the problem by establishing some form of user-support (or information) centre but, sometimes, these have not been very successful. Others have supplemented their traditional mainstream systems with ad hoc solutions based on microcomputers, departmental minicomputers, communication devices, etc. As a result, the approach to, and responsibility for, user access to data is often fragmented and uncoordinated.

A few organisations told us that they planned to continue channelling all requests for access to corporate data through the central systems function. Their approach was to provide analysts and programmers with the necessary tools so they could write bespoke systems or produce ad hoc reports for business users. We believe that, although this approach may be acceptable in the short term, the pressure from users for direct access to the corporate files will become irresistible. Our view is that, within the next year or so, these organisations will have to abandon their centralised approach to providing users with access to corporate data.

One way to create the optimum technical environment for facilitating user access to data would be to rebuild the existing applications. However, this course of action is unrealistic for most organisations. We therefore assumed that users would require access to data held in existing systems.

Most information systems managers have recognised that the growing demand for access to corporate data will present severe problems because existing systems and procedures are not geared to meeting ad hoc requests from a large population of users. These systems managers are seeking advice on the best way of meeting the increasing demand whilst retaining control over that part of the corporate data entrusted to their care.

We conclude that organisations need to take a fresh look at their whole approach to providing business users with access to corporate data. This report sets out a methodology for identifying and resolving the organisational and technical issues concerned with unlocking the corporate data resource. By using this methodology, we believe that Foundation members can create organisational and technical environments that will enable the data locked up in corporate systems to be accessed by business users for the benefit of the whole organisation. The methodology has four steps:

- Adopt a new approach to data management that makes a clear distinction between 'authenticated' corporate data and 'unauthenticated' local data.
- Plan to migrate to a two-level technical infrastructure, with a core systems environment for the mainstream computing activities, integrated with a series of subsidiary systems environments, each of which is optimised for a particular type of systems requirement.
- Establish a user-support centre, paying particular attention to the initial applications and to the way in which the centre charges for its services.
- Select the data-access tools and techniques (directories, database products, data-retrieval products, microcomputer-mainframe links, systemlinkage products and user-interface techniques) that will be used to provide business users with direct access to the data held in corporate, and external, databases.

The report describes each of these steps in detail.

CHAPTER 1

ADOPT A NEW APPROACH TO DATA MANAGEMENT

The amount of data being stored in computer systems is growing significantly, and this data is being used by more and more business users for a greater variety of purposes. In particular, senior managers are making increasing demands for easier access to existing computer-based data. The data required by an individual user may well reside in several different systems, each of which has been developed for a specific systems environment (that is, for a specific combination of hardware and software). This situation will not change in the foreseeable future, although, as we show in Chapter 2, the overall technical infrastructure can be rationalised by coordinating the different systems environments.

The increasing demand for access to data highlights the need to make data accessible, accurate, consistent, and secure — in other words, the need to manage data. Many organisations have recognised that data management is an essential part of their overall information systems policy. (The views expressed at our focus group meetings on the importance of data management are summarised in Figure 1.1.)

Figure 1.1 Focus group views on the importance of data management

"There is a critical need to clarify who controls and owns data."

"There is a need to ensure that common data is used effectively for a whole range of applications."

"Data must be recognised as one of the most important resources of an organisation – although the nature of the business determines the degree of importance."

"Data management should begin with a data definition exercise."

"The rapid increase in the number of microcomputers and the growth of distributed processing have increased the need for data management."

"The responsibilities of the information centre should be extended to include aspects of data management."

During our research we found many examples of the difficulties that can occur if data is not managed properly. One company told us, for example, that data relating to its Norwegian operations appeared in its European and overseas analyses, which led to inaccurate information about the company's overall trading position being presented to management. In another case, inconsistent data about the discount structures of a grocery chain's major suppliers led to incorrect calculations of gross profit margins. Since buying decisions were based on these calculations, the data inaccuracies could have caused considerable damage.

The usual approach to data management has been to establish a centralised data administration function, regardless of the policies for centralising or decentralising company resources and operational database systems. For example:

- Chase Manhattan Bank has invested heavily in setting up clearly defined data-service roles at the centre and in its regions. The central support function establishes data access and communication routines, and the regions are responsible for their own data security and data administration.
- Rolls-Royce Limited has established a central data administration team that is responsible for corporate data modelling. The development of all 'core' systems (that is, systems that are central to Rolls-Royce's business) is validated by this team.

Some organisations have moved away from a centralised approach to data management, however. A large oil company has recently discarded its central database administration role. This move was justified on the grounds that systems development was being slowed down. In our view, the problem arose because a single person had to vet all new applications for data consistency. We believe that providing additional resources for the central function would probably have removed the bottleneck that was being created.

Other organisations have been reluctant to invest in a central data management function because of the difficulty of measuring the success of such a function. The lack of measurement criteria has meant they have not been able to justify the expenditure required to set up a central function.

In general, our research confirmed that there is a need for a central, coordinating data management role. However, the traditional approach to data management may not be able to cope with some of the new demands being made by business users for access to corporate data. Conventional data management techniques are based on the concept of 'accurate' data, where a specific element has a specific unambiguous meaning and a single 'correct' value. But this concept begins to break down when data that was collected for one purpose is used for a different purpose. For example, data that is 100 per cent accurate in a transaction-recording sense may need to be modified before it can be used for marketing purposes. (Marketing staff will always want to present the organisation and its products in the most favourable light.) We believe that the approach to data management needs to be modified to take account of this type of situation.

Providing business users with access to corporate data introduces yet another complication for conventional data management. User demands for access to data will increase as the personal use of business microcomputers passes through the four stages of initiation, contagion (growth), control, and maturity. (We discussed the characteristics of each of these stages in Foundation Report No. 43 - Managing the Microcomputer in Business.) Our research suggests that, from the users' point of view, data management is not important during the initiation stage. During this stage, the bulk of the data is input by the users themselves. This procedure does not cause any great difficulty because most Stage 1 business microcomputer systems are standalone applications, and they address specific tasks.

During the second (growth) stage of business computing, users begin to want to access centralised operational data and to share applications and computing resources. But, because data is processed in different ways in a variety of different environments, problems associated with coordinating the use of data; and ensuring that it is accurate, consistent, and secure, come to the fore. At the same time, the advantages of capturing a data element only once become increasingly apparent. As the growth of business computing continues through the control stage and, finally, the mature stage, the importance of data management will become clearer to users. Moreover, the complexity of carrying out data management will increase in parallel with users' perceptions of the importance of data management.

Although the importance of data management has long been recognised by systems staff, most users are only dimly aware of the need for, and the advantages of, data management. For many business users, data management is, at present, seen as an unnecessary encumbrance that hinders their progress with personal computing.

Thus, any new approach to data management must also be able to cope with users' changing perceptions of the importance of data management. As a consequence, we believe that a two-level approach is now required to encompass both central and local aspects of data management. In order to adopt this approach, it is necessary first to recognise the importance of data, and then to distinguish between authenticated and unauthenticated data. Next, the appropriate degree of decentralisation for data management must be determined, and central and local data management responsibilities must be clearly assigned.

RECOGNISE THE IMPORTANCE OF DATA

Successful data management requires recognition at a high level in the organisation that data is an important asset and that it should be managed accordingly. Data should therefore be managed in a way that is consistent with the organisation's structure, style, and objectives. Only in this way will data management be able to support future information needs and allow more effective use to be made of business information.

Different organisational structures have different implications for data management. A strategic business unit approach that encourages autonomy will, for example, give rise to data management requirements different from those produced by a functional organisation structure. The data management approach adopted should also reflect the corporate 'style' (or philosophy). In particular, it should take account of the organisation 's technical maturity and whether the organisation operates in a constant or changing environment.

If data is perceived as a strategic asset, its management must take account of future information requirements, and these may be markedly different from today's needs. The data management approach must therefore be aligned closely with the system and information planning processes. The approach must also recognise that the strategic use of business data includes the need to access data about competitors, and data about industry-specific factors (regulations, for example).

DISTINGUISH BETWEEN AUTHENTICATED AND UNAUTHENTICATED DATA

The conventional approach to data management assumes that all data is 'authenticated' — that is, controlled and validated by a central data administration function and made available via a coordinated database. Certainly, the need for authenticated data processed within the organisation's data management guidelines will continue and increase. (The data associated with the mainstream systems — order processing, payroll, etc. — will always be authenticated data.) But we believe that the data management approach should also recognise that users will want to use and have access to unauthenticated data as well. Some of the data used by local systems may be input by the users or may be extracts from the authenticated data that are manipulated and interpreted locally. This type of local data should be classified as unauthenticated data.

Authenticated data will always require the involvement of the systems department, either to verify the accuracy and consistency of the data, or to verify the processes by which it is manipulated. However, in some circumstances, business users may deliberately modify (or even corrupt) authorised data for perfectly valid commercial reasons. For example, we know of an airline that chose to show in its public timetable a flight as operating on every day of the week, whereas in the operational database, the 'accurate' data defined two flights, each operating on some days of the week. Both had the same flight number, but on some days the service operated with a smaller aircraft and slightly different timings. However, the airline's marketing department preferred to show in its public timetable a daily flight operating with the larger aircraft. Thus for competitive or marketing purposes, the 'accurate' data from the users' point of view may be somewhat different from the accurate data held in the corporate (authenticated) database.

Moreover, users often do not require the absolutely correct and up-to-date version of the authenticated data. Instead, they require for analytical purposes an accurate version of the data as it stood at a particular point in time — month end, year end, etc.

Thus, business users of microcomputers or workstations require access both to authenticated data and to their own unauthenticated data. The unauthenticated data may have been created by extracting data from the corporate (authenticated) files and modifying it to suit a particular purpose, or it may be 'personal' data created by the user, or it may have been extracted from an external data-supply source. The question arises as to whether unauthenticated data should be brought under the control of the corporate data management function.

Provided that a particular set of unauthenticated data is used only by its originator in standalone applications, there is no need for the corporate data management function to be involved in managing the unauthenticated data. But as soon as unauthenticated data begins to be used by other users, both the originators of the data and its potential users need to be made aware of their obligations and responsibilities.

Some organisations told us that they will prevent this problem by prohibiting the use of unauthenticated data by anyone other than its originator. We believe that this is not a practical approach. Business users will want to use whatever data they feel is necessary to support the particular task in hand — as they have always done when preparing reports or plans. If one of their colleagues has some relevant data available in a machine-readable form, it is inevitable that they will want to use it.

Instead of trying to prohibit the use of unauthenticated data, we believe that the corporate data management function should adopt a 'soft control' approach that encourages business users to take personal responsibility for the unauthenticated data they create and use. Specifically, the data management function should identify those sets of unauthenticated data that are being used (or are likely to be used) by people other than their originators. Such data sets should then be included in the directory of data that is available, but they should be clearly identified as unauthenticated.

Thus, the data that is made available to business users should be clearly marked as being authenticated or unauthenticated, and users should be made aware of the differences between the two types of data. Whenever a business user accesses an unauthenticated set of data, he or she should be reminded that this data does not have the corporate 'seal of approval', and that reference should be made to the data originator for clarification about the accuracy, consistency, and meaning of the data.

Creators of unauthenticated data that is likely to be used by others should be encouraged to adopt good data management practices. The central data management function should not attempt to impose corporate data management practices, however. Instead, it should provide advice and should encourage business users to adopt a basic level of data management that is consistent with their own business requirements and with those of other business users who may use the data.

In many organisations, the best place from which to exercise this soft control of unauthenticated data will be the information centre.

DETERMINE THE MOST APPROPRIATE DEGREE OF DECENTRALISATION

The variety of different systems environments found in most large organisations makes it impractical to think in terms of a single central data resource that is available to all users throughout the organisation.

CHAPTER 1 ADOPT A NEW APPROACH TO DATA MANAGEMENT

For example, personal computer users may have their own systems based on new generations of languages that require bespoke database structures. This type of constraint argues in favour of storing at least some of the data locally. In turn, local data storage also raises questions about whether it is possible (or desirable) to decentralise some of the data management responsibilities. A degree of decentralised responsibility for data management may be appropriate in those organisations where some data processing activities have already been decentralised, or where a centralised data processing structure is considering distributing some of its systems.

In considering whether to decentralise some of their corporate data, most organisations have focused on operational and technical issues such as the relative costs of data storage and data transmission. Although technical concerns are clearly important, we believe they are subsidiary to wider organisational issues, the most important of which are:

- Would decentralised data management responsibilities be compatible with the existing corporate structure and style of management?
- Would a decentralised data management structure be consistent with the organisation's overall policy for managing the corporate information resource?
- Would decentralised data management create insurmountable problems for data consistency, security, and privacy?

Only after these issues have been considered should attention be given to the practical operational aspects of a decentralised data management approach. At this stage it may be necessary to modify the ideal level of decentralisation in the light of technical constraints such as:

- The inability of most current database software to maintain distributed databases.
- The effect that the likely patterns of use will have on storage and communications requirements.
- The extent to which distributed data processing and end-user computing are already used within the organisation.

One approach to decentralising data is to replicate the database at several sites. Creating multiple copies of the data simplifies the communications requirements (and reduces communications costs), increases the availability of the data, and reduces response times. Nevertheless, there are still several technical problems that affect an organisation's ability to decentralise data. In particular, most database software products cannot yet adequately handle distributed databases, and this can create problems, particularly if data consistency is of paramount concern and if the database update rate is high. Some of the factors that should be considered in deciding whether and how to distribute a database are:

- The frequency with which the database is updated, the complexity of the updates, and the way in which updates are initiated.
- —The use to which the database will be put. The number of users, their location, and the frequency with which they will use the database will all help to determine which parts of the database could be replicated or distributed.
- The experience that staff already have with operating and using distributed computing systems. Computer networking systems can be complex to design, and they require complex software. Specialist local staff may be required if these networks are to operate successfully.
- —The costs of local storage compared with transmission costs. Careful analysis will be required to find the optimum balance between transmission costs and storage costs. The most significant parameter is the amount of data that will need to be transferred.

ASSIGN DATA MANAGEMENT RESPONSIBILITIES

In assigning the responsibilities for data management, a clear distinction needs to be made between the data administration and database administration roles. Data administration is concerned with the 'ownership' of data, its meaning, its relationships with other data, its integrity, and so forth; database administration is concerned both with the way in which the abstract data model is physically implemented in a computer system and with controlling the procedures for accessing the data. (Foundation Report No. 32 — Data Management — discussed in detail the management responsibilities of data administration and database administration.)

In this report we are advocating that some data management responsibilities should be devolved to local sites, particularly responsibility for the management of unauthenticated data. We believe, therefore, that management responsibilities can be divided into a central coordinating role and local responsibilities.

Central coordination

The main role of the central data management function is to plan and coordinate the organisation's total data resource and identify its overall data-access requirements. Usually, the requirements are represented in the form of data models, whose purpose is to ensure that the overall data needs are recognised and integrated with the information systems environment. The central role also encompasses the definition of data and the keeping of records about its location, the systems that maintain and use the data, and the users who maintain and rely on the data. The overall aim is to ensure that the data is used effectively throughout the organisation.

Such a coordinating function must, of necessity, be centralised, and it must be given the appropriate authority to enable it to carry out its responsibilities. Figure 1.2 lists the main responsibilities of a centralised data management function.

A centralised data management function also has important responsibilities for data security, particularly in preventing fraud and ensuring the confidentiality of data. Nevertheless, where some of the data management responsibilities are delegated to local sites, some aspects of data security will need to be managed locally.

Local responsibilities

In most large organisations there is a strong trend towards distributing as much computing power as possible throughout the organisation. This trend is evident in the growing use of departmental minicomputers, information centres, local area networks, and personal computers. The result is that users will increasingly want to take more direct responsibility for planning the information systems and the data they need. Indeed, local planning is essential if users are to make the most effective use of their local systems environment.

The local responsibilities for data management must be well defined, and their relationship with the central coordination role must be clear. We believe that the most crucial issue is to draw a clear distinction between the responsibilities for authenticated and unauthenticated data. The responsibility for authenticated data clearly rests with the central data management function, although the central function may choose to delegate some of its responsibilities to the local data management function. The responsibility for controlling the use of unauthenticated data will rest with the local function. The interface between local and central responsibilities must be clearly defined because unauthenticated data will often be created by extracting and manipulating data from the authenticated databases.

In many respects, local data management responsibilities mirror those carried out centrally. The local responsibilities include:

- Ensuring that local data management practices conform with the central policies for data access, security, recovery, and privacy, and applying these in the light of local requirements.
- Defining and maintaining local authenticated data, and understanding its importance.

Figure 1.2 Responsibilities of a centralised data management function

- Monitoring the overall data needs.
- Building and maintaining an overall model of the data needs.
- Supporting overall planning of the systems environment.
- Approving systems development plans with regard to data needs.
- Ensuring that data management roles and responsibilities are identified.
- Specifying approval criteria for authenticated data.
- Liaising with divisional and regional data management functions.
- Approving suitable methods, tools and techniques.
- Ensuring that data analysis follows set standards.
- Providing expertise as appropriate to support local data analysis exercises.
- Ensuring that data definitions are standardised, accurate, appropriate and properly documented.

- Carrying out data audits.

- -Resolving disputes about the ownership of data.
- -Selecting and introducing tools and techniques.
- Providing advice about the impact that changes in business requirements could have on data management.
- Encouraging users who create unauthenticated data to adopt good data management practices, and making them aware of their responsibilities to other users of their data.
- Ensuring that users of unauthenticated data are aware of the limitations of that data, and providing a focal point for resolving disputes about the meaning of unauthenticated data.

SUMMARY

In an environment where more and more users will need access to data, we believe that the key to successful data management is a clear distinction between authenticated and unauthenticated data. In order to achieve this, it is necessary to begin with an understanding of the importance of data. The appropriate degree of decentralisation should then be established, with central and local management responsibilities clearly defined.

By assigning responsibility for the various data management activities to the appropriate level, we believe that organisations can choose the most appropriate approach for providing users with access to corporate data.

CHAPTER 2

ESTABLISH THE TECHNICAL INFRASTRUCTURE

Most organisations believe that their existing technical systems infrastructure is inadequate to meet the changing demands for access to data. In the past, systems were not designed to meet the unstructured ad hoc requests for data that users are now demanding. To meet this demand, many organisations have adopted ad hoc solutions in an attempt to solve immediate data-access problems.

Sometimes, these solutions have been implemented on mainframes and sometimes on microcomputers. Often, the solutions adopted for microcomputers have not been compatible with those adopted for mainframes, and it is proving to be very difficult to interlink the two types of system. However, users will not tolerate the limitations imposed by the complexity of the linkages that are necessary today.

A representative selection of comments made by the focus group participants on the inadequacy of today's technical infrastructure is shown in Figure 2.1. These show clearly that there is an urgent need for a new technical infrastructure that can cope both with the existing applications portfolio and with the changing

Figure 2.1 The inadequacy of today's technical infrastructure as identified by focus group participants

"There is a lack of a single interface for users."

"Today's database management systems are not suitable for unstructured requests."

"The current systems environments often stifle creativity for end-user computing."

"Data design (especially data definition) and data management are not given the attention that they merit."

"There is a disproportionate investment in maintenance compared with the benefits."

"The introduction of new end-user computing facilities are usually not planned with the right attention to detail."

"Today's user demand for access to a variety of databases requires greater controls than currently exist."

requirements for access to data. Before discussing the ways in which this new infrastructure could be established, we first describe briefly the problems created by the technical infrastructure typically found today in Foundation member organisations.

THE PROBLEMS WITH TODAY'S INFRASTRUCTURE

The technical infrastructure comprises the various systems environments found within an organisation. (A systems environment is a specific combination of hardware architecture, operating system, file or database management system, teleprocessing monitor, and development process.) Today, most Foundation members have a variety of systems environments. These environments have evolved to meet the differing requirements of particular types of applications — order processing, office systems, videotex, computer-aided design and manufacturing, and so forth. In order to make it easier for users to get at the data they require, problems will have to be overcome in the hardware, software and telecommunications components of the technical infrastructure. Data-access problems are also created by the variety of systems environments found in many organisations.

Hardware

One of the driving forces behind the increasing demand for access to data is the large number of microcomputers now being installed in most organisations. Spending a few thousand dollars on a standalone microcomputer seems at first sight to be a low-risk investment. But, as we predicted in Foundation Report No. 43 — Managing the Microcomputer in Business — users of standalone machines begin to want access to corporate data after a year or so, and this demand can cause difficulties if microcomputers have been installed in an uncoordinated way.

The experience of one organisation we talked with during the research, although extreme, is not untypical. Originally, there was no formal policy for the acquisition of microcomputers, and users were allowed to purchase whatever equipment they wanted. The rationale for allowing this to happen was that the individual needs were so diverse that they could not (and should not) be coordinated by a central group. But within two years there was an enormous demand both for multi-user systems and for the capability to transfer data from the central hardware to the 'standalone' microcomputers for local processing. As a consequence, this organisation established centres of microcomputing expertise, which set out clear, consistent guidelines for operating systems, preferred suppliers, and so forth.

When we asked why this firm policy had not been established at the outset, we received a philosophical reply. We were told that a 'learn-the-hard-way' approach would be more effective in the long term and would be an effective means of enhancing the standing of the central systems function.

We believe that the approach to the acquisition of microcomputers should be designed to resolve the apparent conflict between microcomputers and mainframes. Microcomputers provide users with the benefits of independence; mainframes provide the benefits of integration, which conflict directly with the microcomputer's advantage of independence. The hardware acquisition policy should allow the benefits of both to coexist.

However, there are other important developments in the accelerating trend towards desktop computing. Departmental computers serving a local community are also being installed in large numbers, for example. As a result, the day is fast approaching when the majority of office workers in most organisations will have a microcomputer or some other form of 'workstation' on their desks. These workstations will take a variety of forms — dumb terminals, word processors, videotex terminals, graphics devices, and so forth. The policy for acquiring all these types of local computing facilities should follow the same approach used for microcomputers.

The policy should also take account of the likely future demands for access to data implied by these developments in desktop computing.

Software

In the context of user access to data, two types of software are important: database software and software designed for direct use by business users.

Database software

Many of the people we met during the research expressed doubts about the ability of existing database software to meet future data-access requirements. Typical comments were that existing database software was too complicated, too inflexible, and not efficient enough to handle future requirements. These requirements are likely to be for a mix of facilities that can handle the routine requirements of the 'core' administrative systems (sales order processing, invoicing, etc.) and that can allow easy ad hoc access to the data collected by the core systems. The database software will need to provide browsing facilities and must be able to handle complex informationretrieval requests based on compound Boolean operators, and other conventional retrieval techniques.

Another disadvantage of current database systems is that they require both a high level of expertise and much time and effort to establish a coherent database structure. Furthermore, new data elements or additional indexing items may well require the structure of the database to be changed.

Existing databases are also difficult to use because they require the user to interact with them via a formalised, structured language. Although some progress is being made in the area of natural language interfaces by products such as Intellect, most database systems require the use of a highly structured syntax.

The inefficiency of much database software arises from the way in which complex enquiries are handled, and the inability of the software to pre-analyse the hardware resource requirements of a particular request.

Direct-access software

Software designed to provide business users with direct access to data encompasses the whole range of decision-support tools, application generators, report generators, query languages, and so forth. It also includes the software interfaces to external information retrieval systems, and specialised retrieval software such as Stairs. The distinguishing feature of this type of software is that it is designed for direct use by business users, either to create a bespoke application or to formulate a bespoke request for data.

Above all, this type of business software must be easy to use, and it must allow useful results to be achieved quickly. Typical examples of successful business software are the spreadsheet packages that are being bought (and used) in large numbers, and are enormously popular with users. Products such as Lotus 1-2-3 use familiar concepts (the conventional spreadsheet) and do not require users to master a new approach to problem solving and decision making. Regrettably, much of the software designed for direct use by business users lacks this essential characteristic and, as a result, is unacceptable to the majority of business users. We also found that the use of external information services was being inhibited for similar reasons. Although the use of such services is growing steadily, we found that users were reluctant to use them because:

- -The interfaces were too complex.
- The meanings of keywords and dialogue terms were obscure.
- The cost of using the services was perceived as being too high.

Telecommunications

Telecommunications facilities are fundamental to providing users with access to corporate data. At present, however, many microcomputers are not linked either to other microcomputers or to the mainstream systems, and many organisations are currently grappling with the issue of how best to link 'standalone' microcomputers to the corporate mainframes. (We describe in Chapter 4 the various techniques that can be used for microcomputermainframe links.)

In many organisations, implementing microcomputermainframe links is made more difficult by the incompatible software at each end of the link. We heard of one extreme example where data is extracted from a Total database (running under the IBM MVS environment) and is passed as a serial file to another machine running under VM/CMS. The data is then either manipulated for the user by a systems analyst using APL, or it is loaded into a mainframe spreadsheet package ready for manipulation by users. Sometimes, the data is further downloaded to an IBM PC for use with Lotus 1-2-3.

Problems created by a variety of systems environments

The need to support several systems environments makes it difficult to create a technical infrastructure that can cope with the demands imposed by user access to data. Ideally, the technical infrastructure should provide a unified interface for users that will enable them to access, from a single workstation, the growing portfolio of applications available within each of the systems environments. Maintaining several systems environments makes it difficult to integrate systems, impedes the provision of integrated user interfaces, and generally obstructs the transfer of information between systems and between users.

Should the technical infrastructure be based on a single systems environment? Regrettably, the answer has to be no, because the requirements, both of users and of different types of systems, are so diverse that no one systems environment could cost-effectively meet all of the requirements. We pursued the

question of the optimum number and nature of systems environments with Foundation members at the focus group meetings. This question proved to be the most difficult of all the issues discussed. A representative sample of the views expressed by the focus groups is given in Figure 2.2.

We are convinced that the key to creating the most appropriate technical infrastructure lies in determining the optimum number of systems environments, and this means that a planned approach to the systems environments that will be used throughout the organisation is essential. The need for coordination becomes paramount as computer networks spread throughout the organisation and as distributed systems, databases, office systems, etc. become more common. If these developments are allowed to take place in a haphazard, unplanned way, they will (in the words of James Martin) "grow like weeds and take over the garden". And, like weeds, incompatible systems are extremely difficult to get rid of once they have become established.

Many of today's problems with providing users with access to corporate data arise because enhancements have been 'bolted on' to existing systems in an ad hoc manner as new facilities have been required. However, the benefits of a planned approach have been recognised for many years, and during our research we met with organisations that were deliberately moving to a greater degree of systems integration both at the corporate level and at the divisional or operating-company level.

Figure 2.2 Focus group views on future systems environments

"There is a need to encompass all requirements in one framework — but beware of compromising the individual components."

"The chosen approach must be practical, and must match a given manufacturer's environment."

"A two-tier system, composed of an integrated core and enquirybased systems, is required."

"There is a need for a strategic approach that emphasises improved access and the provision of data, and that provides rules and guidelines for achieving this."

"The best way forward is to evolve the operational systems so they make provision for future likely demands for decisionsupport-type systems."

"There is a need for a single user interface, irrespective of the technical infrastructure."

There is still	a need for	a central	database.	with the	emphasis
being on cor	itrols and c	data integ	grity."		

Some very large multinational companies are planning a much greater level of systems integration throughout their operations. They see this as a necessary response to changing circumstances, caused partly by increasing competitive pressures and partly by a need to respond better to the market. Those organisations that have opted for an integrated approach at the divisional or operating-company level see it as the optimum method of matching the information systems architecture with the needs of the business.

MIGRATE TO A TWO-LEVEL SYSTEM ARCHITECTURE

We believe that the most appropriate system architecture for the future comprises two levels: a 'core' systems environment that supports the conventional transaction-recording systems and management reporting systems, and a series of subsidiary environments each supporting a specialised systems requirement. (These specialised systems will include office systems, microcomputer-based businesssupport systems, and videotex systems). However, the overall architecture should be designed in a coherent way, because there will be an increasing need for communication between the levels and between the systems environments supporting the specialised requirements. In a conglomerate organisation comprising a head office and several disparate businesses, the head office and each business unit could have its own separate core and subsidiary systems environments.

The core systems environment

The core systems environment will support the mainstream transaction-recording and management-reporting systems. It also needs to provide the means for interlinking the subsidiary systems environments and, in so doing, it will act as a central data manager and as the manager of shared resources.

At present, the core systems environment will often be based on a conventional database management system (that is, a system based on a hierarchical or network structure) geared specifically to the needs of transaction recording and managerial reporting. The types of data-access requests that will be received from the subsidiary systems environments can cause great difficulties for these conventional database management systems. In particular, the data required for decision-support and analytical purposes can create large and unpredictable processing loads that degrade the system's performance for regular online users.

From now on, most of the growth in the systems portfolio will be in the area of decision-support and analysis systems. This leads us to conclude that the core systems environment should be chosen primarily to meet the emerging requirements for data retrieval and manipulation. To do this, it will be necessary to integrate local microcomputers with the core systems environment, and it may be necessary to move to a relational database management system.

The benefits of using a single systems environment for the core applications (transaction processing and management reporting) include:

- Fewer terminal-compatibility problems.
- Easier control of central data, which means it is easier to provide consistent facilities for accessing and manipulating the data.
- -Better integration of decision-support facilities.
- Improved facilities for transferring data (and text) between central systems and users.
- Standardised dialogues between the users and the various systems.

Subsidiary systems environment

There will also be a need for facilities that cannot easily be satisfied by the core systems environment. These include:

- -Word processing facilities.
- Local processing facilities with sufficient computing power to provide fast response times.
- Software facilities tailored to the needs of individual users.
- -Voice recognition facilities.

These facilities will be provided by subsidiary systems environments, the most common of which are office systems, business microcomputer systems, operational support systems, and videotex systems.

Office systems

Many of the present generation of office system products were designed primarily to meet word processing requirements. As a result, they are based on local processing power, usually located in the workstation itself. However, the trend is towards much greater integration between office systems and the core systems. As a consequence, office systems will need to be able to operate both in their own specialised systems environment and as part of the core systems environment. In particular, workstation users will want to access large data files and databases held within the core systems environment.

The office systems environment will be required to provide specialised facilities that it would be inappropriate to provide via the core systems environment. Each of these facilities requires considerably more processing power in the workstation than is provided by current systems. Their availability in commercial products will require future generations of business microcomputers to have substantially more processing power than presentday PCs. However, the continuing developments in microelectronic technology mean that lack of processing power will not be a constraint.

Business microcomputer systems

At present, most business microcomputer systems are used mainly for decision-support purposes. One approach to providing microcomputer-based decision-support tools has been to create a microcomputer version of an equivalent mainframe package. The disadvantage of this approach is that the microcomputer package merely replicates the mainframe version and fails to take advantage of the special features of a microcomputer. In particular, this type of package usually does not provide the highly interactive user interfaces available with microcomputers. The most popular microcomputerbased decision-support tool so far has been the ubiquitous spreadsheet specially designed for use on microcomputers.

However, microcomputer users are increasingly aware of the advantages of accessing data stored on corporate files and manipulating it with more sophisticated decision-support tools. To access the core systems environment, it is necessary to provide some form of communication link between the microcomputers and the mainframe. (The tools and techniques that can be used for achieving such a link are described in Chapter 4.) Microcomputer-mainframe links, as currently implemented, are often cumbersome to use. For example, a microcomputer user wishing to access data held in a mainframe and transfer it to his microcomputer may well have to go through the following procedure:

- -Initiate the terminal emulation facility.
- Use the mainframe command language to retrieve the appropriate data.
- Initiate a mainframe program to convert the data to a new format.
- -Transfer the retrieved file in the correct format.
- -Close down the emulation facility.
- -Initiate the appropriate microcomputer software.
- Instruct the microcomputer software to process the transferred file.

There is a clear need for greater integration between local microcomputers and the core systems environment, but the local systems environment must still be able to meet the specialised local processing needs.

Operational support systems

Besides office systems, other types of operational support systems will need to interface with the core systems environment in order to gain access to corporate data. Such systems include factory automation systems (including flexible manufacturing systems and computer-aided engineering facilities), systems used in a laboratory environment, designoffice systems and executive support systems. Each of these will require its own specialised systems environment, but they will also need to be able to interface with each other and with the core systems environment.

Videotex systems

Videotex can be an effective means for providing users with access to a large volume of information. We believe that, in the future, the best approach will be to integrate videotex with the office or business microcomputer systems environments. Thus, users would not need to know whether a particular menu item invoked a videotex access or a bespoke dataaccess system.

INTEGRATING THE SYSTEMS ENVIRONMENTS

Several software suppliers have introduced relational database products that are designed to provide a flexible response to unstructured requests for data. Adabas from Software AG and IDMS-R from Cullinet Software Inc. are two typical products, and IBM has introduced DB2 as part of a dual database strategy for combining (for decision-support purposes) its hierarchical IMS product for core-system processing with a relational database. These products go some way towards meeting the needs both of core systems and end-user computing. Nevertheless, we believe that true integration between the core and subsidiary environments requires a new approach to database software and overall systems architectures.

Figure 2.3 illustrates the kind of overall architecture we envisage. Essentially, it has three components: a mainframe database management system, a microcomputer-based database management system, and a permanent link that joins the two.

When a microcomputer application makes a call to the local database management system, one of two things can happen. Either the data is present in the local database and is retrieved for the user to process, or the request is passed on to the mainframe database management system. In the latter case, the mainframe data directory is checked to see if the required data is present in the mainframe database. If it is, the appropriate calls to the mainframe system are generated automatically, and the data is retrieved and passed back through the microcomputer database management system to the application. Thus,



Figure 2.3 Overall architecture for an integrated two-level systems environment

the mainframe database management system acts as an extension of the microcomputer's own database management system.

The significance of this type of architecture is that the user has a unified view of data, regardless of whether it is held locally or in a remote system. In addition, the same architecture can be used to provide access to the core systems by a facility in the microcomputer that bypasses the local database. All the facilities would be accessed by the user via a common interface manager, which would provide consistent, user-friendly interfaces.

At the present time, it may be possible to achieve a reasonable level of integration between specific systems environments, but this will usually require the development of bespoke software, or will require the choice of equipment and software to be restricted.

It will be some time before products are available that will allow any subsidiary systems environment to be integrated fully with any core systems environment. Before this can happen, it will be necessary for several significant developments to take place:

- More interface manager software will need to become commercially available.
- Microcomputer applications software will need to be written, or rewritten, so that it can operate both with the microcomputer database management system and with the central database system.
- Mainframe database management systems and query languages will have to be modified to allow them to transfer existing files to the subsidiary systems environments.
- Powerful workstations, with significantly more processing power than those currently available, will be required.

Most of these developments depend on the definition of acceptable international standards for interconnecting different types of computer systems. For the foreseeable future, the development, and adoption, of independent standards will lag behind the 'de facto' standards being promoted by individual suppliers. Thus, although a fully integrated architecture of the type discussed above may not be achievable in the near future, it is essential that this ultimate goal be kept in mind when decisions are made about the type of technical infrastructure that is required in order to provide users with access to corporate data.

Once the appropriate technical infrastructure has been established, it is then necessary to provide the tools and techniques that will be used to access the data held in corporate databases. Many organisations now have some form of information centre (or usersupport centre) and believe that this is the most appropriate means of helping users to get at the data they require. This is a view that we endorse, and in the next chapter we provide advice on how best to set about establishing a user-support centre. In Chapter 4 we review the main types of data-access hardware and software products that are available today and will become available in the near future.

CHAPTER 3

ESTABLISH A USER-SUPPORT CENTRE

The increasing demand for access to data, and the growth of personal computing in general, highlights the need to change the relationship between users and the systems function and to define clearly their respective roles and responsibilities. This topic was discussed by the focus group participants, and a summary of the views expressed is given in Figure 3.1. The participants stressed that users are uncertain about who is responsible for data, documentation, maintenance, and so forth. They believed that a clear, unambiguous statement of the respective roles and responsibilities is required, and that without such a statement, the increased user demand for data access cannot be satisfied.

Figure 3.1 Focus group views on the roles and responsibilities of the systems function and end users

"The systems function should be responsible for end-user computing."

"Irrespective of the way in which the responsibilities are divided, organisations should ensure that the rules are unambiguous."

"The growth of end-user computing is causing the boundaries between users' and the system function's responsibilities to become blurred."

"The systems function should be responsible for routine, highvolume systems and for controlling data."

"The systems function should control systems that share data, and be responsible for security, accuracy, recovery and testing of those systems."

 $``The role of the systems function should evolve to become more of a consultative role.'' <math display="inline">\;$

- "Users should be responsible for:
- Ensuring that business requirements identified as critical by top management are included in the systems plan.
- Reviewing the viability of existing applications.
- Providing the justification for new systems.
- Operating single-user systems and office automation systems.
- Managing their own data."

Many of the organisations we talked with during the research for this report were seriously questioning whether their existing structures could cope with the changing relationship between users and the systems function. The structure of conventional systems functions has evolved in an environment where users were largely passive recipients of the services provided by the systems department. But business users of computer systems are now demanding an increasingly active role, and any new organisational structure must resolve the conflicts and reduce the resulting tensions between the systems function and the users.

These tensions and conflicts manifest themselves in the general feeling of dissatisfaction that users often express about the systems function. The comments we heard included:

- The systems department is inflexible, takes a long time to respond to requests for new work, and cannot deliver what is needed. Nor can it make simple modifications to existing systems.
- The systems department cannot accomodate ad hoc requests for data. In particular, there are too few facilities for accessing and analysing historical data (internal and external) for decision-support purposes.

Comments such as these illustrate the frustration that many users feel about the systems department, but they also highlight the challenge that needs to be faced. The challenge is to create an organisational structure that can accomodate the apparently irreconcilable conflict between the way the systems function would ideally like to operate, and the way the business requires it to operate.

One consequence of users' dissatisfaction has been the rapid growth of personal computing. Users have attempted to overcome the frustrations of using the services of the systems function by taking more responsibility for some of their computing activities. However, many information systems managers and line managers are concerned that personal computing costs are rising too fast and are out of control. They are equally concerned that insufficient attention is being given to the justification of personal computing applications and that unreliable development methods are being used.

We believe that it is most important for the systems function to provide a framework for the overall control of personal computing and user access to data. This framework must be established before the third stage (control) of personal computing is reached, because this is when the importance of data management will finally become clear to users. The best way of providing a controlling framework is to establish a user-support centre.

ROLES AND RESPONSIBILITIES OF THE USER-SUPPORT CENTRE

Many large companies have now established some form of user-support centre (often called an information centre). During our research, we heard about several such centres that had not been entirely successful, mainly because they were staffed by the wrong type of people or because their tasks and goals were not clearly defined. It was clear from our research that in many organisations the role and responsibilities of such a centre were still evolving. Participants in our focus group meetings discussed the changes that are likely to occur, and their views are summarised in Figure 3.2.

Similar trends were identified in the annual survey of its members carried out in 1984 by the Information Centre Managers Forum (ICMF). (The ICMF is a 'club' for about 100 information centre managers in the United Kingdom.) This survey asked about the functions performed by information centre staff in 1983 and 1984 and identified the areas where there had been significant changes. The results of the survey are summarised in Figure 3.3. The survey showed that the number of information centres concerned with marketing and cost-justification issues increased between 1983 and 1984, and (suprisingly) that in 1984 fewer information centres were responsible for data management than in 1983.

In order to establish a successful user-support centre it is essential at the outset to state clearly what its role and responsibilities are. In our view, too many user-support centres have been established without sufficient thought having first been given to the 'mission' of the centre. (The ISTAT case history in the appendix describes some of the problems that can result from insufficient planning of the centre's future role.) The ideal mission statement would define:

- The relationship of the user-support centre to the rest of the organisation.
- The tasks that the user-support centre will have to perform in order to fulfil its mission.

Figure 3.2 Focus group views on the likely changes in the roles and responsibilities of the usersupport centre

"There will be a need to provide a consistent data access method."

"Increasingly, the user-support centre will have to market its service to the rest of the organisation."

"There will be a trend for user personnel to move into the central user-support function."

"The user-support centre will need to audit end-user applications, particularly those that impact 'hard' data (especially accounting data)."

"There will be a greater emphasis on the ergonomic implications of designing the most appropriate working environment."

"The post-implementation benefits of end-user systems will be monitored so that future systems requests can be more easily justified or rejected."

-The criteria for deciding which type of computing work is suitable for personal computing.

The first level in the mission statement defines areas such as the link between the user-support centre and the traditional systems development project teams and the way in which the centre will market its services to the organisation.

The second level describes (for example) how the centre will set about providing appropriate hardware and software products, and the types of usage-measuring systems that will be developed.

Figure 3.3 Changes in information centre responsibilities between 1983 and 1984

Responsibility	Percentage of centres with the particular responsibility		
	1983	1984	
Increases			
Software selection	87	93	
Support for PCs	54	92	
Marketing the centre's services		72	
Workstation selection	-	57	
Cost justification		45	
Decreases			
Applications documentation	82	71	
Data management	62	55	
End-user application development	64	50	
Software modification	54	47	

The percentages are the proportion of respondents to the annual survey of members of the Information Centre Managers Forum who said that their responsibilities included the particular item.

The third level details the criteria (such as the amount of development effort, complexity of the project, level of user participation, and frequency of operation) that will be used to decide whether particular applications should be developed by the central systems function or as personal computing applications.

From such a statement, a detailed operational plan can be formulated, describing the way in which the user-support centre will fulfil its mission. One of the details will be concerned with the availability of data. The use of corporate data by business users is unlikely to be improved unless the user-support centre's operational plans include provisions for:

- -Facilitating user access to data.
- -Ensuring data security.
- -Encouraging and supporting user intermediaries.
- -Providing bespoke user education and training.

Facilitating user access to data

The user-support centre's facilities should provide users with access to in-house corporate data and to external data services as well. Users will require access both to 'hard' data (such as that residing in the organisation's mainstream systems, and product or market data available from external services) and to 'soft' data. Internal soft data would usually be in the form of electronically stored text, and external soft data would take the form of economic forecasts, opinions and speculation about future trends.

The user-support centre therefore needs to establish procedures for ensuring that users know what data is available and for coordinating the ways in which that data is provided. The focus should be on providing access to data that will be used for personal computing activities, rather than on satisfying information needs that can best be met by the traditional applications systems.

Ensuring data security and integrity

Many organisations are concerned that unrestricted user access to data could lead to unauthorised access to (and use of) certain classes of restricted data. The user-support centre must therefore establish procedures to ensure that the organisation's data is secure and is accessed only by those with authority to do so.

In some cases, a personal computing application developed by a business user may be critical to the success of that part of the business. Indeed, the failure of such a system could significantly impair the efficient functioning of the whole organisation. More often than not, these applications are shared by several user communities, which means that data integrity and validity are of paramount concern. The user-support centre must insist that such developments adhere to the same standards as the centrally developed systems.

The next Foundation report (Threats to Computer Systems Security) will address the subject of data security in detail, so we will restict ourselves here to two general remarks based on our research findings:

- Adequate data security can be provided within an information-centre environment.
- All data security procedures depend for their success on the conscientiousness of the people implementing and using them.

Encouraging and supporting user intermediaries

One organisation (Philips in Eindhoven, whose experience is described in the appendix) had identified an 'organisational gap' between its formal, centralised user-support centres and the users themselves. By their nature, user-support centres are usually centralised functions, often located in offices near the organisation's computer centre and therefore remote from most of the user community. Philips (and some of the focus group participants) had identified a serious drawback to this arrangement: users much prefer to have a single, local reference point for advice and support relating to their dataaccess requirements. As a consequence, many user departments in Philips now have their own unofficial 'user intermediary', who is the local expert consulted by his colleagues. Usually, this role is performed by someone from a user department in addition to his or her normal duties

In most organisations, the role of local expert is not explicitly recognised in the organisational structure. Philips, however, has recognised the existence of the user intermediaries and has set about integrating them organisationally with its overall personal computing strategy. The user intermediaries are invited to informal gatherings, where they meet with other people performing similar roles, to discuss new products, common problems, and the educational needs of their colleagues in the user departments. Philips believes that the key to unlocking the corporate data resource is to encourage and support the local user intermediaries.

Providing bespoke user education and training

The education and training requirements for business computing users will vary enormously. The requirements of an infrequent user who will only ever use preformatted commands for straightforward data retrieval will be very different from those of a user intermediary, for example. The former simply needs a comprehensive set of user procedures, preferably describing a menu-driven system, whilst the latter requires more detailed information about various tools and techniques. Our research shows that many organisations believe that their user-support centre's most urgent task is to educate the user community about what data is available and to train users in how to access and use it. The underlying belief is that user-support centres can make a real contribution to organisational effectiveness by encouraging users to access data and convert it into useful information.

There is also a need to provide basic personal computing education for business managers at all levels so they can judge better which systems their staff should develop and the likely costs of such developments.

Education and training will also be required for information systems staff, particularly to help them understand better the capabilities and uses of personal computing software. Historically, some systems staff have believed the systems function is the custodian of the organisation's information. We believe, however, that it is important to make a clear distinction between data and information. As we pointed out in Foundation Report No. 40 (Presenting Information to Managers), the information content of particular data depends on the context in which it is received and the background and experience of the person receiving it. Often, systems staff are not in a position to know what information is important and how it can be used effectively by the business. The role of the systems function is to provide business users with easy access to data that they will convert to information in the context of their particular business environments.

Education and training, both of users and systems staff, therefore require a long-term approach because it is necessary to change fundamental attitudes and preconceptions about the ways in which computer systems are used.

ESTABLISHING A SUCCESSFUL USER-SUPPORT CENTRE

Some user-support centres have not been particularly successful. In one organisation, for example, the user-support centre had effectively been rendered impotent by certain policy decisions:

- The business units' information systems budgets provided only for traditional core systems development and enhancements, and the performance of the local data processing managers was to be measured strictly against these budgets.
- —No processing facilities were to be given to the user-support centre until it could identify a known demand that could be translated into a specific machine size and operating environment.
- -No application development tools were to be used.

—Staff allocated to the user-support centre were employed by the systems development department, and they were switched to conventional applications work if priorities were changed.

Our research suggests that there are two areas that are critical to the success of a user-support centre. The first concerns the way in which the new centre is introduced, and the second concerns the way in which it charges for its services.

Selecting the initial application

Once a manager has been selected and appointed and a clear 'mission' and reporting lines have been agreed, it is important that the new user-support centre quickly establishes its credibility. This can best be achieved by selecting a high-profile application that will be used to demonstrate the capabilities of the new centre. The application should be:

- -Clearly defined, but not too complex.
- -Important in the eyes of top management.
- -Able to be solved by an ordinary end-user product.
- -Shown to have measurable benefits.

Ideally, the initial application should also possess the following characteristics:

- -The requirement should be stable.
- More than one person in the user department should use the application.
- -It should not replace an existing computer system.
- -It should be used frequently.
- It should require access to a limited amount of data.
- -It should be sufficiently important for other user departments to want to use it.

During the implementation of the initial application, the staff of the user-support centre should be trained and the users should be involved in awareness seminars, and the like. At the end of the initial project, a management presentation should summarise the experience and the costs and benefits of the exercise. If the project has been used as a pilot trial, plans for developing a full-scale user-support centre should be proposed, and authority to proceed should be sought.

Charging for user-support services

A successful user-support centre can be expensive to run and will require a substantial initial investment in people, hardware, and software. We believe that the centre should not be a free service, but that the initial users should not be expected to pay for all of the start-up costs. The mechanism for recovering the costs of the user-support centre will vary from

CHAPTER 3 ESTABLISH A USER-SUPPORT CENTRE

organisation to organisation, but we recommend as a general rule that users should pay for the facilities they use. Some organisations base their charging structure on a fixed fee per terminal; others use disc space required as the basic unit of cost; one organisation we met charged an arbitrary percentage of the users' current external timesharing rate.

Our research showed that the most successful usersupport centres use simple, easy-to-understand charging formulae that minimise the chance of users receiving an unexpectedly large bill. The elements of such a formula would typically include:

- -Fixed charges for training courses.
- Fixed rental or purchase charges for any hardware and software supplied by the centre (including telecommunications facilities).

---Variable charges for the use of the centre's computing resource. The scale of charges should relate to the amount of resource used.

All other costs should be considered as overheads and recovered in the above charges.

The recharging mechanism should be designed to make users aware of the cost of the computing resources they are using, but it should not discourage them from using the centre's facilities. This may mean that the costs of running the centre during its first few months of operation may have to be subsidised by the information systems function. Nevertheless, we believe that the user-support centre should aim to break even after about one year.

CHAPTER 4

SELECT THE DATA ACCESS TOOLS AND TECHNIQUES

1 1 N

Part of our research effort was directed at investigating the suitability of current tools and techniques for meeting the data-access requirements of users. The general view of Foundation members was that there are still substantial technical difficulties to be overcome before user access to data can become the norm. Most of the technical difficulties are associated with 'building bridges' between old and new technology. A summary of the views expressed at the focus group meetings is shown in Figure 4.1.

One organisation (Digital Equipment in the United States) presented us with a very different view. This computer supplier told us that, where extensive networking and a common systems environment exists, there are no longer any major technical constraints to providing users with access to corporate data. According to Digital, the main difficulties are concerned with education and training, and these difficulties will be resolved in an evolutionary way.

We also believe that the required technology exists now, and that it is available in products. Unfortunately, its effective deployment is inhibited because some of the best products are not compatible with each other or with existing data management systems (including file handlers). Suppliers are moving rapidly to fill these gaps by extending their products so they can be used with additional systems environments.

Figure 4.1 Focus group views on current data-access tools and techniques

"Data dictionaries need to be improved by incorporating metadata and by using information directories."

"Query languages need to be made easier to use."

"There is a need for better log-on and access tools."

"Common communication standards are essential for future compatibility in a multi-vendor environment."

"I am sceptical about the use of methodologies that are supposed to identify end-user requirements."

"It is essential to begin end-user planning with a thorough data definition exercise."

There is a steady stream of new products, particularly for microcomputer-mainframe links, where we found references to more than 400 products. A detailed evaluation of all of the products is clearly beyond the scope of this report. Instead, we provide general guidelines for selecting products, and we predict the way in which they are likely to evolve during the next three to five years. The views of the focus group participants on the likely future developments in dataaccess tools and techniques are summarised in Figure 4.2. In our view, future developments will focus on reducing the complexity of the interaction between users and the systems they are accessing.

The data-access facilities that are required can be divided into six types, but some products will provide several types of facility. The six types are:

- Data dictionaries, which contain details about the data that can be accessed.
- Database products (software and hardware).
- -Data-retrieval products.
- -Microcomputer-mainframe links.
- -System-linkage products.
- -User-interface techniques.

Figure 4.2 Focus group views on the likely developments in tools and techniques for making data available to users

"More intermediary software to facilitate log-on, access, and search procedures will become available."

"There will be better tools for identifying what data exists."

"Database machines will become more evident."

"Distributed databases will begin to make a real impact, and there will be a corresponding improvement in distributed systems interconnection."

"There will be substantial performance improvements in the operational performance of software products."

DATA DICTIONARIES

From our research, it is clear that users are now demanding that the systems function should provide facilities that help them find out what data is available, how it can be accessed, and what tools are available for processing it and combining it with other data. To answer these questions, some form of data dictionary is required.

The use of data dictionaries has been advocated for many years. The dictionaries that have been constructed have been used primarily as an aid to the system development and implementation processes and not as a practical tool that enables users to locate and access the data they require. Our investigations suggest that the principal reason for data dictionaries not being used to their full potential has been a reluctance to allocate scarce manpower to an area that will not produce immediately visible benefits.

We believe the data-access demands being generated by the growth in personal computing (and in office automation) now make a comprehensive data dictionary essential. In particular, we believe that a centrally organised repository of 'data about data' (often known as metadata) reduces both the effort required for data management and the incidence of errors in describing data.

Current data dictionary systems are complex to use, however, and are not suitable for use by nonspecialists. They can also require substantial computing resources, particularly if some form of associative retrieval is to be carried out, which means that their response times are often too long to encourage routine use. Also, the metadata facilities they can provide are limited. For example, existing products are unable to answer key questions, such as whether the data in two different files is based on comparable definitions or has the same precision and timeliness.

Many of the commercially available data dictionary products are substantial systems in their own right. The selection of an appropriate data dictionary product is a crucial issue, and the choice must be based on the requirements of the individual organisation. Selecting an inappropriate data dictionary system can prove to be a costly mistake because of the data administration and system development effort that will have been wasted before it is abandoned.

We reviewed a new type of framework (known as the environmentally dependent framework) for selecting a data dictionary system. (This framework was described by B K Kahn in 'MIS Quarterly' in September 1985.) In our view, the framework overcomes the shortcomings of a conventional featureby-feature comparison of alternative data dictionary products because it takes account of the organisation's requirements and systems environment. However, before taking account of the organisation's environment, the framework first considers a checklist of basic data dictionary facilities, and whether the alternative products provide these facilities. The checklist of facilities is shown in Figure 4.3.

The environmentally dependent framework relates the choice of a data dictionary system to the organisation's data administration and data management objectives by considering two factors:

- —The requirements for the operational control of data, including system maintenance and documentation, systems and programming standards, data integrity and data security, and systems design and analysis methods.
- The present status of the organisation's database management systems, which may range from no systems in use, through a single database management system, to multiple systems.

Figure 4.3 Checklist of basic data dictionary facilities considered by the environmentally dependent framework

- Data dictionary maintenance facilities
- The method used to input data (online or batch).
- The facilities available for copying data that is already stored in the dictionary.
- The way in which linkages between entities are defined.
- The ease with which the dictionary can be constructed.
- The range of update commands available (add, modify, delete, etc.).
- The method used to update free text held in the dictionary.

Reporting facilities

- The range of standard reports that are available.
- The ease with which report options can be specified.
- The range of report-formatting commands that are available.
- The types of facilities available for scanning the contents of the dictionary and the query facilities available.

Data definition and naming facilities

- The maximum name length that can be accommodated.
- The facilities available for handling synonyms and aliases.
- The facilities for handling duplicate names.
- The facilities for handling free-form text.

Documentation

- The overall adequacy of the documentation supplied with the product.
- The usefulness of the examples included in the documentation.

The framework allows each shortlisted data dictionary product to be evaluated in the context of each of these aspects.

DATABASE PRODUCTS

Database products can take the form of specialpurpose hardware and software systems. The most obvious examples of software products are the various commercially available database management systems. Hardware products, which usually have their own specialised software, include database machines and content-addressable file stores.

Software products

Most large organisations have, for the past few years, based their central core applications systems on 'first-generation' database management systems that is, systems based on hierarchical or network data structures where the relationships between different data elements have to be predefined. We believe that the growing demand by users for ad hoc access to corporate data means that an organisation's database policy should be based on relational database management systems. Relational systems, by their very nature, are more suitable for handling the unstructured ad hoc accesses that users will require.

The implication is that plans should be made either to change the core systems' databases from hierarchical to relational, or to build links between the existing hierarchical databases and the relational databases that will be accessed by users.

For most organisations, the former will be a long-term goal and the latter will be a short-term expedient.

Hardware products

One problem with any mainframe-based database management system is that it can impose substantial processing overheads. As a result, various types of specialised hardware have been developed over the past few years to help ease the problems of maintaining and providing access to large databases. The two most significant developments have been the emergence of dedicated database machines and the content-addressable file store (CAFS) product available from ICL.

Database machines

A database machine is dedicated to the task of database management. The database management software resides in the dedicated hardware, and relieves the host computer of the processingintensive activities normally associated with maintaining and accessing a database.

In the United States, the two leading suppliers of database machines are Britton-Lee and Teradata

Corporation. Britton-Lee has been supplying database machines since the late 1970s, and Sperry is now offering a Britton-Lee system on an oem basis to provide its large 1100-series models with a relational database capability. Britton-Lee's machines range in price from \$57,000 to \$155,000 (although the Sperry version costs \$250,000).

In the United States, there is a Britton-Lee user group with about 450 members. One of them has reported that a Britton-Lee IDB (intelligent database machine) 500 was used to bring together three databases, all in different formats. The alternative would have been to use database software packages that would have required a large host computer. The \$395,000 cost of the IDB 500 was justified because it would have cost \$200,000 more to install the VAX 780 that the software packages would have required. In reality, though, a VAX 780 would have been inadequate for what has been achieved with the IDB 500; this company believes that an IBM 3084 would have been required to do the same job.

Inco, a software system house based in Virginia, is both a user and an oem supplier of Britton-Lee machines. It has developed its own front end, which it calls SQL/Universe. This product provides full SQL capability and can operate with IBM and North Star microcomputers, Prime and DEC-PDP 11/70 minicomputers, and the IBM 4300 series operating under VM/CMS.

After many years of beta-site testing, Teradata is now actively marketing its product. The smallest Teradata system, with six processors and four discs (each with 474M bytes of storage), costs \$320,000. At the other end of the scale, Teradata has demonstrated what it describes as "the largest parallel-processing computer commercially available to the business data processing industry". This database machine has 60 parallel processors working against a multi-million-record database. Teradata claims that the system has the equivalent processing capacity of IBM's 3084Q, which costs approximately \$6.2 million (including storage). The comparable Teradata price is \$1.7 million.

By the middle of 1985, seven large organisations in the United States were using Teradata systems, although all of them were at early stages of implementation. At that time, Teradata had software interfaces to IBM's MVS, VM, and CICS environments. One of the users (Citibank) was considering interfacing the Teradata system to IBM's TPF/11 software. (TPF/11 is an extended, enhanced version of IBM's Airline Control Program, used as the basis for the IPARS seat reservation system; it is designed to handle much higher transaction loads than those typically found with CICS systems.) Until now, database machines have not achieved the market penetration that many industry observers had expected. We believe that the increasing demand for access to corporate data will mean that many more organisations will install these machines to prevent their mainframe systems from being overwhelmed by user-generated processing demands.

The content-addressable file store (CAFS)

The conceptual origin of ICL's content-addressable file store (CAFS) goes back to the early 1960s. The conventional way of randomly searching large files of data has been to construct indexes to the data. There are two drawbacks to this approach: in extreme cases, the absurd situation can arise where the indexes occupy several times as much disc space as the data to which they refer; and it is never possible to guarantee that the indexes will cover all possible search combinations. CAFS overcomes these drawbacks by eliminating the need for indexes; instead it searches all the data, matching its contents against the search keys provided.

The first prototype CAFS machines were built in the early 1970s. The early versions of CAFS searched data on specially adapted disc drives that had multiple read heads and amplifiers. This feature has often been mistakenly seen as an essential characteristic of CAFS. The multiple read heads were adopted purely as a means of increasing the data transfer rate to match that of the CAFS searching and evaluation mechanisms.

During the late 1970s, various commercial trials took place to confirm the market potential of the product. The most significant trial involved a telephone directory enquiry service, where records of about 7,000,000 subscribers were stored on the system. Throughout the trial, the observed average response time was 1.7 seconds. The telephone directory operators also found that the system was much more enjoyable and satisfying to use than the previous paper-based system.

The first commercial version — CAFS 800 — was announced in 1979. However, at £250,000 each, this machine was an expensive piece of equipment. Since then, ICL has progressively incorporated CAFS into its mainstream product line, and the price has fallen substantially. It is now automatically associated with standard file and database structures and is available as part of the standard hardware for ICL's 2966 family of systems.

The CAFS-ISP product was launched in April 1982, and by June 1985 the five-hundredth order for this product was received. Furthermore, CAFS-ISP is now an automatic and integral constituent of ICL's latest Series 39 systems. CAFS is fitted as standard to the first controller on any Series 39 configuration and is optional on second controllers. In 1985, the ICL Computer Users Association prepared a report, 'CAFS in Action', which included seven case histories describing the early experiences of using CAFS. The overall lessons to emerge from these experiences were:

- -CAFS has an important role to play in providing end-user facilities.
- The demand for CAFS-based end-user services has usually exceeded expectations.
- CAFS has increased productivity, both by reducing the amount of development effort required and by making new kinds of analysis possible.

Typical CAFS applications reported in the case histories included:

- A viewdata system installed in a public library that allows the public to interrogate the library's catalogues.
- The Inland Revenue's name and address tracing system, which involves searching 63 million taxpayers' records.

CAFS would appear to be an ideal product for those organisations seeking to provide users with efficient access to corporate data files. It is surprising that it has taken so long for CAFS to become established as a viable commercial product and that no other major computer supplier has followed ICL's lead by building similar products.

DATA-RETRIEVAL PRODUCTS

From the users' point of view, the ideal data-retrieval product would provide a consistent interface regardless of whether the file or database being accessed was stored on an in-house computer or on an external commercial service.

In reality, however, each database management system has its own bespoke command language and dialogue. The scale of the problem can be gauged from the vast number of commercial online databases that are now available. One directory lists 2,764 databases offered by 414 different services. The only area in which any real progress has so far been made is in automating the log-on procedures required to gain access to a variety of databases.

In Europe, the European Commission has been active in trying to define standards for a Common Control Language. Although proposals for a powerful language were made, no progress has been made in developing products that use these standards. More recently, the Commission has sponsored a study, carried out by Butler Cox, that investigated the potential for applying artificial intelligence techniques to information retrieval. Our overall conclusion was that this field is still at a very early stage of development and that it will be several years before products based on artificial intelligence will be available to allow users to access a wide range of databases through a consistent interface.

Nevertheless, a variety of products are commercially available to provide a user-friendly interface for users wishing to access corporate databases. Probably the best-known product is Intellect, developed in the United States by Artificial Intelligence Corporation, and available in Europe through Intellect Software International Limited. Intellect has been designed to provide a user-friendly, natural-English interface to mainframe databases. It also provides statistical, graphical, and modelling facilities for manipulating the retrieved data.

Intellect can handle verbose or terse queries, poor grammar and poor punctuation, it does not require any special commands, codes or characters, and it can clarify the ambiguity typically found in a user's request. The product works by maintaining for each database a knowledge base called a 'lexicon', which essentially is a specialised data dictionary. The lexicon defines the form of the data, how it is known in the database, and the words, phrases and concepts that the user would like to use in referring to the data.

Intellect can also learn from the user. It has a feature called 'English definition' that allows users in the course of a query session to define new meanings for terms or elements that are already known to the system.

At present, Intellect is available for IBM (and compatible) mainframes running under the MVS/VM and DOS operating systems, with TSO, CMS, CICS, and Complete TP monitors. It can be used with Focus, IDMS, VSAM, SQL/DS, Adabas and other file types. Versions for DEC VAX and Wang minicomputers and for the IBM PC will be available during 1986.

Despite the success of products such as Intellect, we believe there is a considerable amount of work still to be done on defining the requirements for naturallanguage (and other user-friendly) interfaces before a general-purpose product can be produced.

MICROCOMPUTER-MAINFRAME LINKS

Systems departments are coming under increasing pressure to establish links between microcomputers and the corporate mainframes, so that microcomputer users may have access to the data in the mainstream files. The proliferation of products claiming to provide such links and the confusing variety of approaches that can be adopted have, we believe, led many organisations into hasty purchases. Most of the products we heard about worked well in specific environments. The major problem we encountered was that users had either overestimated the products' capabilities or expected them to work in situations for which they were not designed.

Our research shows that the principle of caveat emptor (let the buyer beware) applies particularly to the purchase of microcomputer-mainframe links. In particular, we found that:

- The marketing literature for a particular product often makes extravagant claims about the interconnection problems that can be solved by the product. It can be very difficult to determine if the product is genuinely compatible with your particular requirement.
- User experience shows that microcomputermainframe links work best when the software at both ends of the link is provided by the same vendor.
- The standard of documentation is often very poor, and this makes it difficult to track down errors or to find the best procedure for a particular situation.

Ideally, organisations that do not have experience of implementing this type of link, and that do not wish to be pioneers, should wait for market forces to isolate the viable approaches. Unfortunately, this is not an option for most Foundation members, so we now describe some of the features, as well as the advantages and disadvantages, of each of the five approaches to providing telecommunications links between microcomputers and mainframes:

- -Terminal emulation in the microcomputer.
- -Protocol converters.
- Local area networks with a gateway to the mainframe system.
- Incorporating microcomputers into an advanced network architecture.
- -Public network services.

These approaches are not entirely independent: for example, protocol converters are a supplement to asynchronous emulation, not an alternative. Also, public services (in the form of value added networks) require terminal emulation facilities.

Terminal emulation

Terminal emulation is usually achieved by enhancing the microcomputer with an interface board and appropriate software so that it can emulate a dumb terminal such as an IBM 3278 or a DEC VT100. The microcomputer can then take the place of a terminal in the network.

Almost any type of terminal can be emulated, but two cases are particularly common — emulation of an

CHAPTER 4 SELECT THE DATA-ACCESS TOOLS AND TECHNIQUES

asynchronous terminal that has a V.24 (RS232) interface and emulation of a synchronous terminal that has a high-speed coaxial cable interface. Each of these cases presents a different balance of advantages and disadvantages.

Asynchronous emulation

Asynchronous emulation is the obvious (and usually the best) method where the host computer is normally accessed via asynchronous terminals. A microcomputer emulating an asynchronous terminal may be connected direct to the host, or via a switching facility, or via a suitable local area network.

Direct connection is rarely appropriate because it permanently occupies a host port, even though the microcomputer's need to access the host is likely to be intermittent. Connection via a switching facility is appropriate in the following situations:

- --Where a microcomputer is used mainly as a standalone device, but the user requires occasional access to mainframe data.
- Where the microcomputer user requires basic terminal facilities at the lowest possible cost.
- Where a single terminal requires access to several different manufacturers' mainframe systems.

Synchronous emulation

In an IBM environment, synchronous terminal devices (screens and printers) are usually connected to the mainframe via 3274 or 3276 cluster controllers, typically with between 8 and 32 devices per controller. (Similar arrangements occur in some other manufacturers' environments.) Interface boards are now available to allow a microcomputer to connect to a cluster controller, and thence to the mainframe.

In IBM terminology, the add-on communications adaptor must allow the microcomputer to emulate a 3178 or 3278/79 terminal. One of the most commonly used adaptors is the Irma Board (for the IBM PC), supplied by Digital Communications, Inc. More recently, interface boards have become available to allow a microcomputer to connect to multiple sessions on the host computer through a single cluster controller. Each session can be accessed via a separate 'window' on the microcomputer. (The board may also provide other windows for the microcomputer's own applications.)

The typical cost of an interface board and its associated software ranges from \$1,000 to \$1,500. Additional software is usually available for transferring files between the mainframe and the microcomputer via the cluster controller. Without a properly designed file-transfer mechanism (and we know of no such mechanism for 3270-emulation products) transferring files consumes considerable amounts of the host's computing power. Emulating a synchronous terminal that has a coaxial cable interface has several practical advantages:

- -The link operates at high speed.
- The networking arrangements do not need to be changed.
- Products are available from mainframe suppliers.

This approach is expensive, however, and it does have some considerable disadvantages from the overall systems architecture point of view:

- -Switching facilities are limited to those provided by the mainframe and network processors.
- The microcomputer cannot easily access hosts that are not compatible with the terminal being emulated, even if such hosts are connected to the network.
- Interactive terminal protocols are unsuitable for file transfer, and some suppliers' network architectures are too inflexible to allow file-transfer and fileaccess protocols to coexist.

We conclude, therefore, that although synchronous emulation of a terminal that has a coaxial cable interface may have a temporary role to play, organisations should plan to migrate to a more architecturally sound position based on advanced network architectures (see page 23).

Protocol converters

A protocol converter is a 'black box' that sits between a computer system and a terminal and allows them to communicate, with each using its native protocol. Such devices (the Tempus-Link product, for example) have recently become available for microcomputermainframe links.

Protocol converters comprise two communications interfaces linked by a data buffer and processing power for carrying out the conversion. Data from one system, usually a mainframe with a synchronous protocol, is received by the protocol converter, with the converter handling the polling requests from the mainframe. The data is stored in the buffer in preparation for onward transmission to the second system. Using a buffer not only allows data to be stored during the polling sequence but also allows different line speeds to be used on each side of the converter.

Protocol converters typically cost about \$500 (compared with \$1,500 a few years ago). They are a cost-effective approach to microcomputer-mainframe links if there is an existing large investment in microcomputers that now need to be linked to the corporate mainframes.

The most obvious disadvantage of protocol converters is that they are not suitable for handling

CHAPTER 4 SELECT THE DATA-ACCESS TOOLS AND TECHNIQUES

large volumes of data. The problem arises because, typically, the mainframe transmission speed is reduced from 9600 bit/s to 1200 bit/s to avoid overloading the converter. As a result, the transmission of large volumes of data can increase considerably the response times at the microcomputer.

Protocol converters may also cause problems with system security because they usually do not inform the mainframe if a user is disconnected without going through the standard log-off procedure. The next user may then be connected to the uncompleted mainframe session.

Problems can also arise because protocol converters do not take account of the differences in keyboards or screen sizes of the various devices that may be connected to the converter. These particular problems can be overcome, however, and most of the organisations we talked with solved them by employing manual procedures at the microcomputer end of the link.

Local area networks and gateways

The third approach to providing microcomputermainframe links is to add a gateway to a local area network that supports several microcomputers. Gateways translate one network protocol to another, thereby overcoming the problems of hardware and software incompatibility. They usually work by emulating a cluster controller.

The most obvious example of a network gateway is the SNA gateway, which has been developed because of the pressing need of users to interconnect their IBM equipment with other vendors' hardware. The most common type of SNA gateway allows a local area network to connect via SNA links to IBM mainframes. Microcomputer users are then able to access mainframe applications, data, and peripherals at lower cost than with other types of links because no additional hardware is required in the PC. The costs of the gateway, and its benefits, are shared by all the microcomputers connected to the local area network.

The obvious advantages of building microcomputermainframe links via a local area network are:

- Only one connection to the host is required, and no additional hardware is required for each microcomputer.
- Each local area network may have several gateways providing access to several hosts and networks.
- Fewer ports will be required at the host computer.
- No additional cable is required in order to increase the level of connectivity.

Sometimes, a microcomputer-based local area network can provide substantial savings compared to the cost of terminal clusters. Now that IBM has given its 'seal of approval' to local area networks by announcing its own strategic product (the token ring), we expect to see a rapid growth in the availability and use of local area network software and applications.

Incorporating microcomputers into an advanced network architecture

In the past, only larger computers (minicomputers and upwards) could participate fully in network architectures such as SNA and DECnet. Terminals were supported by using small subsets of the full network functions, or by methods outside the architecture, which meant that terminals were regarded merely as adjuncts of the central computer. The increasing power and storage available with today's microcomputers now make it possible to devote substantial microcomputer resources to the provision of networking functions. Thus, personal computers can now operate as full nodes in a network.

As a consequence, the major computer suppliers are incorporating large parts of their network architectures in personal computers. IBM, for example, has been obliged to rethink some of the fundamental concepts of SNA. Recent developments, such as PU2.1/LU6.2 and the APPC (advanced program-toprogram communications) product announced in conjunction with the token ring local area network, have effectively created a new version of SNA. These developments provide peer-to-peer working, interprogram communication, and a layered architecture free of the 'device dependencies' that IBM now admits were included in previous versions of SNA. The new version of SNA can be used to build very small neworks, even networks that do not require a mainframe at the centre.

The implementation of an advanced network architecture that incorporates microcomputers and mainframes will provide several advantages:

- -Support for concurrent communications between a microcomputer and several other machines.
- The ability to use a single microcomputermainframe connection for a mixture of interactive, file access, file transfer, and mailing functions.
- -Support for the management of distributed processing.

Moreover, these advantages do not depend upon the use of a particular means of interconnecting the various computing devices found in an office environment (although local area networks provide a better basis for such interconnections than any kind of circuit switching). An advanced network architecture bypasses many of the problems encountered with emulation and protocol converters because such an architecture has been designed for intermachine communications. All other approaches to microcomputer-mainframe links are attempts to use terminal-support protocols for purposes for which they were never intended.

We believe, therefore, that the use of an advanced network architecture is the best long-term solution to providing microcomputer-mainframe links. Unfortunately, there is, at present, an acute shortage of products, and some important aspects of the architectures are not yet defined. For example:

- OSI is fully defined only up to the transport level.
- —SNA is more fully defined, although it still lacks a standard file-transfer protocol. APPC is not expected to be implemented for some machines until the end of the 1980s. However, it will be available for other machines in 1986, notably for the IBM PC.

Organisations must, therefore, either adopt a proprietary architecture (DECnet or XNS, perhaps) or content themselves with stopgap solutions until OSI and/or SNA is stable and available in products.

Public network services

Some organisations in the United States choose to use the protocol conversion services of third-party network operators (such as Tymnet) in order to provide microcomputer-mainframe links. In Europe, the PTTs (and value added network operators in the United Kingdom) will increasingly be offering this type of network service.

Using a third-party network service for microcomputer-mainframe links can have drawbacks similar to those we identified for protocol converters. In particular, they can cause transmission delays, especially for large volumes of data. There is also the uncertainty about access and performance when a public network is used.

SYSTEM-LINKAGE PRODUCTS

All of the approaches to microcomputer-mainframe links described above provide only the basic telecommunications link between the microcomputer and mainframe. They do not provide any form of systems integration between the two ends of the link. Several commercial products are available to manage the interconnection of the systems at each end of the link. These products usually only provide controlled access to files in a timesharing, rather than a transaction-processing, environment. They may also incorporate directories and provide conversion from mainframe formats to those used by the microcomputer's programs. Ideally, organisations would provide this type of systems linkage by adopting an integrated two-level systems architecture of the type we described in Chapter 2. However, it will be some time before commercially available products provide all of the features that are required to enable such an architecture to be built. In the meantime, several organisations have built their own software to allow a microcomputer or workstation user to gain access to several different mainframe (or remote) services through a common user interface.

ICI, for example, has developed its own product, Conductor, which provides a standard menu framework that can be used to access a variety of applications and utilities. (Conductor was described in some detail in Foundation Report No. 43.) Philips in the Netherlands has also developed its own interface, as described in the case history in the appendix to this report. During our research, we also heard about a similar product (Host Interface Manager — HIM) developed in the United Kingdom by IMI Computing Limited.

HIM enables an IBM PC user to gain access both to PC and to mainframe applications. The user selects the required application from a simple menu, and HIM then either initiates the relevant PC program or automatically sets up the appropriate connection between the PC and the mainframe. The user is therefore relieved of the tedious task of logging on to the mainframe system and initiating the required application.

USER-INTERFACE TECHNIQUES

The most obvious change in the user interface during the past few years has been the development of window-display techniques, accompanied by the use of icons, mouse and pointer. These techniques were pioneered by Xerox and at first were available only with very expensive equipment. They are now finding their way into many organisations through inexpensive personal computers such as the Apple Macintosh, and through microcomputer software products such as Gem from Digital Research and MS Windows from Microsoft. Indeed, the day is fast approaching when all business microcomputers will provide these types of user interface, but only for programs resident in the PC, not for those resident in other computers.

Despite the lack of enthusiasm of many systems staff for these new forms of user interface, we believe they are very significant developments. In particular, they will appeal to business users of microcomputers people who are not computer specialists, but who want to use a personal computer or workstation as a support tool on an ad hoc basis. This type of user will demand to use the same type of interface for all of their computing work, including access to corporate or external databases.

CHAPTER 4 SELECT THE DATA-ACCESS TOOLS AND TECHNIQUES

8 3

The user, or man-machine interface (MMI), is currently receiving a great deal of attention from the research community. Much of this research effort has been in response to the Japanese Fifth Generation project, part of which is aimed at improving the interaction between people and computers. Although this research is not aimed specifically at making it easier for users to gain access to corporate data and external data sources, any progress is bound to help in these areas. We therefore review briefly the direction of one of the research efforts.

In the United Kingdom, the Alvey MMI programme is now under way. The Alvey Programme is a government-sponsored research and development activity aimed at developing new information technology products and techniques.

The Alvey MMI programme covers user-interface design, human factors, input-output devices, pattern analysis, and intelligent front ends. The work on intelligent front ends is aimed at solving two problems, the importance of which is endorsed by our own work. These problems are:

- —The wide differences in hardware, operating software, query facilities, and data definitions that exist in the various databases and services available to users. Ideally, users would like to have a single view of all the facilities, and this view should be clear, consistent, and easy to understand.
- The wide differences that individuals in different departments, or who perform different roles, have in their perceptions about data.

The differences can be accommodated by providing user-specific interfaces, but to do this usually requires complex programming and may require the user to learn unfamiliar techniques. Clearly, there is a need for a general-purpose solution that does not require bespoke programming.

Intelligent front ends are intended to solve the problems by using stored 'knowledge' about the users and about the databases and services. The front end would interpret the user's input, select the required facilities, and then control them on behalf of the user. Artificial intelligence techniques will be used to encode and use the required 'knowledge', and the Alvey research is investigating how this can be achieved.

Intelligent front ends are expected to be able to interact with users in a variety of modes, including text, pictures, and speech. Ultimately, they should be able to 'listen', negotiate, criticise, argue, and explain, as well as respond to requests for data.

At the end of 1984, Dr David Probert of British Telecom, the industrial coordinator for Alvey's intelligent front-end research, described the overall timescale for the development of several 'waves' of intelligent front ends. According to this timescale, no really useful products will be available until 1990, and the most advanced products will not be available until after 1995. Users must therefore plan to solve their immediate data-access and manipulation problems without relying on the availability of the imaginative products that may be developed as a result of the Alvey Programme.

REPORT CONCLUSION

Our research proved that most Foundation members are aware of the growing demand for access to corporate data. They wish to meet this demand, but are seriously constrained by the limitations of their current operating software (including database management systems), development tools, query languages, and network facilities.

Our analysis has shown that a complete solution to the data-access problem will require the definition of new systems environments based on a significantly extended systems infrastructure. Many existing applications will have to be transferred to one of the new environments, or redeveloped completely. In addition, effective user-support (or information) centres will have to be created, and changes will have to be made to the way in which data is managed. We believe that the methodology described here will help members as they plan the appropriate technical and organisational infrastructures that are required to provide access to the corporate data resource.

We have described the data-access tools and techniques that are available as products today, and we have indicated how these are likely to develop in the future. As long as the right technical and organisational infrastructures have been established from the outset, members should be able to provide their managers and professionals with greatly enhanced access to data, and they should be able to exploit the new data manipulation tools as they become available.

APPENDIX

CASE HISTORIES

The six case histories reported in this appendix have been chosen to illustrate the different approaches adopted by a range of organisations, both in Europe and the United States. Each one illustrates one aspect of making data available to users.

In gathering the information for these case histories, we used a standard questionnaire that was designed to explore the following areas:

- -The growing demand for data.
- The most important requirements for personal computing.
- The main benefits to be gained by planning strategically for personal computing.
- The best mechanisms for developing and implementing personal computing.
- Examples and features of decision-support and executive-support systems.
- The requirements for future systems environments.
- The potential problems inherent in a data design exercise.
- The capabilities of current database management systems.
- The main functions of a personal computing facility.
- -The roles and responsibilities of the systems function and of users.
- The ways in which data access should be improved.
- Data-access management issues.
- The developments in personal computing most likely to take place by 1990.

In preparing the case histories, we have focused on the range of data-access options available to organisations, and we have noted the important messages emerging from each organisation's experience. Comments about the impact that user access to data has had on organisational structures have been included in the majority of the cases. All the Foundation members we met with were concerned about the best way of providing users with access to corporate data. In fact, there was a remarkable degree of uniformity in the concerns expressed. However, there were significant differences in the ways in which organisations have attempted to overcome these concerns. These differences arise mainly from the different contexts in which Foundation members are operating and from their different organisational structures and management attitudes. These attitudes are influenced strongly by an organisation's experience of using information technology.

The case histories indicate that organisations are extending their range of systems environments to incorporate increased user participation so that users can more readily access the data they require. The role of the systems function in this situation is twofold: to provide support for all the systems environments (through consultancy, training, etc.) and to build the infrastructure in which end-user developments can take place (for example, by providing tools and developing standards and policies for data management).

AVON PRODUCTS INC.

Avon is a diversified company involved in the manufacture and distribution of cosmetics, toiletries, health-care products, speciality chemicals, jewellery, and clothing. The Avon Division is the world's largest direct-selling business, with approximately 400,000 representatives in the United States and more than 900,000 international representatives. Net sales in 1984 were \$3.1 billion.

End-user computing environment

Visitors to Avon's offices in Rye, New York, are immediately aware that this is a leading-edge company in its delivery of information services; the ratio of electronic terminals to employees is close to 1:1. Avon is a good example of a company that believes in high technology and that has organised itself to support and foster its use at all levels. In the early 1980s, personal computers were delivered to

APPENDIX CASE HISTORIES

the homes of top management so that they could become familiar with their capabilities in private. Today, Apple Macintoshes are used to provide executive-support facilities.

Avon's approach to making data available to users can be summarised as a high-technology environment coupled with comprehensive end-user support services. However, this conceptually simple approach requires a large commitment from those involved in the delivery of services. There are several end-user support facilities and projects at Avon that contribute to making data available to users. The most significant are:

- -Creative learning centres.
- -Atoms Avon's total office management system.
- -Atoms technology education room.
- -Avon Computer Store "solutions unlimited".
- -Proton a computer-based message system.
- -Graphics workstation and PC lab.

Avon is convinced that the key is to support users so that they can utilise currently available technology to access and manipulate the data they need.

Systems environment

The systems function in Avon is highly centralised and is organised to support specific business functions (finance, marketing, international operations, etc.). There are 30 data centres worldwide with IBM 4381s or System 36s or 38s, although there is a multivendor environment at the headquarters office, with equipment from IBM, Digital, Prime, Wang, and Burroughs. The personal computers used in Avon are mainly IBM, Eagle, DEC, and Apple. All systems activities are coordinated by the central systems function. Users are not charged directly for the services of the central function because Avon believes that doing this would inhibit the use of information technology.

We talked with Harry Mah, Avon's information centre manager. He has a staff of eleven and is charged with two major responsibilities:

- -Providing data to users.
- Providing education about technology in general and training for business software in particular.

The information centre uses a mainframe to provide a wide range of services, including a business planning system, database manager, statistics package, modelling, query language, and a fourthgeneration language. If the required data is not available on the centre's computer, the central systems function (after proper authorisation) will write Cobol programs to extract the data so that it can be made available to users via the information centre. Downloading data to PCs is permitted but requires the approval of the vice-presidents of the user department and of the central systems function. Uploading data is not permitted at all. Responsibility for any downloaded data passes to the user, who is required to have his or her diskettes passwordprotected. Data security is an important requirement, and these procedures are enforced.

Avon is firmly committed to the idea that educating and training users is the key to using data effectively and making it accessible, and a wide variety of PC software is available in the learning centres for training and testing purposes. The learning centres can be used by anyone at any time. Help is available during normal business hours. Formal classroom training is also conducted by Mr Mah's staff and covers PCs, word processing, and decision-support software. To date more than 1,000 Avon employees have been trained in-house in the use of personal computers.

Avon is committed to modernising its systems base. As well as developing new applications, it is restructuring existing systems to provide easier access to the data stored in them. Data is organised and controlled through a database administration function comprising a manager and four staff members.

A data dictionary is employed to define and control the data elements.

The Atoms pilot project

Atoms — Avon's total office management system aims to improve communications and administrative productivity by combining the use of minicomputer software for electronic mail with personal computers for word processing, financial spreadsheets, and personal databases. The system has been installed in the Rye facility over the past year, and there are approximately 320 managers, professionals, and secretaries using the system.

The Atoms system consists of two DEC VAX 11/750 minicomputers with All-In-One and DECmail software, together with a network of IBM, Eagle, Macintosh, and DECmate II personal computers. The two minicomputers at Rye can be accessed by dialling up from anywhere in the country. Atoms provides a full range of support services. A telephone hot-line is available for any questions or problems. Training is available in Rye and New York City for DECmate II word processing, All-In-One, DECmail, and personal computing.

According to Mr Mah, Atoms is highlighting the need for an integrated data processing capability within particular departments, and he believes that the

APPENDIX CASE HISTORIES

success of the pilot project has identified two general lessons about the best way in which to provide users with access to data. First, it indicates that, at least for Avon, departmental computing is the way to go. This approach provides user accountability for data and local access to it. Second, the approach eliminates the need for local area networks because minicomputers provide a basic star network for access, and they act as gateways to other departmental nodes or mainframes. Offline local personal computing is provided by using PCs as individual workstations.

Future developments

Mr Mah is surprised that the number of user requests for data is as low as it is. He feels that this is due partly to the lack of good software to enable users to access the data, and to the users' awareness of the need for additional software to meet their data demands.

Furthermore, he feels that the patience and understanding shown by users is due largely to the usereducation programme carried out by the information centre. He believes that dramatic improvements in data-access facilities will occur only when radically different tools are available. In particular, he expects artificial intelligence techniques to play a big part in accessing and analysing data.

Avon's senior management is well aware of the potential for using information technology as a competitive weapon. (The company is building a knowledge-based system with which it hopes to gain a competitive advantage.) Data is an important asset to Avon, and the integrity, administration, and security of data will continue to be important issues, and much of the responsibility will fall on the information centre.

THE CHASE MANHATTAN CORPORATION

The Chase Manhattan Corporation, operating through Chase Manhattan Bank and numerous other subsidiaries, has total assets of over \$85 billion and operates in more than 100 countries. Chase Manhattan provides a comprehensive range of financial services to corporations, individuals, financial institutions, and governments around the world. International operations accounted for 49 per cent of 1984 assets.

Over the past five years, Chase has invested half a billion dollars in information technology. These systems provide customers worldwide with access to Chase's sophisticated banking products. A global electronic banking unit was established in 1984; its purpose was to integrate fully Chase's electronic products and capabilities throughout the bank's wholesale business. We spoke with Egor Andronov, vice-president technical research, corporate management support. Discussions were confined to Chase's internal, noncommercial systems environment, mainly for proprietary reasons.

Data processing environment

Chase's data processing environment is "dedicated to the business unit". This enables the business units to make independent decisions about information processing activities and expenditures. Chase's operations are so extensive that most varieties of data processing exist somewhere in the organisation. Nevertheless, corporate management encourages a controlled multivendor environment, recommending IBM, DEC, and Wang as the primary vendors.

The corporate systems role is one of guidance and monitoring, and it provides support for the decentralised units through the systems management support organisation. Corporate support includes office automation, systems training, planning, technical research and 'end-user computing stores'. Chase operates these stores in New York, London, and Hong Kong. The stores provide advice and guidance to end users, rather than sell products.

Data access

Mr Andronov told us that Chase is reasonably satisfied with its use of existing tools for gaining access to internal and external data. There is no universal access method available that can be established as a standard, however. Chase has accepted that there is unlikely to be a breakthrough that will lead to a general-purpose standard dataaccess tool. End-user computing and data-access facilities will therefore evolve, rather than change dramatically.

At this point in the evolution, Chase is anxious to provide as much individual choice as possible. The corporate systems function does not specify a standard for large-scale database management systems for use in the business units, for example. These systems usually are important and large enough to require their own support staff in the business units, and there is little need to look to, or depend on, the corporate systems function for support services.

A corporate standard (Adabas) has been specified for small database systems, however. These smaller systems cannot justify their own dedicated support staff, and standardising on one product allows dataaccess and communications routines to be shared. Training and installation costs can also be shared among the business units.

Mr Andronov agreed that the introduction of personal computers caused users to demand greater access

to corporate and external data; without doubt, he said, the PC has raised users' awareness of the ability to transform data into information. He made a very perceptive observation: "The PC's impact was to organise data not previously organised. In effect, the data available in the corporate domain was expanded significantly by corporate staff inputting data into spreadsheets and sharing diskettes with each other, mainly for budgeting and planning purposes."

Interestingly, this user-initiated expansion of the available data has increased the pressure on the systems function to provide both access to the data stored in existing systems and facilities for manipulating it. The conventional view is that this demand can be satisfied by downloading data from mainframes to PCs. Mr Andronov takes a somewhat narrower view of downloading, however, and points out that basically it is a method for delivering reports electronically that previously were prepared and delivered on paper. Future software products may broaden this capability, but today's products do not.

Mr Andronov also pointed out the dangers inherent in restructuring existing systems to provide better data availability without also improving the related support procedures. Making the organisation's telephone directory available online may improve its accessibility, for example, but it may also highlight the inadequacies of the data if the support systems are not improved at the same time. People accept the fact that printed telephone directories are out of date by the time they are issued. But they would not tolerate an out-of-date electronic directory. This type of problem can make it very expensive to restructure existing systems to improve access to their data, which explains why this approach is not often used at Chase.

Fourth-generation languages are also seen by many as another vehicle for providing access to data. At Chase, however, the people using fourth-generation languages are systems professionals, not users. Nevertheless, the users benefit from the faster systems development and modification resulting from the use of fourth-generation languages.

Data security is a major concern at Chase Manhattan. In theory, users could be provided with access to data stored anywhere in the world via the organisation's extensive worldwide electronic network. Due to the sensitive and confidential nature of customers' financial data, data encryption and sophisticated security routines are mandatory. In Chase's large IBM systems environment, IBM's Resource Access Control Facility (RACF) package is widely used, with acceptable results.

Control of data, in terms of rules about who can change and access data, is a critical aspect of Chase's business. In many of the systems, this control must go beyond the application, file, and record level to the individual field level. The emphasis on data security heightens the general awareness of the importance of data. In Chase's decentralised environment, users take an active role through their own security administrators. The end result is that the users, rather than systems staff, generally end up being responsible for the data.

The emphasis on data security explains to some extent why Chase is basically satisfied with its use of existing tools for user access to data. The more user-friendly and general-purpose tools, by their very nature, do not provide the necessary data security features. As a result, most access to data at Chase Manhattan is monitored closely and is controlled by proprietary, internally developed software and systems.

Possible future developments

Chase is investigating many different possible approaches to improving access to data. Mr Andronov feels there are three areas in particular that show promise.

The first area concerns the development of vendorindependent standards. De facto, vendor-specific standards may have been acceptable (or even necessary) in the initial stages of data processing, but today such standards are counterproductive and divert effort from the real business problems. Mr Andronov is encouraged by the progress being made by the International Standards Organisation (ISO). The key, he believes, is international cooperation, because modern communications technology has, in effect, removed national barriers. He advocates participation by user organisations in the work of international organisations like the ISO.

The second area concerns the use of reuseable code in the public domain, which would permit common routines to be used by many applications, thereby providing a basis for compatibility and standardisation. Standard access methods could be based on reuseable code developed by vendors whenever they need to change software because of hardware modifications. The same technique could be used by software vendors when they upgrade or change offthe-shelf packages.

The third area identified by Mr Andronov is the one that he feels has the greatest potential for improving user access to data. It is concerned with expert systems and the use of nonprocedural languages. He believes that hardware will be provided with 'intelligence' based on expert system techniques. This intelligence will enable the hardware to interpret input received from the users and automatically access and retrieve the appropriate data. Imbedding these types of expert systems into application gateways, whether to other computers or databases, reduces the risk of the wrong data being retrieved. This approach also has the advantage of presenting users with a common and consistent interface. However, it is not yet clear when such expert systems will be developed and commercially available.

Chase Manhattan is already working on systems in this area, however, because it requires a powerful easy-to-use interface for its retail banking services. In fact, the need to serve the public will hasten the development of an expert system interface because, for Chase, it is a question of competitiveness, rather than convenience. For this reason, Chase is eager to fund this critical activity. But Mr Andronov warned that other organisations should not expect to be able to purchase the resulting systems products from Chase. The results of this development effort are likely to be considered as proprietary by Chase.

ISTAT

This case history describes how one organisation has fundamentally reappraised the way in which it makes data available to users. It also illustrates that the information centre approach is not always successful.

The Istituto Centrale di Statistica (ISTAT) is Italy's official statistical bureau. The organisation has 18 regional offices and employs about 3,000 people. ISTAT's main activity is collecting, processing, and distributing data relating to population, social issues, industrial statistics, domestic economic factors, and foreign trade statistics. ISTAT is also responsible for processing the data collected by the census carried out every ten years. Statistical data is distributed mainly through publications, but it is also available through databases, magnetic tapes, and so forth.

ISTAT's mission is to make statistical information more readily available to all Italians as part of their cultural inheritance. To achieve this aim, it has reappraised the existing systems and has developed a set of strategic objectives. These objectives have been constructed with the help of IBM consultants using IBM's Business Systems Planning (BSP) methodology.

Data processing in ISTAT

Dr F Marozza, Dirigente Superiore of ISTAT, explained that data processing is the responsibility of one of the three major divisions of ISTAT. Computers have been used extensively in the organisation since the 1950s for collecting data and distributing statistical information.

Future systems requirements will be determined by the information system strategy that resulted from the BSP exercise. A major component of the systems strategy will be the specification for database and network development. To enable the appropriate developments to take place, the migration from batch to interactive systems is to be speeded up. This change will allow three levels of database to be maintained:

- A database containing macro-level statistics oriented towards the external network.
- A database containing micro-level statistics mainly for internal use.
- A database containing metadata about the other two databases.

The overall structure of ISTAT's future data processing operation, based on the three levels of databases, is shown in Figure A.1. The figure should be read from the bottom upwards, with the operation envisaged as a triangular shape having a wide base (the enormous amount of raw data) and a summit made up of the relatively small amount of statistical data. The figure also shows the principal functions carried out by the data processing operation.

The whole operation is based on the data suppliers who provide a description of the state of, and changes in, the real world. The apex of the operation points towards the users — the distributors of statistical information.

Making data available to users

Until late 1982, ISTAT's systems were based solely on IBM's IMS. Three years ago, however, Adabas/Natural was introduced, and by the end of 1984 all further development work under IMS had ceased. This change was made because ISTAT believed that its aims of increasing both the amount of data input directly by users and the availability of data would be furthered by using Adabas/Natural. There are now 150 people in the central departments who input data directly, and other users in various public bodies — the Senate, Supreme Court, Deputy Prime Minister's office, universities, etc.

More than 60 of ISTAT's staff now use Adabas, and Dr Marozza told us that their productivity had increased as a result. The relational structure of Adabas has proved particularly useful in carrying out ad hoc modifications to existing systems and in developing new applications. The use of Natural has also made it possible to shift the responsibility for data capture and retrieval away from the central data processing function to the user departments — and eventually to the external users themselves. Dr Marozza believes that ISTAT now has the basic infrastructure in place to allow users direct access to the database so they can retrieve and manipulate the data they require. As a consequence, the need for someone requiring statistical information to



search through volumes of printed data is gradually diminishing.

To improve the direct links to the central departments for key administrative users.

ISTAT's medium-term plans for satisfying the growing user demand for data have three distinct objectives:

 To extend the indirect links to the provincial users both for input and enquiry-based systems. To distribute copies of the database via the PTT network so that the general public may access the data via microcomputers and other end-user devices.

Problems with the information centre

Towards the end of 1983, ISTAT attempted to establish an information centre. The aim was to support the use of end-user tools such as Speakeasy, SAS, APL, and GDDM. The information centre was a failure and was abandoned as a mechanism for introducing end-user facilities within ISTAT. The main problems were that the role of the centre had not been adequately defined, nor had the responsibilities and authority of the staff. The main reasons for the failure were that:

- The same hardware was used both for end-user development and for operational systems.
- There was a lack of organisational structure in the information centre which meant that individuals' roles and responsibilities were not clearly defined.
- There were no clear objectives that mirrored the business objectives.

Dr Marozza now believes that information centres are more applicable in commercial organisations and are not immediately suitable for geographically dispersed public bodies whose diverse range of users do not share a common business aim. The fact that many of the potential users of the information centre did not work for ISTAT made the problems of control even more acute.

Nevertheless, the use of SAS, Speakeasy, APL, and GDDM has had a dramatic impact by reducing the time required to publish the statistics. In addition, the quality of the graphics facilities available with this type of tool has improved considerably the presentation of the statistics. At present, these tools are used by ISTAT's own systems staff. The next step is to make these advanced packages available to the user base.

PHILIPS

This case history illustrates how one company — Philips, the large Dutch-based multinational electronics company — has developed its organisation structures to provide support for a range of development environments and system building tools.

Philips' main business activities are lighting, home electronics, domestic appliances, and professional and industrial products. The company is structured as a matrix organisation with product groups running through each national operation. The company employs about 70,000 people in Holland, 40,000 of whom are based in Eindhoven.

Data processing environment

Data processing functions are organised into several national information systems and automation (ISA) groups. In the Netherlands, the ISA group is called Nederlandse Philips Bedrijven B.V. — Electronic Data Processing (EDP-NPB). The group is composed of a series of functions split approximately between traditional data processing activities (the Data Processing Service) and user-oriented activities (Personal Computers and Information Services).

Personal computing environment

The Personal Computers and Information Services (PCIS) group was formed in 1985 to consolidate the activities of three older EDP groups (Information Services, Micro Applications and Information, and Workstations). The group's role is to:

- -Advise users on personal computing.
- -Promote the use of personal computing facilities.
- Supply products and services supporting the use of personal computing (mainframe and PC-based).
- -Supply PC hardware and software.
- Integrate access to general data and help users to gain access to specialised data.
- Provide guidance and support to those wishing to establish local PCIS organisations (information centres), and coordinate the activities of the local organisations.
- -Provide education and training.

PCIS is organised into seven sections:

- User training (developing training facilities for end users and for information centre personnel).
- Marketing (product management, sales, online help desk).
- -Consultancy.
- Data services and communications (including data storage and retrieval).
- Technical automation.
- Decision support and management information (customising and developing systems, interface management).
- Office systems (supplying hardware and software, providing advice and support).

We met with Mr Mager (Information Centre Manager), who told us that there are now a total of 50 staff involved in supporting business computing in Philips, compared with only 15 in 1982. This increase reflects the growth in personal computing at Philips. Business users are now much less dependent on the Data Processing Service (DPS) for the provision of personal computing facilities although PCIS may in the future purchase computer time from DPS and resell it to user departments via the information centre interfaces.

User access to data

A consequence of the increasing use of personal computing packages and facilities has been the emergence of 'unofficial' information officers. These people have been nominated by user groups as their representatives and, as such, have no direct links with the data processing function. Their role is to help their colleagues make the best use of their personal computing facilities and to keep themselves up to date with information centre and IT developments. There are now more than 50 of these unofficial information officers.

Mr Mager noted that they had appeared in a semiplanned manner, not as part of a formalised strategy. They are regarded as the pioneers of a structured local support organisation (or local information centres).

Philips has developed its own in-house interface manager (called Easy) that allows users to log on and gain transparent access to the packages available via the central information centre. Having logged on, the user is presented with four windows (or choices) — mail, news and information, query, and edit data and systems. Mr Mager regards Easy as primarily for 'button pushers'.

At present, the information centre services are based on a partition of an IBM 3084Q (approximately equivalent to a 3081K). Up to 120 concurrent users can be accommodated, with a workspace allocation of 500k bytes per user, although it is planned to increase the workspace allocation to 1.5M bytes per person.

The increased storage available via the information centre is seen as a way of reducing the emphasis on PC processing. PCIS believes that this shift is necessary because of the better data security and control facilities available on mainframes.

In Philips, therefore, the emphasis on making data available to business users will be on access to mainframe data, not microcomputer data. In fact, the representatives from DPS and PCIS that we spoke with told us that, in their view, positive action was required to prevent data held in standalone microcomputer systems from being made available to other users.

RHONE-POULENC

This case history describes how one organisation has established and developed an information centre to support end-user computing and user access to data. In doing so, this organisation has established an effective organisational infrastructure for personal computing. We believe its success is attributable largely to the central data processing function adopting a coordinating, rather than a controlling, role.

Rhône-Poulenc is a French-based chemical company. It employs more than 80,000 people worldwide and is represented in 60 countries. The company is involved in many aspects of chemical fabrication, covering base and speciality chemicals, health products, textiles, materials for electronics, magnetic media, and packaging. In 1984, the company's turnover was 5 billion francs.

End-user computing at Rhône-Poulenc

Rhône-Poulenc has about 700 people in France actively involved with the data processing function. Approximately a third work in a centralised department, and the remainder are employed in the divisions. There is a large SNA-based network with approximately 1,200 terminals. Two large IBM mainframes (3081Ks) are located at twin computer centres at Lyon and Courbevoie, and a 3033 MP (running under VM/CMS) has recently been dedicated to end-user computing. This requirement for a central host machine for end-user computing has evolved over the last ten years. To understand the reasons for this, it is necessary to step back to the mid-1970s and trace the major events in the growth of end-user computing at Rhône-Poulenc.

In 1974, the IBM GIS product was providing onrequest report-writing facilities for Rhône-Poulenc specialists (mainly in the finance area). Even at this early stage, users were expressing a strong demand for their own facilities and increasingly requiring access to central data.

This arrangement continued to be satisfactory until 1977, when the central team recommended the introduction of APL, mainly as a scientific tool, but also as a query-language tool. The take-up of APL was user-driven, however, with occasional 'champions' emerging from the scientific departments. Even so, the normal process at this stage was for the user to refer a request to a professional programmer.

In 1979, the first major user-driven application was installed. This was a personnel system for manpower planning and strategic personnel planning. The application was written in APL, and it extracted and analysed data from the monthly IMS-based payroll files. It provided access to the records of 40,000 employees and was used directly by the 300 staff in the personnel function. In the same year, an information centre was established to handle user problems and requests. The centre initially was staffed by six people recruited from user departments and the data processing function.

From 1979, there was steady progress in end-user computing, essentially using APL both for prototyping and for building simple enquiry systems. By early 1984, the steady progress became unacceptable because users were initiating more complex, unstructured requests. As a consequence, one of the major divisions (health products) led an investigation into the requirements for decision-support systems, and set about choosing and acquiring appropriate software. The users carrying out the investigation were supported by the central technical specialists when they required assistance.

Three decision-support products were shortlisted — Nomad, Wizard, and Express — and benchmark tests were carried out on each system. Nomad was selected because it provided a more general-purpose tool than the other products. The main selection criteria were ability to handle the type of data structures required and effectiveness in formulating ad hoc enquiries.

Nomad has now been in use for over a year. The main benefit so far has been in reducing some of the applications backlog and in facilitating systems maintenance. In addition, users have been able to reduce the number of periodic reports (even though these were thought to be essential), and have replaced them with more on-request reports. Nomad is also able to handle more complex data structures and data manipulation requirements than APL.

The recent change in operating system environment from TSO to VM has facilitated the use of Nomad, and future growth in its use has been anticipated by the installation of the IBM 3033 MP.

The central team has not imposed complicated, restrictive rules on the use of end-user computing facilities. Instead, it stresses the need for a few unambiguous general rules. In the case of Rhône-Poulenc these are:

- —The information centre databases accessed by users are segregated from the production databases. All access to, and transfers from, the production databases are controlled by the central team.
- Tight control is exercised over the definition of the meaning of data. The systems staff at Rhône-Poulenc are convinced that this area requires a coordinated consistent approach throughout the organisation. This is achieved by encouraging the appointment of a data administrator in each of the divisions.

These rules resulted in an informal strategy for enduser computing. Because of the highly autonomous business environment in Rhône-Poulenc, the central data processing function cannot enforce a mandatory policy. Instead, the rules are disseminated via a policy group made up of senior staff from the major divisions. This group meets six or seven times a year and plays an active and visible role in data processing decisions.

Rhône-Poulenc summarises its approach to end-user computing in the following terms:

- Providing initial intensive support for users, but then pulling back once they have reached a sufficient level of competence.
- Shielding users from the more technical aspects of personal computing: discussions with suppliers and the mainstream development team are handled by the technical team at the information centre.

Future plans

Future plans for end-user computing at Rhône-Poulenc include:

- Providing different types of facilities for experienced and inexperienced users.
- Introducing new user-friendly tools, especially for the more inexperienced users.
- -Introducing mainframe-based spreadsheet facilities, as appropriate tools become available.

The overall aim is to provide an end-user computing environment where users can make the rapid progress that they are now demanding.

SURREY COUNTY COUNCIL

Surrey County Council has used a modern system building tool (Mapper) to promote user-driven development of data enquiry systems as a complement to the traditional systems development process.

The council is responsible for providing local services within its administrative area in Southeast England. It has about 26,500 staff and an annual turnover of about £400 million. The permanent staff of the council report to the chief officers of the service departments. In addition, the Head of Manpower Services and the Head of Computer Services report directly to the Chief Executive.

The Policy Committee is ultimately responsible for the council's computer services, but it has delegated this role to the Computer/IT Subcommittee. An IT Steering Group reports to this subcommittee and has the task of assigning workload priorities.

Data processing environment

There are about 155 computer services staff working in operations, development, technical support, and administration. The recent emphasis has been away from database systems and towards a clearer definition of data items and data structures. As part of this change in emphasis, the computer services department has been restructured into three groups:

- End-user support (including an embryonic information centre).
- Development project teams and operations.
- -Support for existing systems.

The new organisation reflects the increasing importance of end-user computing, office systems, and communications.

The growing emphasis on end-user computing is reducing the need for large-scale development projects. The emphasis is changing from systems that reduce the costs of running the business to the more extensive use of existing data for improving managerial effectiveness and providing information for policy making. The computer services department believes that its long-term role will be to provide information, but it cannot quantify the benefits to the organisation of such a changed role. Consultants are currently helping the council to formulate an information technology strategy, which should clarify the future role of the department.

The use of Mapper to support end-user computing

The change in emphasis and the new organisation structure have been accompanied by the introduction and successful use of Mapper. The use of Mapper will increase in the short to medium term as part of a planned move to specialised databases and related query languages. These moves will be accompanied by a reduction in the use of Cobol.

Surrey County Council sees Mapper essentially as a tool for prototyping and initial development by users. Using Mapper in this way has enabled the rapid development of small systems where the emphasis is on data retrieval, in particular for access to existing Cobol files (although the files do have to be converted before they can be accessed by Mapper).

Despite the difficulties of using Mapper to access conventional files, the council has created several special applications with the product, particularly for the education department. Mapper allows payroll information to be accessed and manipulated by this department, for example. In effect, the system enables staff in the education department to use the applications as a manual filing system.

These applications have been built by computer services staff, although the plan is to train some users so they can develop their own applications.

According to Surrey County Council, Mapper is appropriate for developing small applications that access data files constructed specifically for the applications, and where the data-retrieval requirements are ill-defined. It is also suitable for the rapid development of prototype systems, and the auto-run facility (for repeating a program) has been found to be a useful feature. Furthermore, users can be trained to write their own Mapper routines.

Mapper is unsuitable where large volumes of data have to be transferred from existing systems. Moreover, it should not be used to create larger systems with complex processing requirements and large volumes of data that need to be updated frequently. For this type of system, the inherent inefficiency of programs generated by Mapper would adversely affect the operational performance of the system.

Staff in the engineering services department have been actively involved in designing and implementing a direct-labour accounting application. The whole application was conceived and implemented in less than six months. This application processes 50,000 records relating to jobs-in-hand, and it is used by two departments at four separate locations. Engineering services staff helped to develop enquiries, reports, screen formats, and some processing routines so that estimates for standard jobs can be derived from the rates held in the Mapper databases. Nevertheless, computer services staff retained overall responsibility for data integrity and security.

Surrey County Council has appointed a Mapper coordinator, who provides the technical and business interface between the operational support staff and the end users. At first, this arrangement created some difficulties because the operational support staff felt they should be responsible for the operational use of Mapper.

BIBLIOGRAPHY

Canning Publications Inc. Coping with end-user computing. EDP Analyzer, Volume 22, Number 2, February 1984.

Cash, J.T. and Konsynski, B.R. IS redraws competitive boundaries. Harvard Business Review, Volume 63, Number 2, March-April 1985.

De, P. and Sen, A. A new methodology for database requirement analysis. MIS Quarterly, Volume 8, Number 3, September 1984.

Elsaw, O.A. Implementation by cultural infusion. MIS Quarterly, Volume 9, Number 2, June 1985.

Gandy, M. Microcomputer communications. Manchester: NCC, 1985.

Gerrity, T.P. and Rockart, J.F. Managing end-user computing in the information era. Centre for Information Systems Research Working Paper Number 120. Cambridge, Massachusetts: Sloan School of Management, October 1984.

Hammond, L.W. Management considerations for an information centre. IBM Systems Journal, Volume 21, Number 2, 1982.

Henderson, J.C. and Treacy, M.E. Managing end-user computing. Centre for Information Systems Research Working Paper Number 114. Cambridge, Massachusetts: Sloan School of Management, May 1984.

Henderson, J.C. and Treacy, M.E. Managing end-user computing for competitive advantage. Sloan Management Review, Volume 27, Number 2, Winter 1986. ICLCUA (UK). CAFS in action. ICLCUA (UK) Working Paper, November 1985.

Information Centre Managers Forum. UK ICMF survey. ICMF, 1984.

Kahn, B.K. An environmentally dependent framework for data dictionary systems. MIS Quarterly, Volume 9, Number 3, September 1985.

Kiechel, W. How executives think. Fortune International, Volume 111, Number 3, 4 February 1985.

McFertan, F.W. and McKenny, J.L. The information archipelago - governing the new work. Harvard Business Review, Volume 61, Number 3, July-August 1983.

Martin, J. An information systems manifesto. Englewood Cliffs, New Jersey: Prentice-Hall, 1984.

Meiklejohn, I. Should you be better informed?. Business Computing and Communication, May 1985.

Smith, C.L. A PC support centre. Datamation, Volume 31, Number 5, March 1985.

Wheelwright, G. Micro-mainframe links. Communications/Communications International, November 1984.

Xephon Buyers' Guide. PC/Mainframe links. Newbury, Berks: Xephon Technology Transfer Ltd., 1984.

THE BUTLER COX FOUNDATION

Butler Cox

Butler Cox is an independent management consultancy and research organisation, specialising in the application of information technology within commerce, government and industry. The company offers a wide range of services both to suppliers and users of this technology. The Butler Cox Foundation is a service operated by Butler Cox on behalf of subscribing members.

Objectives of The Foundation

The Butler Cox Foundation sets out to study on behalf of subscribing members the opportunities and possible threats arising from developments in the field of information systems.

The Foundation not only provides access to an extensive and coherent programme of continuous research, it also provides an opportunity for widespread exchange of experience and views between its members.

Membership of The Foundation

The majority of organisations participating in the Butler Cox Foundation are large organisations seeking to exploit to the full the most recent developments in information systems technology. An important minority of the membership is formed by suppliers of the technology. The membership is international with participants from Australia, Belgium, France, Italy, the Netherlands, Sweden, Switzerland, the United Kingdom and elsewhere.

The Foundation research programme

The research programme is planned jointly by Butler Cox and by the member organisations. Half of the research topics are selected by Butler Cox and half by preferences expressed by the membership. Each year a short list of topics is circulated for consideration by the members. Member organisations rank the topics according to their own requirements and as a result of this process, members' preferences are determined.

Before each research project starts there is a further opportunity for members to influence the direction of the research. A detailed description of the project defining its scope and the issues to be addressed is sent to all members for comment.

The report series

The Foundation publishes six reports each year. The reports are intended to be read primarily by senior and middle managers who are concerned with the planning of information systems. They are, however, written in a style that makes them suitable to be read both by line managers and functional managers. The reports concentrate on defining key management issues and on offering advice and guidance on how and when to address those issues.

Previous reports

Developments in Data Networks* Display Word Processors* Terminal Compatibility' Trends in Office Automation Technologies* The Convergence of Technologies Viewdata* Public Data Services* Project Management The Selection of a Computerised PABX* Public On-line Information Retrieval Services* Improving Systems' Productivity Trends in Database Management Systems* The Trends in Data Processing Costs The Changing Equipment Market* Management Services and the Microprocessor The Role of the Mainframe Computer in the 1980s* **Electronic Mail** Distributed Processing: Management Issues Office Systems Strategy The Interface Between People and Equipment Corporate Communications Networks Applications Packages **Communicating Terminals** Investment in Systems System Development Methods Trends in Voice Communication Systems Developments in Videotex User Experience with Data Networks Implementing Office Systems End-User Computing A Director's Guide to Information Technology Data Management Managing Operational Computer Services Strategic Systems Planning Multifunction Equipment Cost-effective Systems Development and Maintenance Expert Systems Selecting Local Network Facilities Trends in Information Technology Presenting Information to Managers Managing the Human Aspects of Change Value Added Network Services Managing the Microcomputer in Business Office Systems: Applications and Organisational Impact **Building Quality Systems** Network Architectures for Interconnecting Systems The Effective Use of System Building Tools Measuring the Performance of the Information Systems Function Developing and Implementing a Systems Strategy *These reports have been superseded.

Future reports

Threats to Computer Systems Security Organising the Information Systems Function Using IT to Improve Decision Making Integrated Telecommunications Networks Planning for the Future Corporate Data Centre The Effect of IT on Corporate Organisational Structure Butler Cox & Partners Limited Butler Cox House, 12 Bloomsbury Square, London WC1A 2LL, England ☎ + 44 1 831 0101, Telex 8813717 BUTCOX G

France Butler Cox SARL Tour Akzo, 164 Rue Ambroise Croizat, 93204 St Denis-Cedex 1, France **2**(1) 4820.61.64, Telex 630336 AKZOPLA

> The Netherlands Butler Cox BV Burg Hogguerstraat 791 1064 EB Amsterdam ☎ (20) 139955, Telex 12289

United States of America Butler Cox Inc. 115 East 57th Street, New York, NY 10022, USA 26 (212) 486 1760

Australia Mr John Cooper Business House Systems Australia Level 28, 20 Bond Street, Sydney, NSW 2000 ☎ (02) 237 3232, Telex 22246

Italy

SISDO^{BDA} 20123 Milano – Via Caradosso 7 – Italy 2498 4651, Telex SISBDA 350309

The Nordic Region Statskonsult AB Stortorget 9, S-21122 Malmo, Sweden 2 46-401 03 040, Telex 127 54 SINTAB