Network Management

BUTLER COX FOUNDATION

Research Report 65, August 1988

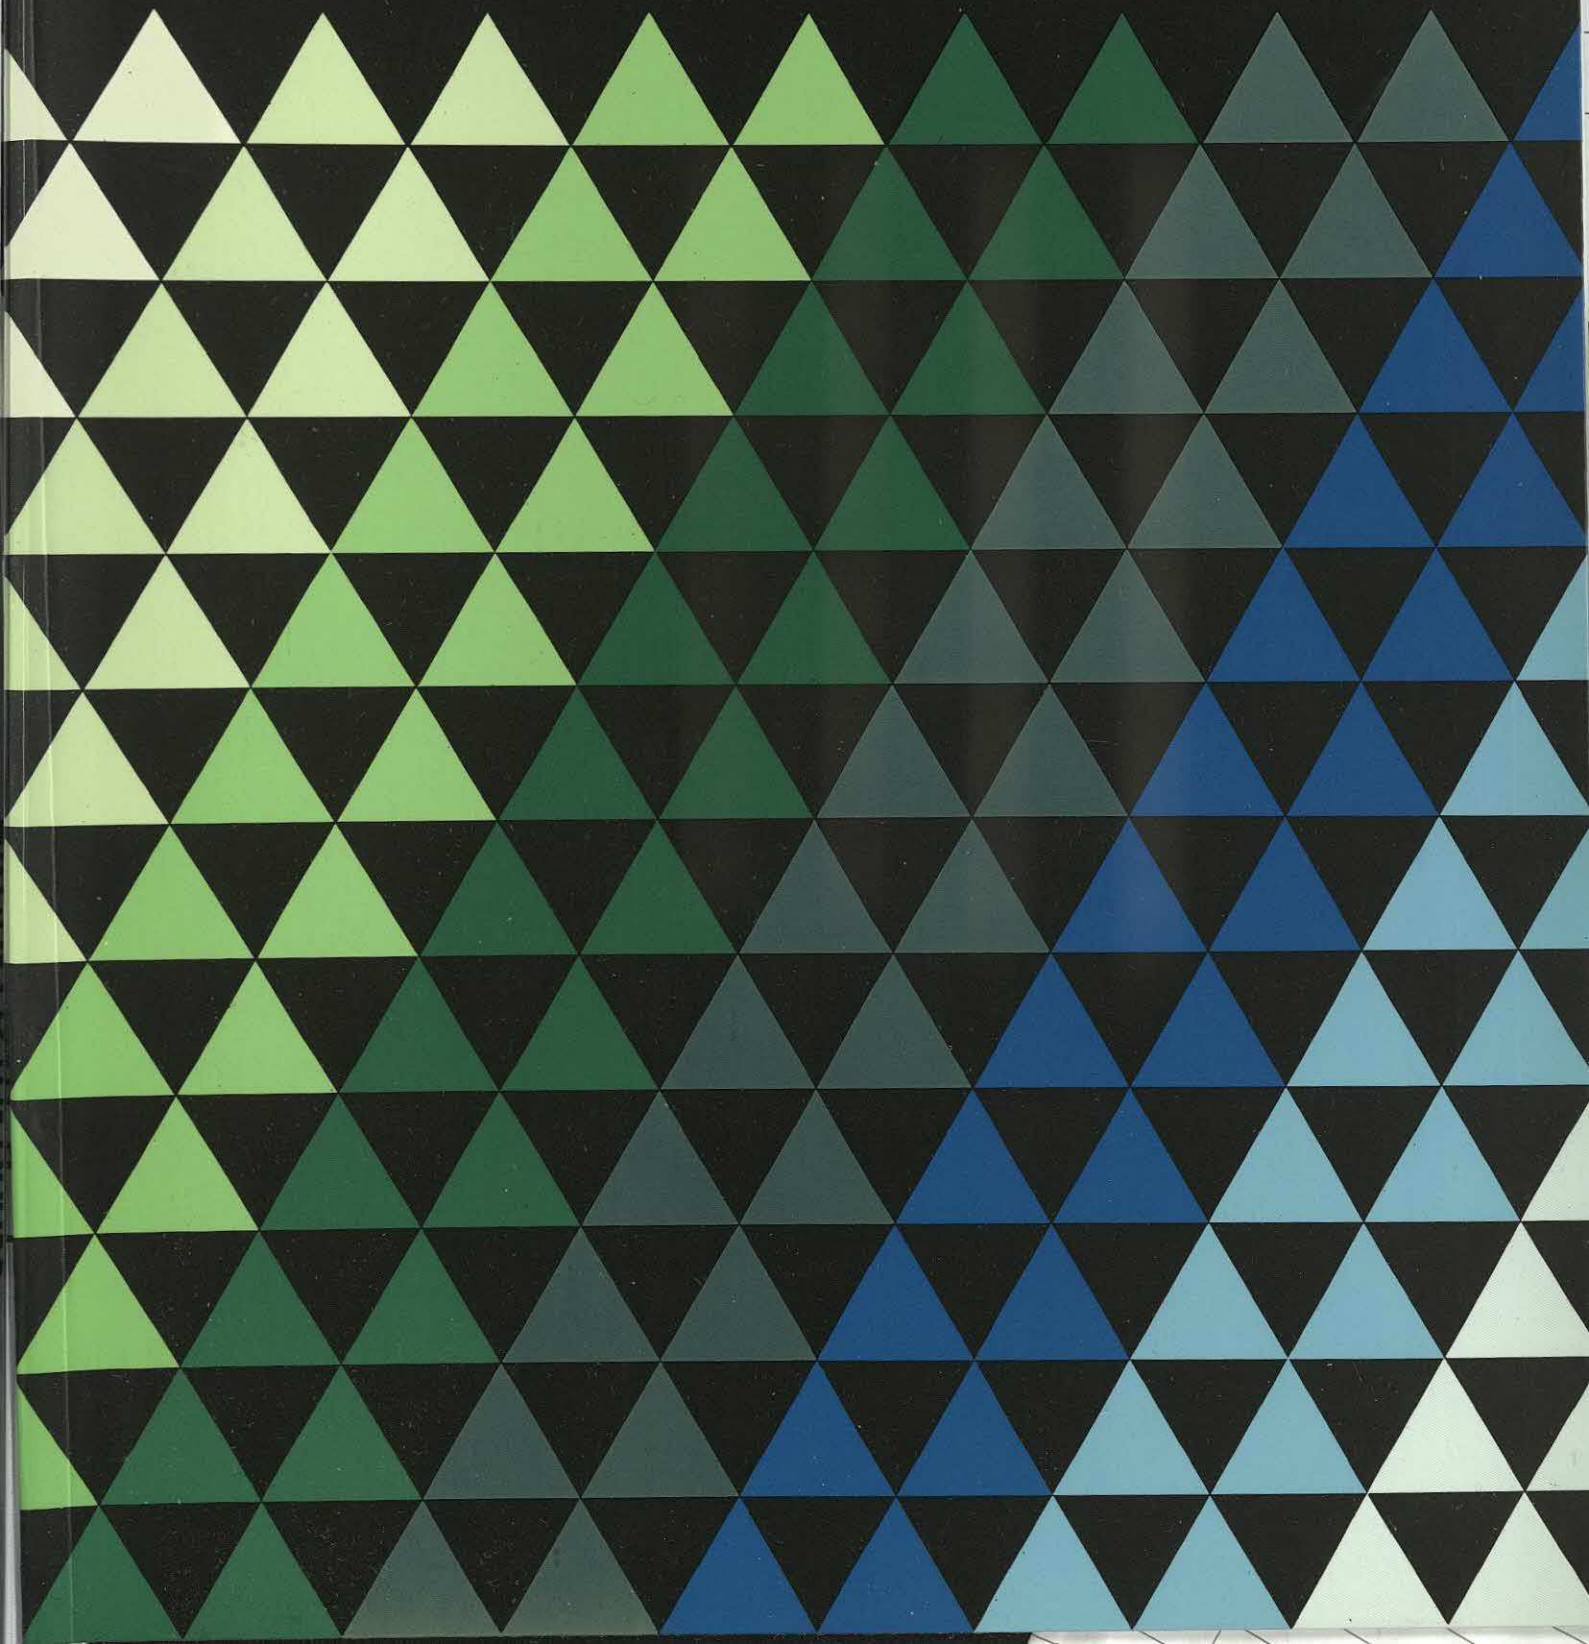

BUTLER COX FOUNDATION

Network Management

Research Report 65, August 1988

Butler Cox & Partners Limited

LONDON
AMSTERDAM MUNICH NEW YORK PARIS

Published by Butler Cox & Partners Limited
Butler Cox House
12 Bloomsbury Square
London WC1A 2LL
England

Availability of reports

Members of the Butler Cox Foundation receive three copies of each report upon publication; additional copies and copies of earlier reports may be purchased by members from Butler Cox.

Photoset and printed in Great Britain by Flexiprint Ltd., Lancing, Sussex.

BUTLER COX FOUNDATION

Network Management

Research Report 65, August 1988

Contents

Management Summary

A Management Summary of this report has been published separately and distributed to all Foundation members. Additional copies of the Management Summary are available from Butler Cox

Chapter 1

The growing importance of network management

The use of corporate telecommunications networks, particularly data networks, is growing rapidly. Many organisations would experience severe difficulties if their networks were not available for an extended period of time. Telecommunications managers (or network managers, as we refer to them in this report) must therefore ensure that their networks provide a reliable, efficient service and are capable of being extended both geographically and to meet demands for new types of service.

The rapid growth in networks, both physically and in the number of services they are used for, means that networks are becoming more complex technically, with components from many different suppliers. In turn, this means that it is increasingly difficult to identify and fix faults when they do occur. Moreover, the problem is compounded by the scarcity of skilled telecommunications staff and the lack of suitable network-management tools. Many organisations regularly experience great difficulties in getting their suppliers to fix faults promptly, or even to admit that a fault exists.

Two examples illustrate the scope of the difficulties that some organisations now face. Our analysis of fault-log data showed that one company in the United Kingdom was experiencing failures on its digital circuits four times more often than other UK companies. The high failure rate was caused by a persistent synchronisation problem. The network-management staff spent months working with several of their suppliers before the problem was solved. The second example concerns a French company that went through a period during 1987 when it experienced an average of five failures a day in its front-end processor. The overall availability of the network was reduced to 80 per cent, but the company thought that the network-management tools available from its computer supplier were too expensive. These tools would almost certainly have helped to identify and solve the faults.

Some organisations have even reached the stage where they put off making changes to their networks. Even though the changes are desirable, they know that by introducing the changes, they stand a very high chance of disrupting the existing network and services.

In the past, the network-management role has been concerned primarily with day-to-day activities — principally with fixing faults on the network. Obviously, ensuring that the network is fully operational is an essential part of the network-management function. However, it is no longer adequate to perceive network management as being concerned just with network operations. If the management of corporate networks is to be improved, it is necessary to widen the scope of the network-management function so that networks are managed in a way similar to other business functions. The purpose of this report is to show how the management of corporate networks can be improved.

THE NEED FOR IMPROVEMENT

Before exploring how network management can be improved, it is first necessary to understand why there is a need to improve and broaden the scope of network management. During our research, we found many network managers who now devote much of their time to activities not directly concerned with the immediate and urgent tasks of keeping their networks running as smoothly as possible. They told us they are approaching a level of activity with which their existing staff and tools cannot cope. They have to find better ways of planning and monitoring networks so that the level of resources required for handling faults and implementing changes can be contained.

There are four main factors creating the need for improved network management:

- The growth of data networking.
- The emergence of networks as a major component of the information technology (IT) infrastructure.
- The increase in user expectations.
- The lack of consistent network-management policies.

Each of these factors is described in more detail below.

GROWTH OF DATA NETWORKING

The tremendous growth in data communications experienced by most organisations is making it much more difficult to manage corporate data networks. Data networks are not only growing in their importance to the organisation but also in their size and complexity. The number of types of equipment connected to networks is increasing, as is the number of links between public and private networks.

Recent Foundation Reports on Communications Infrastructure for Buildings (Report 62) and The Future of the Personal Workstation (Report 63) have highlighted the continuing growth in workstation penetration. Many of these workstations will need to access corporate wide-area data networks. During the research for this report we asked about the growth in the number of users of wide-area data networks and local area networks experienced during the past year and expected during the next year. The answers are summarised in Figure 1.1. The responses indicate that network managers will often have to cope with growth in excess of 20 per cent a year and in some cases more than 50 per cent. Growth rates such as these imply the need for continual addition of network capacity and for high levels of user support and training.

In addition, improvements in technology have enabled the provision of network services that are more cost-effective, but in some cases, this has led to increased complexity. Thus, wide-area data networks must now:

- Provide transparent connections to multiple computers from different suppliers.
- Support protocol conversion.
- Interconnect with local area networks.
- Share high-speed transmission links with voice networks.

This increased complexity makes the operation of data networks more difficult to understand and it means that it takes longer to diagnose problems when they do occur.

As data networks increase in size and complexity, the number of different types of equipment used in networks also increases. Large networks can include analogue circuits (and digital circuits in some countries), modems, multiplexors, protocol converters, packet switches, packet assemblers/disassemblers, local area network gateways, and communications controllers, as well as connections to a variety of types of workstation and computer. The variety of network components means that network managers have to deal with

Figure 1.1 Anticipated growth in the number of users of wide-area data networks and local area networks

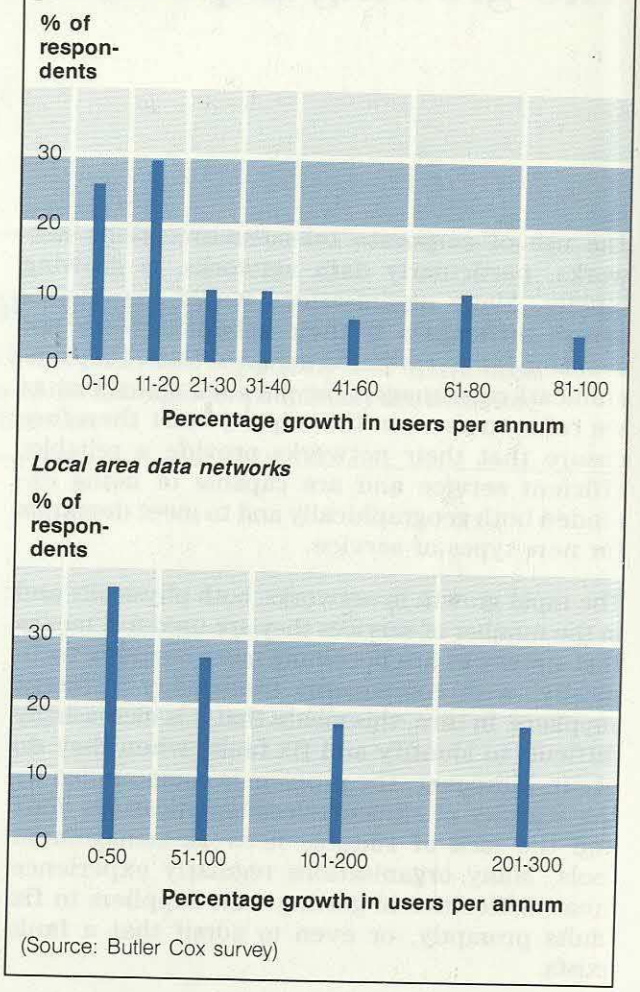

(Source: Butler Cox survey)

a large number of different suppliers, each with different contracts, standard maintenance terms, and ways of conducting business. A survey conducted in the United States by Northern Telecom revealed that, on average, corporate networks use services or equipment from 20 suppliers. In Europe and elsewhere the number of suppliers is likely to be less because the number of public network providers is more restricted, but would probably still be up to 10 for large organisations. Dealing with this number of suppliers, and ensuring that all their equipment works together, can be very time-consuming.

Corporate private data networks (where transmission links are dedicated to the use of the organisation concerned, although the links are usually leased from a public network provider) are also increasingly being connected to public networks. Examples include:

- Connections to value-added networks for purposes such as electronic data interchange (EDI) with suppliers and customers.

– Connections to the public telephone network for access to the corporate network by salespeople, home workers, and customers.

The organisation's network-management function must control the interface between the public and private networks so that users are unaware, as far as possible, of which type of network they are using.

As the number of data network users grows, the importance to the organisation of reliable network performance increases rapidly. For many years, organisations such as banks and airlines have recognised that networks are part of their product-delivery systems. Other organisations, even those that use their networks to carry only administrative traffic, are finding that they cannot conduct their business if the network is inoperative. They cannot revert to the paper-based administrative systems used prior to the introduction of networks and computers. The majority of network managers in our research described their data networks as "vital" or "essential" to the business.

NETWORKS ARE A MAJOR COMPONENT OF THE INFORMATION TECHNOLOGY INFRASTRUCTURE

Five to ten years ago, systems departments were concerned mainly with managing large mainframe computers and the applications that ran on them. As shown in Figure 1.2, the number of external communications links was small because of the small population of remote terminals. The network was relatively straightforward to manage because it was small and all the links terminated in one place.

The situation today is usually more like that shown in Figure 1.3. Physically, there is still one network, although it may now cater for several network protocols and allow access to a variety of computer systems, ranging from mainframes to microcomputers interlinked by a local area network. The number of transmission links in the network has also increased substantially. This means that the proportion of the total IT budget accounted for by network costs has increased considerably because the unit costs of communications have not decreased nearly as rapidly as the cost of computing power.

Figure 1.2 In the past, the mainframe was the central component of the IT infrastructure

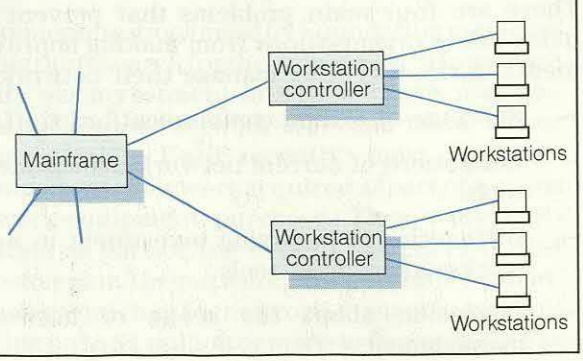

Figure 1.3 depicts the network as a core element of the IT infrastructure. The network is the enabling mechanism that provides users with access to information held on any one of a number of computer systems. For some years now, systems departments have recognised that information is a vital business asset and have developed procedures for managing the storage and retrieval of the information. However, the means of accessing and transferring the information also needs to be managed. The management of networks should therefore now be accorded the same importance by the systems department as the management of computer operations or applications development.

INCREASING USER EXPECTATIONS

The increasing number of business applications that use voice or data networks, or both, means that the use of networks is increasingly becoming an integral part of many people's jobs. Naturally, network users expect the network always to be available for use and always to provide a first-class service. The increasing reliance on networks and computer systems means that any break or fault in the service is very frustrating, particularly when the network is used in a customer-service environment. Customers calling to query their account are quite rightly annoyed if they find that the organisation's telephone lines are constantly engaged or if the customer-service representative tells them that they must call back later because

Figure 1.3 Today, networks are becoming the core of the IT infrastructure

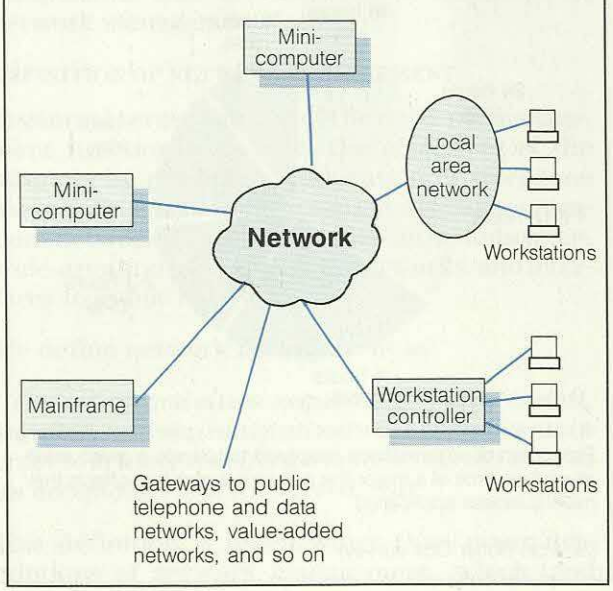

"the computer system (or the network) is down". Equally, the customer-service staff find it tedious to have to listen to numerous complaints from upset customers who want to know *now* what their outstanding balance is. It is not surprising that these frustrations are relayed back to the network manager as demands for a more reliable network.

Most of the network managers we talked with during our research had set a network-availability target of 98 per cent or higher. However, 98 per cent availability still means that the network will not be available for use for approximately 45 hours a year during normal business hours. Consequently, some network managers believe that the current availability of their networks is unacceptable to some network users. (Care needs to be taken when comparing network-availability targets and statistics; for some organisations, 100 per cent availability would mean that their networks were available during normal business hours, Monday to Friday; for others it would mean the networks were available 24-hours a day, every day of the year.)

We asked Foundation members what they believed was the maximum acceptable downtime for the most important application running on their network. The results are shown in Figure 1.4. The majority of members believed that they could tolerate a break in service of up to four hours on a major network link. Four hours seems to be a long time given the increase in user expectations. The majority of suppliers, however, do not provide a guaranteed time to fix a fault — they only guarantee to respond to a service call within a given time, which is frequently four hours. Since the level of service provided by suppliers does not match user expectations for network availability, network managers often find they have to install additional circuits or transmission capacity that will be used only if a link fails.

Figure 1.4 A majority of Foundation members can tolerate a major fault lasting up to four hours

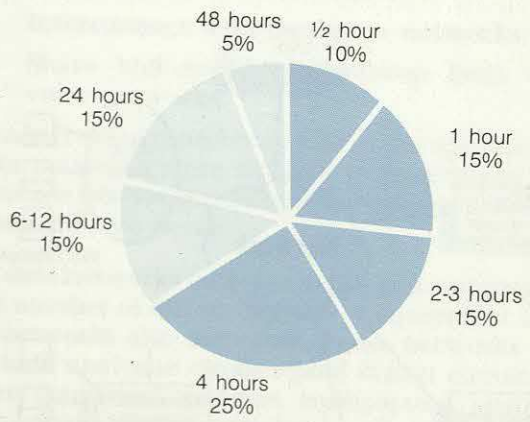

Proportion of organisations prepared to tolerate a given maximum downtime of a major link in their network that affects their main business application

(Source: Butler Cox survey)

Network users can also complain bitterly if they notice an increase in response times, even if the response times may still seem to be within acceptable limits to the network manager. Many network users also expect to be able to give less than a week's notice of the need to move office workstations. Furthermore, they expect the moves to be carried out during a weekend. Long lead times from equipment suppliers and the PTTs often prevent network managers from meeting these expectations.

INCONSISTENT NETWORK-MANAGEMENT POLICIES

Most of the difficulties described above relate to the management of wide-area data networks. Network management, however, also encompasses voice networks and local area networks. We found that several network managers dismissed the management of voice networks as being "not a problem". Furthermore, their network-management policies did not include local area networks because these "are managed by users, not the systems department". The result is that fault logs and records of the changes made to voice and local area networks are not as detailed as they are for wide-area data networks. Hence the procedures and policies adopted for managing the different kinds of network are inconsistent. If inadequate attention is given to the management of voice and local area networks, it will be harder to identify and correct faults when they do occur.

Without doubt, there is a real need to improve the way in which networks are managed. There are, however, important reasons preventing the desired improvements from being made.

BARRIERS TO IMPROVEMENT

There are four main problems that prevent or delay many organisations from making improvements in the way they manage their networks:

- Shortage of skilled communications staff.
- Limitations of current network-management tools.
- Difficulties in justifying investment in network-management tools.
- Confusion about the scope of network management.

SHORTAGE OF SKILLED STAFF

A perennial complaint within systems departments is that experienced staff are in short supply and are difficult to recruit. The growth in the use of data networks and the rate of change in communications technology in recent years have led to an acute shortage of skilled communications staff in many organisations. Inevitably, the increased demand for this scarce resource has resulted in substantial increases in salary levels for this type of staff. It has also led to very rapid staff turnover. As a consequence, many network managers believe that they have insufficient staff with the right skills to manage their networks effectively.

LIMITATIONS OF CURRENT TOOLS

Difficulties in recruiting experienced staff have caused many network managers to seek to enhance the effectiveness and productivity of the staff they do have by using more network-management tools. Unfortunately, the tools available are very limited in scope because, in general, they were developed by telecommunications-equipment and computer suppliers as enhancements to their main products. Each supplier's tools are designed to work only with the supplier's own equipment. Organisations therefore have to purchase tools from most of the suppliers of the equipment used in their network. Tools from different suppliers usually cannot interwork with each other. Also, the majority of network-management tools mainly generate alarm and error messages that are used for fault diagnosis. Performance and usage data provided by such tools is usually limited and is difficult to extract.

In large networks, network-management tools can produce a large amount of basic data, but it requires considerable skill and time to interpret the cryptic and overlapping messages to produce the information required to manage the network. Most network managers believe that the lack of a comprehensive set of network-management tools (or a network-management system) is the greatest barrier to improving the way in which their networks are managed.

DIFFICULTIES IN JUSTIFYING INVESTMENT

A common question raised by Foundation members during the research for this report was "How do we justify our investment in new network-management tools to our network users and to our senior management?" Until recently, most network-management tools were acquired as part of a major network-equipment purchase. They were cost-justified as part of the overall business case for investment in the network. Today, there is an increasing tendency for network-management tools costing up to $1 million or more to be purchased as separate items. Many network managers find it difficult to justify this type of investment because they are unclear how they can identify quantifiable cost savings.

CONFUSION ABOUT THE SCOPE OF NETWORK MANAGEMENT

Most network managers are aware that their responsibilities cover a broader spectrum of activities than identifying and fixing faults on the network. Nevertheless, many of them are unclear as to exactly which activities and equipment should be included in the responsibilities of the network-management function. In particular, there is confusion as to whether network planning, help desks, and mainframe communications software should be included in the network-management responsibilities. In some organisations, for example, disputes have arisen between the network-management function and the computer-operations function about who should run the help desk, who should control the installation of terminals, and who should manage the front-end communications processors.

It will be difficult to improve the way in which networks are managed until such disputes are resolved and the scope of the network-management activity has been defined clearly.

THE SCOPE OF NETWORK MANAGEMENT

Much has been written on the topic of network management. Many of the authors assume the topic has clear, well-understood boundaries, but fail to define what they mean by network management. After a literature review, we concluded that there is no generally accepted definition of what the term 'network management' means. Before proceeding further, we need to explain what we mean by the term, and to describe the activities and network components that are included in the scope of network management.

DEFINITION OF NETWORK MANAGEMENT

In general terms, the aim of the network-management function is to meet the objectives of the business by providing workstation or telephone users with access to information via communications networks. These networks can include voice, wide-area and local area data networks, and interfaces to public networks.

We define network management as:

"The set of activities required to plan, install, monitor, and maintain all network components in order to achieve specified service levels reliably, at an acceptable, and an agreed, cost."

This definition is much wider than many descriptions of network management, which tend

to focus on the day-to-day operations of networks, particularly on monitoring performance, identifying faults, and fixing them. For many organisations, our definition of network management equates to therole of the telecommunications department. Although we do not refer to 'the telecommunications manager' in this report, this job title can more often than not be used whenever we refer to the 'network manager'. In some organisations, however, the network-management function will be one part of the telecommunications function.

Our definition includes service levels as an integral part of network management. Without targets for service levels, network managers have no measure of how well their networks are performing and thus cannot be sure whether the service provided is meeting their users' needs.

The definition also refers to costs. Obviously, costs should be kept as low as possible, but the scope for reducing costs will be constrained by the service-level requirements. Network users therefore need to understand the trade-offs that can be made between service levels and costs.

Our definition is deliberately wider than that formulated by ISO (International Standards Organisation) who are defining standards for network management. Their work is based on, and extends, the OSI (Open Systems Interconnection) standards because the relevant standards committees realise that effective interconnection between systems requires common network-management standards. A draft OSI network-management framework standard is now available for comment and should be formally agreed by the end of 1988. However, this standard only defines the framework for the development of network-management standards. The standards themselves are unlikely to be fully defined before the early 1990s.

The OSI definition describes network management as the facilities to control, coordinate, and monitor the resources that allow communications to take place in the OSI environment. There are five functions within the OSI management framework: fault management, accounting management, configuration and name management, performance management, and security management.

The OSI definition is thus narrower than our definition because we include planning and user-support activities in the scope of network management. We also include the process of change management, whereas the OSI definition only goes as far as managing the components that comprise an existing network.

We believe such extensions to the OSI definition are necessary when considering how to improve the network-management function. While the scope of the OSI definition of network management is too narrow for this purpose, it is valuable when comparing the capabilities of network-management tools, and is discussed further in Chapter 5 in this context.

NETWORK-MANAGEMENT ACTIVITIES

There is a broad range of activities included in the scope of network management because the activities cover all the functions involved in providing a communications service. The activities included in our definition of network management are shown in Figure 1.5. We have identified two types of activity — those for which the network-management function is entirely responsible, and those for which the network-management function has a shared responsibility with other groups in the systems department.

The boundary between the two types of activity is unclear and may vary from organisation to organisation. For example, some organisations regard

Figure 1.5 Network-management activities

Activity	Description
Fault handling	Identification, diagnosis, and repair of all faults related to network components. Provision of alternative service, where possible, during network breakdowns.
Change management	Installing and controlling additions, moves, and changes of users, hardware, circuits, and software related to networks. Controlling the configuration of the network, including alternate routeing. Inventory maintenance.
Performance monitoring	Tracking usage of the network(s) to identify requirements for additional capacity. Analysing performance of network equipment, services, and suppliers on a regular basis. Measuring level of service provided to users.
Tactical planning	Planning required to ensure that the network(s) will accommodate growth or new services. Timeframes for tactical planning will be between one month and one year.
Cost control	Monitoring operating costs of the network(s). Reconciling invoices with predicted costs.
The activities listed above are entirely within the responsibility of the network-management function. The responsibility for those shown below may be shared with other groups in the systems department.	
Billing	Billing users for IT services, including network usage.
User support	Help desks, training, consultancy, and assistance in selecting equipment or software. Defining service levels.
Vendor relations	Negotiating contracts for purchase or maintenance of equipment or services. Agreeing discounts and service levels.
Security	Determining best methods of ensuring that access to confidential or proprietary information is restricted to authorised users. Reviewing violations of security procedures. Auditing effectiveness of security measures.

security as the responsibility of computer systems staff and expect them to provide procedures for controlling access to applications. Others take the view that security is predominantly the responsibility of the network-management function. They believe that security procedures should be invoked as users first access the network; the log-on procedure should therefore check that users have authority to access the requested applications.

It is important to realise that some network-management activities involve other groups from the systems department. For example, it no longer makes sense to provide separate help desks for networking problems and for applications problems. Most users are totally disinterested in whether their immediate problem is caused by a network fault or a software bug. A fragmented approach to these types of activity can cause users to have a poor perception of the quality of the service being provided. We return to this topic in more detail in Chapter 2.

NETWORK COMPONENTS

Some of the confusion about the scope of network management occurs because of uncertainties as to which types of equipment are included as part of the network, and which are external to the network. It is therefore also important to define which equipment is considered to be part of the network, because this will, to some extent, define the scope of network management.

Today's complex networks contain a wide range of equipment and a variety of transmission media. Many different types of workstation (including telephones) and computer system are connected to networks. Some network functionality may reside within workstations (a personal computer circuit board with a built-in modem, for example) or within a larger computer system. This means that it can be very difficult to define the boundary of a network and may mean that the network-management function shares the responsibility for a particular item of equipment with another section of the systems department. Sometimes, even the transmission media (circuits and cables) may not be considered as part of the network. For example, some organisations with multiple DEC computers interconnected by an Ethernet system consider the Ethernet to be part of the computer system.

The equipment components considered to be within the scope of the network-management function vary by organisation. Figure 1.6 lists the components that are usually considered to be part of the network, and other components that may or may not be controlled by the network-management function. Figure 1.7 overleaf shows two actual examples of the different ways that the responsibility for managing the components can be assigned. These examples illustrate how the network-management function may have difficulty in determining where its responsibility ends. From the users' standpoint, the divisions can appear arbitrary and illogical.

The decision about which components fall within the scope of network management appears to depend on the structure of the systems department and the skills resident within the telecommunications area. Thus, PABXs and local area networks are often managed locally by the user community because there are no tools to facilitate central management and because it is not possible to employ communications specialists at every site.

The responsibility for terminals, personal computers, and other end-user equipment can also cause difficulties because these types of equipment do not relate directly to the structure of the systems department. (Foundation Report 63 — The Future of the Personal Workstation — contains advice about how to manage workstations.) In some cases, the network-management function takes on the responsibility for these types of equipment to reduce the number of hardware-support groups.

THE PURPOSE OF THIS REPORT

In this chapter we have explained why network management is a rapidly growing problem for most

Figure 1.6 There are considerable variations in the components considered to be part of the network

Components usually considered to be part of the network	Components sometimes considered not to be part of the network
Private circuits	Front-end processors (FEPs)
Multiplexors	Other communications controllers
Modems	Matrix switches
Packet switches and packet assemblers/disassemblers (PADs)	Communications control software (front-end processors or mainframes)
Terminal controllers	Local area networks (cabling and interface equipment)
Tandem switches (for voice networks)	Terminals, personal computers, and other communicating office systems
Telex machines	PABXs
Message switches	Facsimile machines
Videoconferencing equipment	Any interface to a public network

The network-management function will usually be responsible for components listed in the first column, but not necessarily for those listed in the second column

Figure 1.7 Components considered to be part of the network vary by organisation

Company 1: A German manufacturer

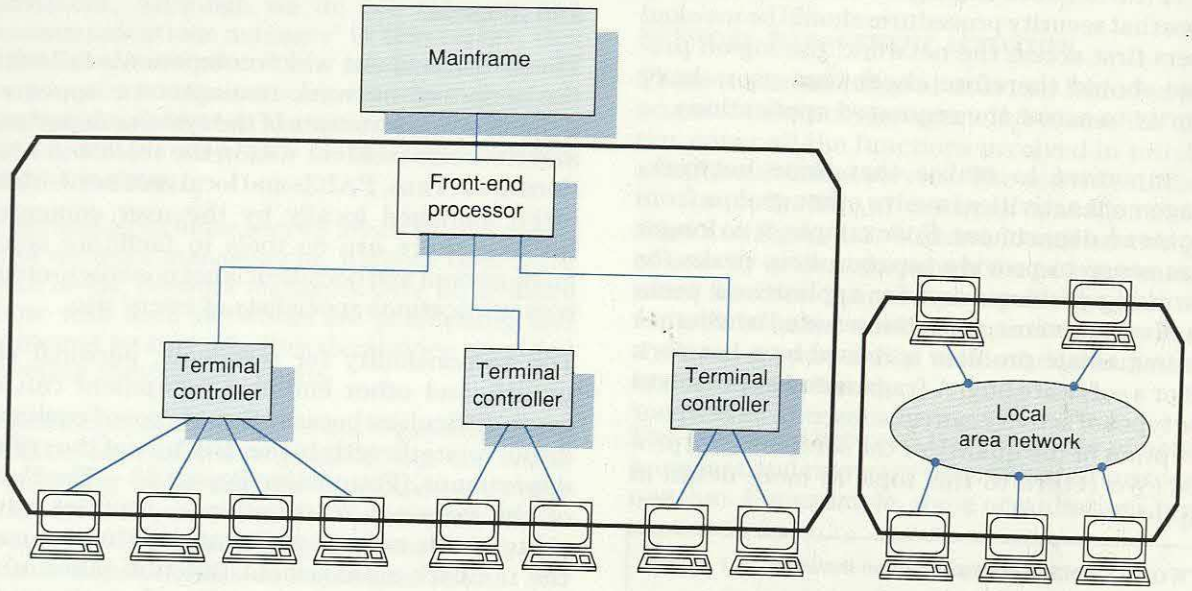

Company 2: A UK manufacturer

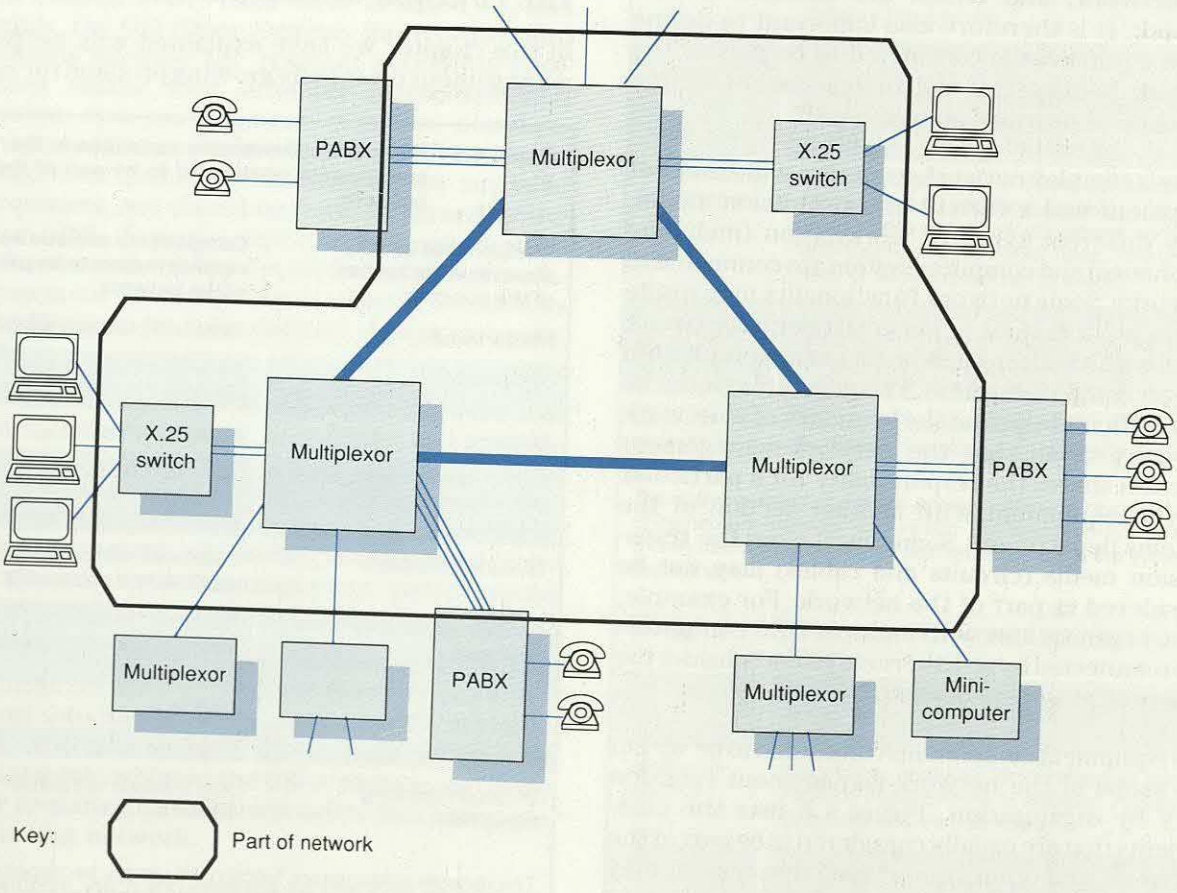

organisations to which there is no easy or quick solution. The outlook appears to be very gloomy. Yet during our research we did talk to organisations that are managing their networks quite successfully with reasonably satisfied users. The difference between the organisations with adequate network management and those that are struggling is not usually more or better tools or staff. It is how well the network-management function is managed.

From our research and consultancy experience we have identified the key indicators of whether the network-management function is performing its activities effectively. These are shown as a checklist in Figure 1.8 and can be used by senior management to assess the effectiveness of the existing network-management function.

The following chapters of this report contain guidance on how to improve the effectiveness of the network-management function and describe the technology developments that will enable the improvements to be made. We concentrate mainly on the management of wide-area data networks since that is the area of most concern to members. The management of voice and local area data networks is referred to where we feel that it is important.

In summary, the network-management function should:

- Be organised to provide the service that users want.
- Be staffed by the right combination of people to handle business concerns, user-support, and technical problems.
- Have procedures in place to handle most routine tasks efficiently, and to identify abnormal situations quickly.
- Understand technology developments and their impact on networks, tools, organisation, and procedures.
- Use tools effectively to assist with a wide range of tasks. In particular, develop a strategy for purchasing *and* replacing tools. The purpose of the strategy is to evolve towards a comprehensive network-management system.

In Chapter 2, we discuss how the network-management function can be organised to provide a better service to users. User support is clearly not just a network-management responsibility but concerns the whole systems department. In order to provide the correct level of service, user requirements must be clearly understood and agreed, and the role of the network-management function within the systems department needs to be clarified. We advocate the use of service-level agreements as a good mechanism for ensuring that the network-management function provides the right service at the right price.

Figure 1.8 A checklist for evaluating the effectiveness of the network-management function

The answers to the following questions will indicate whether the network-management team is performing its activities effectively.
Do users rarely complain about poor response times or service interruptions?
Are most faults cleared in less than one working day?
Is it rare for a fault to re-occur after it has been fixed?
Does the network-management team measure and report on trends in network performance?
Does the network manager have charts showing the actual availability of the network or the amount of downtime for the last three months?
Is the average network availability above 95 per cent?
Does the network-management team obtain a good service from suppliers?
Does the network-management function have good working relations with other groups in the systems department?
Are major network upgrades or changes carried out with little impact on existing users?
Are more than 50 per cent of calls for assistance solved by the help desk?
Are users clear about who to call for assistance with queries or to report a fault?
If most of the answers are 'yes', the networks are probably well-managed.

Chapter 3 describes how to make the best use of scarce technical staff. It also shows that business and people-related skills are required as well as technical skills. In summary, the scarcity of skilled staff can be alleviated by assigning routine tasks to nontechnical staff, by using supplier skills wherever possible, by providing more training for all staff, and by using automated tools to reduce or simplify the workload.

The network-management procedures that can be adopted to handle the changing environment are described in Chapter 4. These procedures cover fault handling, performance measurement, change management, and supplier relations. The chapter also highlights how the changing profile of network faults (a smaller number of faults, but with more of the faults being very hard to solve) will affect routine network-management activities.

Chapter 5 identifies the changing requirements for network-management tools and systems, and describes the types of tools available today and their inadequacies. We explain what a complete 'integrated' network-management system should include, the difficulties in achieving this goal, and the likely future developments.

Even though today's network-management tools cannot be used to construct a complete integrated

network-management system, organisations need to work towards such an integrated system. Chapter 6 therefore provides guidance on building a network-management system in an evolutionary way and on justifying the necessary investment.

SCOPE OF THE RESEARCH

This report is based mainly on an extensive programme of research carried out between November 1987 and March 1988. We received 111 responses to the questionnaire sent out to Foundation members at the beginning of the research. They provided a substantial amount of information about members' current network-management problems and the benefits they expect to gain from network-management tools and systems. The opinions of more than 50 network managers and planners were obtained in interviews and focus groups held throughout Europe and in Australia. Much of the practical advice in the report is based on the insights of network managers who shared their experiences with us. Their contribution is much appreciated.

We also met with 14 suppliers of network-management products and services, both in Europe and the United States. In addition, a telephone survey was carried out to ascertain the range of products offered by a further 30 suppliers. We also drew on our consultancy experience and the extensive body of available literature (some of which is mentioned in the bibliography at the end of this report), and we sought the views of several technical specialists.

The research was led by Janet Cohen, a senior consultant with Butler Cox in London. She was assisted by Simon Forge, a senior consultant with Butler Cox's Paris office, and by Kevan Jones, a consultant with Butler Cox in London. Valuable contributions were also made by several other consultants in Butler Cox's telecommunications consultancy practice.

Chapter 2

Providing a service that meets user needs

The network-management function has a key role to play in the systems department, being responsible for the quality and reliability of the telecommunications service. However, the provision of network services cannot be considered in isolation from the provision of other IT services because many network-management activities have to be carried out in conjunction with other parts of the systems department. These include running help desks, planning, billing, and making arrangements for security. Most of the activities where the network manager's responsibilities overlap those of other managers in the systems department concern the department's relationships with its users. Users frequently do not understand the systems department's division of responsibilities by technology area because they view the provision of IT as a single service. Thus, user support is the responsibility of the systems department as a whole. It no longer makes sense for the network-management function to provide its own separate user-support operation.

Effective user support requires that a business area receives the right type of service for its needs at a price that it finds acceptable. Part of the network manager's role is to make sure business areas are aware of the trade-offs that can be made between cost and level of service. The business areas can then make an informed choice about the type of service that they need.

We believe that the level and type of service to be provided should be formally negotiated between the network-management function and the individual business areas and spelt out in service-level agreements. The network manager then has an effective target for the service needed to meet the business requirements.

In the remainder of this chapter we first explain why the network manager's user-support responsibilities should be part of a wider service provided by the systems department. We then describe the advantages of networking service-level agreements and how to set about constructing them.

ESTABLISH A SINGLE USER-SERVICE CONTACT POINT

Although this report is concerned with network management, one of the most important findings from our research is the trend towards managing networks, computer operations, and system development as a total IT function. This trend is being driven partly by convergence of the technologies, but, more importantly, by increasing demands from the user community for the provision of a unified IT service. The implications for network management are that wide-area networks should be managed centrally, and that network help desks should be combined with other IT help desks to form a single point of contact for the user community. In turn, this means that the systems department will need to move away from its traditional organisational structure based on technical specialisation to a functional organisation.

DEMAND FOR A UNIFIED IT SERVICE

Users of IT services, unlike IT specialists, are generally not interested in technology for its own sake. Instead, they are concerned with how the use of IT assists them in performing their jobs. Their main requirement, therefore, is for reliable and easy access to the computer applications they need to use or to the telephone of the person that they need to contact. In most cases, users do not need to be aware of whether the application or telephone is in the next office or in another country. The network-management team at a major multinational oil company was experiencing difficulties in identifying network problems because the users were not aware of where the applications they were accessing were physically located. Users tended always to claim that an application was running on a computer in the same building. The solution to this problem was not to educate users to recognise which computer they were connected to, but for the network-management staff to obtain current information about application locations from the computer-operations staff.

When users call a help desk with a problem or they require additional terminals to be installed, they do not want to think about which part of the systems

department they should be talking to. Similarly, users who are experiencing increased response times are not concerned whether the problem is caused by overloaded circuits, overloaded processors, or poorly designed applications. They simply want to know when the response times will return to an acceptable level.

Users therefore need a single contact point in the systems department that is able to handle all of their problems or requests for assistance. There are also benefits for the systems department in providing a unified user-support service. Departments providing such a service will have more credibility with their users, and their users are less likely to complain about the service being provided.

We believe that systems directors should review how the various aspects of their user-support services are organised and should reorganise them to provide a 'seamless' user service. This will not be easy to achieve with a systems department that is organised on the basis of technical specialisations because it will require considerable cooperation and interworking between the various sections. Our consulting experience has shown that in some organisations, there can be considerable rivalry between the different sections in a systems department. In particular, different sections will spend time and effort in trying to prove that a fault is the responsibility of another section, rather than concentrating on solving the user's problem.

An implication is that user-support activities, including network support, are likely to be more effective if they are centralised. Moreover, there are other good reasons for centralising the network-management function.

ADVANTAGES OF CENTRALISED SUPPORT

Most Foundation members now have a centralised network-management function. Figure 2.1 shows how the Foundation members who responded to our questionnaire manage their networks. There are several advantages to a predominantly centralised approach to network management:

- Skilled staff can be better utilised and, hence, are more cost-effective.
- Complex tasks can be performed more efficiently because the staff have more opportunity for working on difficult problems.
- Network-management tools and systems can be cost-justified more easily.
- Some suppliers provide network-management tools that are geared to centralised management.

However, there are two situations where distributed network management, or a combination of centralised and distributed management, is essential. First, centralised management is difficult for international networks because of the problems of providing support across multiple time zones, dealing with local PTTs, and being aware of local standards and regulations. These problems require some degree of regional or local network management.

The other situation where some local network management is required is when local computer and telephone systems are not connected to corporate networks. With these types of installation, the network-management tasks are not onerous, and it is often much easier for an individual based at the local site who has some technical skills to carry out these tasks. Where appropriate, the local staff can use the central network-management staff as a source of expertise. The central staff may also specify the standards and operating procedures to be followed by the local staff.

Even where there is a centralised function there are some network-management activities that are best performed by local staff or third-party suppliers. The decision as to which activities should be performed locally involves balancing the expertise required for the activity against the time required to perform it. Thus, it is usually easier for an on-site staff member to move a personal computer connected to a local area network, rather than to wait for a member of the centralised network-management group to carry out the work.

However, some degree of centralised control is required for all services that are connected to a

Figure 2.1 Most Foundation members have a centralised network-management function

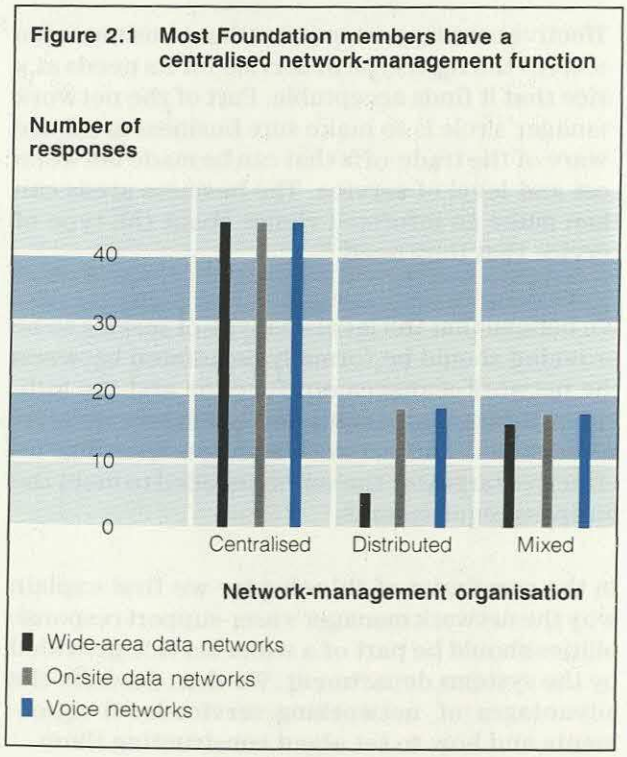

corporate network. Help desks often receive calls from staff who have moved their own workstation or telephone to a different network presentation point and are surprised to find that it does not work at the new location. They are not aware that the network-management centre needs to check that the cabling at the new location is connected to the right type of port, and that the port is recognised by the computer or telephone system.

A SINGLE HELP DESK

Our research showed that some large organisations still provide several help desks for their users of IT services. (The highest number we found in a single organisation was 23.) Most organisations have recognised that operating multiple help desks has several disadvantages. In particular, multiple help desks use staff resources inefficiently, they confuse users, who have to decide which help desk to call, and they can easily result in problems being 'lost' between help desks. The best way of overcoming these disadvantages is to operate a single help desk that is able to handle all user requests for assistance.

However, there are two circumstances where it is necessary to provide multiple help desks. In international networks, where there are time-zone differences and language barriers, it will not be possible to service all requests for assistance from a single help desk. And in networks supporting a large number of applications, a help-desk operator cannot be expected to be familiar with all the applications.

Even so, the multiple help desks should be organised so that an individual user always calls a particular help desk. Other guidelines for running help desks, based on members' experiences, are shown in Figure 2.2.

One of the greatest benefits of a successful help desk is that it improves the relationship between users and the systems department. A well-run help desk persuades users that their problems are understood, under control, and are receiving attention, even if they cannot be solved immediately. Help desks can also identify training needs. For example, if there is an increasing number of calls about the use of a particular application, it might be worth running a training course on that application.

ORGANISATIONAL IMPLICATIONS

In the previous chapter we described how the boundaries between the network-management function and other IT areas are becoming blurred. We have also shown that, in order to avoid user confusion and dissatisfaction, some network-management activities need to be carried out as part of a unified IT service. Without doubt, it will become increasingly difficult to differentiate between computing and communications technologies and there appears to be little reason other than organisational inertia to continue to organise the systems department by technology area. A more appropriate method of organisation is by function. Figure 2.3 overleaf shows in simplified form the structure we believe should be adopted for the future. The functional organisation shown in the figure gives each manager working under the systems director a clear set of objectives and should avoid conflicts between technology areas.

Specialist technical skills will still be necessary in the new organisation but will be spread across the functional areas. This arrangement will encourage interworking amongst specialists with different areas of expertise.

In many organisations, the rivalry between computing and telecommunications specialists is likely to slow down the move to a functionally organised systems department. We believe, however, that the operations area will be organised in this way earlier than other parts of the systems department, mainly because user pressure for a unified service is greatest in this area.

One implication of providing a unified operations service is that common procedures must be adopted both for charging for computing and communications services and for measuring the operational performance of both types of service. In turn, this implies a unified approach to managing the provision of IT services, which means, for example, that it will no longer be sufficient to maintain a separate database containing details of the network

Figure 2.2 Guidelines for running a successful help desk

The help desk exists primarily to support users, not the systems department.
Each user has only *one* help desk number to call.
The help desk 'owns' the problem. That is, the help desk is responsible for ensuring that the problem is resolved, no matter who actually fixes the fault.
The help desk is not merely a clerical function to take details of faults. (Experience shows that where the help desk acts merely as a channel for logging faults and passing them to someone else to action, users will tend to bypass the service.) The help desk should solve the simpler, more common problems. Targets should be set for the percentage of calls for assistance solved by the help desk. Between 50 and 80 per cent is a reasonable target.
The help desk should have standard procedures for passing on problems, reporting back to users, and alerting users about major breaks in network services.
Time limits should be set on attempts to solve a problem before it is passed on to someone with more specialist knowledge or to suppliers. Typically, help-desk staff should pass a problem on if they have not been able to solve it after 15 minutes.

Figure 2.3 In future, systems departments should be organised by function, not by technical speciality

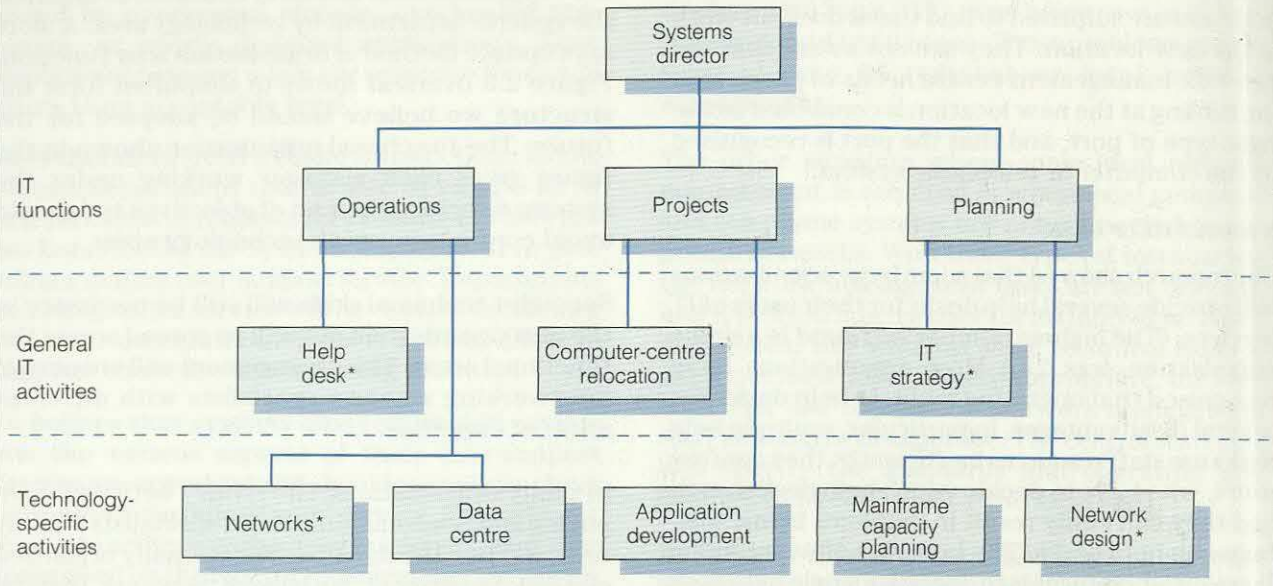

*Network-management activities will be distributed throughout the different functional areas

configuration and equipment. Instead, a total 'IT-inventory' database will be required, containing details about networks, applications, and computer systems.

Although we recommend a unified service approach, the subject of this report is network management, not the management of IT services as a whole. Thus, it is important to bear in mind that, whenever we use the terms 'network management' and 'network manager' in this report, we are describing activities and responsibilities that, in future, will be spread throughout the systems department. There might not necessarily be a separate network-management department carrying out the responsibilities and activities described in this report.

NEGOTIATE SERVICE-LEVEL AGREEMENTS

Systems departments, and thus the network-management function, now have users whose requirements for network services may vary enormously. For example, in an airline the network availability required for a back-office function such as the maintenance of personnel records will be much lower than that for an online reservations system. Network managers must recognise and plan to satisfy the widely different business requirements for each type of application.

Some of the business requirements might necessitate a highly resilient (and therefore expensive) network. Other requirements might well be satisfied by a much simpler, and thus inexpensive, network. It is therefore necessary to find a mechanism of charging users for network services in a way that reflects the level of service they require. We believe that the best method of ensuring that the systems department is providing the appropriate levels of service at a realistic cost is to negotiate service-level agreements with the user community, preferably with the managers who will be responsible for authorising payment of the systems department's charges.

UNDERSTANDING BUSINESS REQUIREMENTS

Given no restrictions, all network users would say that they want:

- Very reliable, good-quality communications.
- Excellent response times for data communication applications.
- No blocked calls on the voice network.
- The ability to move and change equipment at very short notice.
- Low charges for using the networks.

If the charges are to reflect the true costs of operating the networks, it will obviously not be possible to meet all of these requirements at a low cost. An unrealistic 'wish list' such as that above is of no value in trying to determine the required service. In reality, it will always be necessary to make trade-offs between level of service and cost.

In order to understand the real business requirements, and the price that users are prepared to pay, the network manager must sit down with the line managers from each business area who understand

how IT services are used within their business. They should discuss what the business is trying to achieve and how it can make the best use of network services. The two parties should reach an agreement on the types and quality of network services that the business expects to receive. The agreement will be a compromise between cost and quality of service. One of the biggest contributions that the network manager can make to the discussion is to explain clearly and (if possible) in quantified terms the trade-offs between cost and quality of service. A chart similar to that shown in Figure 2.4 can be a powerful aid to persuade line management that 98 per cent network availability is a more realistic aim than 99.5 per cent. The network manager will gain from the meeting an understanding of where each business area is prepared to spend money for a better service. Usually, this will be for the functions perceived as most crucial to the success of the business.

USERS' DIFFERING REQUIREMENTS

It is stating the obvious to say that users have different requirements. However, an often-overlooked implication of this obvious statement is that, not only should systems departments provide different levels of service to different users, but should also provide different billing arrangements. Human nature dictates that those users who complain loudest and most often will usually receive priority service (hopefully, these are the users who really need a priority service). Billing arrangements tend to be standardised, however, and often are based on the simplest method for dividing up the costs. In billing for the use of data-communications networks it is common not to charge by usage and the volume of data transmitted, but to charge on a flat-rate basis for connection to the network.

Figure 2.4 Trade-off between network availability and cost

Increasing network availability from 99 to 99.5 per cent costs about twice as much as increasing it from 96 to 98 per cent.

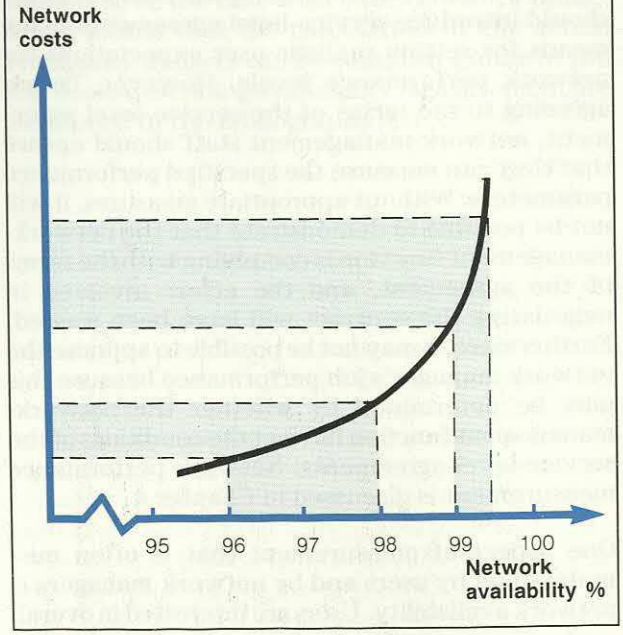

Most corporate data networks are used to provide access to a variety of applications, so a flat-rate billing scheme does not reflect the costs of providing the different levels of service to their respective users. Users with very low service-level requirements have to pay for the higher level of service demanded by other groups of users. The higher network costs resulting from the requirements of some users may not just include the cost of additional circuits for resilience, but can also include the costs of high-security access devices, more sophisticated multiplexors, and staffing the network-management centre for 24-hours a day.

An example of the difficulties that can arise when two groups of users with widely different requirements use the same network is provided by London Regional Transport's Technology and Network Group. This group provides communications services both to London Buses and to London Underground. London Buses is installing a new computer application that allocates crews and buses to routes within the city, taking into account factors such as holidays and scheduled maintenance. If this information is not available at each bus depot, it is much more difficult to run an efficient bus service. Because the information from this application is so crucial, London Buses is prepared to pay for a high level of redundancy and resilience, both in the network and in the computer system.

London Underground uses the same computer systems and backbone network for some of its applications, but does not perceive any of the applications as crucial to its daily operations. The underground is very cost-conscious and is not prepared to pay for resilience that it believes it does not require. The Technology and Network Group is faced with the difficulty of deciding how to bill both groups of users in a way that passes on the economies of scale derived from sharing the same network whilst ensuring that London Buses pays for the entire cost of the higher network resilience it requires.

This example demonstrates that taking account of differing user requirements is likely to lead to very complex billing arrangements for network services. The best way of tackling these difficulties is to negotiate service-level agreements with each group of network users.

SERVICE-LEVEL AGREEMENTS

The majority of organisations participating in our research were implementing (or planning to implement) service-level agreements for the provision of network services. Two main reasons were quoted for introducing service-level agreements:

- The pressure on systems departments from business areas to deliver better service and to demonstrate value for money.
- The need to provide users with information to help them understand what they are being charged for, particularly where they have not previously been billed for network services.

Service-level agreements are contracts between users and the systems department. They specify the types of service that the users want the systems department to provide and the conditions under which the service will be provided. The terms of the contract lay out when the services will be available, the performance levels that are to be attained, reporting procedures, and the costs of the service. The contents of a typical network service-level agreement are listed in Figure 2.5. (A useful source of information about service-level agreements is the article by C N Witzel published in The Journal of Capacity Measurement, Vol 1, No. 4, 1983. Although this article refers specifically to data centres, it contains general advice about constructing service-level agreements.) In France, service-level agreements often include user-satisfaction ratings. Figure 2.6 shows the screen layout used by one organisation to collect monthly survey data on user satisfaction.

Figure 2.5 Contents of a typical network service-level agreement

General clauses
Contracting parties
Period of contract
Provisions for modifying agreement
Performance-reporting procedures
Billing arrangements
Penalties for noncompliance
Performance-related clauses
Availability of service (both normal and contingency)
Maximum number of service interruptions in a period
Response times
Grade of service and percentage of calls successfully completed (voice networks)
Service hours for help desks and other operational support
Mean time to repair faults
Notice required for implementation of major and minor moves and changes
Times of week when changes are scheduled and when unexpected interruptions might occur
User-satisfaction measures
Disaster back-up facilities and service-restoration priorities

Our use of the words 'negotiate' and 'contract' when referring to service-level agreements is deliberate. By perceiving service-level agreements in these terms, the systems department is indicating that it is managed like any other part of the business and can guarantee to provide minimum service levels to its users.

Service-level agreements provide benefits for the systems department as well as for the users. Probably for the first time, users understand the level of service that the network-management function provides and therefore have realistic expectations about what to expect. In particular, user departments are responsible for ensuring that the specified service levels meet their business requirements. Providing that the network manager has explained clearly the implications of the agreed performance measures (such as network availability), users should not complain if the agreed service levels are unsatisfactory in practice. If this proves to be the case, they can then renegotiate the contract.

Network managers can use service-level agreements as a mechanism for understanding user requirements and for demonstrating the high quality of service they deliver. When Eli Lilly, a major US pharmaceutical manufacturer, provided users with monthly reports showing that the actual network-performance levels were within the agreed targets, the number of complaints about poor response times reduced significantly.

We recommend that all large systems departments should introduce service-level agreements as the means for setting realistic user expectations for network performance levels. However, before agreeing to the terms of the service-level agreement, network-management staff should ensure that they can measure the specified performance parameters. Without appropriate measures, it will not be possible to demonstrate that the network-management function is complying with the terms of the agreement, and the effort involved in negotiating the contract will have been wasted. Furthermore, it may not be possible to appraise the network manager's job performance because this may be determined by whether the network-management function has met the conditions of the service-level agreements. Network performance measurement is discussed in Chapter 4.

One aspect of measurement that is often misunderstood by users and by network managers is network availability. Users are interested in overall

Figure 2.6 Screen format used by a French company to measure user specification

```
Date: JJ/MM/AA              MESURE DE SATISFACTION            Terminal: XXXX
                            ======================

Pour les critères proposés. veuillez donner une note aux applications que vous
avez utilisées (de 0 = très mauvais, à 5 = très bon, ou rien = sans opinion).
                                                  D     T     C     A     F
    .........................................................................
    · SIGMA                                    .  __ .  __ .  __ .  __ .  __ .
    · SIROCO                                   .  __ .  __ .  __ .  __ .  __ .
    · BADIN                                    .  __ .  __ .  __ .  __ .  __ .
    · MESSAGERIE INFOCENTRE                    .  __ .  __ .  __ .  __ .  __ .
    · AUTRE (Précisez SVP)_________________    .  __ .  __ .  __ .  __ .  __ .
    .........................................................................

    CRITERES:    D DISPONIBILITE
    ========     T TEMPS DE REPONSE
                 C CONFORT D'UTILISATION
                 A ASSISTANCE, FORMATION ET DOCUMENTATION
                 F FONCTIONNEMENT SANS INCIDENT

Commentaire: ______________________________________________
             ______________________________________________
```

availability at their workstations, but systems departments tend to calculate availability for individual parts of the total IT system. For example, if the network, the computer hardware, the application, and the workstation are each available for 98 per cent of the time, the overall availability to the user could be as low as $(0.98)^4 \times 100$ per cent, or 92 per cent. However, even this calculation is simplistic because it assumes that all failures interrupt the service provided to the user. This is unlikely to be the case with many network faults, which means that the calculation of the actual availability to users can be complex. (Some of the text books describing availability measurement are mentioned in the bibliography.)

Service-level agreements have an important role to play in ensuring that the systems department provides a network service that meets users' needs. However, in negotiating the agreements, the network manager will be under pressure to provide the best possible service at the lowest possible cost. A study carried out in 1987 of network costs in major US corporations showed that personnel costs accounted for 45 per cent of total operating costs and 30 per cent of the total network costs including purchase of network components. Thus, even though skilled network staff are in short supply, network managers need to ensure they are making the best use of the skills that are available. We provide advice on how to do this in the next chapter.

Chapter 3

Making the best use of scarce skills

Our research confirmed that there is still a great shortage of skilled communications staff and that this position will not change in the foreseeable future. In addition to technical skills, the network-management function needs to have staff with user-support skills and managerial skills. Network managers therefore have to come to terms with this situation and make the best use of the skills that are available. In this chapter, we provide advice on how to do this.

First, it is necessary to recognise that many of the tasks now performed by specialists could be delegated to nontechnical staff. We then show how the technical skills shortage can be overcome by using external staff, by developing the skills of existing staff, by providing career-development paths, and by using appropriate tools to support staff. The chapter concludes by describing the situations where it may be possible to subcontract the network-management function to a third party, thereby removing the need to employ technical specialists.

NETWORK MANAGEMENT REQUIRES THREE DISTINCT SKILLS

The network-management function provides a service, and a service-oriented function cannot employ just technical specialists. In addition to technical specialists, staff with managerial and administrative skills will be required, together with staff with appropriate skills for a user-support role. The activities of the network-management function should be allocated so that as many nontechnical tasks as possible are removed from the technical specialists and are performed by staff with appropriate managerial and 'people' skills. The scarce (and expensive) specialists can then concentrate on the more difficult technical problems. Below we describe the responsibilities of the managerial, user-support, and technical staff that will make up the network-management team. We also describe the skills and personal attributes required by each of the three types of staff.

MANAGERIAL STAFF

Successful network managers will, first and foremost, be managers — not technologists. They need to be able to work with business managers and corporate management, and to understand how the use of networks can help meet business objectives. The head of the systems department must therefore ensure that the individual appointed has the right managerial skills, not just the best technical expertise. Members should promote the development of these skills in network-management staff through an appropriate management-training programme.

Responsibilities

The network manager's primary responsibilities include:

- Providing network services that assist the organisation as a whole to meets its objectives.
- Ensuring the performance targets set out in the service-level agreement are met.
- Satisfying the development and support needs of individual team members.

Skills and attributes

The network manager must have a good understanding of the business activities of the organisation and a general understanding of business management. In particular, he or she must be able to translate business plans and objectives into network-management objectives. The case history in Figure 3.1 illustrates the type of problem that can occur if network managers do not understand the business requirement that lies behind a networking requirement.

Nevertheless, the network manager needs to know enough about the technology to understand the work of the technical specialists. A telecommunications background is not essential, however. Some of the most successful network managers we met during the research had a data processing background. Their lack of detailed knowledge about telecommunications technology enabled them to view problems from a different perspective and suggest some new approaches. This viewpoint was endorsed by M Montagnon of France Cables et Radio, who has specified and implemented network-management systems for several major

Figure 3.1 Focusing on technical requirements can cause problems

A major US life insurance company developed very sophisticated software to switch between batch and online data transmission to remote offices. The aim of the communications staff in developing the software was to reduce the number of circuits, thereby reducing communications costs substantially. The resulting software was extremely complex and made the systems quite unreliable because problems were difficult to solve and could be resolved only by the staff who had developed the software.

Branch-office managers found the unreliability of the systems intolerable. One of their major objectives was to process applications for new business as quickly as possible at the end of each month or at the end of a sales-promotion period. (Sales bonuses depended on the volume of new business entered into the system.) More often than not, the network was unavailable for up to a day during these peak periods. It turned out that the branch-office managers were prepared to pay twice as much for a reliable communications service. While designing the software, the communications staff had not discussed the system with the users and were unaware of the requirement for a highly reliable network.

French companies. He confirmed that a manager who concentrates on the technical aspects of network management is often unable to make top management aware of the importance of the network-management function. As a consequence, the network-management function does not receive adequate support from top management.

In summary, we believe that successful network managers will be able to:

- Communicate in business terms with senior management (orally and in writing).
- Balance day-to-day pressures and forward-planning activities.
- Motivate and develop network-management staff.
- Make decisions on priorities under pressure.
- Resist the temptation to become immersed in technical detail.

USER-SUPPORT STAFF

The network-management staff who provide support for network users need a combination of skills that are hard to find in one person. They must be able to relate to user needs and problems, but have sufficient technical understanding both to answer basic questions and to know who to contact for problems that are outside their range of expertise. The activities of user-support staff will usually be performed as part of the help desk, but they may also be used for training activities or to assist in specifying user requirements and co-ordinating changes, and possibly even for installing terminals.

Responsibilities

User-support staff will have a variety of responsibilities, depending on how the network-management function is organised. In general terms, these responsibilities will include:

- Acting as an interface between users and technical specialists.
- Assisting users to obtain maximum benefit from network services.
- Solving routine and straightforward problems.
- Handling requests for moving or changing terminals.
- Acting as 'user champions' within the network-management function.

Skills and attributes

Many organisations that operate successful help desks believe that the best user-support staff have themselves previously been users of the network services. This means that they can better empathise with the users' point of view. Some degree of technical aptitude and training is necessary, however, but this will vary according to the level of support that is to be provided. Most importantly, user-support staff should not be perceived as being technically naive by more experienced users. User-support staff should also have a good understanding of the organisation of the systems department and should know who is responsible for each system, application, or activity.

User-support staff should be chosen primarily on their ability to deal with people, sometimes in tense situations. Other important attributes include:

- The ability to stay calm under pressure and abuse.
- A good telephone manner.
- The ability to recognise when a problem is beyond the scope of their expertise and should be referred to the next level of support staff.
- The perseverance to chase the progress being made by technical-support staff and suppliers in solving users' problems.

TECHNICAL SPECIALISTS

Most Foundation members find it increasingly difficult to recruit skilled technical specialists for their network-management function. The growth in the use of networks, and their increased complexity has outpaced the supply of skilled communications staff. The resulting shortage of skilled staff is now a major barrier to improving the way in which corporate networks are managed. Communications software programmers

and specialists with an in-depth knowledge of X.25 products and protocols are in particularly short supply. Moreover, the high level of demand for qualified staff, particularly from financial institutions, has led to rapid increases in salaries. Organisations whose salary structures do not permit them to pay the 'going rate' experience great difficulties both in recruiting new staff and in preventing existing staff from being lured away by the prospect of higher salaries elsewhere. Some public-sector organisations are particularly prone to these problems. The majority of private-sector organisations, however, do not find it as difficult to retain their existing staff.

Despite these difficulties, network managers should not expend a lot of effort on trying to recruit more-experienced communications staff. Instead, they should seek ways to make more effective use of their existing technical staff. In particular, the work of specialist technical staff should be organised so that they do not have to handle routine faults and administrative work that could just as well be undertaken by user-support staff.

Responsibilities

The main responsibilities of the technical specialists in the network-management function include:

- Providing assistance to user-support staff, in particular by resolving faults that are beyond their incompetence and by liaising with suppliers as appropriate.
- Evaluating the performance-monitoring data provided by network-management tools.
- Planning and implementing changes and enhancements to the network including the selection of new equipment.
- Performing security audits.
- Monitoring developments in network technology.
- Liaising with other specialists in the systems department.

Skills and attributes

Technical specialists will usually have received some form of technical education. The more highly skilled staff, particularly those involved in planning networks, will probably have a technical degree. Many technical specialists will have practical experience, particularly of network operations and fault handling. More experienced staff will have expertise in a particular aspect of communications technology. However, it is desirable that technical specialists understand, at least in general terms, other technology areas both within the communications field and in the wider field of IT. Ideally, technical specialists should also have some experience of working in a user environment and on a help desk.

The personal attributes required to be a good network-management technical specialist are:

- Good problem-solving and analytic skills.
- Ability to work under pressure, particularly in network operations.
- Thoroughness.
- Understanding of the trade-offs that can be made between technical complexity and ease-of-use.
- Ability to understand and apply new communications technology.
- Ability to work with other people including suppliers.

Staff with these attributes and the appropriate level of technical skill will continue to be in short supply. It will therefore still be necessary to find ways of overcoming the shortage of technical skills.

THE TECHNICAL SKILLS SHORTAGE CAN BE OVERCOME

We have identified four main ways of compensating for the shortage of skilled technical staff:

- Make effective use of the skilled staff that work either for equipment suppliers or as independent specialists.
- Provide training programmes designed to develop the skills of existing staff.
- Provide career-development paths.
- Use appropriate tools.

Each of these is discussed below.

USING EXTERNAL STAFF

Many of the most skilled technical specialists work for suppliers of communications equipment or prefer to operate as consultants. Some network managers have found ways of making effective uses of these external sources of expertise to augment in-house technical skills. The case history described in Figure 3.2 illustrates how a major bank uses the specialist expertise of an independent expert to install major upgrades to the networking software. The case history described in Figure 3.3 shows how a leading multiplexor supplier provides the opportunity for its customers' network technicians to work alongside its own technical specialists.

DEVELOPING SKILLS THROUGH TRAINING AND EXPERIENCE

The shortage of network-management skills cannot be tackled without a greater commitment

Figure 3.2 External experts can be used instead of in-house specialists

A major European bank operates a large X.25 network and uses an independent expert (who used to work for its major equipment supplier) to assist in testing and implementing major network-software upgrades. The bank finds it more effective to use the specialist expertise when it is required, rather than employ its own highly specialised staff whose specialist skills would be used infrequently. The independent expert does similar work for several other organisations. He is therefore more familiar with carrying out major software upgrades and is less likely to make a mistake. As a consequence, the upgrades can be implemented more quickly. Another advantage is that the bank's own technical specialists are not diverted from their regular responsibilities in order to carry out the upgrades.

Figure 3.3 Timeplex involves customer staff in final assembly and testing of multiplexors

Timeplex is one of the leading US suppliers of high-bandwidth multiplexors. This company has developed a programme that combines the final assembly and testing of multiplexors with technical training for its customers. After the multiplexors are equipped to the customer's specification, the final stages of assembling the equipment and testing of the multiplexor configuration are performed jointly by Timeplex staff and the customer's network technicians who will operate the multiplexors.

This procedure has many benefits both for Timeplex and its customers. Assembly prior to installation ensures that all the necessary connectors and other components are present and correctly labelled, and ensures that many of the problems that used to occur during installation are found and resolved before the equipment arrives at the customer's premises. This means that Timeplex's technicians need to spend much less time at the customer's premises.

The advantage for the network technicians is that they can see all the multiplexors working together in the same room. Once the multiplexors have been installed throughout their network, it will not be possible to do this. They therefore gain a better appreciation of how the multiplexors interact with each other. They also gain an understanding of fault-finding techniques by being involved with the final assembly and testing.

to training. Organisations are increasingly providing training for less experienced staff rather than attempting to recruit scarce and expensive experts. However, many of the required skills cannot be gained just from formal education and training, but have to be based on knowledge gained from experience. Networks and network-management tools are now so diverse that it is unlikely that a network technician joining an organisation will have experience of all of the network equipment in use. The length of time required for a technician to become fully familiar with a network and competent to handle all network problems and configuration changes is growing. An extreme example is provided by Citicorp's International Communications Center in New York. The network operations manager told us that the time required for a technician to become fully conversant with all of Citicorp's systems and procedures had increased from six months five years ago to two years today. This is because Citicorp has used several customised network-management tools and these have undoubtedly led to the greatly increased training times.

Several organisations have successfully overcome the shortage of skilled staff by recruiting and training graduate trainees to supplement their existing technical experts. They have found that the learning period for graduates is not appreciably longer than that for skilled network technicians, provided that they are introduced gradually to the various types of equipment. Graduate trainees are usually first assigned for several months to carry out tasks that require low-level technical skills such as installing terminals. They may also spend periods working on the help desk, providing operations support, collecting performance-monitoring data, and preparing management reports. As they gain more experience and technical knowledge, they can then be used to provide the first level of technical support, resolving the problems that the user-support staff cannot handle, and passing on the few remaining really difficult problems to the fully trained technical specialists.

Most network managers already provide their staff with some form of technical training. Obviously, the introduction of new equipment requires training for those who use it and the information it provides. Basic technical training is also provided for user-support staff so they can taken on some of the tasks previously performed by technical specialists. However, the past emphasis on technical skills has encouraged the development of highly skilled but very specialised staff who have little interest in nontechnical areas. Such staff are often accused of focusing too much on the technology and of being unable to communicate with users or to understand business issues. In many cases, this is a perfectly valid criticism because most of the individuals concerned have received no guidance or training to cope with anything other than technology problems.

Thus, training is also required to enable technical specialists to understand other aspects of IT and to gain an understanding of business requirements. From now on, network specialists will be involved with highly integrated IT systems, and a detailed knowledge of one particular narrow technical discipline will no longer be sufficient. For example, an understanding of networking, computing, and databases is required to evaluate the potential benefits and drawbacks of a distributed-database system.

Systems departments can help network specialists to develop broader skills by:

— Providing overview courses on areas such as major applications and database technology.

- Encouraging the temporary assignment of specialists to another IT area. (Organising the systems department by function instead of technical speciality will make it easier to do this.)
- Arranging for technical staff to work for a period of time in the business areas.

The training requirements described above can be satisfied only by committing substantial resources to training. Some of the training costs will be offset by reductions in staff costs resulting from the use of less-experienced (and less-expensive) staff. However, the training budgets for some network-management functions will need to be increased significantly. Network managers must recognise that, without the appropriate investment in training and developing people, much of the investment in tools and technology will be wasted.

PROVIDING CAREER PATHS

The high level of specialisation amongst communications staff has made it difficult to provide them with opportunities for career development. Lack of suitable career paths can increase staff turnover because communications specialists may leave if they see no prospects of promotion or career developments. Providing training programmes of the type outlined above will help to open up new career opportunities, either elsewhere in the systems department or in the business.

We recommend that all network managers consider carefully how they could offer more career-advancement opportunities to their staff. Specific actions that can be taken are:

- Create individual career-development plans for those who wish to advance into management positions.
- Adopt salary scales that allow for pay increases without promotion.
- Use specialist staff to monitor technology trends and implement state-of-the-art systems.

The case history in Figure 3.4 illustrates how one company overcame the shortage of skilled staff by changing its recruitment policy and by encouraging existing staff to transfer to different areas of the network-management function.

USING TOOLS TO SUPPLEMENT STAFF

At present, skilled network-management staff spend much of their time interpreting and cross-checking the cryptic messages and reports produced by network-management tools. Over the next few years, there will be gradual improvements in the tools available, particularly in the greater use of automation and expert systems. These developments will mean that less experienced staff will be able to operate the tools. The likely improvements in network-management tools are discussed in more detail in Chapter 5.

We have demonstrated that the shortage in specialist technical skills can be alleviated by using external experts (particularly from equipment suppliers), by providing training to develop the required skills, by providing career development paths for network-management staff, and by using network-management tools that will increase the effectiveness and efficiency of the staff that are available. Another method of overcoming the skills shortage is to subcontract all or part of the network-management function, as we now discuss.

SUBCONTRACT WHERE POSSIBLE

An increasing number of PTTs, computer suppliers, and third-party network suppliers now provide services for managing corporate wide-area data networks, and it is likely that more suppliers will enter this market during the next few years. Many network managers may think that their organisations will never relinquish control of their

Figure 3.4 Career-development paths help to overcome the shortage of skilled staff

In mid-1986, a major multinational oil company was experiencing an annual turnover of communications staff in excess of 30 per cent. Inflexible central grading and salary structures made recruitment very difficult. Skilled staff were leaving because they saw few opportunities for advancement and could obtain higher salaries elsewhere.

The network manager was spending a large proportion of his time on recruitment campaigns aimed at employing experienced, qualified staff. This approach was not successful, so he decided to concentrate on recruiting less-experienced staff. Since then, communications staff have been selected on the basis of their potential skills, rather than their existing skills. About half of them are graduate trainees.

Staff joining the network-management team now have to spend between one to three years in several different sections. All inexperienced staff start in the operations area.

Existing staff are also encouraged to move on to new activities. For example, most network-operations staff are encouraged to move into sections such as network planning and design. Ninety per cent have chosen to do so. Staff can normally only become supervisors in the operations area after they have had planning experience. (In many organisations, network-planning and design activities are performed by staff with higher academic qualifications, usually a degree; network technicians usually do not have such qualifications and so have few opportunities to change roles or advance their careers.).

These policies have led to improved morale and teamwork, and have reduced staff turnover considerably. A somewhat unexpected benefit was the improved relationships between network designers and network technicians. Transferring staff between the two areas has led to a better understanding of each area's goals and priorities. Also, planning staff are able to gain supervisory experience at an earlier stage in their careers.

networks. However, many organisations are already using value-added network services and managed data-network services. They have already entrusted part of their network-management function to a third party. Figure 3.5 describes how the Corporation of Lloyd's in London bases a major network on IBM's Managed Network Service.

There is still a major role for an organisation's network manager even when the network-management function has been subcontracted. Instead of controlling internal staff, the network manager monitors the performance of the service provider and acts as a focal point for translating business requirements into network-service specifications. This latter activity requires both technical and business understanding and cannot be undertaken by a supplier.

Figure 3.5 Subcontracting network services does not mean losing control

The Corporation of Lloyd's provides services to members of the Lloyd's insurance market in London. One of the newest computer and network services, the London Market Network is based on IBM's Managed Network Service. However, the corporation controls the use of the service because it believes that IBM cannot understand the members' needs as well as the corporation can. Thus, Lloyd's staff provide the help-desk facilities, act as an interface for reporting faults and requesting changes to IBM, and gather performance statistics. Lloyd's is investing in network-management tools and is developing its own systems to provide user support.

Use of a third-party network-management service should be considered if the following conditions apply:

- Restrictions on investment in tools and on salary scales for communications staff make it impossible to meet required service levels.
- The limited size of the network means that the cost of meeting the required service levels is excessive.
- The third-party service provider can deliver the required service at a reasonable cost.

The last condition is crucial. Network-management service providers face the same difficulties and problems as an internal network-management function. They too will find it difficult and expensive to provide a high-quality service. Foundation members considering using a third-party supplier must check that the service offered meets the criteria described in this report. Experience has shown that once an organisation has committed itself to a third-party service, it is very difficult and expensive to reverse the process.

In this chapter we have shown how network managers can make the best use of the scarce skills that are available. However, having the optimum mix of skills, training programmes, career-development paths, and use of external resources is not sufficient in itself. It is also necessary to set up network-management procedures that can cope with the constantly changing environment of corporate networks.

Chapter 4

Adopting procedures suitable for a changing environment

Procedures are particularly important for network-management activities. Without adequate and well-documented procedures, the constant pressures to fix faults immediately and implement changes at short notice can lead to silly mistakes being made. Good network-management procedures reduce the likelihood of staff missing out steps or taking them in the wrong order, even when they are very busy. The main benefits of network-management procedures are that they:

- Prevent wasted time. Inexperienced staff are not left wondering what to do next. Also, the amount of duplicated effort will be reduced because it is less likely that two people will work on the same task simultaneously and not know that each other is doing so.
- Provide audit trails. The actions taken will be recorded in a log so that, where necessary, another member of staff can see what has previously taken place. This is vital in shift operations to ensure continuity in following up faults.
- Allow different members of staff to perform the same tasks. Ensuring that all network-management staff follow a set procedure in performing a task means, for example, that they all understand why a piece of equipment has been configured in a certain way. Set procedures will also ensure that someone does not install a performance-monitoring program that no one else knows how to operate.
- Ensure tasks are not neglected. The procedures should specify the frequency at which tasks such as measuring circuit-transmission parameters, reviewing network utilisation and response times, and calculating network availability should be carried out.

The current network-management procedures in many organisations are inadequate, incomplete, or not fully documented. Often, the procedures have been developed on an ad hoc basis and have not kept up with changes in circuit and equipment technology, tools, or network size. However, this report is not the right medium for providing detailed procedures for all network-management activities. (Procedures for most network management activities are provided in Kornel Terplan's book, *Communications Networks Management*, which is listed in the bibliography.) Rather, we wish to draw members' attention to the areas of network management where we believe that many organisations have either not recognised the impact of key trends that are changing the network-management environment or have neglected to implement adequate procedures.

New procedures are required for:

- Resolving faults, particularly to take account of the trend towards a lower frequency of fault occurrence, but a higher proportion of extremely difficult-to-solve faults.
- Monitoring the performance of networks. Without accurate performance measures, it is impossible to manage networks effectively.
- Handling the ever-increasing number of changes that have to be made to corporate networks.
- Managing voice and local area networks. Most network managers realise they need to manage their wide-area data networks. The problems of managing other types of networks can be just as pressing.
- Applying pressure to suppliers to ensure they deliver a better service.

Each area is discussed in turn in this chapter.

NEW TECHNIQUES ARE REQUIRED FOR RESOLVING FAULTS

A common situation that can occur if there are no set procedures for resolving network faults is that one expert carries much of the burden and his or her knowledge is not passed on to other staff. In a major US pharmaceutical company, the use of data networking within the headquarters site had grown rapidly. One particular communications engineer had developed the data-networking plans, installed and tested the initial equipment, and was involved in fixing most of the faults that occurred. As the network grew, this engineer was promoted to the position of data-communications manager

with primary responsibility for planning the future growth of the network. Two new technicians were hired to relieve him of his day-to-day operational activities. Six months later, the newly promoted data-communications manager could frequently be found fixing the most difficult faults. The technicians did their best to learn, while keeping out of his way, but they were frustrated because he did not have the time to explain how he set about solving any but the most simple faults. There were no written procedures to help the technicians, and they began to feel that they were not doing a worthwhile job. Meanwhile, the data-communications manager was complaining that the only time that he could do his planning work was in the evenings or at weekends.

The above experience demonstrates the value of a well-documented set of procedures for resolving network faults. Fault resolution is still one of the most important and most time-consuming network-management activities, so part of our research concentrated on identifying trends in the frequency and nature of network faults. We also assessed the impact these trends would have on fault-handling procedures.

TRENDS IN FAULT PATTERNS

During our research we attempted to obtain quantitative data on trends in network faults and we sought network managers' and suppliers' opinions about this subject. Our overall findings and conclusions are summarised in this section of the report. The detailed results of the survey we carried out are contained in the Appendix. Unfortunately, we were not able to obtain much numeric data because most of the organisations we surveyed do not analyse their fault statistics. (However, several told us that they were in the process of installing a fault-reporting software package that would provide fault statistics.) Two main trends have been identified by the network managers and suppliers we spoke with:

- The number of faults per user is decreasing, and in some organisations even the total number of network faults is decreasing.
- The average time to resolve a fault is increasing.

The first trend was more noticeable in those countries where digital circuits and equipment are now in widespread use. Most network managers believe that digital circuits are more reliable than analogue circuits. They did comment, however that digital circuits seemed to go through a settling-in period after installation during which they tended to fail more often. Another comment made by several network managers was that a digital circuit would operate for months without a single failure, but that several failures would then occur within a short period.

Most respondents in all the countries surveyed had noticed the trend towards an increase in the average time required to resolve a fault. Reasons cited for this trend were greater network complexity and the resultant difficulties of diagnosing the causes of faults. The minority of respondents who said that their average fault-resolution time was decreasing commented that this was due either to the use of better tools or to the fact that the network-management team now had experience of resolving most of the faults that could occur. However, we believe that the trends identified by the majority of respondents will continue over the next few years because:

- Organisations will increasingly benefit from the long-term investments being made by PTTs in newer, more reliable technology such as optical-fibre transmission links and digital switching. As these technologies are deployed in public networks, the quality of network circuits should increase.
- Equipment suppliers and the PTTs will provide more automated fault-recovery facilities, which will prevent some faults from occurring and will allow minor faults to be fixed quickly.
- As a consequence, and as networks grow in size and complexity, the proportion of hard-to-solve faults will increase. These faults will have obscure causes and may take several days to diagnose.

These trends are illustrated in Figure 4.1. Figure 4.2 overleaf shows the way in which we believe

Figure 4.1 There will be fewer network faults, but they will be harder to fix

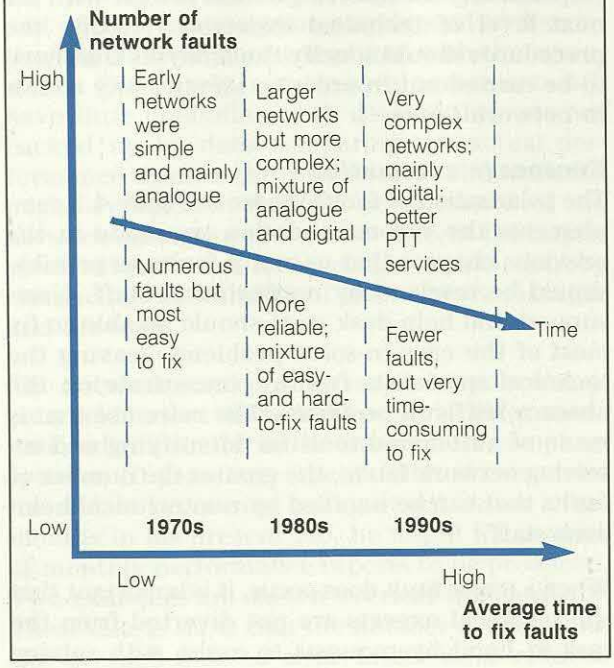

Figure 4.2 Future distribution of fault-resolution times

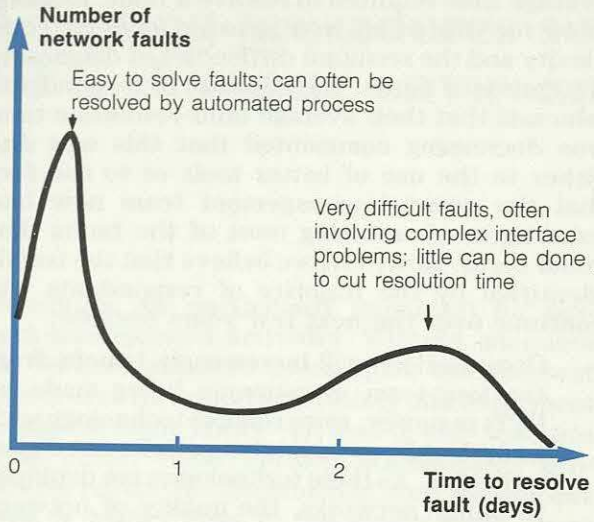

fault-resolution times will be distributed in the future. Faults will tend to polarise into those that are easy and quick to diagnose and repair, and a smaller number that are very hard, and take a long time, to diagnose. The consequence is that most of the fault-resolution effort will go on fixing the difficult faults.

RESOLVING FAULTS

The procedures for resolving network faults should clearly take account of these trends. In particular, the procedures should encourage the use of automated tools and should make provision for the fact that there will, on occasion, be lengthy breaks in the network service. They should also specify clear rules for deciding when to pass the responsibility for resolving a fault to staff with the next level of technical expertise. Finally, the procedures should specify the analyses that need to be carried out in order to identify any trends in network faults.

Encourage automation

The polarisation of faults shown in Figure 4.2 complements the recommendation we made in the previous chapter that as many faults as possible should be resolved by nontechnical staff. User-support and help-desk staff should be able to fix most of the easy-to-solve problems, leaving the technical specialists free to concentrate on the obscure, difficult problems. The more use that is made of automated tools for identifying and resolving network faults, the greater the number of faults that can be handled by nontechnical help-desk staff.

When a major fault does occur, it is important that the technical experts are not diverted from the task at hand by requests to assist with solving other minor faults. Providing nontechnical staff with automated assistance should reduce the probability of this happening.

Make provision for lengthy breakdowns

As the average fault-resolution time becomes longer, network managers should ensure that they have fallback plans for the times when there will be lengthy interruptions to the network services. These plans will specify the back-up procedures in the event of a lengthy breakdown at a network node or on a network link. Most major networks contain alternative routes that can be used, and some of the equipment at the main network nodes is usually duplicated. However, it may be too expensive to duplicate single routes to small sites, and in this situation it may be necessary to provide dial-up facilities for emergency use. The type of back-up provided should depend upon the estimated probability of a failure, the cost of the back-up facility, and the terms of the service-level agreement.

Most systems departments have back-up plans for recovering from a major disaster such as a computer-centre fire. In addition, it is worthwhile spending some time considering the effect of less dramatic network failures on small groups of users, particularly when negotiating service-level agreements. The back-up facilities should then be selected on the basis of this analysis. A major oil company, for example, has a policy that a failure should not affect more than half the terminal users in any department. The only exceptions are at very small sites where this policy would not be economic.

Back-up facilities are only useful if they work when they are required. All standby equipment, circuits, and procedures should therefore be tested regularly.

Specify fault-escalation procedures

There should be clear procedures for deciding if a fault is beyond the competence of the help-desk staff to solve, and for passing it on to the most appropriate technical specialists. The best way of achieving this is to specify time limits for attempts to solve a problem before it is passed on to a higher authority. This is particularly true for the help desk. Help-desk staff should also be given clear rules for deciding who to pass the problem on to. The procedures should also specify when the first level of technical support should pass on the problem to in-house specialists. In turn, they should also be provided with rules for deciding when to involve the equipment suppliers, and the procedures for doing this.

Analyse trends

Identifying the causes of transient and intermittent faults is very difficult because the faults appear to correct themselves and the symptoms

disappear. A similar fault may not occur for days, weeks, or even months, by which time the network-management staff are likely to have forgotten about the previous occurrence of the fault. What is required is a comprehensive problem-tracking and fault-reporting system that allows network-management staff to monitor trends in the occurrence of faults.

In addition, identifying trends in the occurrence of error messages can be helpful in pinpointing potential faults before they actually occur. For example, a network technician might become aware of an increased number of a certain type of error message. By referring back to a similar situation in the past, he or she might be able to inform their manager that, ''The last time we had this number of this type of error message there was a processor overload a few days later''. The technicians in one network-management centre we visited, had asked their systems department to write a program to count the occurrences of alarm messages by type and by network location. The summary reports produced by the program indicate potential faults in the network.

Identifying trends in fault and alarm messages is one of the few methods that network managers can use to anticipate network failures. However, much useful information can also be gained from systematically measuring the performance of the network.

PERFORMANCE MONITORING SHOULD NOT BE NEGLECTED

Many of the automated network-management tools available today provide large volumes of network-performance data. However, the data is produced in a form that is unintelligible to nontechnicians, and requires much manipulation to produce meaningful performance-monitoring information. We were not surprised to find, therefore, that many network managers have neglected performance monitoring because of the lack both of adequate tools and of staff time to process the raw data.

Without an effective means of monitoring the performance of its networks, the network-management function cannot monitor trends in faults or the performance of individual equipment suppliers. The lack of adequate performance measures means that some organisations do not know what the availability of their network is. Indeed, some organisations are unclear about how to measure overall availability correctly. In other organisations, the lack of performance measurement means that the network-management team cannot anticipate the need for additional capacity; their networks are upgraded only after users have complained about poor response times.

THE NEED FOR PERFORMANCE MONITORING

As with any other business activity, a good rule-of-thumb for network managers is 'you cannot manage what you cannot measure'. Without adequate performance measures, network managers do not know if the performance of their networks and staff is improving or deteriorating. There is no quantitative basis for justifying further staff or investment in new tools. Lack of performance data also means that it is more difficult to convince senior management that the network-management team is doing a good job. Senior managers will still hear about complaints from disgruntled users, but will not be aware, for example, that the network-management function is handling x per cent more users and y per cent more equipment changes, even though there are two vacant positions that have not been filled for three months.

Accurate data about response times can also help to resolve disputes between users and network-management staff. Without accurate data, users are inclined not to believe the network-management staff. More importantly, they will not acknowledge when an improvement in response times is achieved. In fact, they will often continue to complain about the performance of the network, even after improvements have been made.

Network-performance measures are essential where there are service-level agreements for the provision of network services because network managers must be able to demonstrate that they are meeting the terms of the agreements. Indeed, part of their remuneration may depend on demonstrating that they are complying with the agreements. Conversely, service-level agreements will have little credibility with users if they are not backed up by data comparing the actual performance against the set criteria. The statistics might even show that the performance of the network-management function exceeds the targets specified in the service-level agreements, thereby enhancing the standing of the network manager and his team.

The network manager at a major insurance company told us that the lack of network-performance data today is similar to the situation for computer systems 10 years ago (the individual concerned has a data processing background). After a few months in his present job, he asked for a series of monthly performance reports to be produced. Two examples are shown overleaf in Figure 4.3. These charts show that the number of faults and the average time to clear faults were reduced

Figure 4.3 Examples of trend-analysis graphs

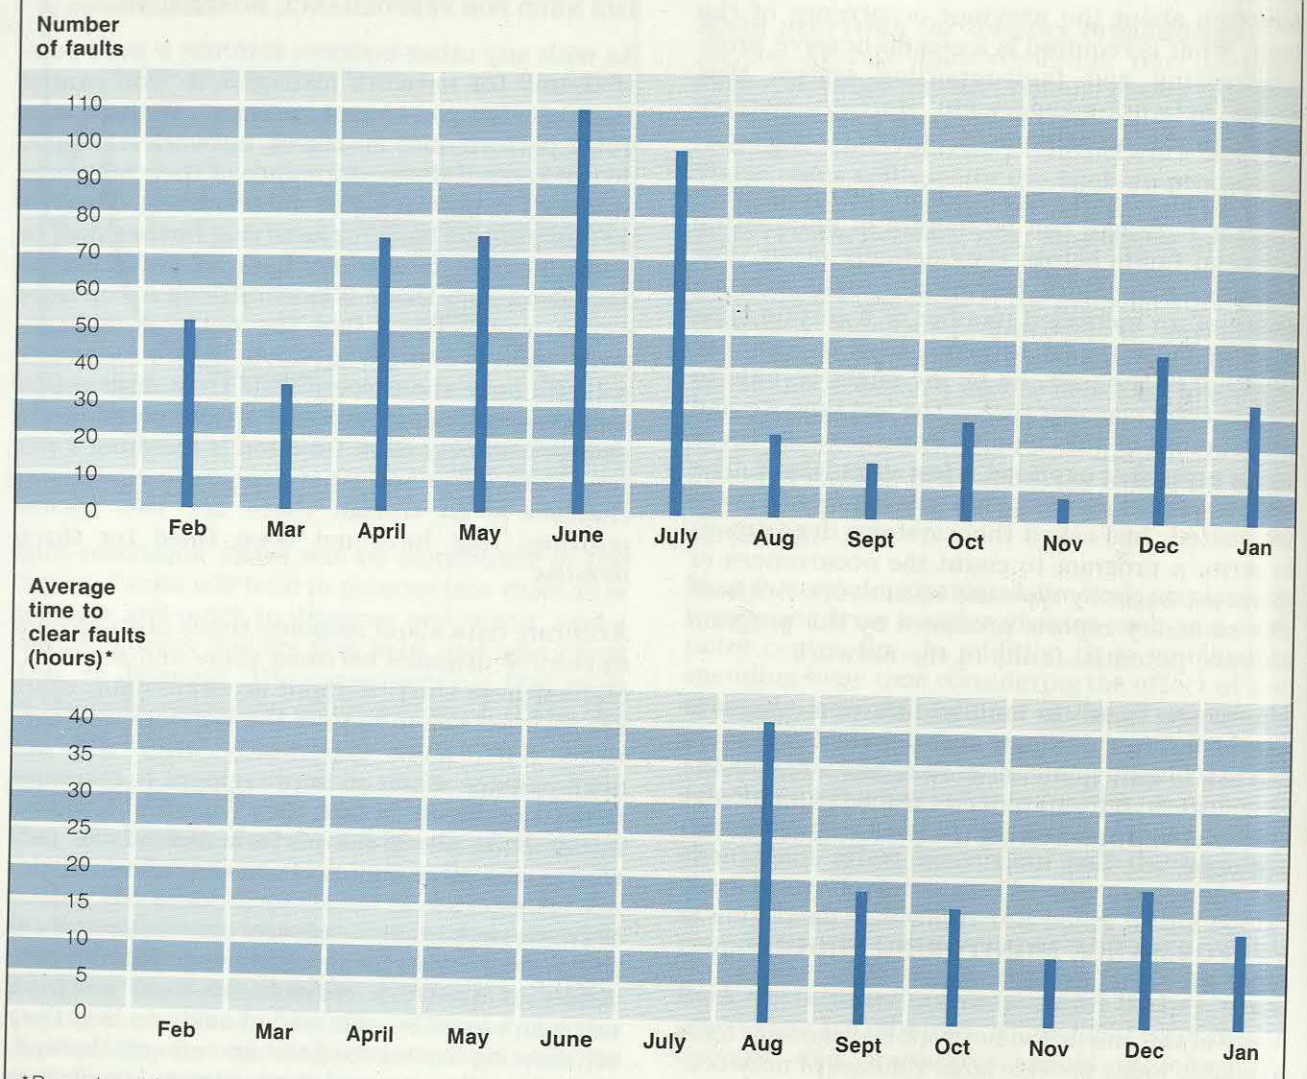

considerably in August and September respectively. The network manager attributed these reductions to his staff's responses to the performance charts. The graphic representation of the number of faults and the time taken to clear faults caused British Telecom to give more attention to fixing faults.

It is important to present performance data clearly. Graphs such as those shown in Figure 4.3 are much easier for users and senior management to understand because they show trends rather than discrete numbers. Trend graphs can also be useful if there are peculiar circumstances that cause the network-management function to fail to meet its service-level agreement targets in a particular month. A graph showing the complete annual picture will place the adverse performance in one month into perspective.

BENEFITS OF PERFORMANCE MONITORING

Few, if any, automated network-management tools produce trend graphs similar to those shown in Figure 4.3. As a consequence, network managers either use manual analyses to produce suitable performance statistics or have developed their own spreadsheet applications. The few network managers we talked with who had made the effort to produce comprehensive performance analyses said that it took several days every month to prepare the statistics. They believed, however, that the effort was well worthwhile because the analyses allowed them to:

- Change the opinions of user and senior management about network performance.
- Identify areas of concern and improve the network-management team's knowledge about the performance of the networks.

- Put pressure on suppliers to improve their performance.
- Set realistic targets for service-level agreements.

IMPROVED TECHNIQUES ARE REQUIRED FOR HANDLING NETWORK CHANGES

One of the most important network-management activities is to manage the changes that take place to the network — adding new users, moving the equipment of existing users, providing additional capacity or features for existing users, upgrading network software or hardware, and so on. Constant change is a fact of life for network managers. The advice of the systems manager at Credito Emiliano (an Italian banking, finance, leasing, and factoring organisation) is that network managers should never regard the network as having a 'steady-state'. Clearly, as the number of network users grows so too will the number of changes to the network.

TRENDS IN CHANGES

As Figure 4.4 shows, some Foundation members reported that the number of changes was increasing at a faster rate than the number of users. Network managers are also being pressured by their users to carry out changes more quickly. Sometimes, network managers receive less than a week's notice of an impending office move and are expected to complete all the necessary work between 5pm on Friday and 9am on Monday. In some cases, the short notice is due to lack of foresight by user departments, although the users themselves do not always receive much notice of an impending move. In general, network-management staff meet the tight deadlines, although it can lead to much overtime being worked.

Figure 4.4 Some organisations report that network changes are increasing faster than network users

Organisation's industry sector	Annual change in wide-area data-network users	Annual growth in volume of network changes
Energy	+20%	25%
Finance	−35%	120%*
Finance	+14%	20%
Manufacturing	+20%	10%**
Transport	+25%	50%

* A decrease in business due to the October 1987 slump in share prices has resulted in a decline in the number of users, but a large increase in network-change activity due to subsequent re-organisations

** This company expects to move between 20 and 25 per cent of all network users each year. Ten per cent growth in change volume, therefore, indicates faster growth in change activity than in the number of users

However, where the changes require new equipment or parts to be ordered, the network-management function often cannot meet the users' required deadlines. The lead times for ordering new equipment, or additional circuit boards for existing equipment, or new circuits or exchange lines, are usually between four and eight weeks — although several months is not uncommon. Network managers told us that their users often find such delays frustrating and difficult to understand.

A large proportion of network faults occur immediately after network changes have been made, particularly major network-configuration changes and software upgrades. The complexity of many networks can mean that a change uncovers a latent fault or creates a problem that is not immediately apparent. (At one of our focus groups, one Foundation member reported that, for this reason, his organisation did not carry out major network-reconfiguration work unless it was absolutely necessary; few organisations have the freedom to operate in this way, however.)

The trend towards changes that need to be carried out at short notice reflects the growth both in network use and the importance of networks to the business. The network manager's objective should be to carry out the increasing number of changes with minimal disruption to users and without an increase in staff numbers. Some techniques for achieving this objective are described below.

HANDLING THE CHANGE WORKLOAD

Managing network changes means that changes are coordinated and planned, rather than implemented hastily to respond to urgent user requests. The management procedures required depend on the type of change and its complexity — moving a terminal from one side of an office to the other requires far less preparation than installing new multiplexors on the backbone network. Based on our discussions with experienced network managers, we recommend the following guidelines for managing network changes:

- Create a database containing a comprehensive inventory of network components and a description of the network configuration.
- Evaluate change requests as they are received and assign them the appropriate priority.

- Assess the potential impact of major changes on the network and its users.
- Always test the changes thoroughly before they are implemented.
- Expect the changes to cause network failures.

The reasons for these guidelines are described below.

Create a comprehensive inventory and configuration database

Without a comprehensive record of the equipment already used in the network, and the way it is configured, it is impossible to carry out changes efficiently and effectively. For example, to determine whether adding a new terminal or telephone extension requires additional hardware in the multiplexor, cluster controller, or PABX, it is necessary to know if there are spare ports available. Similarly, when planning an office move it is necessary to have an inventory of all the communications equipment to be moved and where it is currently located. One way of obtaining the information is to visit each site and carry out a survey. However, site surveys are time-consuming because someone must physically inspect all the equipment and make a written record. In addition, considerable time may be required to travel to remote sites.

However, as the volume of network changes increases, the network-management team will need to ensure that it has ready access to complete and up-to-date inventory information. There will be insufficient time for a quick site visit to check that the inventory records are up to date. An effective way of overcoming these difficulties is to maintain an inventory database. Many organisations have developed some form of network-inventory database, often using a PC-based package such as dBase III or Focus. Others have purchased sophisticated inventory-control packages from computer suppliers or independent software houses.

The inventory database should also contain information about spare equipment and excess capacity. In an emergency, it might, for example, be useful to know that spare equipment in Dusseldorf could be used to solve a problem in the Frankfurt office. For larger networks, the inventory system should be able to record the fact that spare equipment has been reserved for planned changes. This will prevent a member of the network-management team who is not aware of the impending changes from using the spare equipment for another purpose.

A good inventory database will also store information about the network configuration so that network-management staff can easily determine which circuits, equipment, and cabling are linked together. More sophisticated 'change-management' packages can even automatically produce orders for suppliers and weekly lists of tasks to be performed by network technicians.

Although in-house PC-based inventory databases have provided benefits, they may not be able to handle the large amounts of information that network-management teams now need to maintain. We recommend, therefore, that Foundation members should review their present inventory systems. Check how often the inventory information is out of date or how frequently equipment installers find they do not have all the components required to complete a job on the first visit. Assess how the existing inventory system and procedures would cope with a 50 per cent increase in the number of workstations and ancillary equipment attached to the data network, and a 60 to 70 per cent increase in the volume of network changes.

For most organisations, the cost of creating a comprehensive inventory database will be repaid in a very short time. As we show in Chapter 5, the inventory and configuration database will be at the heart of future automated network-management systems.

Assign priorities

Most network managers have limited staff available to carry out changes, so all requests for changes should be evaluated as soon as they are received to estimate the resources required and to assign a priority to the work. If the network-management team believes that the requested deadline cannot be met, the originator of the request should be notified as soon as possible and given a realistic implementation date. Either a paper-based or, preferably, a computer-based change-control system will allow the installers to plan their workload and to check that suppliers deliver on time.

Assess the impact of changes

When planning a major change to a network, it is good management practice to assess the impact that the changes might have on the network itself and on the computer systems that use the network. Questions that should be asked during an impact analysis include:

- Which groups of users need to be notified about the change?
- Will the change affect response times adversely?
- Is the change scheduled to coincide with the running of critical applications?
- Is the change scheduled for a time that will cause least disruption to users?

- Are contingency plans required if the work cannot be completed during the planned period?
- Are there potential interface or software incompatibilities?

Answers to these questions will help to determine the correct timing for the changes and to allocate the resources required to complete all the tasks.

The objective of the impact analysis is to minimise the chance of unanticipated problems occurring when the change is implemented.

Test the changes thoroughly

The need to test major changes thoroughly before they are implemented can easily be overlooked. A complete test plan should be devised before making major changes to the network configuration or software. The test plan will ensure that network-management staff check, for example, that routeing tables have been correctly updated. It is also necessary to check that interfaces, particularly less common ones, work in the same manner after a software or hardware upgrade has been carried out. A special software routine written for a particular interface may not be included in the new version of the software. Although suppliers will have tested any software or hardware before it is supplied, the tests are unlikely to have duplicated the organisation's exact network configuration. Often suppliers can provide routines for testing their products after they have been installed.

Some organisations purchase additional equipment that is used just for testing purposes. One of the top four UK clearing banks, operates a large X.25 network, based on Plessey/Telenet network switches. With the exception of British Telecom, the bank uses more of this type of equipment than any other UK organisation. It has installed a spare network switch which is used to test new software versions before they are implemented. Even with this precaution, there have still been a few occasions when network faults have occurred as a result of a software upgrade.

Expect failures

The above example shows that even the best-planned changes can cause network failures. Network managers should not assume that a change has been completed until it has been used for some time in an operational environment. This means that the network-management team should be ready to respond immediately if faults occur after a change has been implemented. In some cases it will be prudent for network-management staff to remain on-site ready to assist users and resolve any minor problems as they occur.

In addition, the help desk should be aware of the changes that have been made so they can look out for any problems that might be related to the changes.

VOICE AND LOCAL AREA NETWORKS WILL ALSO NEED CLOSE ATTENTION

In Chapter 1 we commented that most network managers tend to concentrate on managing their wide-area data networks. However, considerable changes are also occurring in the areas of voice and local area networks, and the growing problems of managing these types of network can easily be overlooked. In our initial questionnaire we asked Foundation members about their network-management concerns. Their responses are summarised in Figure 4.5. Wide-area data networks are clearly a concern for most members, with more than two-thirds of the respondents rating them as needing improvement or as a major concern. However, only half of the respondents rated the management of local area networks in the same way. There is much less concern about the management of voice networks, with more than half the respondents rating it as not being a problem or a minor concern. (Australian members were much more concerned about voice-network management, however, due to the distances involved.) There are four main reasons why network managers in general are much more concerned about the management of their wide-area data networks:

- The rapid growth in data networking means that data networks currently cause more problems than the other types of network.
- There is almost no growth in the number of voice-network users. Moreover, if the voice

Figure 4.5 Foundation members are more concerned about the management of wide-area data networks than other types

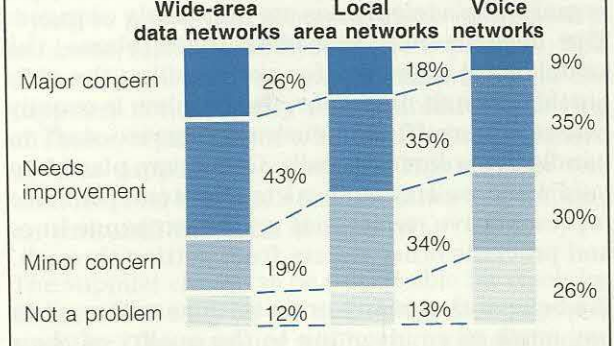

The figure shows the proportions of respondents who expressed each level of concern for managing the different types of network

(Source: Butler Cox survey)

network fails, telephone users can always use the public network. It is therefore not usually necessary to provide a voice network that has very high availability.

- In some European countries, much of the management of voice networks (and PABXs) is still controlled by the PTT.
- Local area networks are only now beginning to be widely used and be connected to corporate networks. Up to now, the performance of smaller local area networks has not been a problem, and they have required little management attention after they have been installed.

We believe that the situation is changing, however, and that network managers will need to give greater attention to voice and local area networks. In our consultancy work, we have identified several voice-management problems that many network managers are not yet aware of. In addition, the management of large local area networks is now too complex for users to handle on their own.

VOICE-NETWORK MANAGEMENT

Many network managers fail to recognise that voice-networking problems exist within their organisations. This is due partly to the lack of suitable network-management tools and partly to the fact that many of the problems are not immediately obvious. A common problem concerns the inadequate provision of telephone-answering facilities in customer-service departments. Organisations are handling an increasing proportion of customer orders and enquiries via the telephone, often with customer-service departments set up specifically to handle telephone queries. In organisations such as insurance companies and other financial institutions, departmental managers frequently underestimate the volume of calls that will be received. As a consequence, the telephone-answering service is poor and no attempt is made to monitor the volumes of calls being received. Inevitably, customers begin to complain when they find that the organisation's telephones are constantly engaged. The departmental managers often blame the switchboard operators for not handling the calls quickly enough. However, the problem is usually caused by insufficient customer-service staff to handle the volumes of calls. Callers are placed on hold while waiting to speak to a customer-service representative, which ties up the exchange lines and prevents other callers from getting through.

We believe that many organisations will need to pay much more attention to the quality of their customer telephone service. Organisations such as airlines and mail-order companies use specialised telephone equipment called automatic call distributors to handle customer calls effectively and to monitor the level of service being provided. Network managers need to provide more advice and guidance to line managers about what is available in this area.

The most common type of voice network-management tools (call loggers) are used primarily to determine how costs should be recharged to users. However, the majority of organisations rarely monitor their telephone usage or even check the PTT's invoices. Sometimes, network managers are not entirely to blame for this situation. In the Netherlands, for example, some companies have to wait for up to a year for a voice-traffic study to be carried out because the PTT does not have sufficient call-logging systems to meet demand.

We find the lack of attention to managing voice networks to be surprising because voice-communications costs are about four times higher than data-communications costs in most organisations. Yet most European organisations spend far less effort on controlling voice costs, even though effective management can yield substantial savings. One UK company installed two new private circuits between two of its offices. Nine months later the company discovered that the new circuits had not been entered in the PABX routeing tables. Thus, the circuits were not being used because the PABX was not aware of their existence.

We believe, therefore, that network managers should review their procedures for managing voice networks and should determine where more effort or new tools are necessary.

LOCAL AREA NETWORK MANAGEMENT

The number of local area networks installed is increasing steadily, as is the size of such networks. Some organisations now connect local area networks to their corporate wide-area data networks, which means that the local network should be subjected to the same control procedures for security and change management as any other network component.

In addition, many user departments are now finding that local area networks are becoming too complex for them to manage themselves, and they are seeking support from the network-management function. It is easy to underestimate the amount of support that a large local area network installation requires. In particular, error messages produced by local area network monitors are difficult for nontechnical staff to interpret. Also, user departments cannot rely on the supplier to support local area network installations because many suppliers are not familiar with all the technical details of the products they sell.

An example of the problems that can be caused by poor management of local area networks is provided by a company that began to experience performance problems with its network about a year after it was first installed. The cause of the problem turned out to be a poor network termination. The fault occurred when the network was originally installed, and was causing the local area network to operate at less than 10 per cent of its design throughput. However, previously the network usage was so low that no one was aware that a problem existed.

Several tools for monitoring the performance and usage of local area networks are now available. Examples include Network General Corporation's Sniffer product and Excelan's Lanalyser. These tools are designed for on-site use — they do not provide remote-diagnostic features. However, there are a few products that do provide these features. Examples include DEC's Remote Bridge Management Software, which is used with Ethernet networks, and IBM's NetView/PC interface for its Token Ring networks.

Until a wider range of local area network management tools with remote diagnostic and control functions becomes available, the central network-management function will find it very difficult to manage remote local area networks. There are no obvious or easy solutions to this problem. However, network managers should adopt the following procedures in order to improve the support they provide for local area networks:

- Make one person in the network-management team responsible for carrying out regular on-site checks of the performance of all remote local area networks.
- Test all new local area network products at the network-management centre before they are installed to ensure they are compatible with existing hardware and software. Do not allow user departments to install a product until the tests are completed.
- Establish in-house standards for interfaces between local area networks and the corporate wide-area data network.
- Create standard procedures to be used by user departments for local area network management. The procedures will cover the maintenance and updating of software-control tables, which contain details about the equipment attached to the network and information about which users are allowed to access which equipment and network services. A nominated user representative should be trained in how to use the procedures.

PRESSURE ON SUPPLIERS WILL IMPROVE THEIR SERVICE

As users' demands and pressures on the network-management function increase, network managers need to increase the pressure on their suppliers to deliver a better service. User demands for guaranteed service levels, faster response times, and faster network changes should be reflected in the network-management function's relationships with its suppliers. Thus, the network manager should be seeking remote-diagnostic facilities, guaranteed repair times, and penalties for late delivery from equipment suppliers and PTTs. We were encouraged to find that several organisations are beginning to obtain a greater service commitment from their major equipment suppliers. Some suppliers appear to have recognised that the levels of service and support they provide will become major criteria in equipment-selection exercises.

In general, the PTTs were identified as the most unresponsive suppliers. Members in France, for example, reported considerable difficulties in getting the PTT to repair faults on transmission lines. France Telecom has no central procedure for coordinating reports of faults so it can be difficult to identify who is responsible for repairing a particular line type. In an extreme case, a failure on a leased line to a plant in Avignon took one month to fix. However, in several countries, including France, members have noticed signs that the service provided by the PTTs is beginning to improve.

The most useful techniques for improving the service provided by suppliers are discussed below.

MAKING ONE SUPPLIER RESPONSIBLE FOR RESOLVING FAULTS

A major difficulty in resolving faults in a multi-vendor network is deciding which supplier's equipment is the cause of the problem. Some suppliers have a tendency to spend more time on trying to prove that someone else's equipment is to blame, rather than on solving the problem. This difficulty can be overcome by making one of the suppliers contractually responsible for resolving all faults, regardless of whether the fault is caused by that supplier's equipment. This supplier will then respond to all calls for assistance and will stay on-site until the fault is repaired.

The supplier chosen to be responsible for resolving faults will usually be the main network-equipment supplier. This supplier will have the largest maintenance contract and therefore has more incentive to agree to nonstandard contract terms. The contracts with other suppliers should require them to cooperate with the main supplier.

TIGHTENING SERVICE GUARANTEES

When a network-management function signs a service-level agreement with its users, it is committing to provide a minimum level of availability or a maximum amount of downtime. However, most suppliers' maintenance contracts only specify a guaranteed time to respond to a fault call — not a guaranteed time to fix the fault. Network managers must therefore avoid being caught between the requirements of their service-level agreements and the terms of their suppliers' maintenance contracts.

To avoid this difficulty, some large organisations are beginning to insist that suppliers' maintenance contracts specify a guaranteed time to fix faults. Suppliers can be pressured into accepting these harsher terms because maintenance contracts are a substantial source of revenue for many of them. (Over 7 per cent of IBM's total revenue comes from equipment maintenance.)

In some countries (notably the United States and the United Kingdom) there is a well-established third-party maintenance market that may provide a further incentive for suppliers to agree to more stringent contract terms.

Some PTTs are now setting internal targets for repair times as a first step in moving towards guaranteed repair times. Telecom Australia, for example, aims to fix 70 per cent of faults within four hours. This PTT is also working with some of its customers to ensure that their internal service-level agreements can be met. Also, British Telecom has recently announced that it aims to repair all business-subscriber faults within five hours.

The service performance of suppliers can also be improved by including penalty clauses in the maintenance contract for not meeting guaranteed delivery dates and by agreeing to pay bonuses when work is completed ahead of schedule.

USING PERFORMANCE CHARTS

As mentioned on page 27, performance charts can be an effective way of goading suppliers into providing a better service, particularly when the charts compare the failure rates and average repair times for different suppliers' equipment. Some organisations display such charts in a prominent place like the lobby of the network-control centre. The immediate visual comparison of the performance of different suppliers gives each supplier an additional incentive to perform well.

HOLDING REGULAR REVIEWS

Many network managers hold regular monthly or quarterly review meetings with their major suppliers. The purposes of these meetings are to:

- Review the progress on outstanding faults, and determine why some faults may have taken longer to repair than was necessary.
- Inform the supplier of anticipated changes and make sure that the supplier is aware of the organisation's priorities for new equipment and network changes.

Usually, both the supplier's main service representative and the account manager attend these meetings. (The account manager will of course, be keen to obtain further sales, so has an incentive to ensure that persistent problems are solved.) They both gain a better understanding of the concerns of their customer from the regular review meetings.

USING ONLINE ACCESS TO SUPPLIERS

Some network-equipment suppliers are now providing their customers with the ability to access portions of their internal computerised service-management system via dial-up communications. Customers can use these facilities to:

- Report faults.
- Order new equipment or circuits.
- Monitor the progress of orders or reported faults .

The benefit to the customers is that they can obtain up-to-date and accurate information without having to call the supplier's customer-services department.

AT&T recently announced that this type of facility, which it calls NetPartner, will be part of the telephone-company management system it is selling to Regional Bell Operating Companies and other PTTs. IBM also plans to provide a similar facility as part of its Managed Network Service. Within the next few years, this type of service is likely to become a common feature provided by value-added network operators.

DEVELOPING INDIVIDUAL RELATIONSHIPS

Sometimes, network managers find that they are unable to obtain an adequate response from their suppliers through the official contact channels. As a consequence, they develop their own individual contacts within the suppliers and use these to stimulate the required action. Some network managers told us that this is the only effective method for obtaining prompt service from their PTT, although the situation is improving as PTTs are placing more emphasis on customer service. However, any such unofficial contact should be used with discretion because it bypasses the normal controls and procedures and can aggravate relationships with the supplier.

A few network managers deliberately set out to recruit people who currently work for their PTT or suppliers. They believe that the personal

contacts these people have will help them to obtain better service in the future.

Having ensured that the basic network-management procedures are in place, it is then necessary to consider the network-management tools that are available. In particular, it must be recognised that existing tools are inadequate and will remain so for some time to come. The next chapter describes this problem in detail.

Chapter 5

Recognising the inadequacies of existing tools

Our research confirmed that network managers are well aware of the inadequacies of existing network-management tools. Despite there being a clear need for better tools, we believe that most organisations will be unable to purchase a comprehensive integrated network-management system until well into the 1990s.

In this chapter, we first set out the requirements for an integrated network-management system and provide a model of the ideal system. We then examine the commercial and technical pressures that have caused suppliers to respond slowly to the need for integrated network-management systems. The chapter concludes with a review of the progress that is being made and some of the likely developments.

THE REAL NEED IS FOR AN INTEGRATED NETWORK-MANAGEMENT SYSTEM

It is important to recognise the differences and the relationships between network-management tools and network-management systems. A tool is a piece of equipment and/or software that automates at least part of one or more network-management activities. Network-management tools are usually associated with particular pieces of communications hardware. A network-management system is a combination of network-management tools that, together, automate a range of network-management activities. Ideally, the tools should form an integrated system and should work across different ranges of hardware. We would classify most of the so-called network-management systems available today as tools, with the exception of some systems available from independent suppliers such as Atlantic Research and Avant Garde.

FUNCTIONAL REQUIREMENTS OF A NETWORK-MANAGEMENT SYSTEM

Most network managers told us they require an integrated network-management system instead of the collection of discrete and incompatible tools they are forced to use today. Our research found that network managers' opinions about the most important features of an ideal network-management system are remarkably consistent. Most of them believe a network-management system should:

- Perform as an integrated whole, even though the system may consist of several pieces of equipment and software from different suppliers. The term frequently used to describe this feature was 'seamless'.
- Collect information from, and control, any network component.
- Support the majority, and preferably all, network-management activities.
- Minimise duplication of information. Ideally, there should be no duplication.
- Automate routine tasks.
- Provide a consistent and easy-to-interpret user interface.
- Display graphically, in realtime, the network's current configuration and status.
- Reduce the expertise or time required to perform an activity.

Figure 5.1 shows very similar requirements emerging from a study conducted in the United States. The only difference from the requirements stated above is that US network managers would like the monitoring facilities extended to applications and systems software. Such an extension is consistent with our view that the network-management function will merge with the management of operational computer systems.

The requirements for network-management systems are bound to evolve over the next few years as advances in technology lead to new network features and new methods of providing better systems.

BETTER INFORMATION IS REQUIRED

Most network-management tools present information such as alarm messages and usage statistics in terse formats that employ incomprehensible acronyms. Moreover, each tool uses its own unique set of formats and acronyms, which means that it is difficult to compare and collate performance and usage data produced by different tools. An added problem is that a parameter (such as

Figure 5.1 Requirements for the next generation of network-management systems

Reduction in technical staff required for network operation.
Reduction or elimination of network downtime by automatic self-healing or bypassing initiated by the network-management system.
Each network element monitored by the system.
Reduction or elimination of network-management system duplication.
Better and more straightforward interface with human operators.
Better compilation, analysis, and presentation of statistical and general performance data.
Monitoring extended to systems software, communications software, applications software operation, hardware, and at least Level 4 of the OSI model.

(Source: International Resource Development Inc.)

response time) can be defined and measured differently by different tools. And many of the warning messages generated when specified parameters are exceeded often need not be acted on immediately. The result is that many of the messages produced by network-management systems are ignored by network operators. Unless the network operator is very familiar with the equipment concerned it can be very difficult to interpret the information produced.

Another difficulty arises because a network fault may cause alarm messages to be created by several tools. Thus, a modem-management system and a multiplexor-management module may each generate an alarm message when a circuit fails. Someone with the appropriate skill has to compare the two sets of messages carefully to identify which of the alarm messages are related. In large network-management centres, there might be eight or more sets of alarm messages being generated, which means there is considerable duplication of the messages. Also, when a major failure occurs a large number of messages are generated, one for each symptom. Network operators have to diagnose the cause of a fault whilst they are being bombarded with a large number of alarm messages. In some cases, a network-management tool can generate up to 100 lines of messages within a minute.

Given the difficulties described above, it is not surprising that network managers are seeking a common and easy-to-interpret format for the output from network-management tools. One approach is to use graphical displays to represent network configurations and indicate (usually in red) where failures have occurred. A pictorial presentation can be interpreted much more quickly and from a greater distance than a text alarm message. Some multiplexor manufacturers now provide quite sophisticated graphics-based network-management tools. One example is the Integrated Network Manager, shown in Figure 5.2, available from Infotron Systems.

The benefits provided by better information from network-management tools are that it takes less time to interpret and cross-correlate information, and that less training is required before staff can use a particular tool. In addition, future network-management tools will provide facilities for allowing the base data to be analysed for a variety of purposes — network design, performance monitoring, and the production of customised reports, for example. Thus, failure messages could be analysed automatically to produce availability statistics, and usage data could be formatted so it can be used by a network-design tool. The result is that network-management staff will be able to spend more time interpreting data, rather than extracting and analysing data as they do at present.

AUTOMATION OF NETWORK-CONTROL TOOLS

A large number of messages produced by network-management tools require a standard response such as resetting a line. Today, the responses are usually actioned by a network operator, although, in theory, routines can be set up to generate the required responses automatically. In the future,

Figure 5.2 The Integrated Network Manager, a graphics-based network-management tool

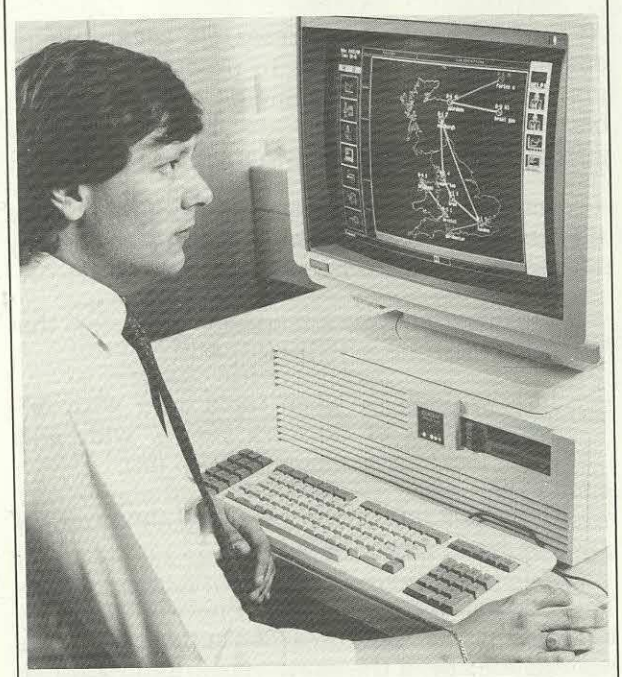

network-management tools and systems will provide much higher levels of automation.

Even today, some degree of automation can be achieved. For example, most network-management tools provide facilities for changing the parameters that trigger alarm messages. The parameter settings should be checked regularly to ensure that the tool does not generate unnecessarily high volumes of messages. And personal computers can be programmed to respond automatically to standard messages.

IBM estimates that more than 60 per cent of the messages displayed by its NetView network-management products could be eliminated if network-management staff developed special routines to filter out unnecessary messages and to generate automatic responses to standard messages. However, these routines need to be written in a low-level language to produce what are known as CLISTs. IBM has recognised that many NetView users do not make best use of CLIST routines and now provides a chart describing the functions that these routines can perform.

The benefits of greater automation are reductions in the routine workload, faster response to standard, easily solved, problems, and fewer messages that have to be considered by network operators. We suggest that network managers examine whether greater automation can be achieved today using existing network-management tools.

INTEGRATION OF NETWORK-MANAGEMENT TOOLS

At the beginning of our research, several Foundation members asked us to predict when integrated network-management systems would become available. There are, however, several different types of integration:

- Interworking between the wide-area data network-management tools provided by different suppliers.
- Providing a network-management system that covers several activities.
- Interworking between tools used with different types of network (voice, wide-area data, local area, image, and so on).
- Gathering information from all the layers of the seven-layer OSI model.

The functional requirements for network-management systems described earlier in this chapter do not exclude any of the above types of integration. In general, however, network managers equate integrated network-management systems with the first three types, and particularly with the first one. Most network managers do not perceive a need to integrate the tools used with different types of network. There appear to be two reasons for the lack of interest in this type of integration:

- A belief that integrating the different tools used for managing wide-area data networks will be difficult enough to achieve, without adding to the problems. Furthermore, the tools for voice networks, wide-area data networks, and local area networks have very different origins. Voice-network tools focus on providing information that can be used for recharging purposes; wide-area data-network tools focus on identifying and handling faults. Local area network tools focus on identifying performance problems. These differences will make it more difficult to integrate the different types of tool.
- With the exception of the basic transmission medium, there has been little progress on integrating voice and data networks. Although the organisation chart may show that voice and data networks are now managed together, there has been little integration of the skills required or of the network-management activities. Thus, there is little demand for integrated network-management tools, apart from those that can be used to control high-bandwidth digital transmission systems.

The final type of integration listed above refers to the OSI framework for network management. This framework envisages the development of network-management protocols that would be used by applications to pass relevant network-management information up through all seven layers of the OSI model. Until the protocols (and applications to use them) are developed, network-management information will be exchanged between each layer on a bilateral basis.

MODEL OF THE IDEAL SYSTEM

In Chapter 2, we defined network management as the set of activities required to plan, install, monitor, and maintain all network components. Figure 5.3 depicts an integrated network-management system that would support most of these activities. The model shown in the figure consists of a series of modules, each of which supports a different network-management activity. Integration between the different activities is achieved by using common databases for the network's inventory and configuration, and for network statistics. These databases are used to provide relevant information to the different modules. All the different network-management tools are connected to the network components via a

Figure 5.3 Model of the ideal integrated network-management system

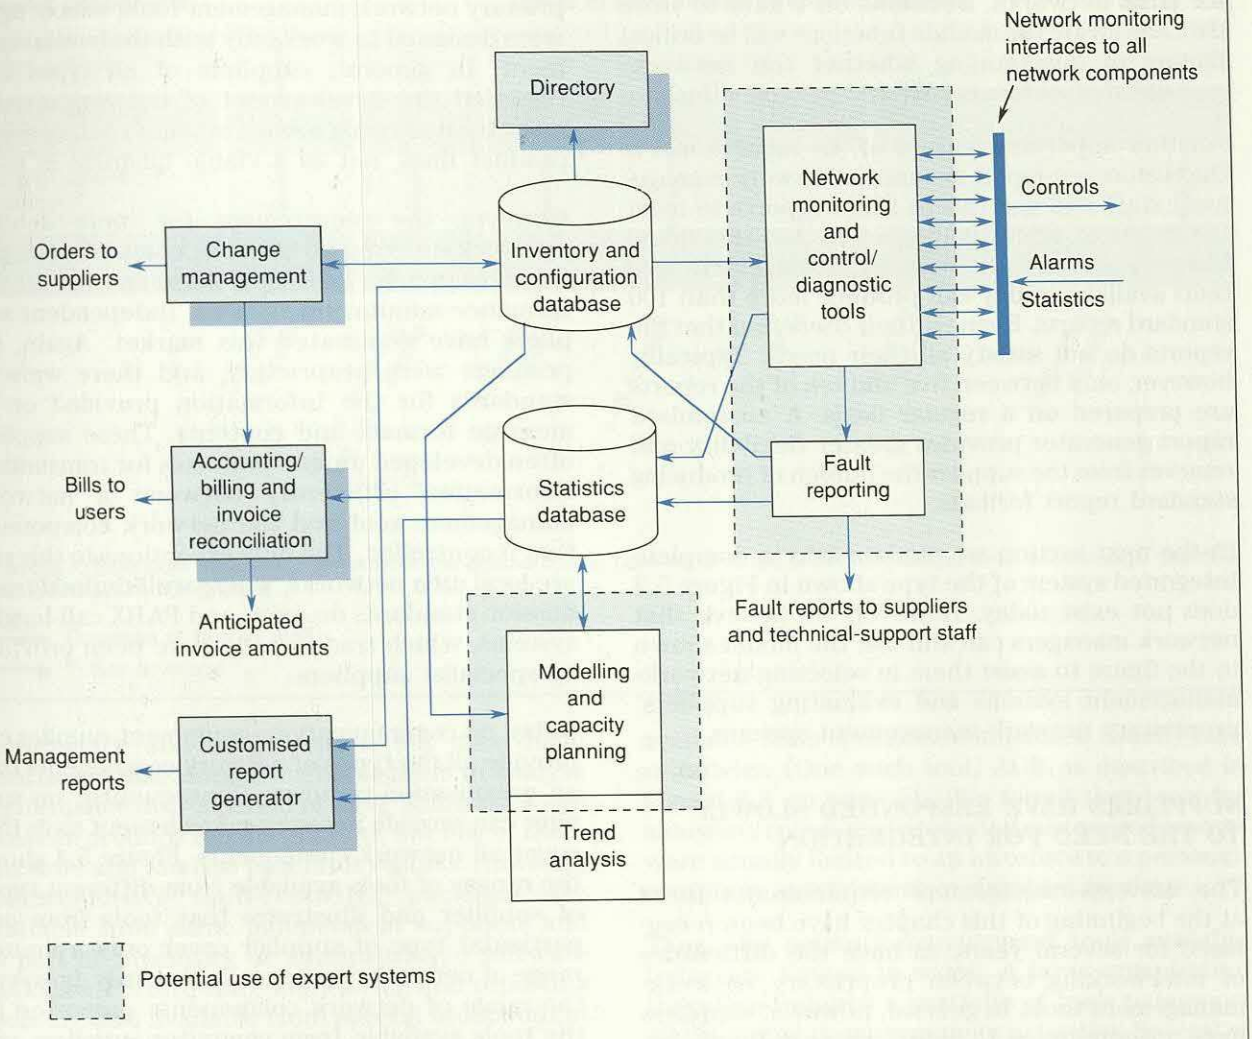

standard network-monitoring interface, thereby ensuring that the tools produce consistent information and can interwork with each other.

Most of the data produced by an integrated network-management system requires interfaces to network components and depends on the quality of the data that these components provide. The network-monitoring and control module is by far the most complex and difficult module to construct because of the wide variety of interfaces it must support and the functions it must provide. The functions include:

- Eliminating duplicate data produced by different network components.
- Converting all the data into a few common formats.
- Filtering out unnecessary data.
- Providing diagnostic tools for use by network technicians.
- Translating control instructions from technicians into the correct command formats for each type of network component.
- Interpreting and responding to certain types of message automatically.
- Providing performance and fault information to other modules.

The remainder of the modules, excluding perhaps the more complex modelling and capacity-planning module, can be specified and developed in a similar manner to any other computer application.

An integrated network-management system and its associated databases may not necessarily exist as a centralised whole, but are likely to be distributed throughout the network. The potential

volume of data needing to be transmitted to the databases, particularly the statistics database, will make it impractical to create centralised databases for large networks. Decisions on where to store data and locate the module functions will be critical factors in determining whether the network-management system is efficient and cost-effective.

Another important feature of the ideal model is the customised-report generator. Network-management staff will use this to tailor reports to meet the needs of senior management and individual business areas. Some network-management systems available today can produce more than 100 standard reports. Even so, their users find that the reports do not satisfy all their needs. Typically, however, only between five and ten of the reports are prepared on a regular basis. A customised report generator provides greater flexibility and removes from the supplier the burden of producing standard report formats.

In the next section we explain why a complete integrated system of the type shown in Figure 5.3 does not exist today. However, we believe that network managers can still use the model shown in the figure to assist them in selecting network-management systems and evaluating suppliers' proprietary network-management systems.

SUPPLIERS HAVE RESPONDED SLOWLY TO THE NEED FOR INTEGRATION

The network-management requirements listed at the beginning of this chapter have been recognised for several years, as have the difficulties of interworking between proprietary network-management tools. In general, however, suppliers have not attempted to rectify these deficiencies. The reason for the lack of progress is that most network-management tools are supplied by communications-equipment and computer-system manufacturers and are designed to control and monitor their own equipment. The manufacturers did not perceive the market for network-management as being large enough to justify the investment that would be required to broaden the scope of their proprietary tools. In addition, interworking between tools from different manufacturers requires the use of nonproprietary standards for exchanging alarm, control, and performance data.

HARDWARE-LINKED TRADITIONAL TOOLS

The first tools for data networks were developed in the early 1970s by manufacturers of modems and multiplexors to provide enhanced control capabilities for their products. The advantage for a manufacturer was that providing a network-management tool tended to lock customers into buying more modems and/or multiplexors from that manufacturer. Equipment purchased from another supplier could not be used with the first supplier's network-management tool. A few years later, computer suppliers also introduced proprietary network-management tools, which, again, were designed to work only with their own equipment. In general, suppliers of all types have regarded the development of network-management tools as being a small adjunct to their overall product lines, not as a viable business in itself.

However, the requirement for more detailed statistics covering all network components led to the development of a small niche market for performance-monitoring systems. Independent suppliers have dominated this market. Again, the products were proprietary, and there were no standards for the information provided or for message formats and contents. These suppliers often developed unique protocols for transmitting information efficiently between a network-management tool and the network components that it controlled. The only exceptions to this rule are local area networks, where well-defined transmission standards do exist, and PABX call-logging systems, which traditionally have been provided by specialist suppliers.

Today, no communications-equipment supplier can provide all the types of network components that an organisation requires. Consequently, no supplier can provide network-management tools that cover all network components. Figure 5.4 shows the ranges of tools available from different types of supplier and illustrates that tools from one particular type of supplier cover only a limited range of network components. Recently, however, the range of network components supported by the tools available from computer suppliers and modem and multiplexor manufacturers has slowly begun to increase. In addition, the PTTs and other network suppliers are beginning to offer more tools for managing the interfaces to public networks. In general, however, the market for network-management tools continues to be fragmented and hardware-specific.

LIMITED SCOPE OF TOOLS

None of the 30 suppliers of network-management tools and systems surveyed for this report offered products that covered the full range of activities shown in Figure 5.3. Instead, the available tools covered a limited range of functions, reflecting the primary interests of the suppliers of the different types of tool. The functions of the tools provided by the various types of supplier are shown on page 42 in Figure 5.5.

The tools available from hardware suppliers usually concentrate on operational control and fault diagnosis. Some of these tools produce

Figure 5.4 Network-management tools have traditionally been linked to specific network components

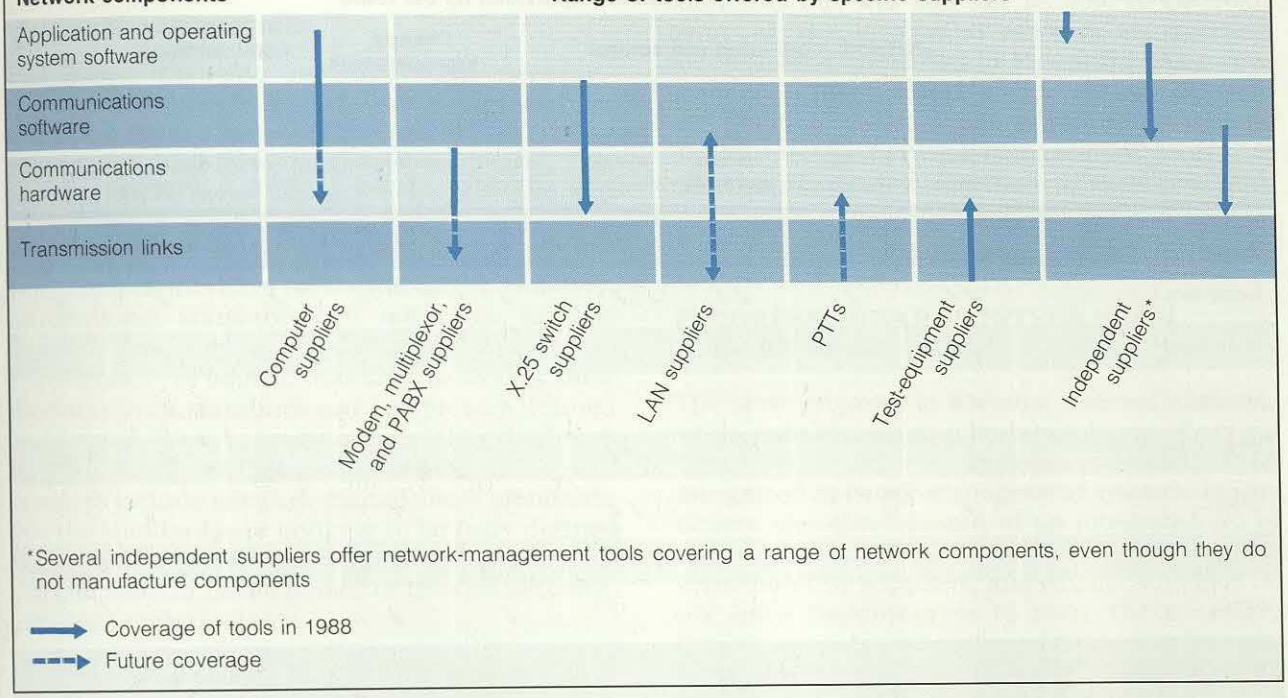

performance and usage reports but most have insufficient disc storage to enable them to analyse performance over a period of time. Similarly, some tools can produce reports of faults, but few of them can store and analyse past fault reports. However, comprehensive fault-reporting packages are available from some independent suppliers and as part of inventory or administrative software packages. Performance-monitoring and reporting tools are also available from several independent suppliers.

We found that the majority of user organisations had developed their own in-house systems for recording the data-network inventory and for recharging network costs, although such systems are available from some computer suppliers (IBM's Information Management software, for example) and from independent suppliers (Computer Associates's NetMan, for example). These proprietary systems also provide some change-management capabilities.

Network-design tools are available from some computer suppliers and also from packet-switch suppliers such as Bolt, Beranek and Newman. These tools range from simple analysis aids to very sophisticated simulation systems. However, these design tools are often available only as part of a service provided by the supplier and can be used only to design networks that will be constructed from the supplier's hardware range. Nevertheless, a few general-purpose network-design tools are available from specialist companies, usually consultancies. (One such tool, ACR, is described in Figure 6.3 on page 51.) We found that tools for analysing trends and preparing customised reports were usually limited to an interface to a personal-computer spreadsheet or database package.

Thus, the network-management tools available today are limited in scope. A large organisation therefore requires a portfolio of tools to support all the network-management activities. Several of these tools will have their own inventory database, their own configuration database, and, possibly, some form of fault-reporting system. The resulting duplication of information causes additional work and can lead to discrepancies between the different inventories.

LITTLE INCENTIVE TO IMPROVE TOOLS

We have already pointed out that most network-management tools are provided by computer-system and communications-equipment suppliers, and that they tend to view network-management tools as a means of enhancing the capabilities of their major product lines (and thus as a means of selling more products). Enhancing their proprietary network-management tools to provide support for other suppliers' equipment was, until recently, perceived as giving other suppliers the opportunity to sell to their customers. Thus, there was a major incentive for suppliers not to extend the range of hardware supported by their network-management tools.

Figure 5.5 Today's network-management tools cover only a limited range of functions

Type of supplier	Functions provided by the tools								
	Network monitoring and control				Change management		Administration		Planning and design
	Alarms and alerts	Fault handling	Automatic and remote operations	Performance monitoring	Auto. reconstruction	Work orders	Billing	Inventory	
Network component suppliers									
Computers	✓✓	✓	✓✓	✓	✓		✓	✓	✓
Multiplexors	✓✓		✓	✓	✓✓			✓	
Modems	✓✓				✓			✓	
X.25 switches	✓✓			✓✓	✓✓			✓	✓✓
LANs	✓			✓✓	✓				
PABXs	✓✓				✓				
PTTs		✓			✓	✓	(1)	✓	
Other suppliers									
Independent[2]	✓✓	✓	✓	✓✓	✓	✓	✓	✓	✓
Test equipment	✓✓			✓✓					
Call-logging systems	✓		✓	✓✓	✓		✓✓	✓	✓
User organisation		✓					✓✓	✓✓	

Notes:
1 Billing of network users *within* an organisation is not offered by PTTs
2 Each individual supplier covers only a subset of the functions shown

✓✓ Tools frequently provide the function
✓ Tools occasionally provide the function

Furthermore, independent suppliers face considerable difficulties. It is not possible to design performance-monitoring systems that work with a range of hardware without detailed knowledge about each manufacturer's proprietary control messages and protocols. The independent suppliers that provide such systems have therefore mostly developed their own monitoring hardware that is located between the workstation and the network. The cost of the additional equipment required makes such tools comparatively expensive and limits the size of the market for them.

In general, most suppliers thought that the market for network-management tools was small, and this means that the range of activities supported by such products has remained limited. In addition, significant investment is required to develop the software required to support network-management tools. Most suppliers have chosen to invest in developing software that is closely linked to their major hardware products. One exception is packet-switch suppliers. Most of these have invested heavily in providing network-design tools because, often, they could not convince organisations to buy their products until they showed them how their networks should be redesigned. Yet again, the incentive for investing in the development of a network-management tool was to sell the main product line.

Compared with the size of the market for their main products, the market for network-management tools appears to many suppliers to be limited. They therefore have little incentive to improve the tools that they provide. For example, market surveys estimate that the annual US market for modems is $2,100 million and the market for

mainframe computers is $14,000 million. By comparison, the US market for network-management tools in 1988 is forecast to be under $500 million.

STANDARDS ARE REQUIRED

The major difficulty in integrating network-management tools from different manufacturers is the proprietary nature of the protocols, naming conventions, and formats used by each tool. This problem can be overcome only once international standards for exchanging management information between network-management tools and network components have been established. However, international standards are notorious for the length of time they take to define and implement. Suppliers are, of course, reluctant to modify their products until standards are completely defined and agreed. Work is progressing within the International Standards Organisation to extend the OSI model to include network-management standards, but the standards are unlikely to be fully defined before the middle of 1991. Figure 5.6 describes the present state of development of the OSI network-management standards.

In the meantime, several major suppliers have announced their own proprietary 'open' standards for exchanging information between network-management systems, and one or two of these will be adopted as *de facto* industry standards. Proprietary standards will develop faster than OSI standards but will still not be able to cater for many of the integrated network-management requirements identified in Figure 5.1. As a consequence, other network-equipment suppliers and independent vendors will continue to provide highly specific network-management tools that offer improved user interfaces or more capability than similar tools available from major computer-system suppliers and PTTs.

Figure 5.6 Slow progress is being made in defining OSI network-management standards

The OSI network-management standards were originally scheduled to be completed by 1990 but it is unlikely that this schedule will now be met. At present (mid-1988), only the management framework is close to becoming a full international standard. The Communications Management Information Protocol (CMIP), which defines standards for exchanging network-management information, is partially defined, although a protocol definition does not yet exist. Working groups are still defining the form of activities for the five areas defined by the management framework.

One of the biggest gaps to date is the lack of definition of the contents of the Management Information Base (MIB). This corresponds to the inventory, configuration, and statistics databases shown in Figure 5.3. Until the standards committees decide what information should be stored, suppliers will not know what data their tools need to provide.

The OSI network-management standards do not cover voice networks, although voice could be included if the signalling and transmission standards conformed to the seven-layer model.

Network-management standards will also probably emerge from the work of the Transmission Control Protocol / Internet Protocol (TCP/IP) network-management task force. The members of this task force have agreed on a single protocol — Simple Network Monitoring Protocol (SNMP) — for exchanging management information, and are now defining the objects that the protocol will manage. The intention is to merge SNMP with OSI's CMIP when CMIP is fully defined.

At current rates of progress, international network-management standards will not be fully defined until the early 1990s. It will then take a further year or two before products conforming to the standards are available.

PROGRESS TOWARDS INTEGRATION WILL BE SLOW

The slow progress in defining international network-management standards is only one of several factors that will retard progress towards fully integrated network-management systems. In particular, the development of an integrated set of network-management tools requires considerable investment by suppliers, and smaller suppliers do not have the resources to make the necessary investment. As a result, smaller suppliers are forming alliances with each other. The different network-management 'architectures' announced by major computer suppliers and some PTTs will also slow down the progress towards a universal integrated network-management system. We now discuss each of these points in turn and conclude the chapter with our views on how progress towards integrated network-management systems will proceed during the next three to five years. We believe that fully integrated network-management systems that manage all the network components within an organisation are unlikely to be available within this timeframe.

THE INVESTMENT OBSTACLE

Integrated network-management systems cannot be developed quickly and require a large investment of resources by suppliers. Codex, for example, states that it has 60 engineers enhancing and maintaining its network-management systems. And IBM is widely reported as having more than 200 analysts working on the development of its NetView network-management products. Smaller suppliers do not have the human and financial resources to enable them to commit to this scale of development effort.

The development of a comprehensive network-management system that can satisfy the requirements listed earlier in this chapter means that the supplier has to employ a variety of new, and perhaps unfamiliar, technologies such as realtime high-resolution graphics and expert-system tools. Furthermore, to produce integrated tools that can be used with wide-area data networks, local area

networks, and voice networks means that the supplier needs a detailed understanding of all these areas. Even when standard protocols exist for exchanging information between network tools, and between network tools and network components, suppliers will still have to write software that:

- Filters out unwanted messages.
- Provides automated control and diagnostic routines.
- Stores, relates, and controls all the data in a network-management database.
- Produces a wide range of reports.
- Provides easy-to-use user interfaces.

Some suppliers may also write software that converts another supplier's network-management information from a proprietary protocol into a standard format.

Another difficulty arises from the large volumes of network-management data that may need to be stored and processed. Some networks contain thousands of components. Suppliers will need to consider carefully how best to construct network-management systems so they can efficiently gather, process, and store these large volumes of data.

SUPPLIER ALLIANCES

Pressure from network managers for integrated network-management systems and the development of network-management standards means that it will become more difficult for smaller communications-equipment suppliers and independent vendors to sell their tools. They lack the resources or the market presence to compete successfully with the PTTs and major suppliers such as AT&T and IBM.

The only way that most communications-equipment suppliers and independent vendors will be able to remain in the market for network-management systems is by forming an alliance or merging with either a computer supplier or a PTT. Computer suppliers lack knowledge about physical network management and the graphics-based configuration tools used by the communications-equipment suppliers. They are also anxious to promote the concept of integrated tools that can be used to manage all aspects of the IT function — computer operations, operating systems, hardware inventories, systems development, and networks.

The PTTs lack knowledge about managing the noncommunications aspects of IT but they can provide expertise in the management of both private and public networks. In fact, the PTTs have more experience in managing networks than other suppliers. It is possible that they could adapt products developed to manage their own networks for use as corporate network-management tools.

Figure 5.4 showed that the network-management tools available from different types of supplier covered different network components. This figure illustrates the logic of alliances between suppliers to obtain a network-management product range that covers the entire systems and networking hardware range. The trend towards such alliances is already clear. Examples include:

- Tandem's purchase of Ungermann Bass.
- Unisys's purchase of Timeplex.
- IBM's purchase of PacTel Spectrum Services (a network facilities-management company) and its collaboration with Network Equipment Technologies (a US-based manufacturer of multiplexors).

An appropriate alliance enables medium-size suppliers to combine development resources, but more importantly encourages network-management products that cover the whole range of network components. We believe that more alliances of this type will occur over the next few years. The right combination of expertise between the companies involved in an alliance, and the speed at which they can integrate their products, will be major factors in the success of their network-management system.

NETWORK-MANAGEMENT ARCHITECTURES

All the major computer suppliers, some PTTs, and some of the larger communications-equipment suppliers have announced that their network-management products will conform to a network-management architecture. In this context, the term 'architecture' usually implies a series of products linked together to form what we call a network-management system. A few of these architectures use proprietary protocols to pass management information between the network components and the network-management system. Most suppliers claim that their architecture conforms to OSI standards, although this is clearly not possible where standards do not yet exist. The characteristics of the main architectures proposed by different types of suppliers are summarised in Figure 5.7.

The group of network-management products that has received the most publicity and support from other suppliers is IBM's NetView products. NetView brings together several of IBM's existing network- and systems-management products. The NetView/PC product allows alert messages to be collected from, and limited control information to be transmitted to, other suppliers' products. At

Figure 5.7 Characteristics of suppliers' network-management architectures

Supplier and major products	Overview of network-management system architecture	Key features
AT&T (PTT and telecommunications equipment supplier)	Unified Network Management Architecture (UNMA) is designed to combine management of an organisation's on-premises network components and its PTT services. PTT management systems and other suppliers' network-management tools communicate with the central-management system using Network Management Protocol (NMP), a published set of standards.	Central-management system can be on-premises or at AT&T's network-control centre. NMP is based on existing OSI standards.
Avant Garde (network performance-monitoring system supplier)	Net/Command accepts alarm information from network components or other management tools. Information is filtered, prioritised, and translated into 'plain English'.	Avant Garde develops its own interfaces to other suppliers' equipment. Some interfaces are provided to PABXs as well as to data-network components. Collection of NetView messages from IBM systems is also provided.
BBN Communications (X.25 switch supplier)	The C/70 Network Operations Center (NOC) provides centralised network-management for BBN Communications' X.25 switches. The NOC covers a wide range of network-management activities. DESIGNet is an expert-system-based network-design tool used in-house by BBN to design X.25 networks.	Graphics screen presentation using colour-coded alerts. Centralised configuration database. Usage-based billing support. NetView/PC support.
IBM (computer supplier)	IBM's goal is to evolve Systems Network Architecture (SNA) network-management products into one integrated network-management system for both IBM and non-IBM components. The overall name for this structure is Open Network Management (ONM). ONM allows centralised management and provides published network-management architectures. Most of IBM's products in this area are grouped under the NetView name and are host-based. NetView/PC is a PC-based product that can gather data from other suppliers' network components, pass alert information onto NetView, and receive commands from NetView.	Suppliers need to write their own software for NetView/PC applications. IBM provides NetView/PC interfaces to IBM Token Ring local area networks and 8750 PABXs.
Motorola Information Systems (supplier of modems and other data communications equipment)	Codex 9800 Integrated Network Management System (INMS) is consistent with currently defined OSI network-management standards and draft standards. The goal is to control any network component using one screen format. Adaptors convert other suppliers' proprietary protocols into Codex's standard.	Graphics screen presentation using icons and windows. Common database containing all component attributes.
Timeplex (multiplexor supplier)	TIME/VIEW is designed to support most network-management activities and can act as either a master or slave network manager. Timeplex has also defined an 'open' specification to enable information to be passed to and from other suppliers' equipment, called TIME/VIEW Open Access. TIME/VIEW will support OSI standards.	Graphics screen presentation using icons and windows. Relational inventory database. Customised reports including graphs. Assistance in automation of routine activities. NetView/PC and UNMA support.

The above companies are a representative sample of the different types of supplier offering network-management systems

least 25 other suppliers have announced that their products will support the NetView/PC interface, making it a *de facto* standard.

However, NetView at present has several limitations, which means that it falls short of meeting all the full requirements of an integrated network-management system. The most important limitations are that:

- The number of control messages is limited.
- It does not provide graphics-based network-configuration facilities.
- Links between the NetView monitoring systems and the inventory and change-management system (which is called Information Management) are inadequate.
- The routines for constructing automated responses have to be programmed as CLISTs (a low-level language). Most network-management staff find it cumbersome to do this.
- The performance-monitoring routines can consume a significant amount of mainframe processing capacity.

As a result, most NetView users still require tools from other suppliers as well, and are likely to do so for several years. Suppliers such as Codex and Timeplex are enhancing their network-management tools while supporting a NetView/PC interface. Their aim is to provide features that the computer manufacturers' products lack, such as graphics-based displays. Other independent vendors offer products that emulate some of IBM's products but which provide improved functionality. One example is Cincom's NetMaster which is a NetView lookalike but uses a fourth-generation language, rather than CLISTs, for programming the automatic-response routines.

Other computer suppliers emphasise that their products will manage both networks and systems. Programs that can automatically distribute software to remote sites are early examples of this class of product (both Siemens' TRANSDATA products and ICL's Community Management systems provide this facility).

Suppliers of call-logging systems are also beginning to extend the scope of their products to include:

- Collection of alarm information from PABXs.
- An inventory and configuration database.
- Change-management facilities.

PROGRESS TOWARDS INTEGRATION

The examples above illustrate that progress towards integrated network-management systems will be steady but slow. Figure 5.8 indicates the timeframes in which we believe the different levels of integration will occur. Products that integrate network-management activities will be available first as the larger suppliers, or alliances between different suppliers, start to produce more comprehensive systems. Systems that integrate the management of different types of equipment will follow as manufacturers implement proprietary and international network-management standards. Faster progress will be made with systems that use proprietary standards, or proprietary versions of OSI standards, than with systems aimed at providing full network-management integration between all the OSI layers. A regrettable fact of life is that suppliers will always want to differentiate their products from those offered by other suppliers. They will do this by incorporating facilities that go beyond the agreed standards.

The integration of systems that manage both voice and data networks will take longer to occur. The reason is not because this form of integration is more difficult to achieve, but because most network managers are not interested in this type of integrated network-management system. The need for integrated voice and data network-management tools will not occur until the use of integrated applications becomes common.

Figure 5.8 indicates that network managers cannot expect to purchase ready-made network-management systems that meet most of their requirements until 1991, or even 1993, at the earliest. However, in previous chapters, we have shown that network managers cannot afford to wait that long before purchasing new network-management tools and systems. In the next and final chapter of the report we examine the steps that Foundation members can take in the meantime to build a network-management system that makes the best use of the tools currently available.

Figure 5.8 Rate of progress in network-management integration

Type of integration	Rate of progress
Between different network-management activities	Quite rapid between 1988 and 1991. Major suppliers will expand the range of activities covered to include better inventory databases, more performance-monitoring routines and design aids, and automatic change-control systems.
Between management of different types of wide-area data-network equipment	Quite slow but steady until the end of the 1980s. Accelerating after 1990 as more installed network components conform to network-management standards. NetView/PC will predominate as a standard until OSI standards are clearly defined.
Between local area network and wide-area data-network-management systems	Almost no integration today because local area network-management tools are in their infancy. However, well established LAN transmission standards will ensure quite rapid progress over the next two to three years. By 1990/91 the level of integration will probably be the same as that between different types of wide-area data network equipment.
Between OSI levels	Slow because of the time required to define and implement international standards. Complete integration at all levels is unlikely to be widely achieved within five years.
Between voice and data network-management systems	Generally very slow because suppliers are taking different approaches and users do not perceive a need. Exceptions will occur in the areas of high-bandwidth multiplexor management tools, PTT circuit-ordering systems, and probably inventory databases.

Chapter 6

Building a network-management system

Most network managers cannot afford to wait until fully integrated network-management systems become available. They need to purchase new tools now, particularly to help overcome the shortages of skilled staff. In this final chapter of the report, we first identify the three different strategies that can be used for building a network-management system. We then emphasise that the network-management capability of network components needs to be a key criterion when selecting equipment. Expert-system techniques will also become an important element of network-management systems and we describe both the benefits that can be gained, and the difficulties that still exist in using expert systems. The report concludes with advice about how to justify the investment in a network-management system. The key is to focus on the business benefits, not the technical merits of the system.

A STRATEGY IS REQUIRED

Even though integrated network-management systems are not yet available, network-management tools should be selected with the ideal model of a network-management system in mind. This means that tools should be chosen to be consistent with each other, for their adherence to standards as they develop, and for their ease-of-use. However, the tools available will develop rapidly over the next few years, so they should be implemented in a way that makes it easy to replace them with later and better products. There are three different strategies that can be adopted for ensuring that the tools selected form a consistent network-management system:

- Adopt a single-supplier policy for all network components and network-management tools.
- Develop a customised network-management system.
- Mix and match products from different suppliers.

In choosing the most appropriate strategy, network managers must first determine where they can realise the biggest improvements in productivity or user service and then aim to satisfy those needs. The two key areas for network-management productivity gains are increased automation and reduced duplication of the data produced by network-management tools. The strategy should be chosen with these two key areas in mind. The benefits and pitfalls of each strategy are described in more detail below.

SINGLE-SUPPLIER POLICY

By restricting its network components and network-management tools to those provided by a single supplier, an organisation can quickly achieve a consistent network-management system. Products from one supplier are more likely to use a consistent management-information protocol, removing one of the biggest problems of compatibility between different tools. Moreover, the supplier has a large incentive to make all its products compatible with an overall network-management system. The recent increased awareness of the importance of good network-management products means that almost all major suppliers will develop new products in this area.

Figure 6.1 describes how one organisation has standardised almost exclusively on the products of a single supplier. However, there are three major drawbacks to adopting a single-supplier policy. First of all, most network managers know that it is unlikely that any one supplier can meet

Figure 6.1 Some organisations adopt a single-supplier policy for all their network-management tools

Sears Communications provides networking services to the Sears retailing organisation. It has chosen to use IBM products almost exclusively for its networking requirements and bases all of its network management on NetView. However, some customised software has been developed to interface Series/1 computers with NetView. The Sears data network connects more than 120,000 terminals and personal computers. NetView runs on a dedicated IBM 3090 mainframe, and a second back-up 3090 is available at another site on a 'warm-standby' basis. Sears chose NetView because it allows the network to be managed from a central site and because standardising on one network-management system means that network technicians have to master only one set of commands. However, this policy would not have been viable if Sears was not already committed almost exclusive to using IBM products.

all their needs, particularly for both network and computer equipment. They also do not like the thought of being dependent on one supplier and most organisations already use equipment from several suppliers. For these organisations, the single-supplier approach is not a realistic option because the computer systems and larger network components cannot easily be replaced. Deciding to standardise on a single supplier is a major strategic decision for the systems department and would usually be based on several factors, not just the need for improved network-management tools.

The second disadvantage of the single-supplier policy is that, even excluding public-network circuits, no supplier can provide a complete range of network components. Even IBM, who at present probably offers the most comprehensive range of products, cannot cover all customer needs. Further alliances between computer-system and communications-equipment manufacturers may increase the number of suppliers who can provide a wide range of equipment. However, restricting the choice of components to those available from one supplier may limit the flexibility of the information systems offered to users.

The third disadvantage is that the choice of network components available from one supplier is unlikely to be ideally suited to all an organisation's communications requirements. Adopting a single-supplier policy could therefore increase hardware costs significantly because one supplier cannot provide cost-effective products across the whole range of network components. The increases in hardware costs must be weighed against the benefits of improved network management produced by a single-supplier policy.

CUSTOMISED NETWORK-MANAGEMENT SYSTEMS

Some organisations with very large networks that are vital to their businesses have chosen to build their own network-management systems. Examples include major financial organisations, airlines, and Electronic Data Systems (EDS). These organisations believe that they cannot afford to wait for suppliers to develop integrated network-management systems and that their needs are so unique that they require a customised system. The base data for such systems is usually provided by existing network-management tools, and is then processed and analysed by customised software.

In the United States, EDS manages networks for several of its customers, the most notable being General Motors. From its network control centre in Plano, Texas, EDS manages 250,000 telephone extensions and more than 100,000 terminals and personal computers. EDS has found that the General Motors' network users are more demanding than they were when the network was run as an in-house function and expect a better service.

EDS's biggest network-management problem is controlling the changes made to the network and ensuring that documentation is kept up to date. The company anticipates a 40 per cent growth in its computer-system and network-management workload over the next few years and aims to accommodate the growth without increasing staff whilst also improving the service provided. To achieve this, EDS plans to invest nearly $10 million over the next few years on developing a better network-management system for administration, change control, and planning. One of the key features of the system will be a centralised management-information database.

EDS is an exception, however. Most organisations cannot contemplate developing their own customised network-management system because their networks are not large enough, or are not sufficiently crucial to the business, to justify the expense involved. A custom-built network-management system is also expensive to maintain. As suppliers enhance their network components or network-management tools, the customised system has to be changed to incorporate the new features. We expect there to be a large number of enhancements to proprietary tools over the next five years, which will lead to substantial maintenance requirements for customised systems. A supplier can spread the costs of enhancements across its whole customer base; an organisation with a customised network-management system must bear all the costs itself.

We believe that most organisations that have chosen to develop their own customised network-management system will eventually have to move to a proprietary system because the cost of maintaining large customised systems will be prohibitive. Only very large network-service organisations like EDS and the PTTs can afford such systems in the long term.

MIX-AND-MATCH APPROACH

For most organisations, a single-supplier policy or a customised network-management system will either be impractical or too expensive. Their only option is to continue for several years to use a variety of incompatible network-management tools. These organisations should purchase new tools that can improve their network management whenever the tools can be cost-justified. In selecting the tools, however, it is important to consider how they will 'fit' with existing tools.

After 1990, those tools that conform to standards, either proprietary or international, will usually provide the greatest improvements because they will be capable of at least partial integration with other tools that use the same standard(s). However, until then other features of a specific tool, such as a better user interface or more automated fault-correction routines, may provide greater improvements.

In addition, there are some actions that network managers can take to enhance their existing tools and make them more useful and to reduce the amount of duplicated data they produce. Areas where a relatively small investment could yield quite significant improvements are:

- Use of small, customised expert systems. (The benefits and difficulties of using expert systems in a network-management environment are discussed later in this chapter.)
- Use of PC-based programs to enable reports to be produced and availability calculations to be performed more easily.
- Enhancements to existing tools to provide trend-analysis reports and better analyses of alert messages.
- Use of personal computers as front ends to existing tools so that responses to the most common and predictable errors can be automated.
- Development or purchase of a database and systems for managing the changes made to the network.

With one exception, the investment required to make these types of enhancement is not large, which means it is easier to discard them when proprietary systems can provide the same facilities. The exception is the change-management database. As we described in Chapter 5, this is a key element of an integrated network-management system, and it could grow to become a very large subsystem. Many of today's network-management tools contain their own inventory and configuration database. Without careful planning, it would be very easy to find that several such databases, all containing the same information, are being maintained. Updating these databases and ensuring that they are in step with each other, is likely to be a difficult and time-consuming task. The difficulties will be compounded when the time comes to move to a proprietary integrated network-management system. One of the biggest implementation tasks will be setting up the inventory and configuration database, and converting the existing databases will not be easy. Network managers will therefore need to think very carefully about how they plan to manage such databases. Questions to be addressed include:

- Should the databases be maintained at local sites or centrally?
- How much duplication of information is unavoidable?
- How can the updating effort be minimised?
- How might the migration to a more comprehensive network-management system be undertaken?

The advice given in the previous Foundation Report — Managing the Evolution of Corporate Databases — will be a valuable source of information in answering these questions.

NETWORK-MANAGEMENT CAPABILITY IS A KEY CRITERION FOR COMPONENT SELECTION

With the exception of external performance-monitoring systems, most network-management systems rely on the data provided by network components. Even the most sophisticated system will be of limited use if it cannot receive information from, or send control messages to, the network components. However, it will be several years before network-management standards reach the stage where all modems or multiplexors provide information in a consistent format. We recommend, therefore, that network managers select new network components on the basis of their present network-management features, particularly their conformance to standards for providing management information.

Network-management capability should be a major criterion in selecting any new equipment because this capability will determine the success of future network-management systems. Some network managers even suggest that, in future, they will select the best network-management system and will then purchase equipment that it is able to manage. This is a complete reversal of the way most organisations currently select equipment and network-management tools.

Most suppliers are promoting a distributed approach to network management. (The OSI network-management standards also imply a distributed approach.) With a distributed approach, a hierarchy of computers is used to support different network-management activities, as depicted overleaf in Figure 6.2. Data-gathering and control activities are mostly performed at the lowest level in the hierarchy, closest to the network equipment that is being controlled. The tools at this level filter the information being provided by the network components and pass relevant information to the central network-management system. The central system is concerned only with high-level activities such as planning, administration, and billing. This

Figure 6.2 A hierarchy of computer systems may be used to support different network-management functions

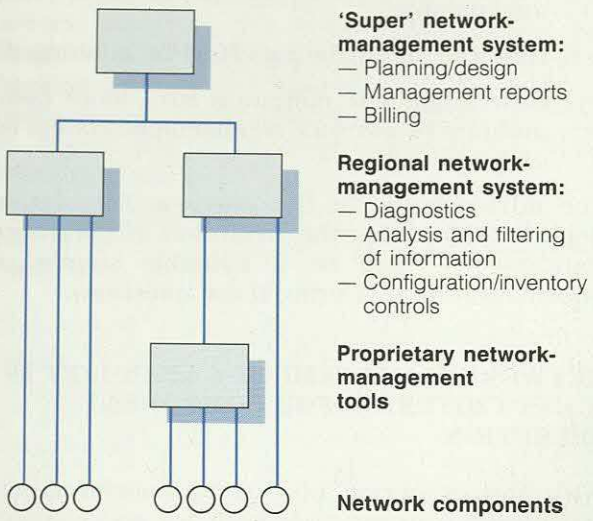

approach is effective because much of the network-management processing load is shared by several smaller computers and the amount of information that must be transferred to the central system is reduced.

It is possible that the proprietary network-management tools provided by individual component suppliers could become the lower level network-control systems. This would reduce the effort required to develop a fully integrated network-management system. But again, the better the individual tools are and the more closely they adhere to standards, the better will be the overall network-management system.

EXPERT SYSTEMS WILL BE IMPORTANT

One of the most significant developments in network-management tools is the use of expert-system techniques. Expert systems and their potential applications were discussed in Foundation Report 60, published in October 1987. One of the application areas identified in that report was network management, which displays some of the most important characteristics of potential expert-system applications:

- There is a shortage of expertise.
- It takes a long time to gain the expertise.
- The information provided by network-management tools often has to be acted on quickly.

Several network-service providers, network-equipment suppliers, and individual systems departments have developed or are developing expert systems for use in network-management applications. In addition, several PTTs, including SIP in Italy and AT&T, have constructed expert systems to assist in the control of public telephone networks. In the United Kingdom, the Alvey Data Processing Expert Systems (DAPES) community club funded research into the use of expert systems for network fault diagnosis. Network management is now widely regarded as an area where systems departments can gain some experience with expert systems without having to involve other areas of the business.

There are three main areas where expert systems could be applied to network management:

- Assisting with fault diagnosis, either at the help desk or as part of the first level of technical support.
- Providing more intelligence for automated network-control functions and configuration management.
- Assisting with network design.

BENEFITS

Expert systems can provide significant benefits for the network-management function because they can be used to:

- Replicate the knowledge of the most experienced network-management staff.
- Analyse the large volumes of data and/or messages produced by network-management tools.
- Interpret information in a consistent manner.
- Allow less-experienced staff to perform an activity.
- Teach less-experienced staff to perform an activity.
- Reduce the time required for experienced staff to perform an activity.

It is unlikely that all the above benefits can be obtained from one expert-system application. Network managers should therefore determine the type of potential benefits that a particular expert-system application could provide before it is developed. Useful advice on how to do this is contained in Foundation Report 60.

Figures 6.3 and 6.4 each describe the benefits being gained by organisations that have successfully used expert systems for network-management activities.

DIFFICULTIES

There are some difficulties in applying expert systems to network management, however. Suppliers will experience difficulty in building a

Figure 6.3 A network-design expert system

France Cable et Radio (FCR) has developed an expert system for data-network simulation called Aide á la Conception de Reseaux (ACR). ACR was developed because there were no suitable design tools available from other suppliers. FCR has used ACR successfully for several projects and claims that it reduces the time taken to design a network by a factor of 10. ACR is primarily an aid for network-design experts, but can be used to train designers. ACR can be used to evaluate:

- Centralised versus decentralised solutions.
- Network costs, including international tariffs.
- Quality-of-service factors.
- Trade-offs between anticipated downtime and increased costs.

FCR believes that problems that subsequently arise from a poorly designed network configuration are very difficult to solve. Using ACR reduces the likelihood of configuration problems arising after the network has been implemented.

Figure 6.4 A network help-desk expert system

A major pharmaceutical company developed an expert system called Rupert (Resolves User Problems Expertly) and introduced it as the network help desk was set up. (Prior to this, network users contacted anyone in the management services division they thought could help them.) Rupert is primarily a diagnostic aid, but also acts as a fault-reporting system. It is also connected to the computer systems so it can automatically disconnect the terminals from the network. Originally, it was not designed to handle user queries about the use of applications, but it is now being enhanced to do so.

The help desk is staffed by two ex-secretaries from the management services division. They told us that Rupert provided them with confidence when they were first answering queries. They rapidly learnt how to handle the most common queries, and now take short cuts when they are using Rupert.

In the first two months, the help desk staff:

- Handled nearly 2,500 enquiries.
- Solved 70 per cent of all queries themselves.
- Found that a number of queries were from users who did not understand the applications (more training courses were arranged).
- Discovered that the maintenance support provided by the terminal supplier was unsatisfactory. (The company asked its supplier to change its support system.)

The major benefits of Rupert to the company are:

- Its role as a training aid for new help-desk staff.
- The ease with which new knowledge can be added to the system.
- The time taken to resolve user problems has been halved.
- The improved 'image' the systems department now has in the rest of the company.
- The better statistics it provides about problems.

The last two benefits could probably be obtained from any help-desk function and fault-reporting software. However, Rupert's excellent user interface has made this a very successful application of expert-system techniques.

generic expert system that can be used for diagnosing faults in a network that uses equipment from a wide variety of suppliers. The difficulty can be estimated from a paper published in Volume 25, Number 2 issue of the IBM Systems Journal.

This paper describes the development of an expert system for automating Multiple Virtual Storage (MVS) operations on IBM mainframes. The authors stress the need to provide customising capabilities in the expert system to take account of the differing operating procedures and priorities in different computer centres. The level of customisation required in a multivendor networking environment would be several times greater.

It therefore appears likely that network fault-diagnosis expert systems developed by suppliers will be restricted to their own equipment in the next few years. Moreover, many organisations will find that this type of expert system is too difficult to be a practical proposition for in-house development. Indeed, at the time of our research, we did not find many in-house network-management expert-systems that had been completed. We believe that members should choose expert-system applications with a clearly defined and realistic focus, as were those described in Figures 6.3 and 6.4.

JUSTIFICATION OF INVESTMENTS SHOULD BE IN RELEVANT BUSINESS TERMS

We stated in Chapter 1 that network managers have difficulty in justifying substantial investments in network-management systems. Usually, the investment in the network itself has already been justified, and the management system is seen as an additional cost that appears to reduce the net benefits from the network. The benefits of the network-management system are not easy to explain in nontechnical terms. Moreover, the benefits appear at first sight to be for those responsible for operating the network rather than for the business as a whole.

When we asked network managers how they intended to convince their senior management of the need to purchase a network-management system, their answers tended to fall into one of three categories:

- "We had a major network failure last year. We should have no problems in obtaining board approval for expenditure that will prevent this situation ever occurring again."
- "Our management understands the importance of networks to the business and will agree to spend money on a network-management system."
- "I don't know."

None of these responses is satisfactory. We believe that network managers cannot rely on top management understanding the importance of network-management systems. Instead, they need to construct a case that demonstrates the benefits of a network-management system in business terms. Only a few of the network managers we talked with in our research had produced business cases to justify their network-management systems. The main reason for this is that they find it difficult and time-consuming to quantify the benefits of better services and performance from the network. There is also a tendency to concentrate on the technology and its technical and operational advantages for the network-management function.

It is easier to justify a network-management system when a new or enhanced network is being planned. In this situation, the management system can be justified as part of the total network cost, as has traditionally been the case for individual hardware-linked network-management tools.

Figure 6.5 contains guidelines for presenting network-management proposals to senior management. One of the guidelines is to gain as many allies and supporters as possible. Network users have the most to gain from reductions in downtime or improved response times. Enlist their help in identifying what the improvements will mean to the business and in presenting the case to senior management. Users will be powerful allies if they are made aware of the benefits that a network-management system will bring to them.

The first step in preparing the business case for a network-management system is to identify the most critical concern business managers have about the applications the network supports. If the applications are online and are essential for taking orders from customers or for delivering products or services, then the main concern is likely to be the risk to the future of the business, rather than the extra costs that may be incurred because the network is out of operation. In this situation, the business case should be based on how the network-management system will minimise the risk of downtime. Alternatively, if the network applications are mainly concerned with back-office support, and network downtime is likely to result only in additional work and expense, and in inconvenience to clerical staff, then the main concern is likely to be cost-reduction. The business case for the network-management system should therefore focus on the scope for minimising costs.

Figure 6.5 Guidelines for presenting network-management proposals to senior management

Sell the business benefits of the project, not the technology. How does the project fit in with other business objectives and approved programs?
Cover people issues, as well as equipment.
Consider alternatives.
Have lower-cost solutions available, even if less effective.
Be accurate in your budgeting, to establish credibility.
Get help from other people.
Get as many allies and supporters as possible.
Be prepared for being turned down.

(Based on a presentation at the Telecommunications Association conference held in San Diego, California in September 1987)

We provide guidelines on how to prepare these two kinds of business case below. A cost-based business case will usually require an estimate to be made of the cost of downtime. We therefore also provide advice about how to calculate this cost. In deciding which type of business case to prepare, the network manager should, however, be alert to his or her own managers' and organisation's specific needs and concerns. They may not always follow the simple logic outlined above.

A RISK-BASED BUSINESS CASE

A risk-based business case is most appropriate where a time-critical business operation requires a high level of network availability. An airline's seat reservation system is an obvious example (although the passenger check-in system is likely to be even more time-critical). In this situation, it is clearly desirable that the network is as reliable as possible, and that if a failure does occur, the fault is cleared as quickly as possible. Although the network will have been designed to be highly reliable, the network-management system should contribute to anticipating possible network failures (by monitoring trends in performance, for example) and to restoring normal operation more quickly when a fault does occur (by providing better diagnostic and remote-control facilities, for example).

With time-critical applications, the cost of a network failure is not simply the amount of business that may be lost whilst the network is out of operation, but includes the impact on the future business and operations of the organisation. For example, Peter Keen, executive director of the International Center for Information Technologies, cites the case of an airline whose reservation system was out of operation for 36 hours. As a

result, the airline lost 5 per cent of its market share and had not fully recovered it a year later Major business losses such as this are more likely to result from prolonged interruptions to network services. A network-management system can significantly reduce the chance of prolonged interruptions because it makes it easier to diagnose and rectify faults.

The above example illustrates that the financial impact to the business of a major network failure can be several orders of magnitude greater than the cost of the network-management system. The business case should therefore emphasise the small cost of the network-management system (say $100,000) and set this against the risk of losing a large amount of business (say $10,000,000). The case should identify very clearly and specifically the nature of the disaster that could occur and should attempt to quantify the financial impact. A precise figure is not required — an order of magnitude is good enough. The relevant business managers should be involved in estimating the financial impact. Senior management will be most easily convinced of the validity of the case if any examples of comparable disasters can be quoted. The presentation should make it clear, in a nontechnical way, how the network-management system will contribute to reducing the chance of such a disaster.

Finally, no attempt should be made to present the potential 'savings' in terms of so much per hour of downtime. This is not an appropriate measure to use when preparing a risk-based business case for a network-management system.

A COST-BASED JUSTIFICATION

A cost-based justification is most appropriate when the network applications support business activities that are not so dependent on continuous operation and immediate access to the computer systems. For example, a back-office accounting system may not hold up any other work if it is out of operation. The only impact may be that some staff have to work overtime to catch up when the network service is restored. In such circumstances, the business case for a network-management system needs to be based on its ability to reduce the costs of any downtime, either by reducing the number of network failures, or by enabling the network services to be restored quickly or inexpensively.

Sometimes the network-management system will also enable other business costs to be minimised or a better service to be provided to customers. For instance, better measurement of the network's performance may allow a business manager to match more closely the number of staff he or she assigns to a particular activity (telephone answering, for example) to the demands for that activity. An example of better customer service was provided by the systems manager of Credito Emiliano. He or she could demonstrate that the investment in NetView has increased the number of hours that the bank's network is available for use by customers.

THE COST OF DOWNTIME

Estimating the cost of downtime for a particular application is not easy. The estimates should be made jointly by the network manager and business-area managers. However, business areas often do not know how much application downtime costs because they have never needed to calculate such a figure. One major airline spent several months trying to calculate the cost of downtime at its check-in counters. It found that network failures rarely caused a complete break in service at all the check-in counters, so most failures had to be assigned a 'passenger-inconvenience cost'. Check-in personnel could cope with short failures but the cost escalated rapidly after the network had been out of operation for 20 minutes.

A study of the costs of downtime performed in the United States by the Yankee Group showed that most estimates ranged between $1,000 to $10,000 per hour. Suppose the network manager estimates that the proposed network-management system will produce a 1 per cent improvement in network availability and that the cost of downtime is $5,000 per hour. On these assumptions, there will be an extra 20 business hours of productive time per year, equating to an annual saving of $100,000. This means that the payback period for many network-management systems would be less than three years.

However, the cost of downtime may not increase linearly with the length of time the network is unavailable. In some instances, such as airline check-in counters, the cost of downtime may increase at an alarming rate after a minimum critical time is exceeded. In situations such as these the calculations need to be modified accordingly.

As well as working with network users to estimate the cost of downtime, network managers should also draw on the experience of suppliers (and the suppliers' other customers) in estimating the increased network availability that can be expected once a new network-management system has been installed. It is in the suppliers' own interests to help network managers convince their managers that a network-management system is a good investment.

REPORT CONCLUSION

In this report we have described the current problems facing network managers — the main ones being an increasing workload, shortages of skilled staff, and inadequate network-management tools. Although the problems of staff shortages and inadequate tools will remain for the next few years, there are several short-term and long-term actions that organisations and network managers can take to alleviate current difficulties and improve the management of their networks.

The short-term actions are to:

- Ensure that the network manager has good management skills; technical skills are of secondary importance.
- Reduce the number of help desks and ensure that no user needs to call more than one help desk.
- Employ 'user-oriented' people to staff the help desk.
- Negotiate service-level agreements between the network-management function and its users. Ensure that the conditions of the agreements can be met.
- Review the current procedures for network management. Include procedures relating to voice networks and local area networks.
- Establish a network-inventory database to assist in managing the changes that have to be made to the network. Avoid creating databases that contain duplicated information.
- Enhance existing tools to increase the levels of automation and to reduce the skill levels required to perform some tasks.
- Gather data on the cost of downtime. This data can then be used to build the business case for investing in a network-management system.

The long-term actions are to:

- Make network-management capability and interfaces a major selection factor when purchasing new network components.
- Monitor developments in network-management systems. Be prepared to replace existing tools as better ones become available. Ensure suppliers conform to the emerging network-management standards.
- Increase the budget for training network-management staff.
- Broaden the level of business and technology understanding of network-management staff. Establish individual career-development programmes.

Appendix

Trends in network faults

In this appendix we present details of the results of our research into trends in network faults. We describe why we believe it is important to identify trends in faults, the nature of the research undertaken, and the results of our surveys.

Our research was hampered by difficulties in obtaining quantitative data. However, we believe that the information contained in this appendix will be of interest to all network managers. The appendix demonstrates the value of performance monitoring and trend analysis, and we recommend that all network managers should undertake these activities.

IMPORTANCE OF IDENTIFYING TRENDS IN NETWORK FAULTS

In the body of the report we emphasised that network management covers more than just fault-resolution. However, as Figure A.1 shows, resolving network faults is still one of the most time-consuming network-management activities for Foundation members. It is likely that some of the user-support activity (which is the other most time-consuming activity) is also concerned with aspects of fault-resolution — running the help desk, for example.

Effective procedures for resolving faults are required to ensure that the network availability meets or exceeds the terms of service-level agreements. We believe that all network managers should monitor trends in network faults so they can improve their fault-handling procedures and determine the most effective ways of providing network resilience. In addition, several organisations have found that preparing detailed analyses of faults on a month-by-month basis enables them to obtain better service from their suppliers.

Figure A.1 Resolving network faults is the most time-consuming network-management activity

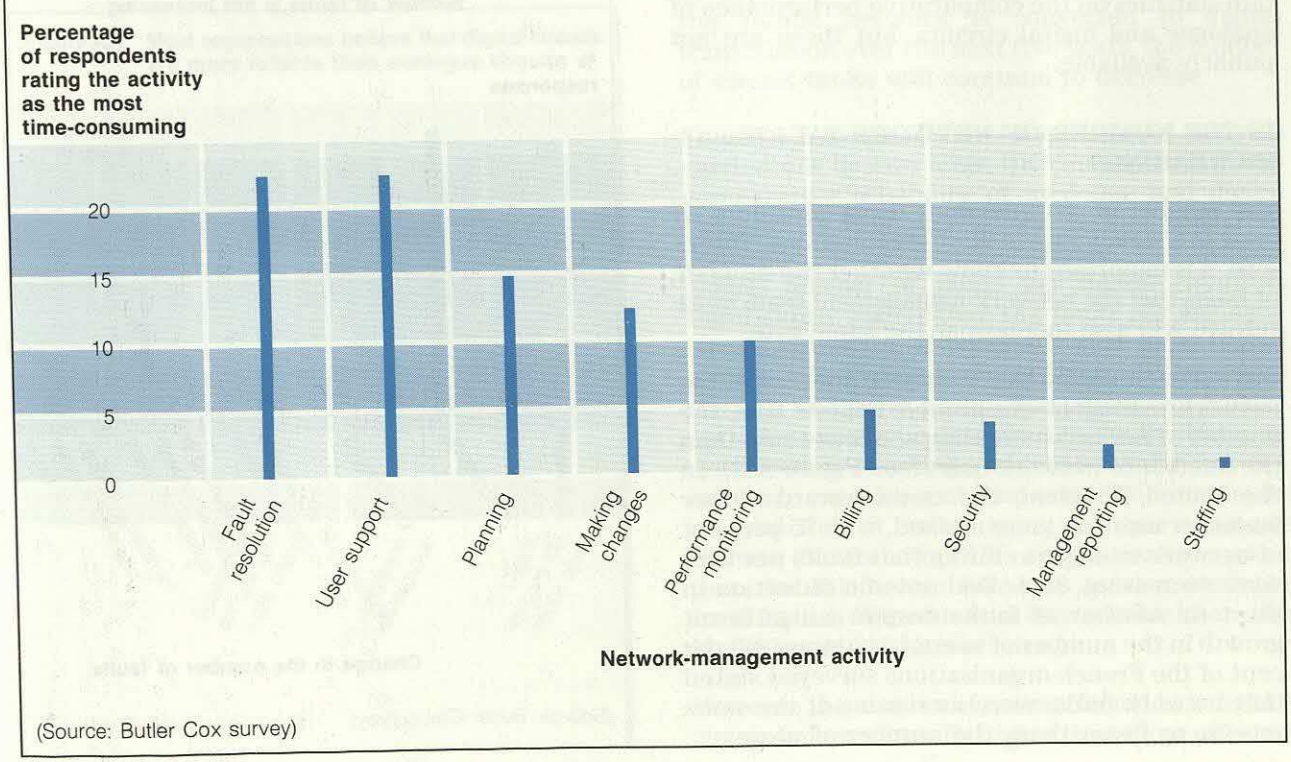

(Source: Butler Cox survey)

RESEARCH UNDERTAKEN

The information on which our analysis is based was obtained from two main sources:

- Interviews conducted with 28 organisations in Australia, Belgium, France, Germany, Italy, the Netherlands, Sweden, the United Kingdom, and the United States.
- Replies from six UK Foundation members to a detailed questionnaire.

Wherever possible, we attempted to obtain data about trends in network faults by asking for any information available from analyses of fault logs. However, we found that most organisations do not analyse their fault data on a regular basis, or have only recently begun to do so. We therefore also asked four questions about network managers' perceptions of faults and their causes:

- Is the number of network faults increasing?
- Are digital circuits more reliable than analogue circuits?
- Where do faults most occur often — in circuits, software, or equipment?
- Are faults becoming harder to resolve?

The responses to these questions are discussed in the following sections.

Since the number of responses obtained was limited, we also asked suppliers about their perceptions of trends in network faults. In general, suppliers agreed with our findings but were unable to provide us with any numerical data to back-up their statements. We understand that PTTs maintain statistics on the comparative performance of analogue and digital circuits, but these are not publicly available.

IS THE NUMBER OF NETWORK FAULTS INCREASING?

The number of data-network users continues to grow at a steady rate in most organisations. Therefore, it is important to know whether the number of faults that the network-management team must resolve will also increase at the same rate.

Our survey revealed that nearly 60 per cent of the network managers questioned believe that the number of faults is increasing at a lower rate than the number of network users (see Figure A.2). In the United Kingdom, the trend towards fewer faults per user was more marked, with 75 per cent of network managers stating that faults per user were decreasing. Some had noted a reduction in the total number of faults despite a significant growth in the number of users. In contrast, 80 per cent of the French organisations surveyed stated that network faults were increasing at the same rate as, or faster than, the number of users.

The reasons for these national differences cannot be determined from the limited data available. However, some UK members remarked that the introduction of digital circuits and the improved service from British Telecom were contributing factors to the decline in the number of faults.

Figure A.3 shows the network downtime experienced by a major Swedish manufacturing company during the period June 1984 to November 1987. The trend to increased availability is clear. Network unavailability has reduced from an average of 5.8 per cent in 1984 to 1.1 per cent in 1987. Unavailability in this context takes account both of the number of faults and of the average fault-resolution time. Thus, it is not possible to state that the number of faults has decreased during the three-year period, because the improved availability could be due solely to improvements in fault-resolution times. However, it appears likely that at least some of the improvement can be attributed to fewer faults occurring.

ARE DIGITAL CIRCUITS MORE RELIABLE THAN ANALOGUE CIRCUITS?

Figure A.4 shows the responses to the question about the relative reliability of digital and analogue circuits. It is clear that most organisations with digital circuits believe that they are more reliable than analogue circuits. This opinion is supported

Figure A.2 **Most network managers believe that the number of faults is not increasing**

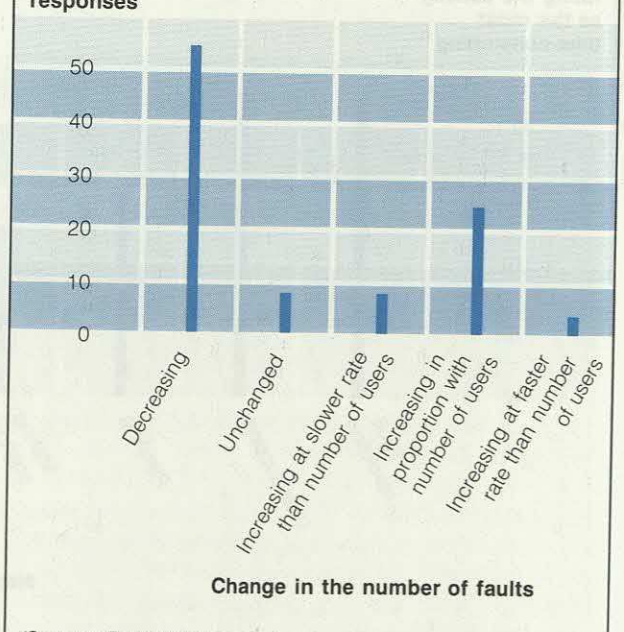

(Source: Butler Cox survey)

Figure A.3 Network downtime has decreased over the last three years in a Swedish manufacturing company

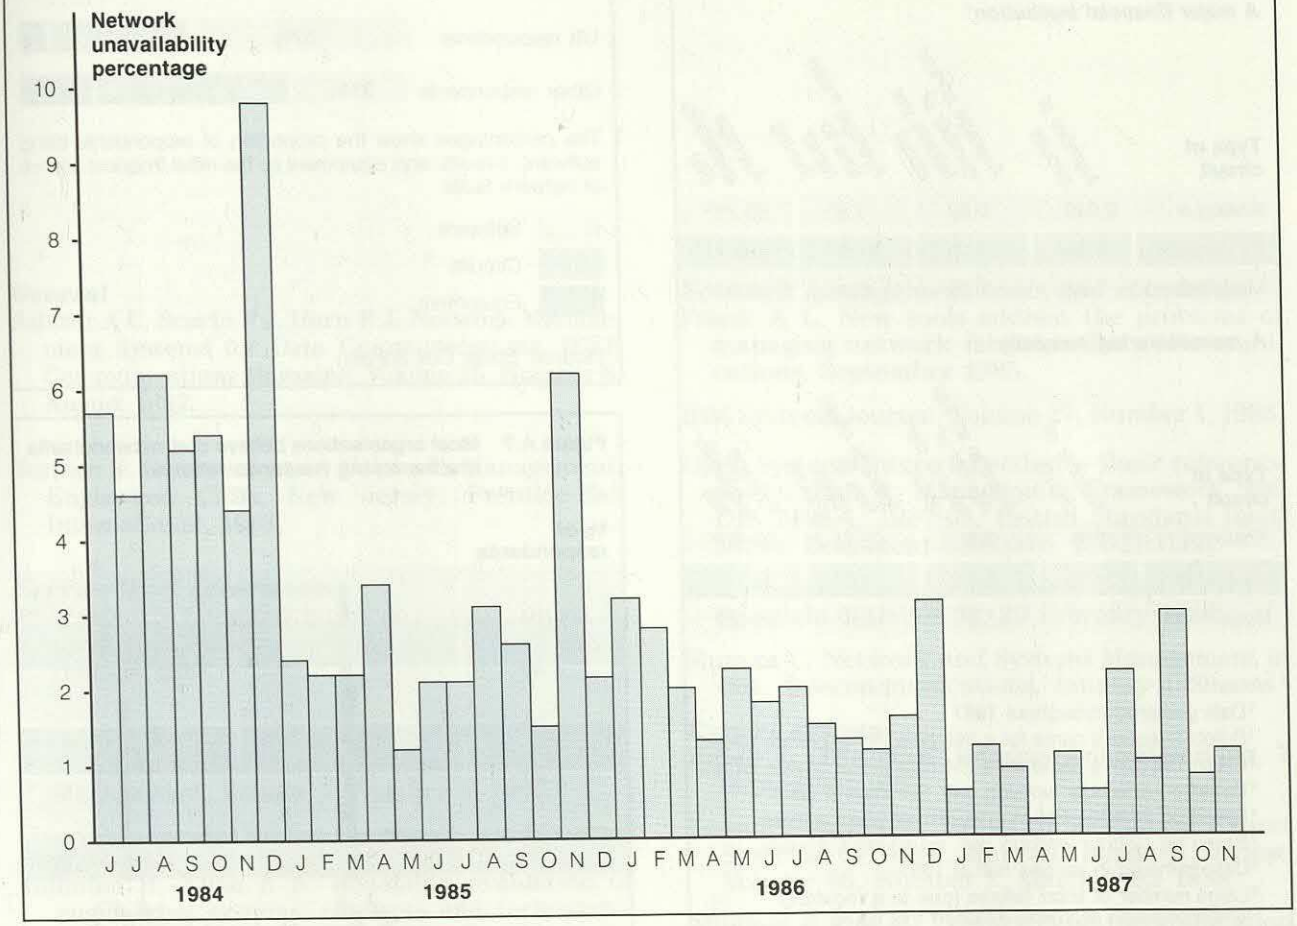

Figure A.4 Most organisations believe that digital circuits are more reliable than analogue circuits

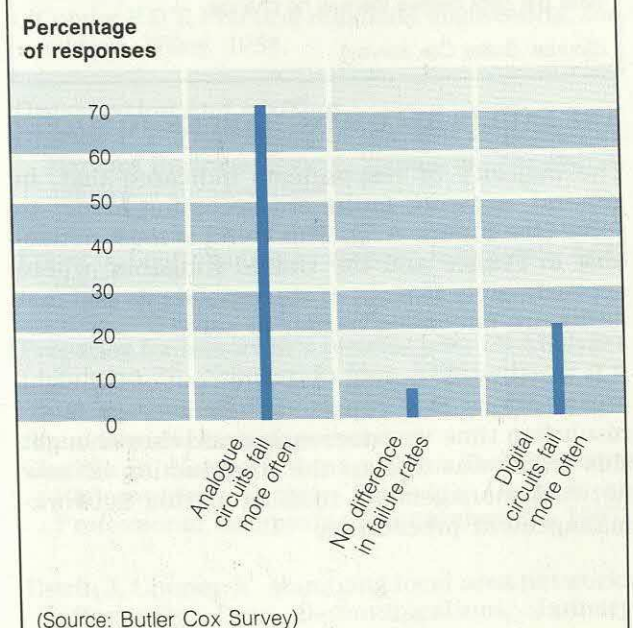

(Source: Butler Cox Survey)

by the PTTs. The implication is that, as more of the public network is converted to digital transmission over the next five years, the number of circuit faults will continue to decrease.

Figure A.5 overleaf shows the results of analyses carried out by two large UK companies into the comparative reliability of analogue and digital circuits. At first glance, the results appear to contradict those described above. The average annual number of faults per digital circuit is significantly higher than the faults per analogue circuit. However, a wideband (2M bit/s) digital circuit usually replaces at least 10 analogue circuits, and a 64k bit/s digital circuit can replace up to four analogue data circuits. On this basis, an organisation installing a wideband digital link can expect to experience at least one-third fewer faults in total than it would using multiple analogue circuits. (Several other UK organisations remarked that the KiloStream 64k bit/s service available from British Telecom was less reliable than the MegaStream 2M bit/s service.)

However, the number of users affected by a digital-circuit failure is usually higher than those affected

Figure A.5 Summary of results of fault analyses performed by two UK companies

A major financial institution[1]

Type of circuit	Number of circuits	Average annual number of failures per circuit	Average mean time to repair (hours)	Average availability (%)
Analogue	2,045	0.69	7.9[4]	99.77[4]
KiloStream[2]	78	1.13	14.0[5]	99.82[5]
MegaStream[3]	24	4.25	3.4[5]	99.84[5]

A manufacturing company[7]

Type of circuit	Number of circuits	Average annual number of failures per circuit	Average mean time to repair (hours)	Average availability (%)
Analogue	248	0.8	23.5[6]	99.63[6]
KiloStream[2]	12	5.2	4.0[6]	99.75[6]
MegaStream[3]	7	57.5[8]	0.4[6]	99.55[6]

Notes:
[1] Data gathered throughout 1987
[2] British Telecom's name for a 64k bit/s digital private circuit
[3] British Telecom's name for a 2M bit/s digital private circuit
[4] Based on nine-hour working day, five days a week
[5] Based on 24-hour working day, seven days a week
[6] Based on 12-hour working day, seven days a week
[7] Data gathered in second half of 1987
[8] Large number of short failures (due to a recurring synchronisation problem) distorted this figure

by a single analogue circuit failure. Network managers therefore need to consider carefully the back-up facilities required when they install digital circuits to replace several analogue circuits.

DO FAULTS OCCUR MOST OFTEN IN CIRCUITS, SOFTWARE, OR EQUIPMENT?

Figure A.6 shows that circuits and software are the most frequent sources of network faults. The equipment itself is, in general, much more reliable.

Again, there is a marked difference between responses from UK organisations and from those in continental Europe and Australia. The majority of UK organisations stated that the largest source of faults was software. Elsewhere, circuits were cited as the major source of faults. The reasons for these national differences may well be the same as those mentioned when describing trends in number of network faults.

Figure A.6 Circuits and software are the most common sources of network faults

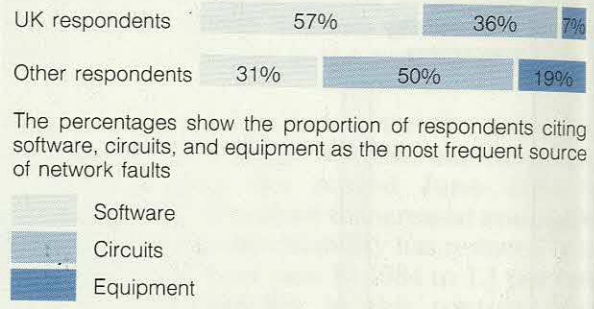

The percentages show the proportion of respondents citing software, circuits, and equipment as the most frequent source of network faults

(Source: Butler Cox Survey)

Figure A.7 Most organisations believe that network faults are becoming harder to resolve

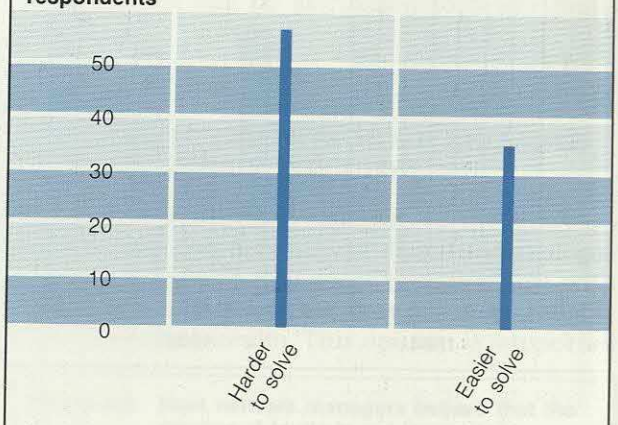

Nine per cent believe there is no change

(Source: Butler Cox survey)

ARE FAULTS BECOMING HARDER TO SOLVE?

The majority of respondents indicated that, in general, network faults are becoming harder to solve (see Figure A.7). This trend is most noticeable in France and the United Kingdom, where approximately five out of six respondents said that faults are becoming harder to solve.

It is worthwhile noting that three out of the eight organisations that reported their average fault-resolution time was decreasing said they thought this trend was due to the introduction of new network-management tools or better network-management procedures.

Bibliography

General

Salazar A C, Scarfo P J, Horn R J. Network Management Systems for Data Communications. IEEE Communications Magazine. Volume 25, Number 8, August 1987.

Terplan K. Communication Networks Management. Englewood Cliffs, New Jersey: Prentice-Hall International, 1988.

Service-level agreements

Ghernaouti S. L'administration de réseaux: un art au service des utilisateurs. Télécoms International — Dossiers 1987.

Witzel C N. Service-level agreements: A management tool for technical staff. Journal of Capacity Management. Volume 1, Number 4, 1983.

Availability

Billinton R, Allan R N. Reliability evaluation of engineering systems: concepts and techniques. Plenum Press, 1983.

Dummer G W A, Winton R C. Elementary guide to reliability. 3rd edition. Pergamon, 1986.

O'Connor P D T. Practical reliability engineering. 2nd edition. Wiley, 1984.

Organisation and staffing

Drucker P F. The coming of the new organisation. Harvard Business Review, Volume 66, Number 1, January-February 1988.

Fraser G. Learning how to manage. Infomatics, April 1988.

Preparing for tomorrow's systems jobs. I/S Analyzer, Volume 26, Number 5, May 1988.

Local network management

Clark R. PC Network Management: Two Opposing Philosophies. Andrew Seybold's Outlook on Professional Computing, September 29, 1987.

Estrin J, Cheney K. Managing local area networks effectively. Data Communications, January 1986.

Network management tools and standards

Frank A L. New tools address the problems of managing network facilities. Data Communications, September 1985.

IBM Systems Journal, Volume 27, Number 1, 1988.

Open Systems Interconnection — Basic reference model, Part 4: Management Framework, ISO DIS 7498-4, 1987-08. British Standards Institution Document 88/60376, IST/21:1115.

Performance management tools. Computerworld, Spotlight Number 39, 29 February 1988.

Sluman C. Network and Systems Management in OSI. Telecommunications, January 1988.

Expert systems

Haber L. The Man. Information Networking, 7 September 1987.

Leonard-Barton D, Sviokla J J. Putting Expert Systems to Work. Harvard Business Review, Volume 66, Number 2, March-April 1988.

Milliken K R et al. YES/MVS and the automation of operations for large computer complexes. IBM Systems Journal, Volume 25, Number 2, 1986.

Cost justification

An Approach to Network Cost Savings. Net/Alert System Handbook. Mt. Laurel, New Jersey: Avant-Garde Computing Incorporated, 1984.

Exploiting Tele-Communications Benefits. I/S Analyzer, Volume 26, Number 4, April 1988.

Johnson B. Network Management Means Selling, Too. Communications Week Information Networking Closeup, 22 February 1988.

Shockley M. Overcoming Challenges and Risks and Reaping the Rewards of Successful Network Management. Communications News, May 1987.

Surveys

Blumberg D F. Private network growth demands new service perspective. Data Communications, August and September 1986.

Network Management Hardware and Software Market. International Resource Development Incorporated, June 1988.

Treacy M E, Nedzel A E. What does it cost to own and operate a corporate network? Index Group, April 1987.

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, Germany, 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 shortlist 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.

Selected reports

8 Project Management
20 The Interface Between People and Equipment
21 Corporate Communications Networks
22 Applications Packages
23 Communicating Terminals
24 Investment in Systems
25 System Development Methods
26 Trends in Voice Communication Systems
27 Developments in Videotex
28 User Experience with Data Networks
29 Implementing Office Systems
30 End-User Computing
31 A Director's Guide to Information Technology
32 Data Management
33 Managing Operational Computer Services
34 Strategic Systems Planning
35 Multifunction Equipment
36 Cost-effective Systems Development and Maintenance
37 Expert Systems
38 Selecting Local Network Facilities
39 Trends in Information Technology
40 Presenting Information to Managers
41 Managing the Human Aspects of Change
42 Value Added Network Services
43 Managing the Microcomputer in Business
44 Office Systems: Applications and Organisational Impact
45 Building Quality Systems
46 Network Architectures for Interconnecting Systems
47 The Effective Use of System Building Tools
48 Measuring the Performance of the Information Systems Function
49 Developing and Implementing a Systems Strategy
50 Unlocking the Corporate Data Resource
51 Threats to Computer Systems
52 Organising the Systems Department
53 Using Information Technology to Improve Decision Making
54 Integrated Networks
55 Planning the Corporate Data Centre
56 The Impact of Information Technology on Corporate Organisation Structure
57 Using System Development Methods
58 Senior Management IT Education
59 Electronic Data Interchange
60 Expert Systems in Business
61 Competitive-Edge Applications: Myths and Reality
62 Communications Infrastructure for Buildings
63 The Future of the Personal Workstation
64 Managing the Evolution of Corporate Databases

Forthcoming reports

The Marketing of the Systems Function
Computer-Aided Systems Engineering (Case)
Mobile Communications
Software Strategy
Electronic Document Processing

Availability of reports

Members of the Butler Cox Foundation receive three copies of each report upon publication; additional copies and copies of earlier reports may be purchased by members from Butler Cox.

Butler Cox & Partners Limited
Butler Cox House, 12 Bloomsbury Square,
London WC1A 2LL, England
☎ (01) 831 0101, Telex 8813717 BUTCOX G
Fax (01) 831 6250

Belgium and the Netherlands
Butler Cox BV
Burg Hogguerstraat 791c,
1064 EB Amsterdam
☎ (020) 139955, Fax (020) 131157

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cédex 1, France
☎ (1) 48.20.61.64, Télécopieur (1) 48.20.72.58

Germany (FR)
Butler Cox GmbH
Richard-Wagner-Str. 13,
8000 München 2
☎ (089) 5 23 40 01, Fax (089) 5 23 35 15

United States of America
Butler Cox Inc.
150 East 58th Street, New York, NY 10155, USA
☎ (212) 891 8188

Australia and New Zealand
Mr J Cooper
Butler Cox Foundation
3rd Floor, 275 George Street, Sydney 2000, Australia
☎ (02) 236 6161, Fax (02) 236 6199

Ireland
SD Consulting
72 Merrion Square, Dublin 2, Ireland
☎ (01) 766088/762501, Telex 31077 EI,
Fax (01) 767945

Italy
SISDO
20123 Milano, Via Caradosso 7, Italy
☎ (02) 498 4651, Telex 350309, Fax (02) 481 8842

The Nordic Region
Statskonsult AB
Stora Varvsgatan 1, 21120 Malmo, Sweden
☎ (040) 1030 40, Telex 12754 SINTABS

Spain
Associated Management Consultants Spain SA
Rosalía de Castro, 84-2°D, 28035 Madrid, Spain
☎ (91) 723 0995