

# BRITISH TELECOMMUNICATIONS ENGINEERING

## *Included in this Issue*

*Integrating the Internet with the  
Voice Telephone Network  
Callscape*

*CTI for the Small Business*

*How 'Smart' Can We Get?—*

*Technology and Business*

*Opportunities for Smart Cards*





# BRITISH TELECOMMUNICATIONS ENGINEERING

## Contents

VOL 15 ■ PART 4 ■ JANUARY 1997

<b>Editorial: Welcome to Our New Readers</b> Bill Mills	289
<hr/>	
<i>Perspectives</i>	
<b>Integrating the Internet with the Voice Telephone Network</b> Fred Briggs	290
<hr/>	
<i>Computer Telephony Integration</i>	
<b>Callscape</b> <b>Computer Telephony Integration for the Small Business</b> Graham Hillson, Chris Hardcastle and Marc Allington	293
<hr/>	
<i>Operational Support Systems</i>	
<b>Work Manager</b> Gerry Garwood	300
<b>A Service System for BT's Major Customers</b> Paul White	309
<hr/>	
<i>Telecommunications in the 21st Century</i>	
<b>Superhighway... Supergridlock?</b> Chris Gibbings, Dave Newson, Ian Henning, Jeremy Barnes, Dave Pratt and Peter Cochrane	315
<b>Advocating a Managed Optical Platform</b> Ewart Lowe and Ian Hawker	322
<hr/>	
<i>Process Engineering</i>	
<b>Using Software Agents for Business Process Management</b> Paul O'Brien, Mark Wiegand, Brian Odgers, Christos Voudouris and Donald Judge	326
<hr/>	
<i>Making it Happen</i>	
<b>Christmas in BT's Network Management Centre</b> Stephen Morgan and Graham Barratt	334
<hr/>	
<i>Human Factors</i>	
<b>Learning to 'Work Smarter' in a Virtual Business</b> Roger Martin and Peter Hornsby	340
<hr/>	
<b>How 'Smart' Can We Get?</b> <b>Technology and Business Opportunities for Smart Cards in the Age of Electronic Commerce</b> Mike Arnavutian	346
<hr/>	
<b>Organisational Learning</b> <b>Early days in a change process</b> Mick Cope	356
<hr/>	
<b>Workplace Assessment for National Vocational Qualifications</b> Alan Martin	363
<hr/>	
<i>Regular Features</i>	
<b>pcochrane@btlabs</b>	369
<b>Book Reviews</b> 370, <b>BT News</b> 372, <b>Industry News</b> 374, <b>IBTE Contacts</b> Inside Back Cover	

### Board of Editors

**Chairman** Bill Mills, M.A., C.ENG., M.I.E.E.

**Managing Editor** Paul Nichols, B.SC.    **Secretary** Les Brand, M.ENG., M.B.A., D.I.C.    **Treasurer** Denis Taylor, F.C.M.A.

**Assistant Editor** Richard Kemp, B.SC., C.ENG., M.I.E.E.; Roger Blake, C.ENG., M.I.MECH.E.; Gavin Costa, C.ENG., M.I.E.E.;  
Martin Evans, B.SC., PH.D.; Warwick Flury; David Greenop, B.SC., M.SC.; Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E., M.I.D.M.;  
Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT.; Keith Ward, C.ENG., F.I.E.E., M.I.MGT.; David White, M.I.P.D., M.B.I.M., L.I.A.A.

*Bill Mills*

# Welcome to Our New Readers



*Our aim is to provide you with information about developments and trends within our dynamic and rapidly changing industry.*

This issue of the *Journal* marks a significant stage in the evolution of the new IBTE. Last November, the first Conference of the new, combined IBTE took the key decision to distribute the *Journal*, together with its associated *Structured Information Programme (SIP)*, to all Full Members as one of the universal range of benefits of membership. This issue is the first to be distributed in this way. It is therefore with great pleasure that I welcome more than 4500 new IBTE Members to these publications and I trust that you will draw increasing benefit from them.

Our aim is to provide you with information about developments and trends within our dynamic and rapidly changing industry. Our objective is to increase your knowledge and professionalism and equip you with a set of transferable skills and background expertise to enable you to be successful within telecommunications and its related fields.

In each quarterly edition of the *Journal* you will find a number of high-quality readable articles covering a wide range of technical and business issues, together with news and other topical items relating to BT and the wider telecommunications industry. You will find that much of the *Journal* content is divided into 'themes', whereby a series of articles on one particular topic is spread across several editions. Themes currently being planned include the Internet, Multimedia Communications, Optical Developments, Broadband and External Plant and Practices.

Professor Peter Cochrane also contributes a regular column intended to stimulate you into thinking about the awesome possibilities of the future.

The *SIP* is IBTE's valuable reference work on telecommunications. It is a partwork that has been building over several years. Units for the *SIP* are issued with the *Journal*. Its 16 chapters cover everything from the Market to Networks, from Customer Interaction to Future Trends. New readers can purchase back issues of the *SIP*, but it is also our intention to reissue earlier units, as they are revised, with future editions of the *Journal* so that new readers can gradually build up their full collections.

Another important development is our plan to provide information supporting people undergoing telecommunications National Vocational Qualifications (NVQs). A number of articles have already been published providing background information about the nature and structure of NVQs and we intend to continue and extend this theme, as well as providing typical samples of evidence to assessors.

We are endeavouring to provide publications that are attractive, informative and of real benefit to Members. With Members coming from such a broad range of disciplines, this is a challenging task, but a challenge we are determined to tackle with vigour, and with your help. The Board of Editors would therefore be very keen to hear from anyone who has suggestions for topics that they feel should be included or indeed any suggestions for improvement. Equally, they would be interested to hear from any reader wishing to offer an article for publication. Please write to the Managing Editor, Post Point G012, 8-10 Gresham Street, London EC2V 7AG, or e-mail to [nicholsp@grsec2.agw.bt.co.uk](mailto:nicholsp@grsec2.agw.bt.co.uk). Alternatively, call the IBTE HelpLine on 0171 356 8008.

*Bill Mills***Chairman, IBTE Board of Editors**

*Fred Briggs*

# Integrating the Internet with the Voice Telephone Network

*Continuing the Journal's series of 'Perspectives' articles, Fred Briggs, Chief Engineering Officer of MCI, discusses the emergence of Internet telephony and examines the opportunities it brings.*


## A US History Lesson

Imagine a technology that transforms the way you live and work right before your eyes. Picture a global network that delivers content faster than you've ever dreamed, using the most advanced technology available in the world. A technology that creates a vast array of new markets and products, and can give a business a national presence, virtually overnight.

No, I'm not talking about telephone networks, or even the Internet. I'm actually looking back to the 1860s, when the first US network was being constructed mile-by-mile across this young nation from steel rails and wooden ties. The new network was the railroad, and if you were 'on-line' in Chicago, Illinois, your business experienced phenomenal growth. But if the technology passed you by, as in Tombstone, Arizona, your business withered and died. This is a US history lesson to be sure, but one with global implications. It is a lesson we cannot afford to forget as we move into the new millennium and as we continue the historic merger of BT and MCI.

## Growth of the Internet

In the US and around the world, the telecommunications industry is experiencing the same phenomenal force of change today that our predecessors experienced in the 1800s with the advent of the railroad—only the stakes are much higher and have grown to global proportions. Instead of shipping cattle across the country in weeks down railroad tracks, we are carrying data around the world in seconds via



*Fred Briggs Chief Engineering Officer of MCI, Chief Technology Officer of Concert*

the Internet. Traffic on MCI's Internet fibre-optic backbone network is growing at a monthly pace of 15 percent, as nearly 350 terabytes per month of data flow at speeds of 622 Mbit/s.

Once driven by the telephone, demand for network capacity is now determined by the rapid proliferation of fax machines and computer modems, as more and more companies figure out how to do business on the Internet and on their private intranets. We currently project that, just after the year 2000, the Internet will rival the telephony network in size and scope.

It was just a decade ago in the United States that the break-up of the world's largest monopoly launched a revolution in the telecommunications industry, bringing greater choices and lower prices to consumers, while sparking waves of unprecedented technological innovations.

## Emergence of Internet Telephony

Today a second revolution is at hand, as the Internet evolves from a curiosity into a valuable business tool.

As we begin to consider how best to link the Internet and the switched network in ways that will bring an exciting array of highly advanced broadband services directly into the home and workplace, Internet-based telephony has emerged as one of the most intriguing technologies and perhaps the logical next step. Internet telephony, of course, involves the transmission of voice traffic over the Internet (packet switched network) as opposed to the traditional switched telephone network. The growing interest in this application is simply a symptom of a larger phenomenon: a desire to use the

communicate using the IP network with the same ease as they do today via the telephone network. Consumers will be able to send and receive information using computers, traditional telephones or a combination of the two, via cable-data modems, asymmetric digital subscriber loop (ADSL) and even wireless service. We'll soon be offering consumers services that used to be seen only in science fiction movies.

There are many technological hurdles to overcome, and until the quality of Internet telephony improves significantly, the vast majority of businesses and consumers will continue to pick up the familiar telephone when they need to ring up someone down the street—or around the world. There are about 660 million conventional telephone line terminations worldwide, compared to about 13 million Internet hosts, and

web site. Clicking a voice button on the site, they could be directly linked to a mortgage counsellor who could assist with filling out a loan application. With another click, a real estate agent could take those same home buyers on a virtual tour of available properties in their price range.

An editor could splice tape for the evening news in full view of a producer located in a newsroom thousands of miles away. A travelling executive could check e-mail while returning telephone calls from her hotel room. Multicasting, video conferencing, and document sharing all become easier through the use of Internet-enabled communication.

Internet-based applications will make use of the voice network by putting its services under software control.

And just as the personal computer has moved from the business into the home, consumers will want access to these high-speed, broadband services 24-hours a day. High-speed technologies like ADSL, that once required fibre, will now deliver communication capacity over the twisted copper pair. Speeds in excess of 50 Mbit/s have been achieved over distances of 1000 feet on twisted copper pair. Such speeds will bring interactive services of all kinds within the economic reach of homes and small offices, helping to accelerate the number of people who telecommute.

The good news for proponents of Internet-switchnet technology is that by the beginning of the next century, use of the Internet will be as common as the telephone system. For most people, using the multimedia applications via the Internet and the switched network will be an extension of a technology that is already a familiar part of their daily lives.

## Issues

Clearly, many issues must be resolved before these two powerful networks can be truly linked. Current technology doesn't bring users the clear, crisp, readily avail-

## *Internet telephony will create new value-added services that will drive new business and new markets*

Internet for a full suite of integrated services including voice telephony, 'radio' broadcasting, network video, interactive multi-party computer games, and a host of other computer-mediated applications.

### Exciting Opportunities

Internet telephony will fuel demand for the creation of complementary products, evolving in a manner similar to cellular and paging.

But that's just the beginning. As the lines continue to blur between computers and telephones, broadcast and telephony, long distance and local calling, the integrated switched and Internet protocol (IP) networks will play a key role, bringing us closer to a time when communication becomes seamless across all media, revolutionising the way we do business. We envisage the day when consumers and businesses will

54 million global Internet users. Plainly, the Internet has some distance to go before it is as widely deployed as the telephone network.

But more importantly, we believe that the wide variety of business applications made possible via Internet telephony will create new value-added services that will drive new business and new markets. And we believe Internet telephony will complement and add value to our existing voice and data products.

Internet telephony could be used, for example, in a company's customer service centre. A customer service representative would be able to talk to the customer while both are also interacting with their computers—without the need for separate telephone and computer access lines.

The technology would be similarly useful in other business settings. A prospective home buyer might research mortgage rates on a lender's

able communication they enjoy via the switched telephone network. The Internet capacity to carry real-time services in quantity will depend on some serious upgrades in its implementation (requiring new features in routers) and in its total capacity. In most cases, Internet 'phone calls' must be arranged in advance and still must take place via personal computers equipped with specialised software. Interoperability standards among vendors and service providers will play a key role in the success of this medium and are still to be developed. Common databases that can operate using either a uniform resource locator (URL) or a traditional telephone number would be a key element to a truly integrated network.

There are other issues that must be faced. How well will the current fibre-optic-based Internet backbones, such as the ones operated by MCI and BT, handle the new surge of traffic? It is clear that as traffic grows on the Internet, backbone providers will have to boost capacity while utilising new technologies to achieve real-time Internet service. The Internet's demand for bandwidth is growing far faster than the switched telephone network. MCI has begun taking necessary steps, recently upgrading its Internet optical-fibre backbone to OC-12, or 622 Mbit/s, as we make ready for the next wave of Internet applications. Increases to OC-48 (2.4 Gbit/s) and beyond will be required over the next few years.

There are also many regulatory uncertainties. Will Internet telephony and the rest of the Internet enjoy freedom from governmental interference, or will the entrepreneurial spirit of the Internet ultimately be stifled by restrictive regulatory measures, as advocated by some of our competitors in the US and abroad?

## MCI and BT

MCI is already the leading US carrier of commercial international Internet

access service, providing connections in nearly 60 countries.

Through Concert InternetPlus (CIP), announced in June 1996, MCI and BT are creating the first globally managed IP network, providing businesses with consistent performance standards, enabling companies to utilise the Internet as a virtual private global network for applications such as groupware, messaging and electronic commerce. CIP offers high-speed dedicated access up to 45 Mbit/s to the Concert backbone from locations around the world, and will increase the overall international capacity of today's Internet by 30 per cent. Through CIP, MCI and BT will continue to build the world's most technologically advanced communication network and take advantage of the functional commingling of conventional voice services and Internet-based applications.

## Conclusion

As in the past, the winners in the new telecommunications era will be those who grasp the new opportunities it creates, to remake themselves and their businesses, to take advantage of new ways to augment older services and to discover new services not achievable in the past, through the combining of the traditional with modern methods.

The choices that vendors, IT managers, and businesses make, as we enter a new millennium, will be crucial to their success. They must carefully consider if their decisions will result in a skyscraping monument to the future or a tombstone commemorating the past.

## Biography

Fred Briggs is the Chief Engineering Officer of MCI Communications Corp. The engineering division encompasses such key areas as network systems engineering, global network services, strategic accounts engineering, network systems development, intelligent services platform, wireless engineering, network architecture and other advanced technology. When the BT/MCI merger is complete, he will become the Chief Technology Officer of the new company, Concert Communications. He began his career at MCI in 1983 with the International Network Planning organisation, and in 1984 directed the operations department of Western Union International. Later, he was named Vice-President of Data Services for MCI International, responsible for international messaging and private-line services. Most recently, he served as Vice-President of Data Services Engineering, and then senior Vice President of Network Services Engineering. Before joining MCI, he held operations management and engineering positions with AT&T and Mobil Chemical Company respectively. He served in the US Navy as an engineering officer on board nuclear submarines, qualifying as Chief Nuclear Engineer. He holds a Bachelor of Science degree in Electrical Engineering from Oregon State University.

Graham Hillson, Chris Hardcastle and Marc Allington

# Callscape

## Computer Telephony Integration for the Small Business

*BT's Callscape product integrates a basic range of telephony features with an IBM compatible PC. This article examines the features of the product and the philosophy behind its design. It describes the hardware and software systems, and closes by looking at future developments.*

### Introduction

The July 1995 issue of *British Telecommunications Engineering* introduced the theme of computer telephony integration (CTI)<sup>1</sup>. The article discussed the key types of CTI, including desk-top integration, PBX-based integration and network-based integration. Desk-top integration is the term used when a computer application is linked via a physical interface at the desk top to a single telephone and/or exchange line. This type of integration is also known as *first-party CTI*, as the application has the same view of the network as the telephone it replaces.

This article describes BT's Callscape—an example of desktop integration realised using the external telephone interface approach discussed in the original article.

### Product Overview

Callscape's design integrates a basic range of telephony features with an IBM compatible PC meeting a defined performance specification. The key feature of the product is support of the Caller Display service; this enables the PC to display information about the caller before the call is answered.

The product consists of two main elements: the hardware and the software. The hardware interfaces the PC to the telephone line and provides a socket for the connection of any standard telephone. The hardware has a range of detection circuits to enable the application to know what is happening on the line, and can send dual-tone multifrequency (DTMF) and




Figure 1—Callscape connection arrangement

timed-break recall signals to the line. The software, supplied on several floppy disks, consists of the application itself and several other components needed to support Callscape's features and different versions of Windows.

The minimum PC specification required to run Callscape is a 486SX 25 MHz processor with 8 Mbyte random-access memory (RAM), running Windows 3.1 or above. The application requires 8 Mbyte of hard disk space. Figure 1 shows the general connection arrangements.

### Features

#### General

Callscape runs in the background mode while productivity software such as word processors and spreadsheets run in the foreground. The application may be set to open automatically when the handset of the associated telephone is lifted or when an incoming call arrives. As with any Windows application, Callscape can also be opened at any time by clicking its icon.

#### Directory and quickdial lists

Callscape features a large directory, which holds name and number

Figure 2—Callscape case design

information about the customer's contacts. In addition, users can create multiple quickdial lists to group their contacts logically under user-defined headings. Callscape includes pre-programmed quickdial lists for the most commonly used BT national and international service numbers.

### Caller Display


The key benefit of Callscape is the way in which it handles incoming calls. Callscape detects an incoming call by the arrival of a Caller Display message; this in turn causes Callscape to screen pop details about the caller. If the Caller Display message contains the number of the caller, the number and geographical location are displayed. At the same time, Callscape searches through its directory looking for a matching entry and if one is found also displays the details of the caller's name and company. If information about the caller has been stored from previous telephone calls using the Callscape notes facility, then this information is available instantly, even before the user answers the call, by pressing the Notes button.

### Select Services and Featureline

Callscape supports most current Select Services, such as divert and reminder call, which have common exchange enabling codes. A set-up screen option allows the user to customise Callscape to show only the services to which the customer subscribes. Where Callscape is connected to Featureline, the buttons normally used for Select Services can be reconfigured for direct access to Featureline facilities.

### Calls log

The calls log keeps a record of all incoming and outgoing calls made from either Callscape or the associated telephone. In addition to the name and number details, each record in the log gives the time, date and duration of the call, and for an incoming call whether the call was



answered. This information allows users to manage their calls and to return those left unanswered.

### Outgoing calls

Callscape supports DTMF signalling and timed-break recall only. This provides 100% compatibility with the public switched telephone network (PSTN), as well as with most PBXs that support two-wire equipment. Calls may be directly dialled from the associated telephone, from the PC keyboard or from the dial keys shown on the application, although it is expected that most users will use the sophisticated directory and quickdial features to dial from memory.

### Importing data

Callscape can import data from external databases into its directory, and can also save its directory in formats suitable for use by other databases. Callscape also supports Open Database Connectivity (ODBC), which allows it to use data from remote databases.

### Powering

To enable Callscape to continue to log calls when the PC is switched off, the hardware module is separately powered by a plug-top transformer. Details of line activity are stored in the hardware's memory and automatically downloaded to the application the next time it is opened.

### Design Philosophy

From the outset, Callscape was designed to be a volume product addressing the needs of small

businesses and home-office markets. Ease of installation and use were considered to be two key requirements that would enable the product to successfully address the chosen markets. Although these requirements are simple common-sense statements, they greatly influenced the direction taken by the design.

As discussed in the original article in the July 1995 issue of *British Telecommunications Engineering*<sup>1</sup>, the hardware could have been implemented in several ways. These include a PC-hosted telephony card, data connection from a telephone and a telephone interface unit. Anyone who has installed a sound card in a PC will know that the installation of a PC-hosted telephony card is likely to cause severe difficulty for many customers. Therefore, although this method would have offered many advantages, it was rejected. It was decided to implement the hardware using the external-interface approach rather than a telephone-based approach, as this enabled customers to use the telephones that gave them the features they required and, importantly, kept the cost down. The hardware casing was designed to reflect modern BT telephone design, and has several design features, such as the wavy loudspeaker grill, readily associated with BT products. The case colour is intended to harmonise with the colours used on most business PCs. A socket was included in the case for connecting the associated telephone, and directly connected cables were used for the PC and line cord connections. Figure 2 shows Callscape's case design.



Figure 3—Main Callscape window

Although Callscape provides a range of sophisticated features, and control of complex Select Services, it was imperative that the features were accessed in a logical and easily understood way, that encouraged the user to explore the product without recourse to complex user documentation. Additionally, although the product operates within a standard Windows environment, it had to be given a look and feel that could readily be associated with the BT brand. The way in which these requirements were met are discussed in the following section.

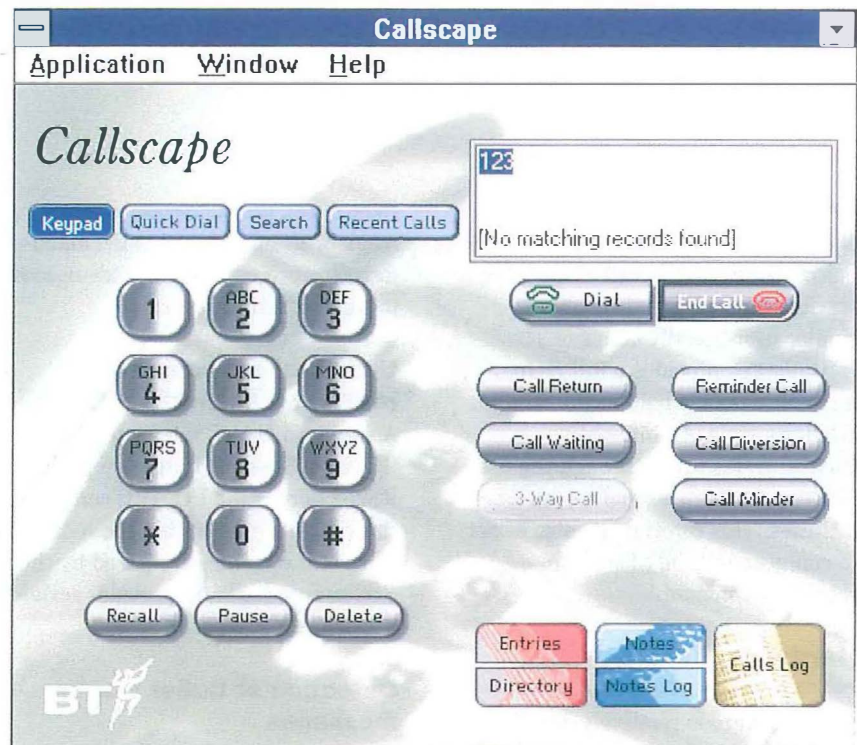
### User Interaction Design

It would have been relatively easy to design Callscape's user interface using standard Windows controls and colours. However, this would have resulted in a design that looked like any other Windows interface, and would not be readily associated with the BT brand. To help BT realise its requirements, a company was appointed to investigate the design options, and to model the chosen option so that BT could confirm customer acceptability before the real design work commenced.

Initially, three design options were considered, and each of these was modelled on a story board to see how they would appear on a PC. The option chosen used the metaphor of a telephone for the design, as it was felt that it would enable users to quickly feel comfortable with the product. The development of this concept was continued using Microsoft Visual Basic, and various options were experimented and tested until the final design was agreed.

Figure 3 shows the main Window of Callscape, and shows several of the key features of the design that help the product reflect the BT brand:

- The button shapes reflect those used in BT's telephone ranges, rather than the normal rectangular shapes used in Windows applications. Each of the button



shapes uses two bit maps, one for the UP position and one for the DOWN position. The text on the buttons was carefully edited to ensure maximum clarity.

- The graphics use 256 colours from a non-standard palette rather than the 16 colours used by standard Windows applications. This enabled the high-quality background graphic, depicting a telephone, to be designed. Sixteen-colour bitmaps are provided for PCs that do not support 256 colours.
- Separate background pictures are used for the main function keys to indicate the functionality of the key.

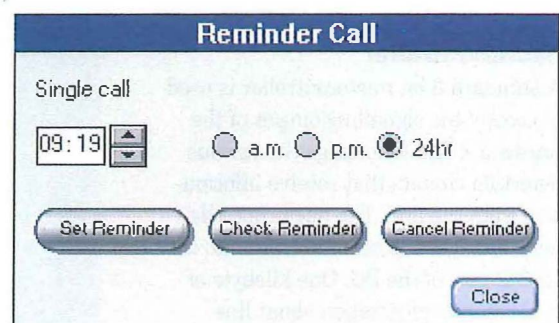
Figure 3 also illustrates an important concept used in the design—context sensitivity. The feature keys and other functions of Callscape are shown only if the feature controlled by that key is available.

Figure 3 also shows the 3-Way Call key greyed out. This tells the customer that, although Callscape has this particular feature, it is unavailable at this time because either Callscape is not in the correct phase of call set-up, or the customer does not

subscribe to this service. A set-up screen shown on initial installation, or from the menu at other times, allows the customer to tell Callscape to which services he/she subscribes. Callscape automatically arranges the positions of the feature buttons and status based on this information.

Figure 4 shows the dialogue box that appears when the Reminder Call key is pressed. This enables the customer to set, check and cancel reminders using the 12 hour or 24 hour clock, without needing to remember any of the Select Services codes. When the customer presses one of the keys, Callscape automatically seizes the line and dials the enabling code for the feature; the user can confirm that the service has responded correctly by listening to the exchange voice prompts through Callscape's call-progress monitor. The

Figure 4—Dialogue box for setting reminder call



Reminder Call dialogue box is typical of the way in which Callscape controls Select Services, and each of the other services controlled is through a similar user-friendly dialogue.

Callscape provides a full Help menu and bubble help in the familiar Windows style. An interactive tutorial walks users through the use of the main features and set-up facilities. This tutorial can also be set to run automatically for use in a retail environment.

**Hardware Design**

The hardware is relatively straightforward and makes use of off-the-shelf integrated circuits where possible. The hardware meets both the statutory UK approval requirements and any enhanced internal standards where considered appropriate. Some key performance areas are outside of the scope of the statutory approval requirements however, and these are designed against BT's own internal standards. Examples of these are the off-hook detection of the associated telephone, and the surge protection circuits built into the hardware. Figure 5 shows a block diagram of the hardware.

**Exchange line interface**

The line interface circuit ensures that the unit presents the correct DC conditions to seize the line properly, and the correct AC conditions to match the circuit to the line and exchange impedances. The line interface circuit also includes high-voltage surge protection, which protects Callscape from voltage surges that could damage its circuits.

**Microcontroller**

A standard 8 bit microcontroller is used to control the signalling output of the hardware, and to manage the various detection circuits that receive information from the line. The microcontroller also manages the communication to the RS232 port of the PC. One kilobyte of RAM stores information about line activity when the PC is turned off. To

reduce cost, the microcontroller is mask-programmed during the manufacturing process with the code needed to run the hardware.

**Universal asynchronous receiver transmitter (UART)**

A standard UART integrated circuit (IC) interfaces the hardware to the RS232 port of the PC. This converts the voltages used within the hardware to the voltages required by the RS232 standard and provides serial-to-parallel conversion.

**Detection of Caller Display messages**

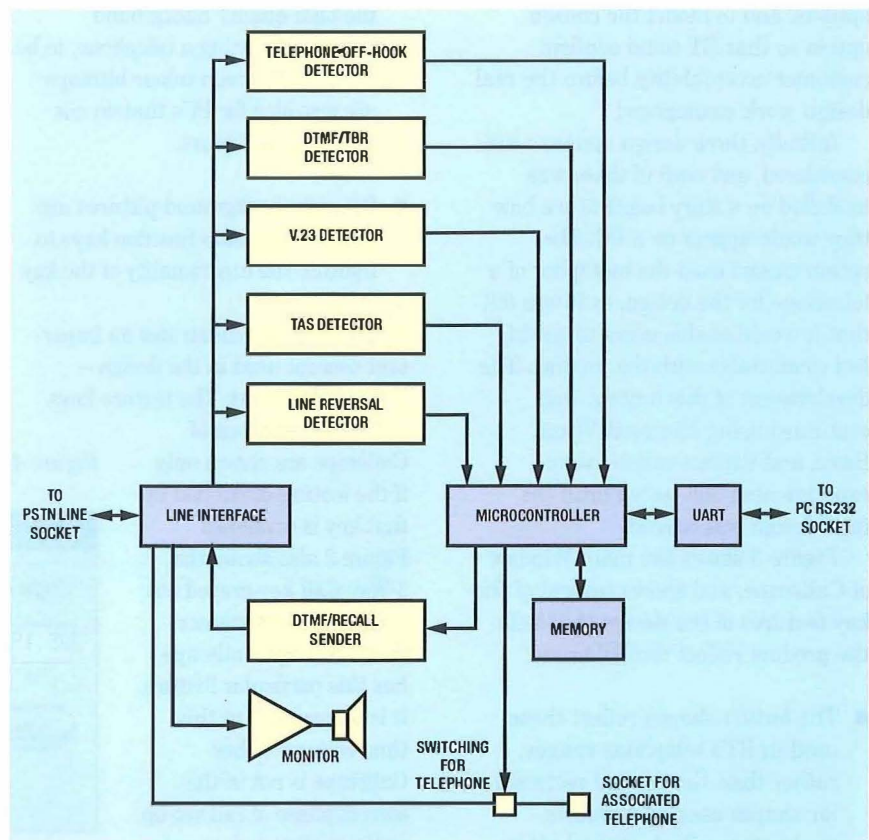
The line reversal detector, tone alert signal (TAS) detector and V.23 detector are required for the detection of a Caller Display message. Figure 6 shows the sequence of signals involved in the delivery of a Caller Display message. All the Caller Display detection circuits are realised using a standard integrated circuit from Mitel.

As shown in Figure 6, Callscape provides a current wetting pulse just after the TAS is received. The purpose of this pulse is to clear any high-resistance paths in the connection, to improve the reliability of message detection.

**Detection of associated telephone off-hook condition and signalling information**

The telephone off-hook detection circuit detects the change in voltage at Callscape's terminals when the associated telephone is taken off-hook. The detection of this event can be used to open the application from its minimised mode. Certain features and delays built into the circuit and controlling software prevent false triggering. A standard DTMF receiver circuit detects the outward signalling from the associated telephone, so that the information can be echoed on Callscape's feedback screen and entered into the call log.

Figure 5 – Callscape hardware



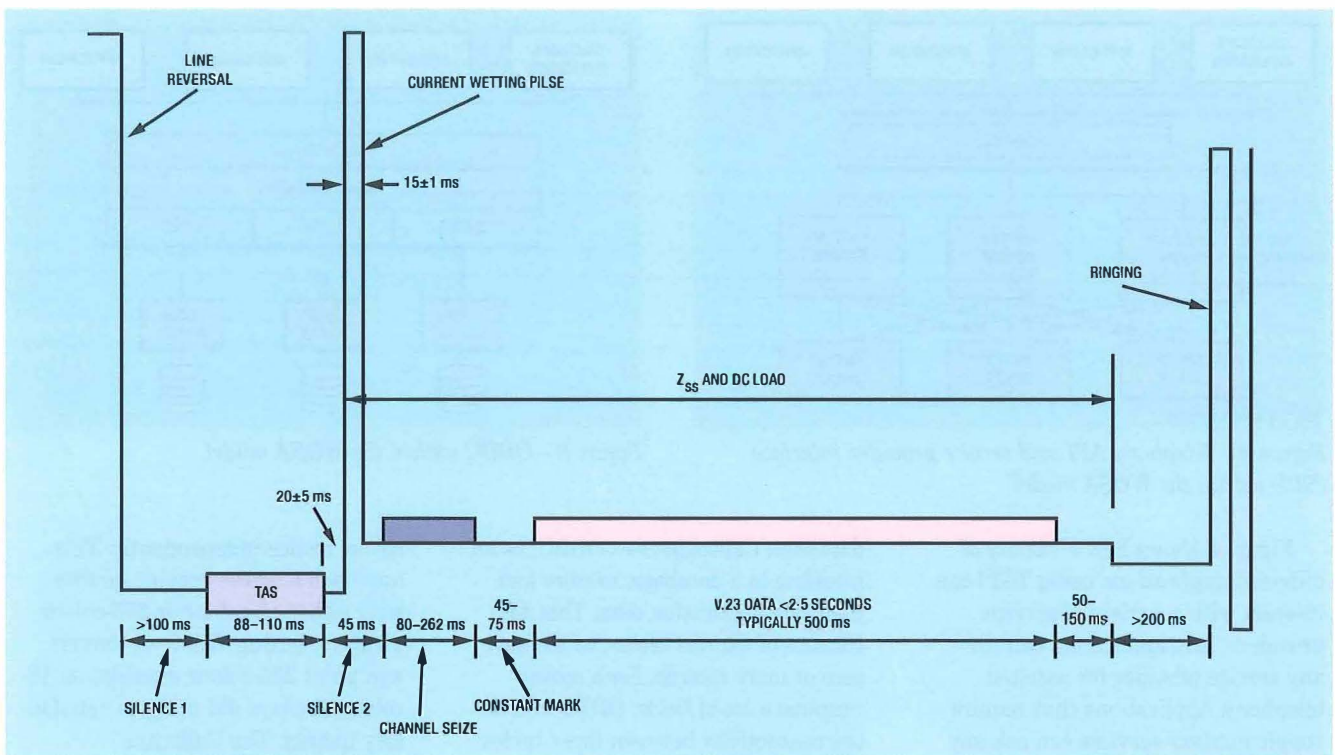


Figure 6—Calling line identity signalling sequence of events

### Outbound signalling

Outbound tone and timed-break recall signals are generated by a standard dialler IC under the control of the microprocessor.

### Call progress monitor

A call-progress monitor enables customers to monitor the progress of outgoing calls so that they can take control of the calls at the appropriate time. The level of the monitor, which can be set low or high, is controlled from the application.

### Telephone socket

The telephone socket on the hardware unit is simply connected in parallel with the line connection, but Callscape can switch the telephone in and out of circuit when required. For example, when the hardware is being used to send dialling information to the line, Callscape switches the telephone out of circuit to prevent accidental corruption of the signalling information.

### Software Design


Callscape was written using C++, which was considered to be the best language for writing commercial Windows applications using an object-oriented methodology. Wherever possible, common Microsoft

technologies were used to help the development process. These technologies can be split into three main areas which are implemented on the Windows Operating System Architecture (WOSA):

- the application framework, using Microsoft Foundation Classes (MFC);
- telephony, using Telephony Applications Programming Interface (TAPI); and
- database, using Open Database Connectivity (ODBC).

The Venn diagram in Figure 7 depicts the Microsoft technologies used in the development. These are now described in further detail.

Figure 7—Callscape foundation in the Windows 16 bit architecture



### The application framework

Callscape uses MFC as its application framework. An application uses a framework as an architectural base. It includes standard features such as creation and control of the graphical user interface (GUI), and handling user input from the keyboard and mouse; MFC handles many of the routine chores on behalf of the application. The MFC is popular because of its *document/view* architecture where a *document* is a container which stores and maintains an application's data, and a *view* is a method for displaying that data in a fashion suitable to the user. For example, Callscape's Note dialogue displays a note associated to a telephone call—the note data is stored in the document container, and the note view knows how to display that note data to the user.

### Telephony technology

Microsoft's TAPI enables developers to support telephony in applications. TAPI allows three types of telephony interaction: assisted, supplementary and extended. Callscape implements all of the assisted telephony features, a subset of the supplementary features including hold and transfer, and the extended features that allow Caller Display messages to be captured by Callscape.

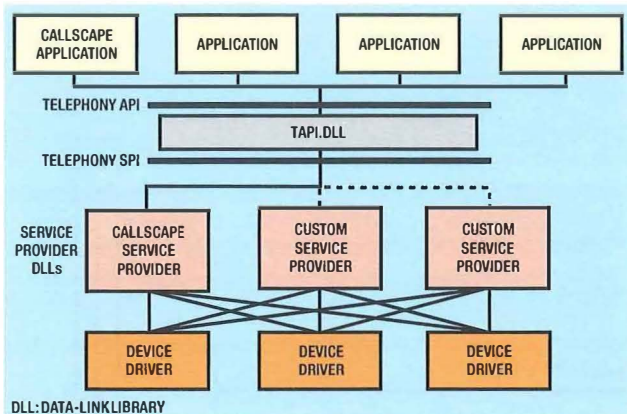


Figure 8—Telephony API and service provider interface (SPI) within the WOSA model

Figure 8 shows how a variety of different applications using TAPI can interact with a variety of service providers. Any application can use any service provider for assisted telephony. Applications that require supplementary services can ask any service provider if it supports that service. Other applications which require extended service functionality require specific service providers to support them. A service provider is used to communicate with the hardware, via the device drivers; Callscape implements its own service provider when extended telephony services are used.

The Callscape application uses TAPI to communicate with the Callscape hardware via the TAPI.DLL. Each telephony function is passed to the TAPI.DLL, and examined to determine which service provider should handle it. It is then passed on for further handling by the designated service provider. For example, when Callscape makes a call, the 'make call' function is passed to the TAPI.DLL. If the Callscape hardware is ready, then its service provider handles all the aspects of making the call to the destination, and then returns a result of success or failure.

TAPI allows other applications access to use assisted and supplementary telephony services via the Callscape service provider. For example, the Microsoft dialler application can use the Callscape service provider for all of its requirements.

### Database technology

Although many different mechanisms and technologies are available to store an application's data, Callscape's data is particularly well suited to that of a

database. Callscape uses ODBC, as an interface to a database, to store and maintain application data. This data consists of various tables, which have zero or more records. Each record contains a set of fields. ODBC handles the connectivity between these tables and the Callscape application.

Figure 7 shows a region occupied by both ODBC and MFC. Microsoft has wrapped ODBC within the MFC framework architecture, and Callscape uses this encapsulation of ODBC wherever possible.

### Potential uses of the ODBC technology

Any application which uses ODBC can extract data from the Callscape database (Figure 9). This allows businesses to generate reports which integrate data from both legacy databases and their Callscape database. An advantage of using the described Microsoft technologies is that other applications may benefit from Callscape's data either dynamically or by batch processing.

For example, a small business running Callscape and Excel can produce a bar chart which displays the busiest part of each month for its inbound customer calls against the outbound calls made by staff. This can be achieved by writing a simple Excel macro and associating the macro to the 'Tool' top-level menu in Excel. The business may also wish to use Excel (or similar) to study customers' or its own business telephony practices by using this method.

### Callscape components

Callscape consists of several program components. Firstly, Callscape supports both 16-colour and 256-

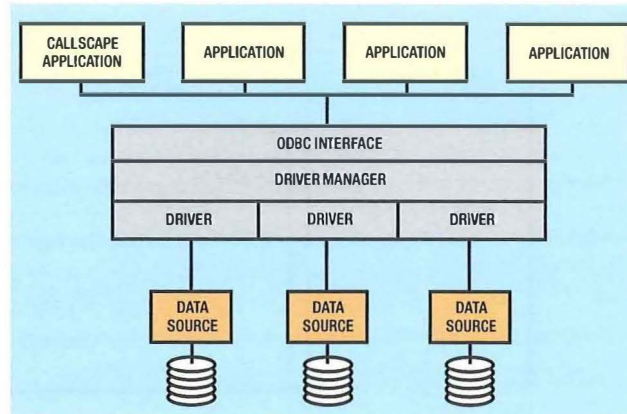


Figure 9—ODBC within the WOSA model

colour modes independently. This approach ensures graphic quality with either 16-colour or 256-colour displays; letting Windows convert and paint 256-colour graphics on 16-colour displays did not give satisfactory quality. The Callscape application uses one executable and two resource dynamic link libraries.

The service provider forms another important component which has to be installed onto the PC using Callscape. The service provider can be configured via the 'Telephony' applet found in the system's Control Panel. The installation for the Callscape applications, service provider and ODBC configuration is carried out initially by the Callscape set-up application, another component. The set-up application, written using the Microsoft Set-up SDK, also installs the initial database and any necessary ODBC drivers, as well their configuration. The customer can configure them manually, using the 'ODBC Administrator' found in the system Control Panel.

### Future Developments

Network enhancements will soon provide Caller Display with Call Waiting; consideration is being given to allowing Callscape to support this. Although this may seem a relatively minor change, it represents a considerable engineering challenge, as the TAS must now be detected in the off-hook mode. In this mode, speech may be simultaneously being transmitted and the detection circuits must now be able to reliably detect the tone in the presence of this speech. Two well understood problems that result from this are:

- The near-end or far-end speech may simulate the TAS, and cause the product to respond as if a legal signal had been received. This will result in the transmission being muted and a signal being sent to the far end. This effect is known as *talk-off*.
- The opposite problem is that if the speech at the particular instant that the Caller Display message is received is so loud that it masks the reception of the TAS, the message will not be received. This effect is known as *talk-down*.

Key to the success of any product that supports this service is that talk-off and talk-down problems are reduced to an absolute minimum. Bellcore has defined tight performance requirements for these effects and it is BT's intention to meet them. *New Suppliers' Information Notes* (SIN) being published by BT Regulatory Affairs may make some recommendations on performance but these are unlikely to form any part of future mandatory approval requirements.

Developments to Callscape's functionality could include making the method used by Callscape to communicate with the TAPI extended services publicly available. This would allow bespoke applications to be produced to meet specific requirements with minimum effort; advanced solutions could be implemented by small businesses themselves. For example, telephony functions may include power dialling, screen popping, or agent routing. Database functions may tie together legacy databases with Callscape's database. These solutions would be made available via component software implemented in Visual Basic, Delphi, Visual C++ or similar languages. Other future developments may include service providers for various PBXs and ISDN allowing Callscape to run from a variety of telephony networks and switches. These and several other options are currently being considered for what is expected to be a range of highly successful products.

## References

- 1 CATCHPOLE, ANDREW; CROOK, GARY; and CHESTERMAN, DOUG. Introduction to Computer Telephony Integration. *Br. Telecommun. Eng.*, July 1995, 14, p. 98.

## Biographies



**Graham Hillson**  
BT Networks and  
Systems

Graham Hillson joined BT in 1968 as a TTA in London South West Area. After

basic training he worked in the external planning office, subsequently moving to Northampton to join the Customer Works Group, planning large PBX and data communications systems. After several years in this role, he was promoted and moved to Anzani House in Felixstowe to work on the development of switching systems, and was involved in the development of several very successful systems including the Ambassador ESS and the Senator. Another promotion took him to BT Laboratories in Martlesham Heath, where he was the team leader for answering machine development and was responsible for the launch of many successful products. He currently manages the Telephone Product Support team, which is responsible for the development of most of BT's customer premises equipment, ranging from simple telephones to the CTI product described in this article.



**Chris Hardcastle**  
BT Networks and  
Systems

Chris joined BT Laboratories in 1975 after working in the electronic component industry. Initially he was involved in studying the reliability of

reed relays and then designing and running test systems for optical receivers and transmitters. In 1984, he joined the Telephony division with responsibility for hands-free telephony products, the launch of the highly acclaimed Prelude featurephone and subsequently the successful Duet range of telephones. More recently, Chris has been responsible for the development of screenphones based on the use of Bellcore's ADSI protocol. This led to the launch of the Mondex screenphone, currently being trialled in Swindon. This involved not only the development of the product, but also the specifications and UK requirements with which the product had to comply. He now leads the team responsible for the development of Callscape.



**Marc Allington**  
BT Networks and  
Systems

Marc Allington started in BT as a sponsored student in 1990, programming in 68000 for

Standard C mobile satellite communications. Starting his graduate BT career in 1992, he was involved in the development of local area planning tools for radio spectrum management. In April 1993, his role changed to a team member developing demonstrators of intelligent network services. From April 1995 he joined the Distributed Computing Group, and helped provide network services based upon computing architectures using various distributed languages. He has held his current position in the CTI group since December 1995. His work has been to develop Java CTI demonstrators which can be used within a call centre environment. He has been a technical consultant for the Callscape product, helping resolve various technical issues.

Gerry Garwood

# Work Manager

*A real-time control system has been developed and deployed nationally to support BT's Work Management Programme. This article traces the history, system architecture, development, deployment and service aspects of this very large programme. Many issues covering technical, procedural and operational aspects had to be tackled and overcome.*

## What is Work Management/Work Manager?

BT employs a large workforce to install and repair telecommunication equipment both at customers' premises and within its own telephone exchanges and transmission stations. Many are mobile and the challenge is to be able to assign all required work (upwards of 150 000 jobs per day) in the most effective manner to all available people (20 000).

BT's Work Management Programme provides the company with the ability to assign work to people in an optimal manner and to measure the performance. This involves understanding, defining and implementing work-flow processes in both reactive repair and provision across both customer access and core networks. Further, it is responsible for defining any development required by operational support systems (OSS) to support these processes.

Work Manager is a collection of interconnecting computer-controlled systems that support this programme by automatically collecting, allocating and distributing work to BT's field technicians, while providing management information to monitor the progress and measure the effectiveness of the task. The central computer is the Work Manager system (WMS) itself.

## Background

In the mid-1980s, as part of the expert system/artificial intelligence research programme<sup>1</sup>, investigations at BT Laboratories were looking at how the application of artificial intelligence technology could be applied to BT's operations. The focus of attention was developing an allocation algorithm that could match people to work. By using such

languages as Prolog and Poplog, demonstrator systems were produced and underwent trials with the cooperation of BT's operational division, now known as *Zones*<sup>2,3,4</sup>. This very early work in such areas as appointment/diary management, fault analysis, and job allocation to people in both repair and provision, together with a proposal on how to achieve repair service centre logistics<sup>5</sup>, laid the technical foundation for today's Work Management Programme. As a result, BT was able to show, at Telecom 87 in Geneva, a knowledge-based management system, with the ability to deliver a work allocation system matching available people to required work.

By 1988, a personal computer (PC) single-user production system, known as the *district control support system* (DCSS)<sup>6</sup>, was launched in conjunction with the Midlands Zone, with some 40 systems delivered nationally by 1989. This allowed each operational unit to manage out-of-hours call-out by matching the call-out condition to the most appropriate person, based on required skill, availability and location.

Recognising the success and importance of DCSS to BT, the company then embarked on a major development programme (starting 1989), to deliver the network operations management system (NOMS2) for the core network domain. This multi-user application supported network operations units (NOUs) and network field units (NFU) by assisting control officers to manage work and people and manual allocation.

Being a much larger application, one system per NOU was deployed using open system platforms based upon the UNIX operating system and the ORACLE database manager.

The Work Management Programme was born during the late-1980s, based on the early success of

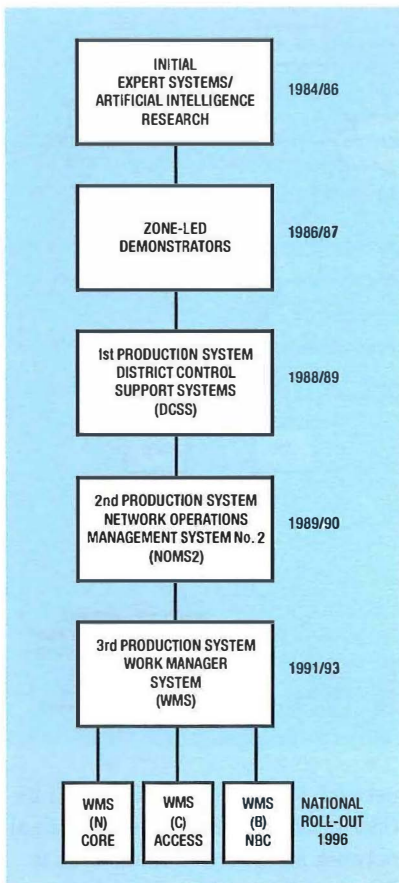


Figure 1—Historical development on Work Manager

DCSS and the development work of NOMS2. This programme focused on the development of Work Manager as a company-wide solution for all tasks related to work management. The key driver was to automate many of the previous manual allocation tasks. Figure 1 traces this development evolution path indicating that today not only are there systems for both the customer access and core networks but also a third for National Business Communications (NBC) to support its workforce within the large business systems community.

## Requirements

The Work Management Programme had to define and develop not only a system but also the operational processes to support the business requirements for the next 10 years, laying the foundation for BT's operations for the 21st century.

The three principal aims were:

(a) to achieve significant improvements in operational productivity by reducing unit costs in reactive repair and provision,

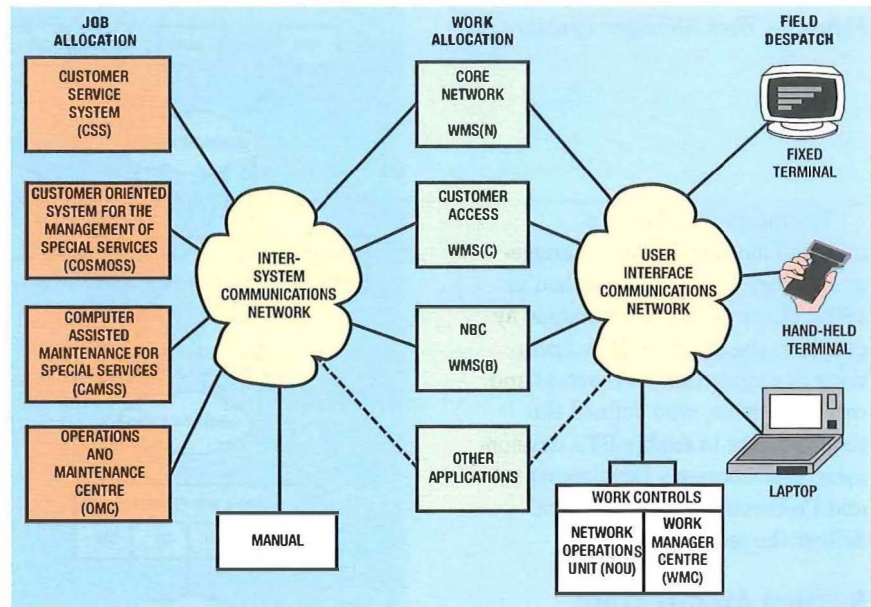


Figure 2—Overview of work management system

(b) to maintain the required quality of service for customers, and

(c) to provide a singular common approach to work management that was the basis for all future applications.

These were a significant challenge to the company because (a) and (b) are orthogonal to each other: generally, one can be achieved only at the expense of the other, and improving both together is another matter. Achieving (c) requires a very flexible system architecture that could be adapted to the many related work management tasks and evolve as the company's organisation changes.

To encapsulate the above aims into the programme's main objectives, the following simple statements were originally stated and are still appropriate today.

- get the right person
- in the right place
- at the right time
- with the right equipment/information
- in an optimum and cost efficient manner.

Thus Work Management in business terms is one of the company's key enablers.

The programme had to recognise and align itself with other strategic initiatives such as the Network Administration Implementation Programme (NAIP)<sup>7</sup> and the Strategic Systems Plan (SSP)<sup>8</sup> and more recently with the direction of the Breakout programme. All three are company-wide programmes to improve the effectiveness of BT.

The key requirement—to provide a company-wide solution that would allow the move away from manual allocation to an automatic facility—represented a significant cultural shift because it would allocate work on a 'just-in-time' basis. Thus individuals would no longer be able to determine and plan their own daily work pattern. This change is as significant as that which took place when the company moved towards automatic telephone exchanges and away from manual boards.

In technical terms, the solution had to collect jobs automatically from the appropriate systems, allocate work and despatch this to field technicians via a range of terminals. Figure 2 gives a simple schematic overview of the arrangement. It had to monitor progress and effectiveness, and close the work on completion of the jobs. In total the overall solution had to handle more than 150 000 jobs each working day, supporting 20 000 field technicians and 2000 control officers. Consequently it had to be flexible, reliable, serviceable, responsive and enduring.

Figure 3 – Work Manager structure

To support the above<sup>9</sup>, a multimillion-pound Work Management Programme was launched in 1989. This was led and managed by people at the centre of BT's operations in conjunction with access and core field users, who defined the requirements to enable BT's development unit (formerly Development and Procurement) to design and deliver the technical solution.

**System Architecture**

**Overview**

A national work management system had to accommodate both the core and access domains across both the reactive repair and provision workforce. Given the range and size of the tasks and current computing technology available, a single physical system was neither practical nor desirable. The proposed architectural solution was to develop a logical structure (Figure 3) that could be realised as a number of separate and independent physical systems.

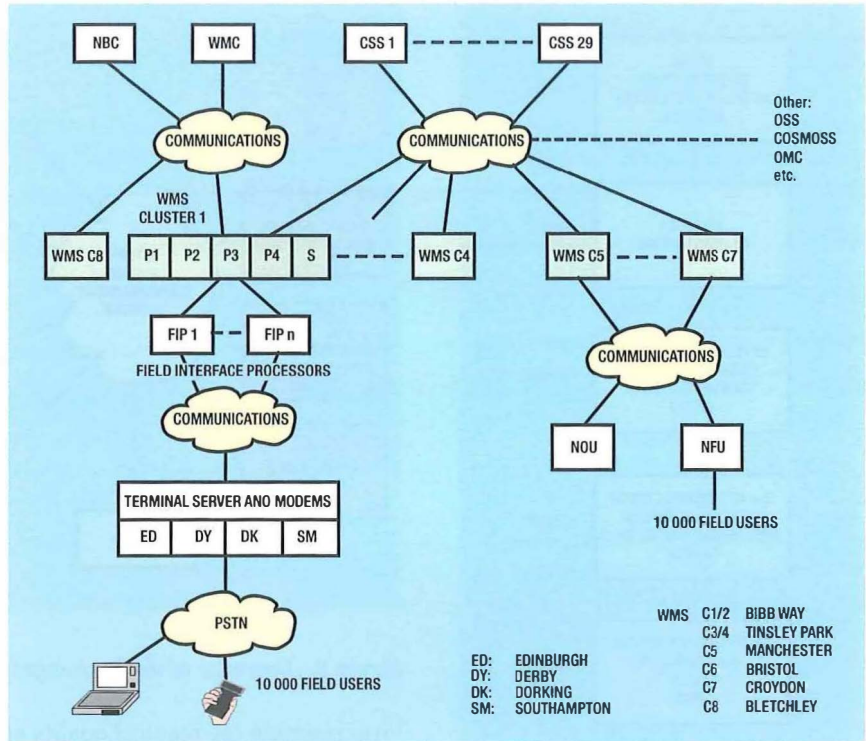
At the centre of the structure would be the Work Manager system (WMS) itself. This would be interconnected to other existing support systems (for example, customer service system (CSS)<sup>10</sup>), where appropriate, along with any new systems that are required to be developed.

The design objective for the WMS was to realise a single logical solution (Figure 4), but to meet operational requirements this would be available in two variants known as *WMS(N)* and *WMS(C)*.

In the core domain, which already used NOMS2, to support the NOU and NFU operations, *WMS(N)* would be the direct replacement<sup>11</sup>, using the existing infrastructure.

In the access domain, where no such system existed, *WMS(C)* and the required supporting infrastructure had to be provided anew. This embraced:

- the existing CSS,



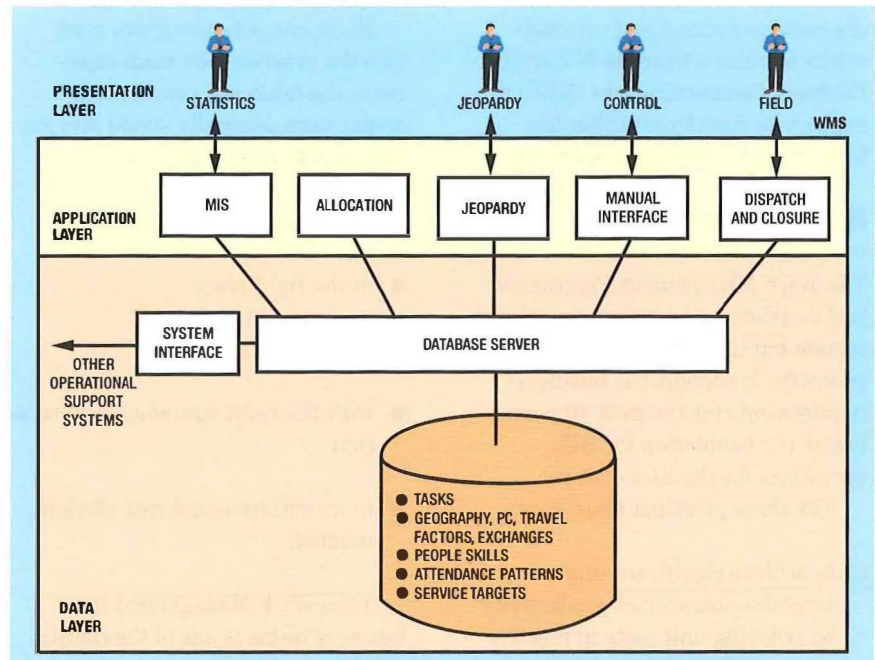
- WMS,
- a field access interface system (FAIS) to provide access to/from WMS from field terminals, and
- a hand-held terminal (HHT) for use by the mobile access workforce.

The total work management solution is shown in Figure 3, which gives an indication of the size and complexity of the arrangement. It uses the company's main computer

network for interconnection to all 29 CSS databases, the existing terminal network for local communication to PCs, and the public switched telephone network (PSTN) for HHTs. In the future, the WMS can be connected to other systems such as COSMOSS, CAMSS and OMC.

To give the initial national coverage, the design had to provide for 10 *WMS(N)* and 16 *WMS(C)* systems, but the design also provided the flexibility for additional *WMS*s to be accommodated, as demonstrated

Figure 4 – Work Manager logical design





by the later addition of WMS(B) for NBC. (Also, see the section on deployment.)

Within the architecture, three basic tasks are undertaken: job collection, work allocation and field dispatch.

**Job collection**

The company collects work in the form of jobs to be undertaken. These can be held on paper but are generally stored on a wide range of computer systems such as CSS. WMS scans the support system's output files and transfers details periodically via a bespoke interface. A manual interface allows those jobs not collected automatically to be entered. The frequency of transfers is variable but a suitable time interval is every five minutes.

Initially, a manual link was used to transfer job closure information from WMS, but this has now been superseded using the system interconnect service shell (SISS) technology.

A generic interface has now been designed to overcome the need for further bespoke interfaces.

**Work allocation**

The WMS is the central system for Work Management. Indeed it could be described as undertaking the combined heart and brain function, being a real-time allocation *control* system rather than a support system. To function, it relies on its own intelligence rather than that given by manual input.

The key functions are:

- to collect jobs automatically from other support systems such as CSS;
- to allocate work automatically to the most appropriate technician in a business-optimum manner;
- to automate and deliver the dispatch and closure process on demand;
- to maintain full visibility of the status of both jobs and workforce; and

- to generate, on demand, comprehensive management information statistics (MIS).

The functional outline is shown in Figure 5.

The system is required to support residential and business repair and provision across customer access and core network workforces. The resultant workload is highly variable. Some activities are known in advance and can thus be planned and scheduled, but others are unsolicited (for example, customer fault reports) and have to be dealt with as they arrive.

Having collected the jobs from its various inputs, WMS uses its real-time algorithm (RTA) to match all the required work to be undertaken to all the available technicians within business imperatives. This task is a continuous process, as all the allocations have to be constantly

readjusted as changes to real-time events occur.

The various attributes that have to be matched are:

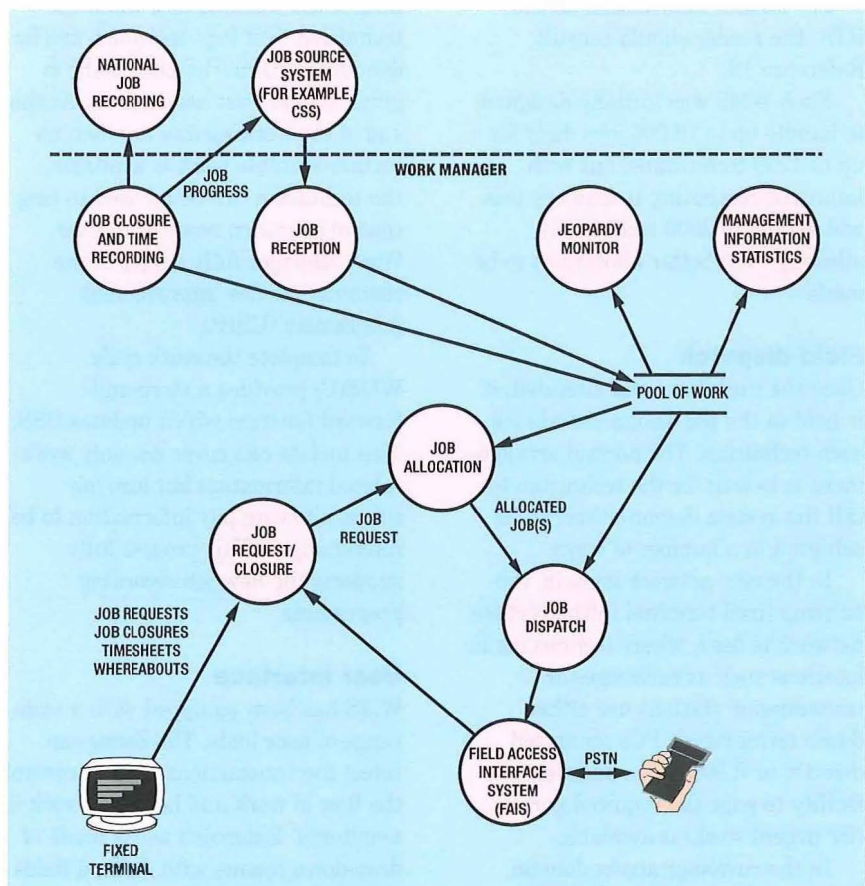
**Job:**

- customer's number,
- location,
- type of equipment,
- type of work,
- skill level required,
- estimated duration,
- target completion time,
- appointment time,
- number of people required.

**People:**

- identification
- start, current and finish location,
- current work-load,
- skill available,
- planned absences,
- overtime availability,
- preferred work,
- number of people available.

Figure 5— Work allocation and dispatch systems



**Business:**

productivity targets,  
 quality-of-service and service-level-agreement targets,  
 allowed overtime,  
 geographic information,  
 travel time,  
 risk of failure,  
 priorities,  
 job bundling,  
 job dependencies.

The RTA balances the many real and virtual costs for each job and produces the 'best' overall cost when assigning the work to each technician. Further, account is taken of those people who are about to finish, and could handle the job in the remaining allotted time. However, if unforeseen events prevent a person completing his/her current job on time, then his/her next assigned job is reallocated to an alternative technician and a new job found for the original person when he/she does complete the task.

For further information on the RTA, the reader should consult Reference 12.

Each WMS was initially designed to handle up to 10 000 jobs daily for up to 1200 technicians, but with improved computing technology this will extend to 2000 technicians, allowing even better allocations to be made.

**Field dispatch**

Once the work has been allocated, it is held in the pre-assigned state for each technician. The normal arrangement is to wait for the technician to call the system (log-on); this can be achieved in a number of ways.

In the core network domain, the existing fixed terminal infrastructure network is used, where technicians in locations such as exchanges or transmission stations use either dumb terminals or PCs connected directly to WMS(N). In addition, a facility to page the required person (for urgent work) is available.

In the customer access domain, where the vast majority of techni-

cians are mobile, visiting customers' premises in vans over a large geographical area, technicians are equipped with an HHT. This device, which is in fact a miniature PC, calls the front-end processor FAIS via the PSTN. A connection is made either directly via a telephone socket, via a telephone acoustic coupler or via a mobile (GSM) connection, using a non-metering pre-allocated telephone number. Current work is now introducing a laptop PC as the technician's tool which is likely to progressively replace the HHT.

The task undertaken by the FAIS is to manage the load, monitor and control the HHT. Further, it manages the access rights, security aspects, line protocols and exchange of data.

Once the link is established, details of current work, closure information and progress is up-loaded. The option to undertake remote line test is now available. At this time, details of the next assigned job are down-loaded. At the beginning of the working day, when the technician first logs on (which can be done from his/her home), he/she is given his/her first assignment. At the end of the working day, or when no further suitable work is available, the technician can be advised to ring control or return home or to base. Work Manager fully supports the customer service improvement programme (CSIP).

To complete the work cycle, WMS(C) provides a store-and-forward function which updates CSS. This update can cover not only work-related information but also job times, allowing pay information to be interchanged. This process fully supports the new job recording programme.

**User interface**

WMS has been equipped with a wide range of user tools. The Zones can select the transactions used to control the flow of work and how that work is monitored. Extensive use is made of drop-down menus with defined fields that have to be filled in.

Data build tools also exist to allow users to create the working environment they require. Parameters on people, jobs, locations, geography constraints, and travel factors can all be created, modified, or deleted.

In the customer access domain, where WMS(C) is utilised, Work Manager controls have been established. Here people and their managers monitor and control the whole operations. For core network domains, the already existing NOUs have accommodated the Work Manager control function.

**Jeopardy monitoring**

The key to meeting customer and business targets is jeopardy management. This is undertaken in real-time to provide a continuous picture of the progress of the total work and the availability of the resource to undertake that work. Early warning of any work likely to fail its agreed completion target is flagged in sufficient time to allow alternative action to be undertaken. Indeed vital work can be reassigned to safeguard customer agreements and contracts.

**Management information**

Management information statistics (MIS) is a module within Work Manager and provides a raft of real-time data and trend information.

The main tool available within the standard software has restricted access to safeguard system performance and generates wide-ranging statistics on anything from productivity performance to quality of service. Statistics can be measured on a zonal, patch or even individual basis. Examples range from the number and types of jobs done each working day to management targets and failure reports.

To allow complete flexibility on a zone-by-zone basis, privileged individual users, under controlled access, can create their own statistical requirement.

This MIS is a powerful management tool, able to provide virtually any type of report, either as a

spreadsheet or directly on screen. It can be run in background or on demand, providing up-to-date real-time information. MIS fully supports the field effectiveness and quality programme.

### **WMS(C) and WMS(N)**

Principally, Work Manager was designed as a single logical system, but to support both core and access domains with their widely different roles and management organisations, two variants have been delivered. In reality it is a single delivery with each variant being separately configured at installation. Currently, approximately 70% of the total software is common to both variants.

Nevertheless, the design is such that as the operational requirement migrates towards a single common workforce, the software can be modified to produce one single common Work Manager system.

### **Fast-track solutions**

Because of the size of the development programme and the complex nature of some operational requirements, the Work Management Programme accommodated a fast-track process. Under this arrangement, individual Zones could specify outline requirements which a small team translated into a PC-based solution within a maximum of 90 days at low cost.

This had the advantage of delivering an urgent requirement quickly, provided it was self-contained and could extract its information via current available interfaces. The disadvantage is that, as they are outside the main system developments, fast-track solutions are vulnerable to any change. This imposes a significant change management and testing overhead. Although several fast-track solutions have been delivered, future use of this arrangement needs to consider the whole-life costs.

Nevertheless, the great benefit of this approach is to bring clarity to user requirements and ascertain

benefits which, in the fullness of time, can be embedded into the main application software.

### **Development**

As stated earlier, the development of NOMS2 was based on the open-systems approach. This strategic decision was vital in that it allowed many parts of NOMS2 to be reused in the development of the WMS. However, the most fundamental benefit to accrue from this decision was the ability to be able to port software from one hardware platform to another.

To support the open approach, the implementation was undertaken using the standard software languages C and C++ running under the UNIX operating system and using the ORACLE relational database. The application software, which is large with over 2 million lines of source code, has been developed on a different platform to that used in the field. Each day a port takes place to help to ensure constant portability. The WMS can be ported to any large UNIX platform. During the course of NOMS2 and WMS development, no less than five major UNIX platforms have been used.

This programme was one of the first large systems to have its production application software ported to a completely different manufacturer's platform. The saving in procurement by having this ability was significant.

Extensive use of software development tools such as Software Through Pictures (StP) greatly enhanced the development in both quantitative and qualitative terms.

### **Deployment**

WMS was designed as a system for national deployment. Its design had to be flexible to evolve with the changing organisation of BT. With Project Sovereign it aligned itself easily with the respective parts of Personal Communication Division (PCD) and Worldwide Networks (WN). The result

was 16 machines for PCD and 10 for WN. While each was basically standalone, they were delivered in clusters of up to four machines, plus one extra machine per cluster as a standby. The technology and cost did not permit the deployment of the full worker/standby arrangement: however, it did provide a level of fallback in the event of machine failure. The programme also provided a complete cluster at Harmondsworth with the dual purpose of being a reference centre and a fallback in the event of a site disaster.

After a successful pilot using the WMS(C) configuration, NBC now requires its own cluster to support its entire BT operations, due in 1996. This clearly demonstrates the flexibility of the original design and capability since a new business domain can be accommodated without the need to build a new system from scratch. Not only is WMS portable but it is also reusable.

At present, eight clusters, in addition to the fallback centre, are deployed.

### **Service and Support**

The WMS has been designed to deliver a service availability of 99.8% over the working day (06.00–22.00) across all systems. For the 1995/96 financial year, the Computer Service and Operations (CSO) organisation independently measured the WMS(C)'s service availability at 99.84%. Unfortunately, for organisational reasons it was not possible to measure the WMS(N) in the same manner but the indications were that this achieved similar performance.

During 1995/96, CSO progressively introduced its open system management framework (OSMF) for its computer systems. Work Manager was modified to fit into this structure. Performance agents are now embedded with both WMS and FAIS to enable a central control at the Ipswich Computer Centre to monitor the state of all systems at all operational sites. Progress on using the on-line backup

Figure 6—Management structure

tool (known as *ADSM*) and the upgrade transfer facility (known as *Open X-FER*) is being made during 1996.

In the past, Work Manager was supported by two different organisations in two different types of computer centre. During 1995, management of both service and support was brought under one single control, namely CSO. The processes involved and their implementation are now acting as a role model as other systems are being integrated into CSO during 1996/97.

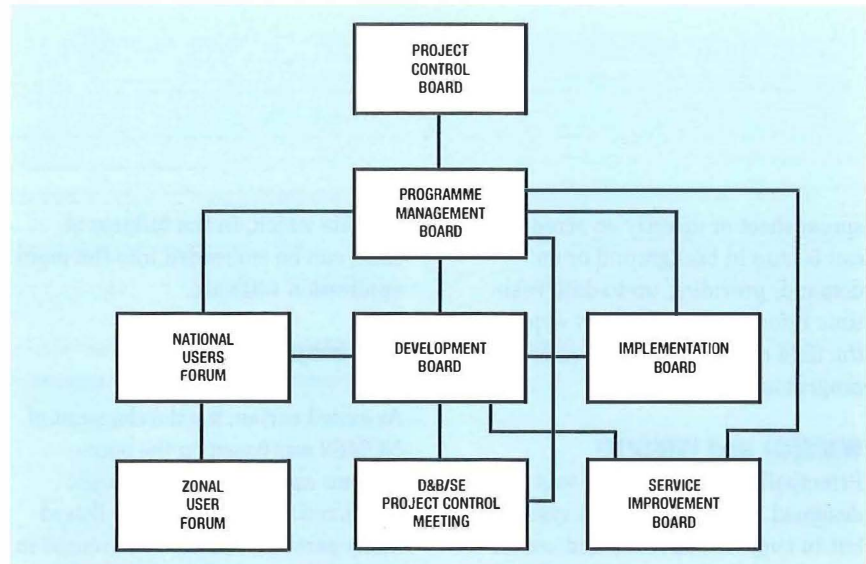
Providing an acceptable level of service is paramount to a successful programme. WMS has made good progress but more can be achieved. Once the systems fully conform to all aspects of the OSMF, the business can expect even higher levels of service and availability. This approach must now be applied to the communications infrastructure, which at present is the weakest aspect of the Work Manager service.

### Programme Management

The Work Management Programme is one of the largest and most complex programmes undertaken by BT. It has brought together many parts of the company each with its own distinctive roles and responsibilities. Over the years, the management organisation of the programme has evolved to its current structure as given in Figure 6.

At the most senior level the project control board (PCB) steers the direction of the programme chaired by the Director of Network and Systems Operation's personal representative. The programme is governed by the programme management board (PMB). At the appropriate management level, both boards are represented by field users from access, core and NBC, BT's development organisations—Design & Build (D&B) and Systems Engineering (SE)—and the customer central project team from Manchester.

At the lower levels, the national and Zonal user forums debate and



scope user requirements, which are then specified by the Work Management Programme development board in conjunction with D&B and with support from SE and CSO. The implementation board manages the delivery of the function to operational systems via pilot and then roll-out. The service improvement board chaired by CSO plays the vital role of monitoring and improving service.

The key aspect of the structure is that all parts of the company are represented and contribute to the end solution, with the end users providing the most vital role.

### Potential

The Work Manager technical vision is to provide the company's operational managers with a technically independent solution that gives them the ability to organise and manage all their engineering workforce as a *single integrated unit*.

The initial architecture of Work Manager has provided a sound foundation on which to build a company-wide solution with sufficient flexibility and potential to evolve towards the above vision. By the year 2000, with the migration to a three-tier architecture, the WMS could evolve from its present singular logical design to a multi-application structure that would support all aspects of work management including the proposed access operations unit and its NetMaster system (Figure 7). Such an architecture, which anticipates pragmatic enhancement of computing technology,

will deliver many aspects of Work Management covering:

- work allocation,
- work scheduling,
- work collection and distribution by standard interfaces,
- network information,
- fault diagnosis and location,
- people information, and
- resource management.

This virtual machine would have sufficient capability to support the largest current operational Zone, negating boundary conditions, and thus enabling the company to move towards a single managed workforce, (sometimes referred to as a *multidomain* workforce), within a very large geographical area, covering many types of work; for example, reactive repair, provision, and capital. The potential for Work Manager as an integral part of work management is enormous. Indeed it provides the foundation to manage any type of workforce.

### Lessons

Any large programme will identify many lessons through its life cycle, and the Work Management Programme is no exception. The following are some of the key lessons covering technical, procedural and operational aspects:

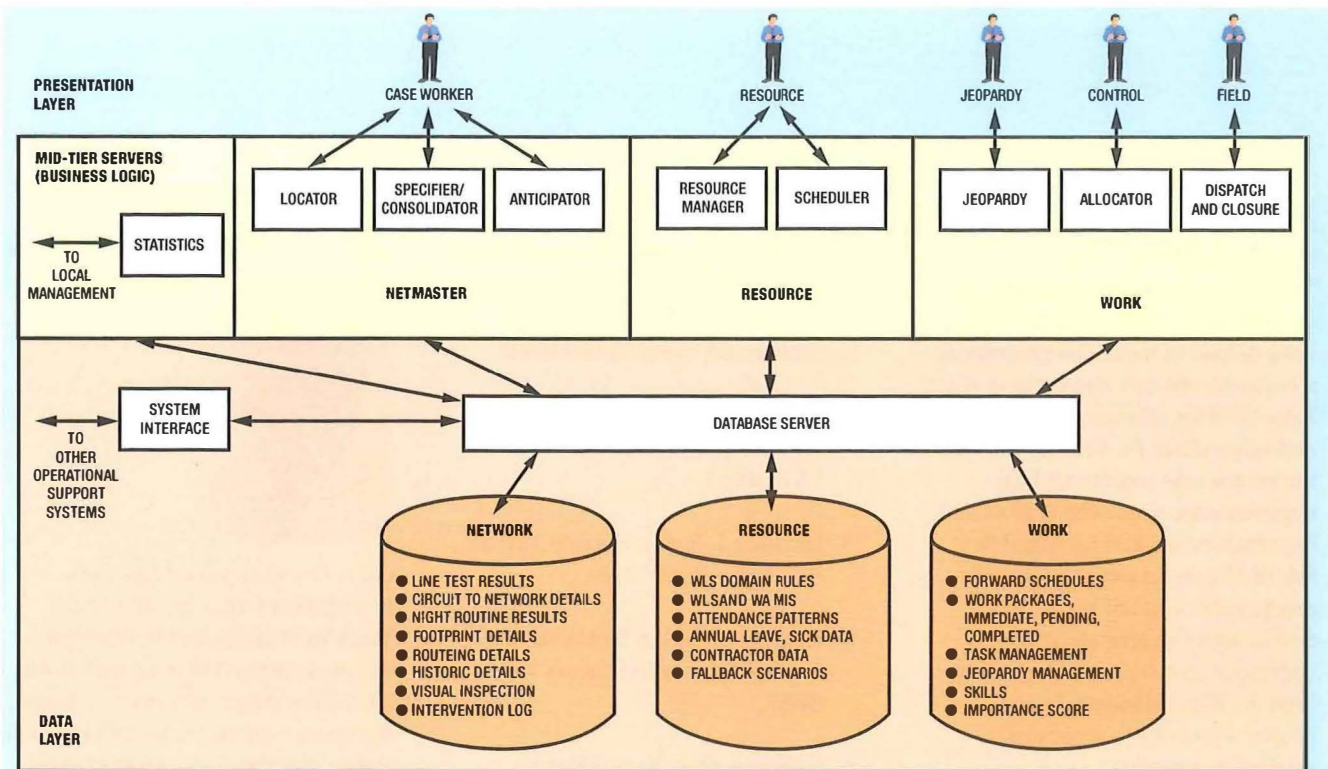


Figure 7—Work Management logical design—potential vision

### Technical

- Assess technical complexity and never underestimate the task. Emulating real-world attributes is difficult. It will not be achieved first time and designers must be prepared to experiment and trial in operational situations. Allow sufficient time to understand the requirements and where possible use fast-track systems to demonstrate the proposed solution.
- Recognise that a national system is a much larger version of a demonstrator, and attention to scale is vital. Consider carefully scalability issues across all hardware and software domains.
- Service operability and availability can easily be underestimated. These attributes must be designed in at the beginning. Supporting standard service arrangements like OSMF is vital.
- Recognise the capability and limitations of the communications infrastructure. This aspect is not in the direct line of control of the programme, which has to influence its evolution rather than direct it.

- Use state-of-the-art and proven technology but ensure that the future path is evolutionary.

### Procedural

- All technical solutions must support company processes. Do not expect technology to provide the process. Understand the business requirement being addressed and explain proposed solutions to clarify the need. A simple demonstrator can help enormously.
- Establish a company champion at a very senior level to support the programme.
- Establish team working across field users, central management and developers. Consult, advise and liaise on all issues at all levels. A sound programme management structure is vital.
- Develop the programme vision, identify the benefits/costs and explain these in both business terms and local needs.
- Expect problems to occur, and be prepared to provide a fast/temporary fix, while an enduring solution is developed.

### Operational

- Take notice of end-user views and needs.
- Never underestimate the culture change. Expect user resistance to the new system as perceptions at local level never initially align with national requirements. Never rush from pilot to national roll-out without first understanding all operational and procedural issues.
- Find a local manager to champion the system which will help smooth the implementation.
- Communicate to all what is happening.

### Conclusion

The Work Management Programme has succeeded in designing, developing, delivering and deploying a key operational support system for BT. The original objectives have been met, indeed WMS is now vital to a significant part of the company's operations. Where WMS has been deployed, both productivity and quality of service targets have been met.

## Acknowledgements

The author gratefully acknowledges the many hundreds of people who have helped to make the programme a company success: the users in every Zone for their tolerance, perseverance and suggestions for change; those in the centre who coordinated the requirements, drove the deployment implementation and managed on behalf of the company the overall programme; and not least his development colleagues whose dedication and technical ingenuity brought Work Manager from a laboratory model to a national production system.

## References

- 1 GARWOOD, G. J. A R&D Programme for an Intelligent Maintenance System. BT memo TA12-01/86.
- 2 HYDES, N. S. The Repair Service Appointment System. BT memo TA12-23/87.
- 3 DENMAN, J. Repair Service Adviser. BT memo TA12-07/86.
- 4 HERBERT, J. M. An Installation Control System. BT memo TA12-24/87.
- 5 GARWOOD, G. J. Repair Service Control Logistics: How to Achieve. BT memo TA12-36/87.
- 6 LAITHWAITE, R.; and THOMAS, A. N. District Control Support System. BT memo TA12-37/87.
- 7 Network Administration for the 1990s. *Br. Telecommun. Eng.*, Oct. 1990 (special issue).
- 8 BROWN, D. W.; BELL, R. C.; and MOUNTFORD, J. A. Strategic Planning of Business Operations and Information Systems. *ibid.*, April 1990, **9**, p. 16.
- 9 GARWOOD, G. J.; and ROBINSON, A. C. Work Management System. *ibid.*, Oct. 1991, **10**, p. 204.
- 10 GORTON, R. K. Customer Service System. *ibid.*, April 1986, **5**, p. 18.
- 11 MORRIS, BILL.; and KELLY, LIAM. Workforce Allocation in the Core Network: NOMS2 and WORK MANAGER. *ibid.*, July 1995, **14**, p. 106.
- 12 LAITHWAITE, R. Work Allocation Challenges & Solutions in a Large-Scale Work Management Environment. *BT Technol. J.*, Jan. 1995, **13**(1), p. 46.

## Biography



**Gerry Garwood**

Gerry Garwood joined the then British Post Office in 1960 as a Youth-in-Training in the Norwich Telephone Area. He then undertook technician duties covering, test desk, linesman, and exchange and repeater station maintenance. After graduation in 1969, he spent four years in computer development within the Postal Mechanisation Department. He moved to Research Department in 1973 and became involved in software systems development for stored-program control exchanges and the associated support. After some research work on the use of expert systems to support maintenance activities, he developed a number of software systems which demonstrated that software technology could be applied to work management. He was responsible for the development of the AI demonstrators, DCSS, NOMS2 and WMS. As the D&B delivery programme manager he was responsible for further development of the WMS and its deployment nationally. Gerry has an Honours Degree in Electrical and Electronic Engineering from Brunel University, holds a Diploma in Management Studies, and was a Chartered Engineer and a Member of the Institution of Electrical Engineers. He retired in September 1996.

*Paul White*

# A Service System for BT's Major Customers

*Customer service excellence is regarded as one of the key differentiators of BT's product and service portfolio. This article discusses, from a technical perspective, the major customer service system (MCSS), and how it is helping to achieve excellent customer service for BT's global and major customers.*

## Introduction

This article focuses on the technical challenges associated with the development and deployment of the major customer service system (MCSS)—a system designed to support BT's top 4500 global and major customers, who between them account for nearly one third of BT's total revenue.

The article discusses some of the unique features of the MCSS architecture and illustrates how a diverse range of technologies from 4GL tools, three-tier client/server and intranet-based network computing have been combined to provide a key competitive edge to BT's customer service capability.

The following three key requirements have dictated the architecture of MCSS.

- **Electronic customer access (ServiceView)** MCSS users include both traditional service agents within the customer service centres responding to telephone enquiries, but also includes direct electronic customer access (with the attendant security issues). This allows customers to raise and track orders and problem reports without the need to directly involve BT agents.
- **Product diversity** The range of products and services supported by MCSS ranges from the public switched telephone network (PSTN) and integrated services digital network (ISDN), through private circuits, asynchronous transfer mode (ATM) and 0800, and encompasses global products

such as Concert frame and packet services. Central to the external market positioning of ServiceView is the ability to support all of these services through an integrated management system, rather than building stand-alone systems to support specific products.

- **Performance and scalability** MCSS is deployed in a large number of separate and distinct environments, from small systems supporting BT's joint ventures up to the main UK customer deployment with the requirement to support 1700 users.

## The Business Context for MCSS

MCSS is part of the major customer service excellence programme, the scope of which encompasses building uplifts, business process re-engineering and much more besides. However, from a systems perspective the business benefits of MCSS are delivered through three key capabilities.

Firstly, BT and its major customers can share information electronically as the basis of forging an outgoing customer relationship. This is realised through an integrated system supporting a wide range of information and transaction types (from order entry to bill enquiries), and encompassing the vast majority of BT's products and services portfolio. This electronic interaction and the associated software applications significantly improve the accuracy and timeliness of information flows between the customer and BT.

Secondly, customers in this market segment typically have over 5000 distinct BT products and services. By building a customer-centric information repository, BT has moved away significantly from the traditional support systems strategy where customer data has been distributed through a range of product-based systems.

This repository can now serve as the customer's main inventory system and is fundamental to supporting the major MCSS applications, which customers can access through ServiceView. It also empowers the BT account support teams with the key information they need to start to become more proactive in anticipating and meeting their customers' needs.

Thirdly, there are operational cost savings, since a side effect of providing electronic customer access is that information can be automatically interchanged with the customer. However, if MCSS was just a stand-alone system this benefit would still need to be entered into other BT systems; therefore, a series of automated interfaces has been built that allow for this interchange of data. These interfaces will also reduce the need for training on a wide variety of complex legacy systems, with users instead needing to learn only the single graphic user interface for the MCSS.

**Systems Architecture**

Figure 1 illustrates the logical architecture of the MCSS, which is based on a 'traditional' three-tier model; that is, Windows or Web clients communicating with an open systems mid-tier server, which in turn interacts with a variety of BT and global mainframe servers.

Customers interact with the global and major service centres via a variety of routes. Early customers used the Windows-based MCSS manager client; more recently, customers are starting to use a World

Wide Web (WWW) intranet client (Bali). This latter direction has been inspired by the need to reduce the operational cost of delivering electronic access to an ever-increasing customer base. Of course, many customers still contact the service centres via telephone, but in addition to voice communication ServiceView customers are provided with a video link, to enable face-to-face dialogue with the service-centre agents. Within BT the service agents currently use the MCSS manager client.

The MCSS server is based on an Oracle database, which in high-end deployments is configured as a parallel database running on clusters of multi-processor SPARC Enterprise servers. Currently MCSS has its own in-built operational customer and product databases, but the intention is to migrate towards the emerging corporate systems as they become available.

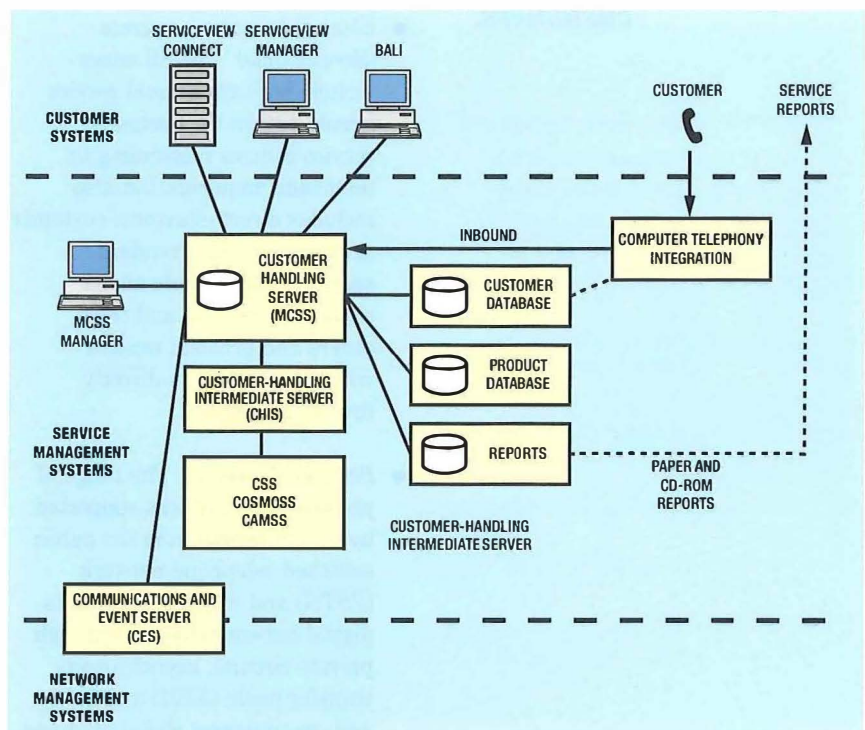
MCSS supports integration with a large number of other systems, including the customer service system<sup>1</sup> (CSS), the customer-oriented system for the management of special

services<sup>2</sup> (COSMOSS), the computer-assisted maintenance for special services (CAMSS), OMNI and CTMS via its external-systems interface architecture. A combination of server-based robotics and, increasingly, customer-handling intermediate server (CHIS) services<sup>3</sup> are used to facilitate this communication. The same architecture also supports the ability of separate MCSS systems to interwork; this multiserver working is used to provide scalability for large installations, plus the ability for different MCSS-based organisations to interwork, enabling capabilities such as 'follow the sun' working.

MCSS also handles feeds of network events from a variety of BT's network services, thus enabling proactive reporting of service-affecting faults to customers and service agents. A standard component (communications and event server (CES)), based on HP's OpenView platform, is used to provide this capability.

Regardless of the client 'type' or the back-end systems supported, a standard application programming

Figure 1—MCSS logical architecture





interface (API) provides access to all of the core server-based functionality of the MCSS system. This API is now also being used by third parties (for example, BT's joint ventures and remote development teams) to extend and enhance the core application and interface set.

## Functional Architecture

MCSS applications have been developed using a highly modular and extensible approach, with a variety of distinct applications supported by a common infrastructure and data model (see Figure 2). The descriptions below provide a brief overview of some of the key components together with an insight into the approach taken for their construction.

### Infrastructure and data

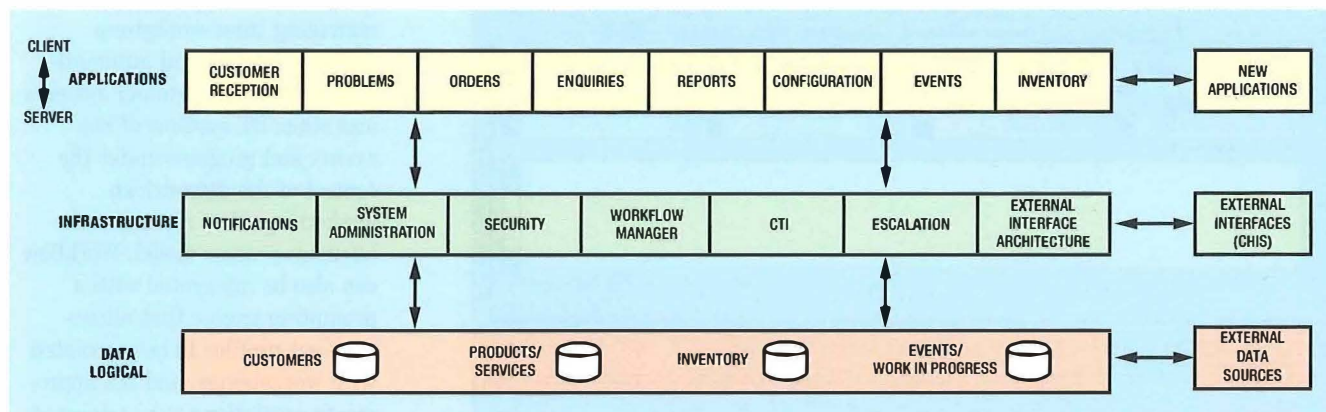
There are five key components of the MCSS infrastructure and data modules:

- **Security** Because of the need to provide external customer access, and to support the fact that, operationally, MCSS users are typically organised into distinct customer teams, security and data partitioning considerations have been paramount in the design. MCSS supports strong end-user authentication mechanism, full data encryption across the client-server links and low-level data

partitioning based on an MCSS-developed (BT patented) algorithm. A primary characteristic of the design has been to provide the security features transparently to the application developers, who therefore do not have to worry about the mechanics of the security regime. MCSS applications have been certified to European community standard E4.

- **Workflow** MCSS was always designed to support a variety of operational deployments and a wide range of products. An early design decision therefore was to build all the applications around a finite state model, where the available states and the legal state transitions are defined by data-driven workstrings. This approach has enabled the tailoring of MCSS around existing business processes; for example, it is possible to define in data the order-taking process from initial data capture, through order validation, authorisation, and interaction with the appropriate order fulfilment system.
- **External interfaces** External systems are treated just like any other client of the MCSS server; however, a series of API calls is available to facilitate integration with distributed computing environment services such as CHIS.
- **Computer/telephony integration (CTI)** The MCSS programme currently uses CTI to enable its customer-reception application. This facility causes an automatic screen pop of relevant customer details based on the calling and called telephone numbers. The application is based on the Genesys T-Server CTI product, and is currently on trial as part of an extension of the scope of MCSS into the extended business customer base.
- **Data** One of the fundamental design decisions in MCSS was to create a customer database at the mid-tier that replicated much of the data on the back-end systems. The key rationale for this approach was to create a consolidated customer-centric view of information that could be provided to both customers and to the customer teams. This information is a key part of the sales proposition of ServiceView, and is the most frequently cited reason for attracting customers to the

Figure 2 – MCSS functional architecture



ServiceView portfolio. It also begins to empower the customer teams with the information they need to move from a reactive to a proactive approach with their customers.

The customer data includes customer names, hierarchies, contacts and addresses, sites and locations, full product and service inventory, and details of work in progress details; for example, outstanding orders and faults.

### Applications

The MCSS applications can be divided into both client and server functions, the server functions typically being realised as reusable business objects implemented in Oracle's PL-SQL programming language.

Although the client applications have traditionally been Oracle forms-based applications running on Windows PCs, many of these applications have now been rewritten as WWW applications in a combination of HTML, CGI and Java. There is an ongoing development programme to accelerate and extend the availability of the MCSS applications on the WWW.

- *Customer reception* Customer reception is one of several entry point applications; it supports initial customer profiling and data capture, and then enables rapid navigation to the most appropriate application, based on the initial customer dialogue.

- *Inventory* The inventory application provides the ability to browse and configure a customer's portfolio of products and services. Numerous tools are provided to report interactively on the inventory, using techniques such as query by example and user-definable filters. Typically, users might want to ask questions of the inventory database such as: 'How much bandwidth is available between sites A and B?', where that bandwidth may be realised by a variety of different services (for example, flexible bandwidth service, KiloStream and MegaStream).

All service inventory object classes contain a standard header which identifies the object to the MCSS system. The details of a particular service are normally

contained in an inventory body, supplied as a distinct executable (analogous to a Windows .dll), or alternatively (and more typically) by one of several meta service bodies. In this latter case, all of the attributes, screen layout and validation rules are defined in a server-side product database, with no need to modify the client software.

- *Order processing* The order-processing application is responsible for capturing order details, updating the MCSS inventory and then, via the workflow engine, forwarding the order requests to the appropriate external system. Figure 3 shows the order entry screen.

The same concepts of order header and order body described above are reused in this application. In addition, an order installation routine can be supplied which determines how fulfilled orders are added to the inventory database.

- *Fault processing* The fault-processing functions also operate in a similar manner. It should be noted that faults can also be automatically raised via CES or an external system and can then be automatically forwarded to the customer's ServiceView terminal.

The workflow engine is sufficiently powerful to control all the key state transitions within the fault and order-processing applications, and is responsible for managing inter-workgroup communications and automatically informing customer systems and other BT systems of key events and progress under the control of the data-driven workstrings that represent the business process model. Workflow can also be integrated with a prompting engine that allows time-out profiles to be associated with workstrings, and the appropriate escalations to be triggered;

Figure 3 – Order entry screen

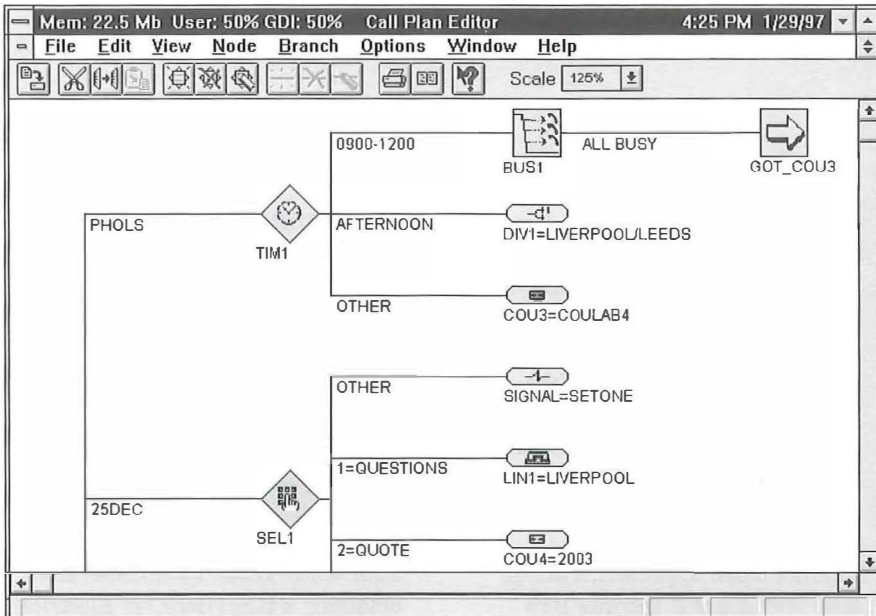


Figure 4—Call plan editor screen

for example, if an application remains in any given state too long.

In both the fault and order-processing applications it is possible to track progress, within MCSS and on the external systems as key events are delivered across the automated interfaces.

- **Configuration** Configuration tools are a classic example of the

occasional need to supply product-specific applications to supplement the more general-purpose ones available from the core MCSS product. One of the more ambitious of these applications has been developed for BT's telemarketing services (0800, 0345, etc).

Figure 4 illustrates a call-plan editor application that allows BT's telemarketing customers to manipulate their call routing

plans with an almost immediate response time, enabling campaign routing to be tuned graphically against a wide range of possible decision points dependent on numerous factors from date and time of day to which region of the country is generating the call.

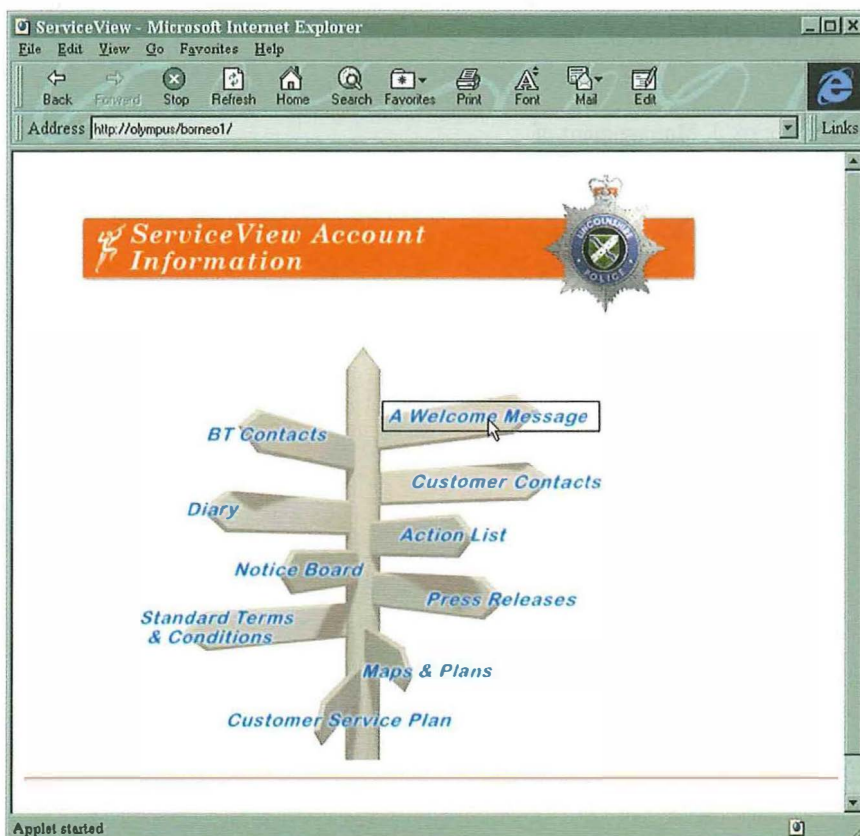
- **Shared workspace** Shared workspace is a feature supported only by the WWW clients; it basically provides a jointly owned and maintained information pool shared between BT and the customer. Account team contacts, key diary events, customer service plan details are held on-line; content management tools are supplied to enable end users to manage the information directly and therefore to keep it up to date. Figure 5 shows the shared workspace signpost.

**Key construction challenges**

At the outset of the development stages for MCSS in 1994, it was decided to use a 4GL toolkit rather than an in-house developed platform, using the Oracle Designer and Developer 2000 products with the Oracle database. This has enabled a significant amount of the standard functionality to be exploited, such as encryption, multiple-platform support, scalability and password management. Further evolution will make use of the Oracle Web Server, Sedona and Oracle 8 products, as well as the common object request broker architecture (CORBA).

Of paramount importance in the deployment of a system is performance, and the key area is the finite network pipe between the client and the server where indiscriminate and unnecessary data transfer costs user response time. This desire to minimise network bandwidth between client and server, especially for WWW clients, led to the decision to encapsulate functionality behind a server-side API, which produced all the results prior to returning them to the client.

Figure 5—Account information signpost screen



The decision to maintain a customer-centric database has been described earlier in this article. This feature is central to the MCSS functionality; however, replicating and managing up to 30 Gbytes of data sourced from at least six external system sources leads potentially to a variety of data quality and integrity issues.

The MCSS solution to managing these issues has been the creation of a customer data management centre (CDMC), where the approach has been the development of a data-cleanse tool to enable the correction of data at the source systems, using techniques such as PAF (Post Office address file) matching and de-duplication, prior to loading on to the MCSS. This also has the advantage of improving the data for the native external system users. While the MCSS order and configuration applications should then maintain synchronisation with the external systems via automated interfaces, the CDMC also supplies a routing function to periodically check data integrity with MCSS.

Deploying and managing large client-server systems is always going to be a massive undertaking. Now there are the necessary tools for software distribution and systems management in the client-server environment; for example, the NetView/6000-based enterprise systems management facility, into which MCSS has been integrated via the generation of SNMP traps.

Finally, another fairly unique requirement of MCSS is multi-language deployments. The strategy to date for multiple-language support is restricted to help file changes or a specific language port of the software. Here, again, the WWW variant of MCSS may provide a true solution to multi-language support without the need for client-side software changes.

## Future

Intranet-based applications are the key to the rapid and widespread deployment of MCSS to the high-end business customer market place, and the MCSS programme is already delivering this type of application to external customers and starting to reap the benefits in operational cost savings and speed to market.

Internet-based applications are on their way and will enhance BT's ability to extend the functions of MCSS to a much broader customer base than would otherwise have been cost effective.

## Acknowledgements

The work of the MCSS design team over the last two years, in particular Steve Comish, Martin Roberts and Paul Buckley, is gratefully acknowledged.

## References

- 1 HARRISON, P. F. Customer Service System—Past, Present and Future. *BT Technol. J.*, Jan. 1997, 15(1).
- 2 CHANDLER, J. Management of Special Services—Designing for a Changing World. *BT Technol. J.*, Jan. 1997, 15(1).
- 3 MATHIESON, G. Customer Handling Intermediate Server—An Architecture-Led Project. *BT Technol. J.*, Jan. 1997, 15.

## Biography



**Paul White**  
BT Networks and  
Systems

Paul White joined BT in 1984, after obtaining a B.Sc.(Hons.) in Electrical and Electronic Engineering from Loughborough University. In the early part of his career he worked on the development of operations and maintenance systems for BT's core network (for example, OMC2). More recently he has worked on the design of customer management solutions which have culminated in the development of MCSS. He is currently the chief designer of BT's portfolio of service, marketing and sales systems.

Chris Gibbings, Dave Newson, Ian Henning, Jeremy Barnes,  
Dave Pratt and Peter Cochrane

# Superhighway... Supergridlock?

*The global broadband network of the future will increase still further our reliance on telecommunications, but the way it will react to heavy loading is still a matter of conjecture rather than solid prediction. Will our electronic money be stranded in a network traffic jam?*

## Virtual Rubbernecker

In the late-1950s, the hard shoulder of the newly opened M1 motorway was punctuated by overheated cars, unaccustomed to dizzying speeds of 70 mph and more. Today, the M25 features overheated cars, but they are more likely to have been sitting in one of the traffic jams that have become the trademark of London's orbital motorway. Often, there is no obvious reason why cars are stationary. The crashed cars, the victims and the police might have gone, but they still remain as a memory in the mysterious dynamics of the overloaded motorway. The opposite carriageway is not immune, as 'rubberneckers' slow down and precipitate tailbacks.

Is a similar shift from predictability to chaos about to affect the world's data networks? Adding more bandwidth to a network is analogous to adding more lanes to the motorway, but with the uncertainty that predicting future traffic levels on the Internet is vastly more difficult than forecasting the number of cars on the roads: nobody can drive more than one car, but Internet users can

simultaneously run several multimedia applications, with no limit on the maximum amount of network usage. Furthermore, different applications have different requirements—the steady load of an Internet telephone or videoconference, or the exceptionally bursty nature of World Wide Web (WWW) traffic.

Will the archetypal butterfly wing beat give rise to huge peaks in demand? Will the software 'rubberneck' and amplify the problem still further? We do not know the answers to these questions, and it appears that current modelling is not in any position to supply them. However, it is vitally important to find out soon. Brownouts cause disruption, even with our current level of dependence on networks, and this dependence will be much greater in 10 years time. Bank holiday jams are a tradition on real highways. If chaos on the information highway, from whatever trivial cause, gridlocks our electronic money then perhaps the banks might just as well be on holiday.

BT's *futures test bed* is one attempt to answer these and many other questions, mixing together the

Figure 1 – M25 traffic jam



TELEGRAPH COLOUR LIBRARY

## *Predicting the future behaviour of wide area broadband networks is especially difficult because they represent the coming together of many different networking and operational cultures.*

latest hardware, software and communications protocols, a wide range of external inputs, 700 researchers and the mission to 'do and discover'. In the absence of modelling techniques for future networks, we have to create the network and the broadband traffic. A diverse range of traffic is encouraged by remote working, multimedia and group working, broadband connections to universities and schools, and the use of satellite, wavelength-division multiplexing and conventional fibre links to transport the data. The challenge is: can the technology be bent to our needs, or do we have to bend to the limitations of the technology?

### **Cultural Exchange**

Predicting the future behaviour of wide area broadband networks is especially difficult because they represent the coming together of many different networking and operational cultures. Telephone companies have traditionally sold point-to-point connections, either on the basis of connection time (telephone calls) or as a permanent private line. Whether customers use this data capacity fully is not usually of any great concern to the provider. In contrast, local area network (LAN) and Internet users have access to a shared medium where the cost is normally independent of, or decoupled from, usage. This is a sensible solution, given the bursty

nature of data communications, as all users benefit by sharing costs and pooling capacity which is used when needed.

Real-time applications provide a challenge for 'best-effort' networks such as the Internet and LANs. Without any means of reserving bandwidth the only way of guaranteeing that users continue to get an unbroken video stream is to make the data capacity so high that congestion never occurs. Although this is feasible on a LAN, the history of the Internet suggests that to keep capacity this far ahead of user demand is impossible. Applications always expand to use the available capacity, and video users, or anyone else with a real-time application, lose service to congestion.

Resource reservation is thus becoming an important issue for the Internet world, and a new reservation protocol (RSVP) is being developed. Meanwhile, telephone companies can see the advantages of 'best-effort' data transfer in offering higher bandwidths at lower cost. The standard for future broadband telephony networks, asynchronous transfer mode (ATM), aims to combine the best of both worlds, with data carried in streams of small cells in multiple virtual connections over the same data link. The cells can be classified according to whether they can be discarded or not if congestion occurs.

Table 1 summarises, somewhat unfairly, the differences between the principal players. In reality, many

stereotypes, such as the non-commercial academic nature of the Internet, are no longer valid. However, there is a fundamental question that has to be asked: 'Is it better to put resources into getting the last few per cent of utilisation out of a link, or to buy a bigger link? This control strategy differs widely. Telephone companies have been keen on high utilisation (>80%), with a significant fraction of their expenditure going towards the people and software involved in charging, billing and network management. Although network control centres provide the most striking image of the efforts involved in maintaining telephone network quality of service, the fundamental mechanisms to avoid congestion are pricing by distance, by time of day and call duration. Meanwhile, Internet providers struggle to increase capacity to deal with more and more voracious applications, and usage is visibly limited by congestion.

Generally speaking, the various parties involved agree that the future will involve a mixture of reserved bandwidth data streams, 'best effort' data and management messages. The wide area broadband network is therefore moving into uncharted territory, and chief among our concerns should be the stability of this combination under heavy and dynamic loading.

### **Data Tailbacks Likely**

Using the stereotypes described, it might be thought that telephone network engineers are cautious, conservative folk, and if they can get into trouble with a particular innovation then everyone should worry. They have indeed run into highly publicised problems with switches. Network brownouts have occurred when inter-switch signalling was modified. Switches sent out messages saying how overloaded they were. These messages then overloaded the other switches in turn, which then sent out more messages saying how overloaded they were...

**Table 1** Stereotypical network cultures

Network	Control Strategy	Standards Process	Future Evolution
Telephone Networks	Charge for connection time	Glacial timescales. 'Kitchen sink' mentality	Need flexible bandwidth
Internet	Live with congestion	Pragmatic approach. 'Anoraks'	Need ability to reserve bandwidth
Local Area Networks	Throw money at problems	Many de facto standards	Need to be optimised for speed and manageability

Brownouts are common across the Internet, but they have not received the same publicity as the telephone network's mishaps, probably because the Internet is not yet as central to professional and social activities, and many uses are less time sensitive (for example, mail). The transition from 'best effort' to 'no chance' is less clearly defined than in the telephony case, as a user might find it harder to notice the difference between very poor service and no service. As real-time Internet applications become more significant, these brownouts will become more newsworthy.

Clearly network components should be designed to cope with disruption in a way which does not result in storms of management messages. However, this represents a difficult design target if this stability is required when they are running in already heavily congested conditions.

Like our motorways, broadband data networks will not be operating in isolation from the rest of the world. External events can precipitate a large increase in traffic. Telephone networks are used to responding to events such as phone-ins, which generate a lot of traffic in a short time. Control centres can also help the network cope with disasters, where a lot of people want to call a particular number at the same time. In the same way, broadband networks will have to cope with surges in demand precipitated by external factors. However, management may be a great deal more complex, and the presence of certain sectors of activity on the network will introduce complications, as outlined later.

Internet traffic is highly bursty, with 'fractal' behaviour. If a WWW user clicks on a 'hot spot', it is highly likely that a further 'click' will come within several seconds, adding to the traffic load. If the network usage of many users is aggregated, the network load becomes a chaotic distribution, rather than averaging out. Adding an external factor—such as a breaking piece of financial

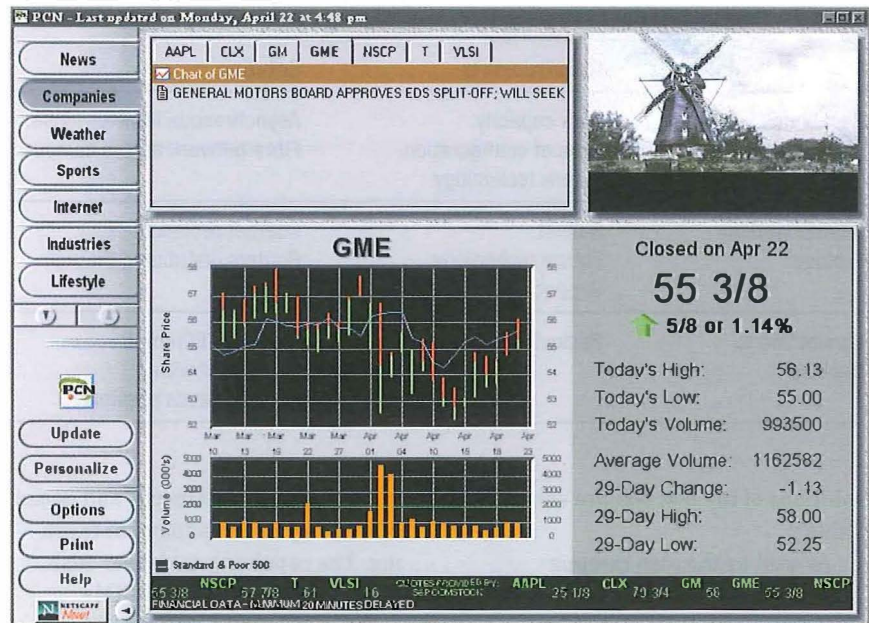


Figure 2—Financial news using the PointCast service

news—will spur many users, or their software agents, into activity. The flurry of mouse clicks in different parts of the globe may well be dealt with quite happily by the network. However, if some parts of the network are overloaded, many people will find their news supply cut off. This is analogous to the communications problems during the Wall Street Crash of 1929. Enormous trading volumes meant

the effects on shares, currency values and, eventually, ordinary peoples' lives should this network fail. This underlines the importance of studying the stability of the network against the various threats—interactions with the outside world, broadcast storms, and, where congestion is high, virtually any small perturbations. In the absence of a successful model, we must construct a broadband test network and fill it

## *The futures test bed is BT's window onto the new opportunities and challenges*

that information on the ticker tape lagged many hours behind the real position. Investors were left to worry and panic without up-to-date information<sup>1</sup>.

Figure 2 shows one of the applications that are available on the Internet today<sup>2</sup>. Many other screenshots could also be included at this point, showing the pioneering applications in electronic commerce over the Web. It is not a great extrapolation from the current position to imagine a significant fraction of economic activity taking place on global broadband networks, and it is natural to be worried about

with data. The futures test bed is BT's window onto the new opportunities and challenges described above.

### **The Futures Test Bed**

The futures test bed<sup>3-5</sup> is not an isolated high-speed network with artificially generated traffic. It provides the network connections for over 600 researchers at BT Laboratories. Given the task of filling large data pipes, a collection of imaginative people working on advanced applications with leading-edge hardware are well placed to fulfil this mission. The general specification and design

**Table 2 Design decisions**

Network Components	Requirements	Solution
Backbone	High capacity. Ease of configuration. Future technology	Asynchronous transfer mode. Fibre network across campus
Internetworking Protocol	Robust. Future technology. Support high data rates	Internet protocol. Routers distributed through network
Connection to Desktop	Backwards compatibility	Switched Ethernet. Structured wiring. Higher speeds possible

solutions of the test bed are shown in Table 2.

As well as the core campus network there are integrated services digital network (ISDN), public switched telephone network (PSTN) and global system for mobile communications (GSM) connections for teleworkers and people away from base. Specifically, Peter Cochrane's experiences with hotel telephone connections have been described previously in the *Journal*<sup>6</sup>.

Finally, there are a range of test bed connections to the outside world, ranging from data lines to local schools up to SuperJanet, the switched multi-megabit data network that links universities and other research centres.

### Phase 1

The first phase of deployment was completed in April 1995 and involved the establishment of a fully interconnected router-based network (Figure 3). Two hundred users were given their own dedicated router port, providing 10 Mbit/s of uncontested bandwidth and supporting current applications (e-mail, file sharing, group working WWW), and protocols (IP, IPX, Appletalk). This combination of good access, high performance routers, and fast optical links for interconnection demonstrated the capabilities which could be delivered by an advanced Internet. Currently, the Internet suffers from poor performance both in the backbone and in the local area. The test bed

removes these bottlenecks and opens up the way for broadband networking. The capability to deliver real-time services such as MPEG1-encoded video and associated audio was demonstrated at an early stage, and it has since formed a platform for the investigation of new emerging protocols such as real-time protocol (RTP), resource reservation protocol

(RSVP), and multicasting (using protocol independent multicast (PIM) and DVMRP). Indeed, the implementation of PIM using IP over ATM on the futures test bed formed the basis of the City Media Network described later.

### Phase 2

The second phase of the network was completed at the end of 1995. The network was expanded to cover the whole research department, and switched Ethernet was introduced as the 'access' technology. As in the first phase, ATM was used as the bearer. Each of 700 users was connected via switched Ethernet, but in a much more cost-effective manner (Figure 4). The network supports LAN emulation (LANE), which is a new capability developed by the ATM community, providing a LAN-like environment over ATM for seamless support of legacy applications. It is

Figure 3—Phase 1: The testbed as a fully interconnected router network

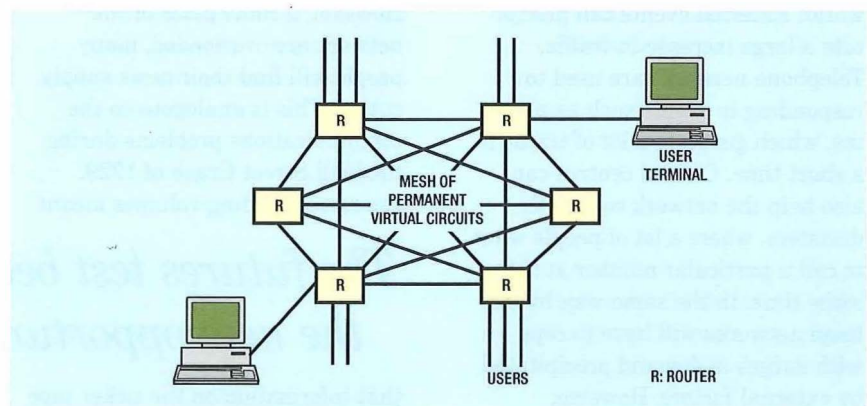
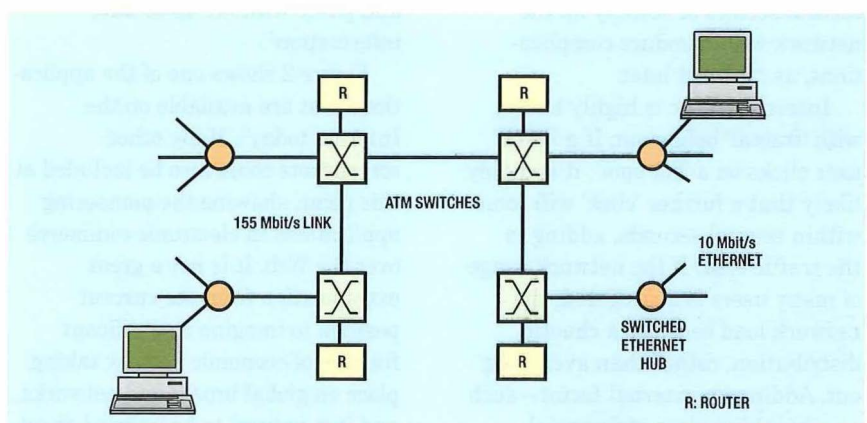


Figure 4—Interconnection of testbed elements





often seen as a stepping stone along the way from the 'shared medium, shared resources' world of the traditional LAN towards the 'dedicated connection, resource on request' world of ATM, which is widely regarded as too expensive for wide-scale deployment at present.

155 Mbit/s ATM desktop connections were introduced in this phase to form small communities of 'power users' who typically use high-end workstations. Shared servers were connected with both ATM and 100 Mbit/s Ethernet; their performance is being assessed and compared. Applications running include shared computer-aided design (CAD) environments and virtual worlds.

### Phase 3

The final phase is now being realised, introducing wide area switched virtual circuit (SVC) switching across ATM using the newly installed second-generation ATM switches. Supporting the PNNI and ABR protocols (recently agreed by the ATM Forum), these switches permit advanced and efficient connectivity and traffic management. The number of ATM connections to the desktop is continuing to increase.

For most users, electronic working has meant the use of non-real-time applications (mail, file sharing, WWW etc.), and while this has improved the way we work, there are still further gains to be had. On-line video and audio services such as training, television briefings, departmental seminars, lectures and the adoption of desktop conferencing to partly replace real meetings are currently the subject of trials and experiments.

Education is seen as a prime application. People increasingly need to update their skills, and on-line training facilities provide a very effective delivery mechanism.

Advanced wavelength-division multiplexing transport technology is being used to extend the test bed off site. The flexibility inherent in WDM allows alternative routing and

easier management. A further advantage is that capacity can be upgraded without putting in new fibre.

## Applications

### City Media Network

The London City Media Network Service (CMNS), to be launched commercially in early 1997, results directly from the experience gained during operation of the futures test bed network. CMNS is a multimedia network delivering 2 Mbit/s video over Internet protocol (IP), with core transmission over ATM.

CMNS has been engineered for the media and financial sectors and comprises two distinct services. The first, IMM, offers download of interactive multimedia material from a central server capable of dealing with over 50 concurrent sessions. It should be stressed that CMNS is carried over IP, and is thus tackling multimedia from a 'computer' viewpoint. BT's video-on-demand (VoD) trials in Kesgrave, Ipswich and Colchester approach the problem from the 'TV' perspective. Both approaches have their own relative advantages, but IP may be better for the corporate market place where computer awareness is higher, and because it allows the PC to be used for a variety of services, such as word processing, e-mail, WWW-browsing, videoconferencing and viewing of multimedia material.

The second CMNS service offers multicast real-time video downstreamed from a server, such as from a specialist financial information supplier. The advantage of multicast is that it greatly reduces the barriers-to-entry of sending material to a number of different users. It matches both the low initial cost of a point-to-point solution, with the flexibility of a broadcast; that is, many viewers can watch the same stream.

Multicasting can be used for a variety of other applications; for example, stock market or sales feeds,

videoconferencing, reducing the network load when updating 'mirror' WWW sites, or simultaneously upgrading the software of PCs on a network.

### Facilityline 15

Digital video post-production is increasingly used in the media industry for everything from TV commercials to Hollywood blockbusters. Within the UK, this market segment is concentrated around central London.

BT's Facilityline 15 allows post-production houses to pipe D1 uncompressed (270 Mbit/s) data. ATM networks are attractive for such a market because of the high bandwidths they can offer and the ability to multiplex a range of services over a common infrastructure. In the future, this service will include a fast link to Hollywood, allowing data to be piped worldwide for around-the-clock post-production.

## The Future—a Holy War ?

The rapid development of bandwidth-hungry applications, the continuing increase in desktop PC power and the rise of the Internet show that, however much we upgrade local and wide area links, the problem of congestion will not go away.

The ability of future networks and systems to deal with real-time applications and congestion will depend upon the protocols they run. Having started as a relatively simple robust connectionless protocol offering 'best effort' quality of service, IP is evolving towards managing real-time streams of data with guaranteed bandwidth. This is the role for which ATM was designed, and therefore one must ask which approach is better suited to carry future broadband services.

Alternatively, the best of both worlds may be possible: IP as people navigate the Net, ATM when they click on a real-time application. This could be seen as an evolution away from the traditional flawed 'stove

pipe' approach of intelligent network servers that perform both service functions and network functions. These could be replaced by platforms built using a more layered structure, with servers dedicated to running services and servers dedicated to providing connectivity. These connectivity servers will be able to take a connection request from the service, specifying the required bandwidth, quality of service and cost, and choose the most appropriate network technology to suit those needs. In this way, ATM and IP can be complementary.

Brownout resistance and manageability are key factors in deciding future network architectures. However, the intrinsic merits of IP and ATM alone will not be decisive; the choice will also be influenced by the evolution path from current networks and services. Clearly these important decisions should be based upon careful analysis and working experience, rather than the entrenched attitudes of the participants in a holy war. The primary aim of the futures test bed programme is to provide an earlier and more appropriate reaction to global gridlock than slamming on the brakes at the last minute.

## References

- 1 GALBRAITH, JOHN KENNETH. *The Great Crash*, 1929. André Deutsch, London, 1973.
- 2 <http://www.pointcast.com/>
- 3 BARNES, J. W. R.; CHALMERS, J.; COCHRANE, P.; GINSBURG, D.; HENNING, I. D.; NEWSON, D. J.; and PRATT, D. J. An ATM Network Futures Test Bed. *BT Technol. J.*, July 1995, 13(3).
- 4 BARNES, JEREMY; GINSBURG, DAVE; NEWSON, DAVE; and PRATT, DAVE. IP Multicast of Real-Time MPEG over ATM. Proc. 7th Joint European Networking Conference (JENC7), p. 163-1, 1996.
- 5 BARNES, JEREMY; HENNING, IAN; NEWSON, DAVE; and PRATT, DAVE. BT Laboratories' Futures Testbed. Proc. IEE Colloquium 'Towards Gigabit Networking', Manchester, UK, May 1996 (Institution of Electrical Engineers, London).
- 6 COCHRANE, PETER. The Great Socket Hunt—From 12 Days to 12 Hours. *Br. Telecommun. Eng.*, Jan. 1996, 14, p. 362.

## Biographies



**Chris Gibbings**  
BT Networks and  
Systems

Chris Gibbings gained a Ph.D in Semiconductor Physics from the University of Cambridge in 1986. He joined BT in 1985 to study the growth of silicon-germanium strained layers using molecular beam epitaxy. In 1992 he moved to the security research area and in 1996 joined the futures test bed team. He has recently been working in the areas of IP multicast and PNNI.



**Dave Newson**  
BT Networks and  
Systems

Dave Newson graduated from Cambridge University with B.A. and Ph.D. degrees in Physics. He spent 2 years working in the research and development center of Toshiba in Kawasaki, Japan, researching optical phenomena in III-V semiconductors. After joining BT, he worked on the development of optoelectronic and millimetre wave electronic components. Part of this work was directed towards systems requirements for ATM. He then joined

the team to design, install, commission and run the futures test bed ATM-backed network at BT Laboratories. More recently he has been heavily involved with the network design and testing of the City Media Network Service, a BT service which delivers digital video over an IP and ATM network. He has authored or co-authored over 50 papers or articles, and three patents.



**Ian Henning**  
BT Networks and  
Systems

Ian Henning gained a B.Sc. (Hons.) in Applied Physics from the University of Cardiff. He remained there for a further four years, three of which were spent gaining a Ph.D. on defect levels in GaAsP LEDs, and one as a fellow extending this work on to InGaAsP LEDs and lasers, and InGaAs photo detectors. In 1989, he joined BT Laboratories at Ipswich and began working on the design, modelling and characterisation of all aspects of semiconductor laser performance. He led a group whose responsibilities have covered the design, fabrication and characterisation of optoelectronic integrated circuits and photonic integrated circuits and encompassed the development of novel components for interfacing between radio and optics. In April 1994, he changed job to head up a group responsible for designing, building and running an advanced broadband multiservice network for 650 users. In September 1996, he made a further change and is now working on setting up a longer-term activity on the management and support of advanced networks and services. He has authored or co-authored over 50 publications, three books, and is currently a Visiting Professor at the Department of Electrical Engineering, Sheffield University.

**Jeremy Barnes**

BT Networks and Systems

Jeremy Barnes is a senior engineer with Concert Communications' Internet Plus and ATM team. Previously, he worked in the Advanced Applications and Technologies department at BT Laboratories. Most recently, this included three years in a team analysing information networking technologies and realising their potential over the futures test bed. During this time, he also represented BT in the LAN Emulation and Multi-Protocol groups of the ATM Forum and at the Internet Engineering Task Force. His earlier work included modelling using system dynamics techniques to investigate traffic scenarios for broadband networks and services. He holds a degree in Physics from Durham University, is a member of the IEEE and an IEE Associate. He has published a number of papers in the field of information networking, and was a member of the lecture team for the IEE's 1995-97 Silvanus P. Thompson UK tour.

**Dave Pratt**

BT Networks and Systems

Dave Pratt has a B.Sc., M.Eng. and Ph.D. from Southampton University and, after a period with Plessey Research where he developed fibre-optic sensor systems, has worked since 1987 at BT Laboratories. Initially, he was in charge of advanced function laser source development including transfer into production. He later led a team developing various optical wavelength-division multiplex (WDM) components. In April 1994, he joined as a founder member of the futures test bed team. He set up a network of Unix workstations and regularly gives presentations to important visitors. He deployed multicast (MBONE) on the futures test bed, on the City Media Network, and for Globecom 96. He represents BT at the IETF and on the ATM forum, covering the LANE, MPOA and PNNI areas.

**Peter Cochrane**

BT Networks and Systems

Peter Cochrane joined the British Post Office in 1962 and is a graduate of Trent Polytechnic and Essex University. He is a fellow of the IEE, IEEE, and Royal Academy of Engineering, a visiting professor to Essex and Kent Universities and University College London. He joined BT Laboratories in 1973 and has worked on a variety of analogue and digital switching and transmission studies. He has been a consultant to numerous international companies on projects concerned with systems, networks and test equipment developments. In 1978, he became manager of the Long Lines Division and directed the development of optical-fibre systems, photonic amplifiers and wavelength-routed networks for terrestrial and undersea applications. His team received the Queen's Award for Technology in 1990 for the production of optical receivers for TAT-8 and the PTAT-1 undersea cable systems. In 1991, he was appointed to head the Systems Research Division at BT Laboratories, which is concerned with future computing and communications developments. He was further promoted in 1993 to head the Advanced Applications and Technologies Department with 620 staff.

# Advocating a Managed Optical Platform

*A managed optical platform for the BT UK core network would reduce capital costs significantly and enhance resilience of the basic infrastructure. This becomes increasingly important as traffic demand grows substantially with the introduction of broadband services.*

## Background

Telecommunications network operators are designing broadband networks to deliver interactive multimedia services to business and residential customers. This will increase the traffic carried on core transport networks. Today's modern transport networks are based on electronic synchronous digital hierarchy (SDH) cross-connects and high-capacity single-carrier optoelectronic transmission line systems supporting various switched platforms: asynchronous transfer mode (ATM), switched multimegabit data service (SMDS), frame relay, etc. Tomorrow it will be feasible and cost-effective to build the transport infrastructure by using layers of electronic and all-optical cross-connects (Figure 1), interconnected with very-high-capacity wavelength-division multiplexed (WDM) carrier transmission systems<sup>1</sup> with significant capital cost savings.

A managed optical platform will also simplify resilience management. When breaks occur in optical cables, routes can be reconfigured rapidly by using optical add-drop multiplexers (OADMs) or optical cross-connects (OXC) without the need to reconfigure large numbers of 150 Mbit/s (VC-4) and 2.3 Mbit/s (VC-12) blocks.

With the introduction of broadband services, the amount of traffic requiring transport across the core network will reach a level where it becomes cost-effective to route large blocks of traffic at VC-4 and above. Consider, for example, the current SDH layer (Figure 2). All traffic entering an SDH node is processed at 2.3 Mbit/s and 150 Mbit/s, whereas much of it could pass through without grooming or switching at this level. An OXC (or OADM) could route tandem traffic effectively and provide additional resilience options.

Hence there exists an opportunity to re-address fundamentally both PDH and SDH resilience strategies while utilising the currently installed fibre infrastructure for future growth. WDM is the enabling technology.

This article first considers a high-level network design for a managed optical platform, and issues of functionality, dimensions and cost. It goes on to examine a migration strategy, using

Figure 1—Network platforms

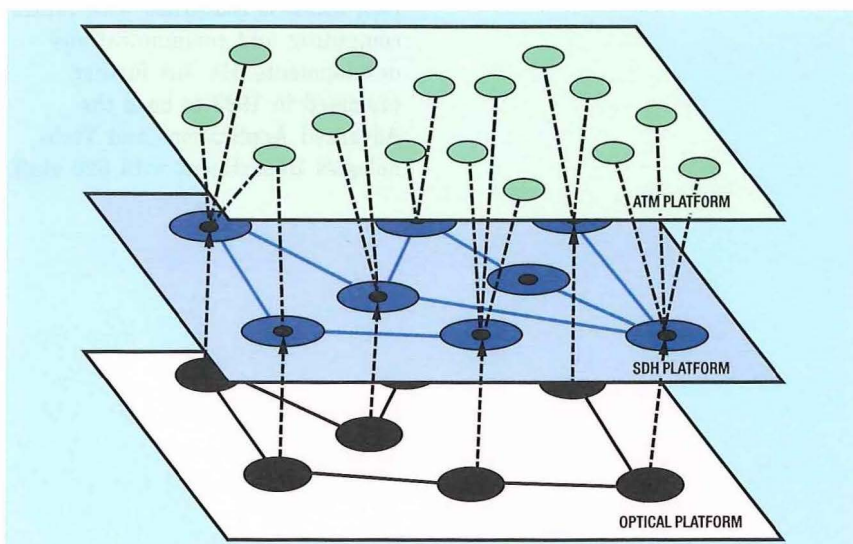


Figure 2—Additional flexibility using the optical layer




Figure 3—A core network hierarchy

optical technology that will soon be available from suppliers, and to link this to current network needs.

### High-Level Network Design

Let us dimension a core network for 2, 5 and 20 times the current traffic demands (for example, bulk services migrated from 64 kbit/s telephony circuits to 2 Mbit/s video services require a 30-fold increase in transport capability) and estimate the cost benefits of providing a managed optical platform<sup>2</sup>.

Consider a network (Figure 3) with a series of Tier 2 rings with STM-*n* ( $n \times 155$  Mbit/s) opto-electronic transmission systems and SDH ADMs (add-drop multiplexers). These rings are dual-parented onto an SDH mesh network at Tier 1 composed of high-capacity STM-16 (2.44 Gbit/s) transmission systems and a mixture of low-granularity-routeing VC 4/1 digital cross-connects (DXCs) and high-granularity-routeing VC 4/4 (VC-4) DXCs providing multilayer routeing and grooming within the Tier 1 transport layer.

With the introduction of broadband services, the amount of traffic requiring transport across the core network reaches a level where it is cost-effective to route blocks of traffic at 150 Mbit/s and above. The high-granularity routeing can be carried out by either an OXC<sup>3</sup> (Figure 4) or a high-capacity electronic DXC (SDH, ATM, etc.) with traffic streams routed at  $n \times 150$  Mbit/s.

### Network and Traffic Model

A template traffic matrix based on demand for bidirectional 2 Mbit/s private circuits illustrates the traffic flows between exchanges. Different scaling factors can then be applied to this template to represent future traffic forecasts. An example UK network structure showing node and link locations is shown in Figure 5.

The network design was performed by a multilayer tool using simulated annealing<sup>4</sup> to minimise

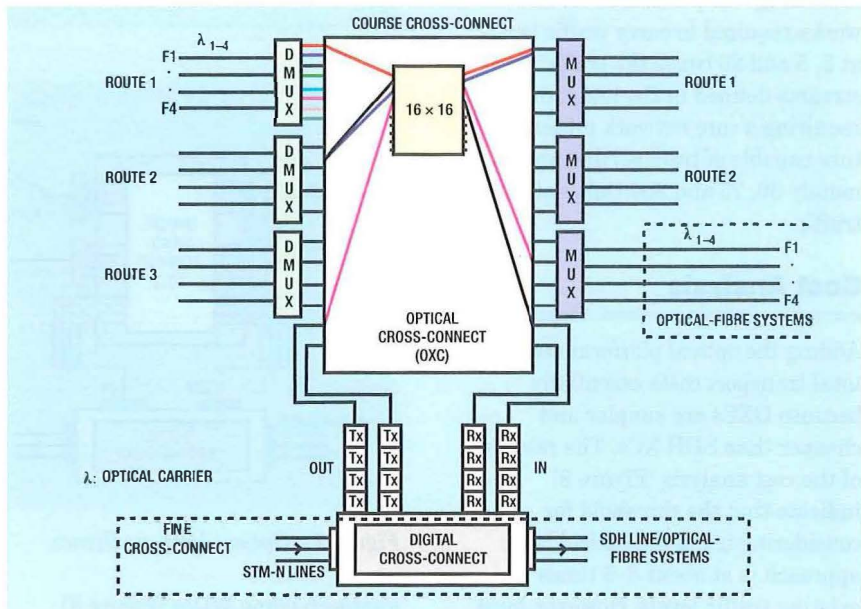
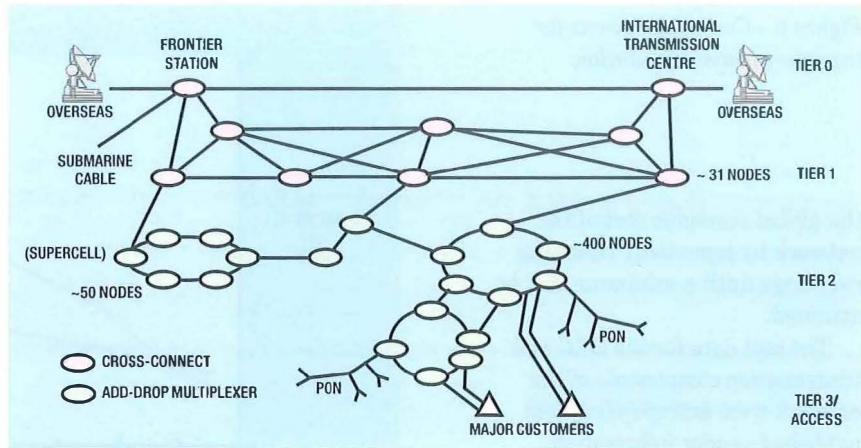


Figure 4—High granularity routeing multi-layer node structure

Figure 5—Selected model topology




Figure 6—Cost comparisons for transport network evolution

the global economic cost of the network by repeatedly changing routings until a minimum cost is attained.

The cost data for the OXC and transmission components of the network were extrapolated from published vendor information. Modelling was performed for networks required to carry traffic levels at 2, 5 and 20 times the traffic streams defined in the template, requiring a core network infrastructure capable of transporting approximately 30, 75 and 300 Gbit/s of total traffic.

**Cost Analysis**

Adding the optical platform reduces total transport costs essentially because OXCs are simpler and cheaper than SDH XCs. The results of the cost analysis (Figure 6) indicate that the threshold for considering using the optical layer approach is at about 5–6 times existing traffic levels. However, local hot spots would see immediate benefit.

**Network Resilience**

Cable cuts continue to be a major source of revenue loss to telcos. Generally, these are restored by centralised restoration techniques using operational support systems to identify spare capacity and re-route traffic.

A clear advantage of using OXCs is the ability to reconfigure connections at 150 Mbit/s and above after a cable break (Figure 7) without the need to reconfigure a large number of VC-12 and VC-4 blocks in the SDH platform. This will reduce the complexity of resilience management and speed up the restoration process.

**A Migration Strategy**

In 1998, OADM rings will be the first optical networking products available from manufacturers, several years ahead of full OXCs. The OADM

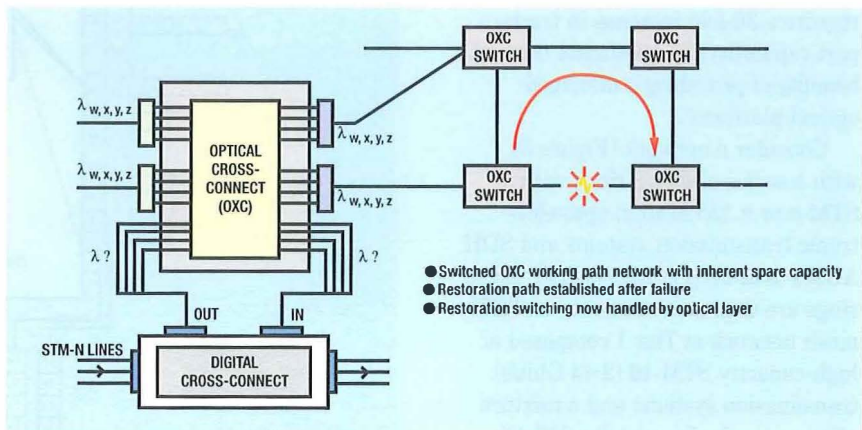
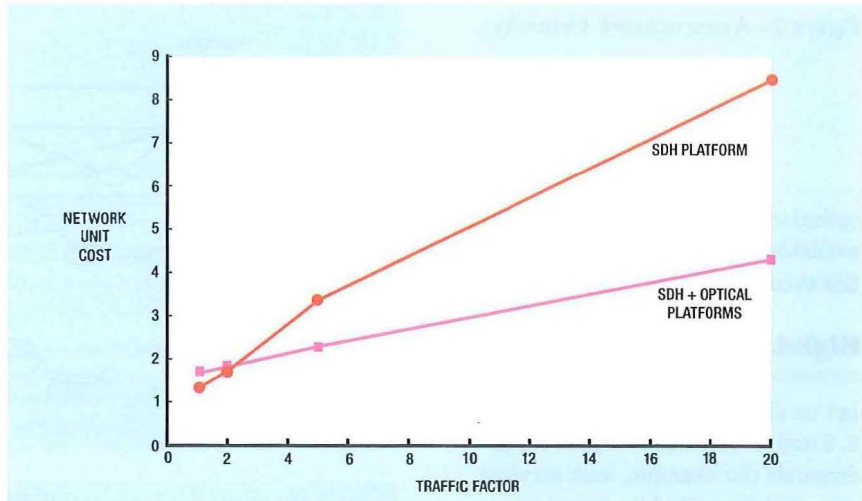


Figure 7—Optical layer resilience

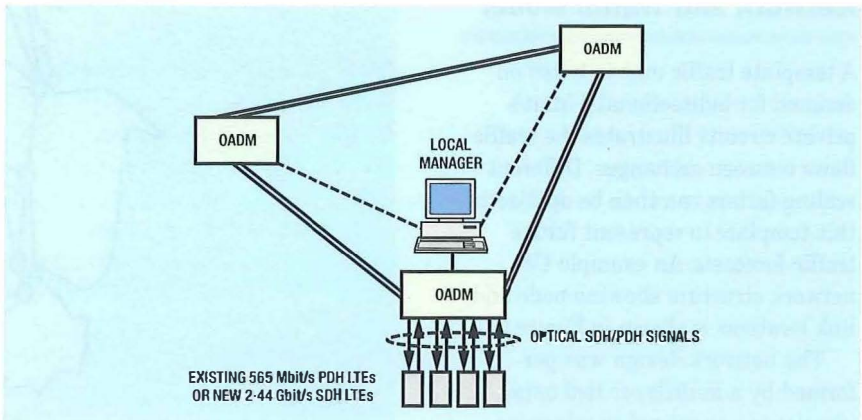
approach using WDM (Figure 8) increases both network capacity and resilience in a cost-limiting and cost-incremental fashion. Thick PDH and SDH routes (carrying upwards of four bidirectional 565/2440 Mbit/s systems) can be provided with millisecond restoration for duct, fibre or line-system outages, while spare fibre or WDM channels are provided in large blocks at incremental cost as opposed to an immediate high capital outlay for 48 or 96 fibres in a cable. SDH

equipment savings (and repeaters) will be another major economic benefit.

**Conclusions**

- Soon after the year 2000, the volume of traffic requiring transport across the core network is expected to increase to a point where it becomes economic to consider high-granularity routing at the STM-16 level (2.44 Gbit/s routing granular-

Figure 8—OADM rings



ity) or above using high-granularity electronic-optical cross-connects or OADMs.

- Bandwidth transparency should allow large units of capacity to be routed cheaply around networks with the possibility of flexible provision in the face of uncertain market forces. On a purely capital cost basis, a 50% reduction could be accomplished by using the optical layer, when core traffic exceeds about 20 times current levels.
- After a major outage, such as a cable break, there are benefits in restoring in the optical layer since this eliminates the requirement to reconfigure large numbers of 140 Mbit/s blocks independently. This in turn should reduce complexity and the impact of a major break on the end user.
- In 1–2 years, OADM rings will be available, providing a medium-term option to upgrade the capacity and resilience of the basic infrastructure and the first stages in realising a managed optical platform.
- Optical networking systems (OADMs and optical line amplifiers (OLAs)) will have associated element managers which will be remotely managed and configured. These functions can then be integrated into the network control layer allowing configuration, fault and performance management of both the optical platform and the platforms carried on it.

### Acknowledgements

The authors would like to thank Paul Botham, Graham Reeve and Alan McGuire of BT Laboratories for their contributions to this work.

### References

- 1 HILL, G. R.; CHIDGEY, P. J.; KAUFHOLD, F.; LYNCH, T.; SAHLEN, P.; GUSTAVSSON, M.; LAGERSTROM, B.; GRASSO, G.; MELI, F.; JOHANSSON, S.; INGERS, J.; FERNANDEZ, L.; ROTOLO, S.; VEZZONI, E.; TESTA, F.; SCAVENNEC, A.; O'MAHONY, M. J.; ZHOU, J.; YU, A.; SOHLER, W.; RUST, U.; and HERMANN, H. A. Transport Layer Based on Optical Network Elements. *IEEE J. Lightwave Tech.*, May/June 1993, **11**, pp. 667–679.
- 2 LOWE, E. D.; BOTHAM, P.; and HAWKER, I. An Upgrade Route for Core Transport Networks. GLOBECOM '96, London, November 1996.
- 3 JOHANSSON, S.; LINDBLOM, M.; GRANESTRAND, P.; LAGERSTROM, B.; and THYLEN, L. Optical Cross-connect System in Broadband Networks: System Concept and Demonstrators Description. *IEEE J. Lightwave Tech.*, May/June 1993, **11**, pp. 688–694.
- 4 KIRKPATRICK, S.; GELATT, Jr., C. D.; and VECHII, M. P. Optimisation by Simulated Annealing. *Science*, 13 May 1983, **220**(4598), pp. 671–680.

### Glossary

- PDH** Plesiochronous digital hierarchy
- SDH** Synchronous digital hierarchy
- SMDS** Switched multimegabit data service
- ATM** Asynchronous transfer mode
- DXC** Digital cross-connect
- OXC** Optical cross-connect
- OADM** Optical add-drop multiplexer
- OLA** Optical line amplifier
- LTE** Line terminating equipment
- VC-n** Virtual container  $n$

### Biographies



**Ewart Lowe**  
BT Networks and  
Systems

Ewart Lowe received a B.Sc. in Applied Physics from UMIST in 1988 and a Ph.D. in Electrical Engineering from Essex University in 1995. He joined BT Laboratories in 1988 and was involved with the development of BT's transmission surveillance system in the Network Management Department. For the past five years he has been closely involved with the development of future optical transport networks and is currently responsible for a small team of engineers and researchers engaged in planning and designing BT's advanced transport network requirements.



**Ian Hawker**  
BT Networks and  
Systems

Ian Hawker received B.Sc. and Ph.D. degrees from Leicester University. He joined BT Laboratories and began work in the design of high-speed integrated circuits for TAT-8 submarine systems and for inland 565 Mbit/s transmission systems. He then undertook studies in design, control and quality-of-service for SDH networks. He currently leads a team researching future telecommunications networks including issues in resilience, advanced network design and optical networking.

# Using Software Agents for Business Process Management

*Organisations increasingly are becoming distributed. Agent-based process management systems (APMSs) combine the latest distributed computing techniques with autonomous software-agent technology to provide a service-oriented view of a business process in order to meet the requirements of open distributed enterprises, now and in the future. This article outlines the key benefits of adopting an APMS approach. It describes the APMS architecture and an example in the ADEPT system.*

## Introduction

The problems of defining and maintaining business processes within a large organisation, such as BT, can seem to have much in common with the famous story about painting the Forth Rail Bridge: no sooner do you think you have finished than you must start all over again! This was one of the salutary lessons learnt in the late-1980s from the work of BT and NYNEX<sup>1</sup> that led to the initiation of project Sovereign. This created a five-layer BT business model which could not easily be stabilised. That we live with constant change is a platitude that has never been more appropriate than of telecommunications in the last 15 years.

A great many advances have been made in recent years within BT in preparing a culture of dynamic improvement. The Breakout programme in particular has served as a beacon for this effort, and this has been well supported by other initiatives, such as the development of the telecommunications requirements, impact analysis and decision support (TRIADS) system<sup>2</sup>. From the Total Quality Management initiatives, through regulatory and competitive pressures, and on to ISO 9001-conformant quality management systems, the emphasis has been placed increasingly on the notion of *service*. This embraces the service that BT provides to its customers, and the service that trading units within the company provide to each other.

Business process definition and maintenance requires considerable effort to consult everyone involved

and to capture the requirements and constraints of all parties to the process. This involves not only the individuals contributing to a process but also the underlying information systems which facilitate and often have processes embedded within them. Unless this process is coupled with ways of engineering for future changes, the end result may serve simply as a model of how the process should have worked one or two years ago. The process support infrastructure must include effective tools for monitoring (and measuring) the performance of the process and for re-engineering its design, as in the vision of process integration suggested in Reference 1.

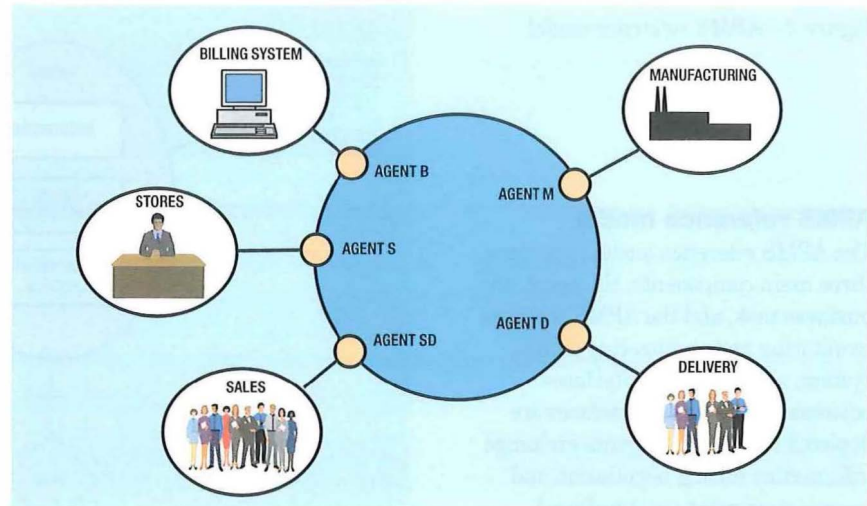
This article describes a business process support infrastructure that is inherently service-oriented. This environment builds upon the latest techniques from distributed computing and intelligent systems techniques<sup>3,4</sup>, and utilises software agents to bring on-line aspects of service provision. We call this new approach *agent-based process management systems* (APMSs).

## Agent-Based Process Management Systems

APMSs offer an alternative vision of how organisations can be structured and managed. They take a decentralised service-oriented view of business process management, where the resourcing and coordination of activities to support an end-to-end business process involve negotiation and collaboration between customer and provider agents. APMSs combine the



Figure 1—Agents and the virtual enterprise



latest distributed computing technology with agent-based techniques.

The service-oriented approach of APMSs can reflect the inherent distributed nature of large organisations and make the management of an organisation transparent to its logical or physical structuring. Similarly, this approach allows an organisation to adapt and evolve with minimal disruption so that new services or tasks can be defined incrementally, without the need to redesign an entire distributed system. This empowers local semi-autonomous groups to define how they will perform and manage tasks and processes; Reference 5 discusses in detail the empowerment of personnel in large organisations.

An important feature of APMSs is that dimensioning of business processes is brought on-line and integrated with process enactment, resulting in improved redeployment of resources and increased flexibility during exception handling. Therefore, APMSs have two objectives: firstly, the timely execution of business functions, and, secondly, the efficient use of resources.

### APMS: An overview

A business process is made up of several functional activities or tasks, such as capturing a customer's details. A task represents a primitive functional component of a business process. Dependencies exist between tasks, so they have to be executed in a controlled and ordered way. The execution of a task may consume certain resources. Resources in most organisations are grouped into business units that have semi-autonomous control over the way in which those resources are deployed.

Autonomous software agents represent these groups of resources. They communicate with each other over a network and negotiate over how they can collaborate to support an overall business process. Each agent offers 'services' to other agents. A service is a packaging of tasks and other (sub)services that allows an

agent to offer or to receive from another agent some functional operation. A service can be reused as a component of another service, and agents can take the role of provider (server) or customer (client) for services. There are parallels with some aspects of the organisation of large enterprises, where service-level agreements may exist between groups and departments (for example, BT's trading units).

Figure 1 shows an example of five agents managing a business process. Each agent provides service(s) in support of a typical customer sales business process.

Agent SD, representing the sales department, would negotiate for services from other agents to support the sales function. This could involve negotiating for manufacturing capacity with agent M, and delivery services from agent D. These agents in turn might require other services; for example, the agent M might negotiate for services from stores through agent S. As agreements are reached between the agents (as both client and server agents), contracts are established.

Each agent has to ensure that it has sufficient resource capacity to provide a service before it commits itself to delivering that service. This requires the agent to allocate sufficient resources to support those tasks under its direct control, as well as ensuring it has access to sufficient component services offered by other agents.

When a customer requests a service from the sales team, an instance of this business process is enacted. The agents manage the execution of the business process instance ensuring that each task and

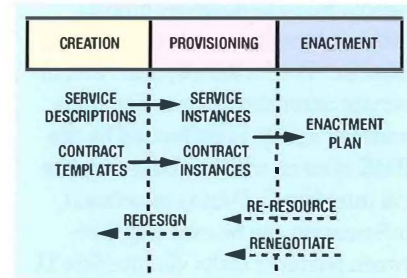


Figure 2—The APMS life cycle

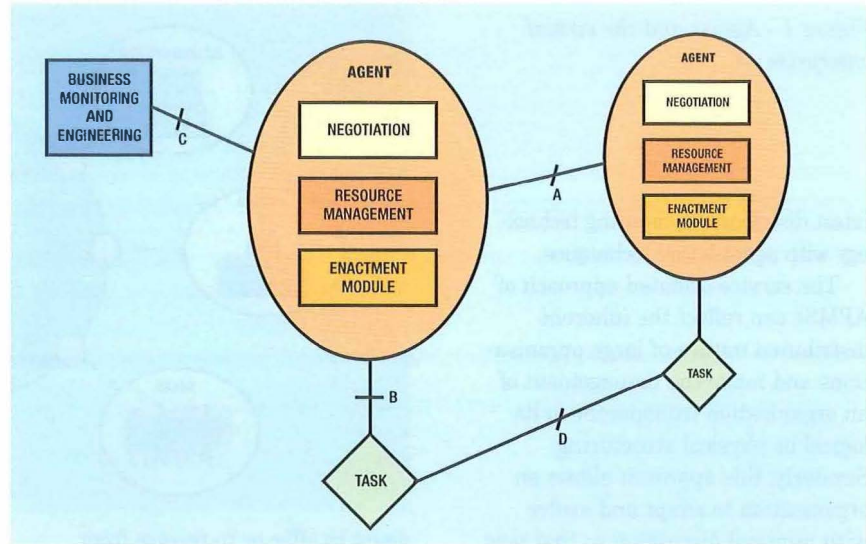
service is performed in a timely and efficient way within the constraints of the agreed contracts. Agents interact to coordinate enactment, facilitate the exchange of information and handle exceptions (this might require rescheduling and/or renegotiation).

### APMS life cycle

The APMS life cycle consists of three key phases: creation, provisioning and enactment (Figure 2).

The service creation phase involves the definition of service descriptions, contract templates and the selection of negotiating strategies for the agents in the APMS. In existing systems (see project ADEPT later), this is a manual activity, but ideally some form of decision support would be provided to ease this process. The service provisioning phase is performed automatically by agents. This involves agent negotiation for services, the agreeing of contracts between agents and the assignment of resources. The delivery phase is the enactment of those agreed services. If exceptions occur which cannot be resolved by the enactment module, then the service reverts to the provisioning phase for re-resourcing or renegotiation.

Figure 3—APMS reference model



### APMS reference model

The APMS reference model comprises three main components: the agent, the business task, and the APMS business monitoring and engineering (BME) system, as well as the interfaces between them. These interfaces are depicted in Figure 3. Agents exchange information during negotiation and service enactment via interface A. Agents manage business process tasks, information and resources via interface B. The development and in-service administration and maintenance of agents is performed by the BME system, which accesses agents via interface C. During enactment, information can be exchanged between business tasks via interface D.

An APMS agent comprises three core modules<sup>6</sup>:

- negotiation module,
- resource management module, and
- enactment module.

#### Negotiation module

Negotiation is the process whereby two or more agents seek a mutual agreement and commitment on the delivery of a service. Quality, time and cost are three typical parameters that would form the basis for negotiation. As is the case with all negotiations, agreement cannot always be reached.

A contract is the result of an agreement between client and server agents during the negotiation process. The contract contains a list of agreed values for parameters, establishing the terms and conditions for the delivery of the service. For example, the agreed time(s) at which the service will be available and/or activated, the maximum duration of the activation, the minimum quality of the service, etc. will be defined.

In a realistic business process, an agent will be required to negotiate for multiple contracts simultaneously. Therefore the negotiation management module of an agent must be multi-threaded so that it can

support multiple negotiations for different services concurrently.

During negotiation, an agent correlates and balances multiple criteria both within the negotiation for a single service, and across all the negotiations the agent is involved in. The criteria can be modelled as (partial) ordinal value spaces that represent parameters such as quality, time, cost, etc.

#### Resource management module

An important aspect of an APMS is its ability to perform direct management of resources: the systems, databases, equipment and people that make up an organisation. Resource management is one of the key advantages of the APMS approach.

When a set of resources is under some form of semi-autonomous common ownership, there is usually a need to control the commitment of those resources, to maximise speed, efficiency, etc. and to minimise cost, waste, etc. Delegating some form of executive responsibility to a single agent for the set of resources is a way of achieving this. Resource management functionality (in particular scheduling) is either implemented in, or available to, the agent.

#### Enactment module

Enactment involves the activation of tasks and agreed (sub)services in order to meet the obligations established in a contract. Server agents activate tasks and services when triggered by client agents. This can involve, for example, the execution of software, or the sending of a work schedule by fax to an operative. An

agent executes multiple services and tasks simultaneously. This module must be multi-threaded to allow the agent to activate concurrent operations.

When tasks and services fail, agents can perform corrective actions and try to resolve the failure. The enactment module receives exceptions from the tasks within its agency and from other agents (as servers). Exceptions can be resolved by either:

- restarting the task/service,
- re-resourcing/rescheduling the task/service by the resource management module,
- renegotiating the terms of the contract (as the defaulting server agent),
- relocating the service with another agent (as the aggrieved client agent), or
- ignoring the exception and accept a penalty (if appropriate).

#### Business monitoring and engineering (BME)

The APMS BME system supports developers of organisations to engineer and monitor agent-based business process management systems.

Within an APMS, agents offer services to each other based on process activities (tasks). These services are combined to realise business processes. The techniques for defining services that may be useful in such a scenario correspond

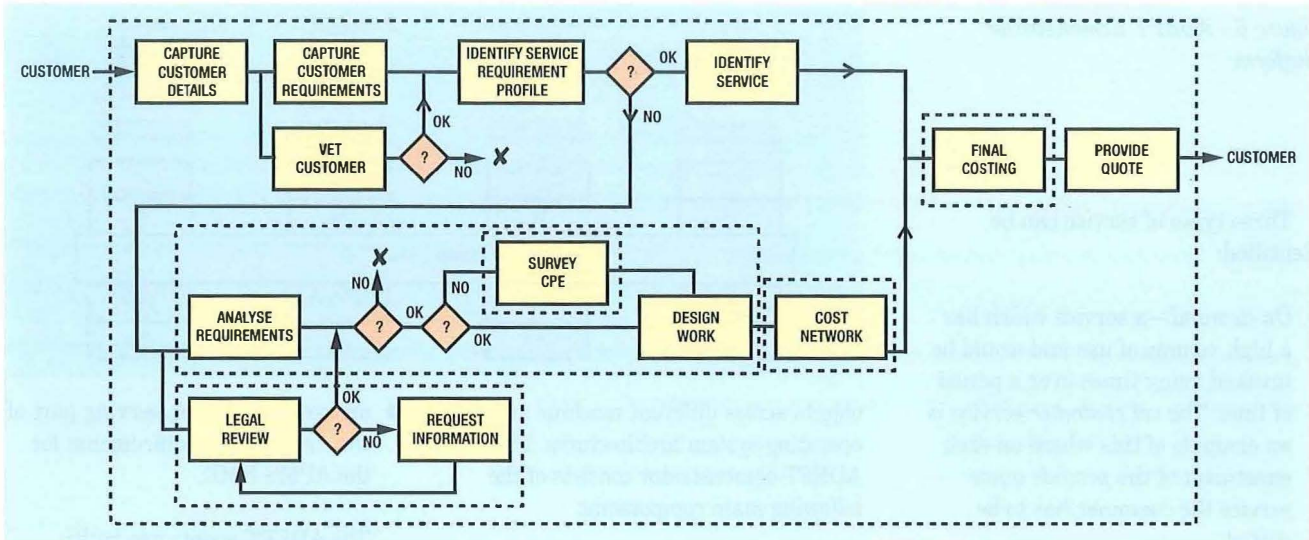


Figure 4—Provide Quote business process

to business process definition tools, although it is important to note that these services are not expressing the complete process. The service definition phase must be linked to a methodology that ensures that the defined services are useful to the enterprise (and will therefore be used) and that they will together realise an end-to-end business process. These conditions are referred to as *necessary* and *sufficient* for service definitions.

The structure of an integrated system should reflect the structure of the existing organisation(s) rather than be imposed. Each agent manages business functions that are under its direct control, and the realisation of an end-to-end business process could involve contributions from many different agents. The monitoring and administration of a business process based on the APMS approach demands the ability to collate information from all the agents involved, to filter this information, and to present it in a way that allows a business process owner/manager to understand the contribution from all agents in the enterprise, and to see where problems might occur. The BME has to strike a balance between agent autonomy and the overall business requirement.

**ADEPT: An APMS Prototype**

Project ADEPT has developed a prototype APMS which demonstrates multiple autonomous agents managing over 200 simulated business tasks which make up a typical BT business

process. The ADEPT infrastructure consists of a community of nine agents that can negotiate concurrently with one another in order to reach agreement on how resources are to be assigned to support a business process. The ADEPT agents are used for resourcing business processes, coordinating process tasks and exception handling when business processes break down.

Figure 4 illustrates a typical Provide Quote business process. This business process consists of those tasks which need to be performed when providing a customer with a quotation. The process is initiated by a customer requesting a quotation and terminates with a quote being provided to that customer. The process is partitioned into groups of tasks and services that are performed by particular business units.

The ADEPT agents negotiate with each other to agree which unit will perform which service and when (see Figure 5).

The top-level service *provide quote* is offered by agent CS (customer sales), which represents the interests of a customer servicing department. The *provide quote* service comprises the tasks that agent CS can perform itself and has complete control over; combined with the services that can be negotiated from other agents. Customer vetting is offered by agents CC, VC, and XC. These could represent external companies offering credit checking facilities. The network design department represented by agent ND offers both *design network* and *cost network* services. The *design network* service uses an embedded service, *survey cpe*, offered by agents S1 and S2. On-line legal advice is provided by agent LA.

Figure 5—ADEPT solution

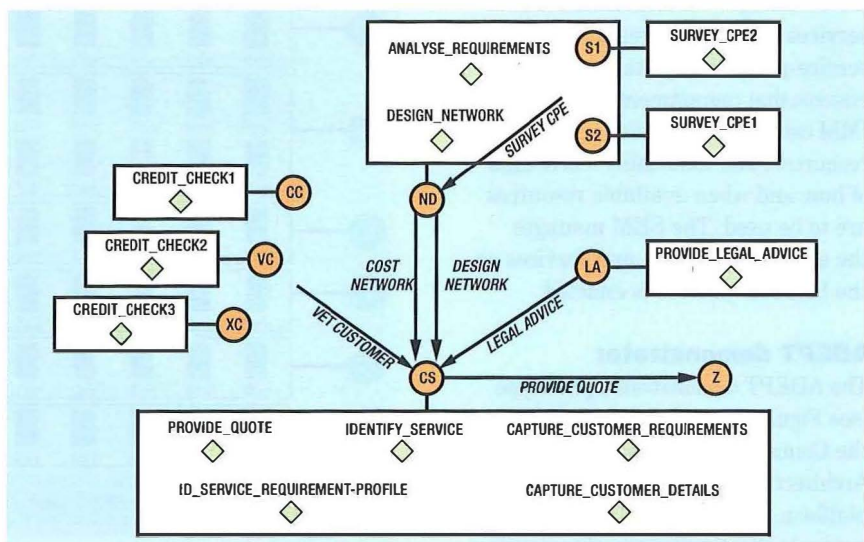


Figure 6—ADEPT demonstrator platform

Three types of service can be identified:

- *On-demand*—a service which has a high volume of use and would be invoked many times over a period of time. The *vet customer service* is an example of this where on each enactment of the *provide quote service* the customer has to be vetted.
- *One-off*—a service which is used occasionally and consumes a resource of high value. An example of this is the *legal advice service*, which would be enacted only under particular circumstances.
- *Regular*—a service which is used regularly at a particular time. An example would be a database back-up service.

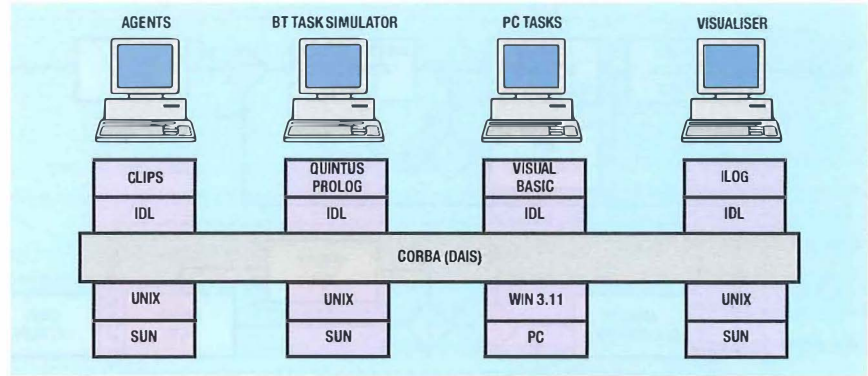
The particular type of service is reflected in the type of service level agreement made between agents. In addition to this, the agents would agree quality, timing and cost of services.

**ADEPT Agent**

Each ADEPT agent consists of three core modules: the interface management module (IMM), situation assessment module (SAM) and service execution module (SEM). These map respectively into the reference model outlined in the previous section. The IMM manages the negotiation for services with other agents during the service-provisioning stage. The SAM ensures that commitments made by the IMM can be discharged using available resources, and maintains a schedule of how and when available resources are to be used. The SEM manages the execution of tasks and services as the business process is enacted.

**ADEPT demonstrator**

The ADEPT demonstrator prototype (see Figure 6) is a full integration over the Common Object Request Broker Architecture (CORBA)-compliant platform, DAIS, from ICL. CORBA supports the distribution of computing



objects across different machine and operating-system architectures. The ADEPT demonstrator consists of the following main components:

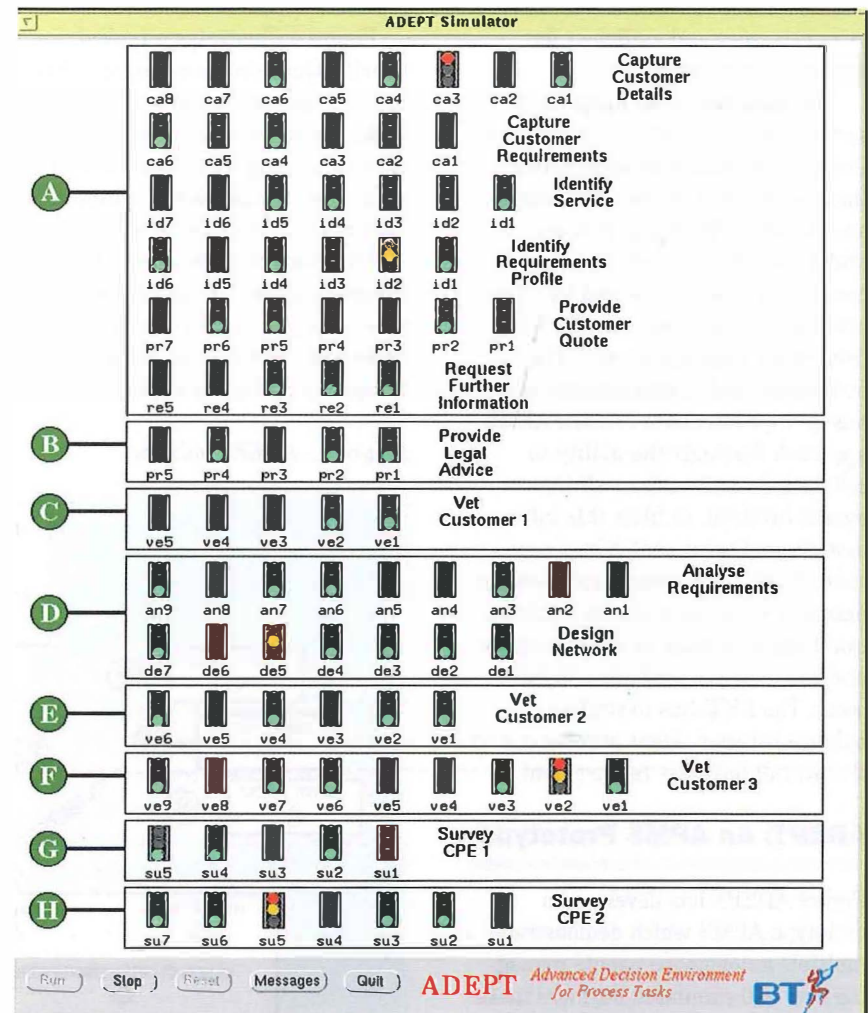
- a community of nine agents using the ADEPT agent architecture,
- a BT task simulator that emulates the relevant behaviour of over 200 distinct business tasks,
- a simulation of three PC-based business tasks including front-ends, and

- an agent visualiser, serving part of the monitoring requirements for the APMS BME.

The ADEPT agents are multi-threaded, supporting the concurrent execution of several business processes. The CLIPS rule-based language has been used in the implementation of the agents.

The BT task simulator (see Figure 7) emulates the relevant behaviour of distinct business tasks, providing a simulation environment for the agents which is used both for testing and demonstration. Simulated

Figure 7—ADEPT task simulator



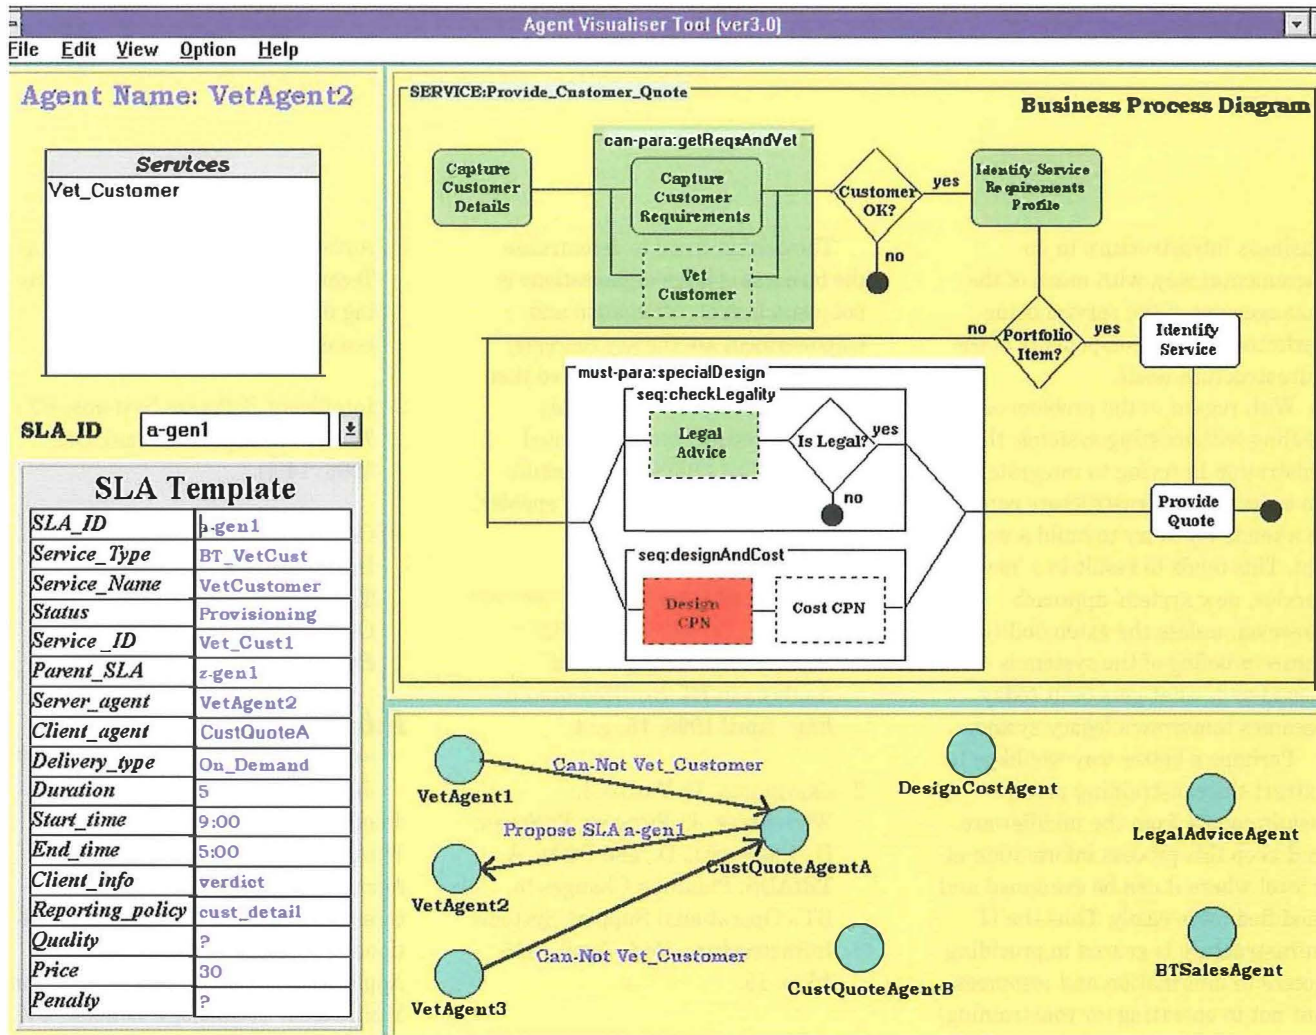


Figure 8—ADEPT visualiser

tasks consume and generate information and have varying execution times and different probabilities of success or failure. Additionally, three PC-based tasks are represented. These provide examples of how an ADEPT infrastructure would impact on those actually performing business tasks as part of a business process.

Agent visualisation tools (see Figure 8) provide the means of displaying agent communication and interaction. The visualiser can be seen to include functionality that would be included in an APMS BME. It performs the role of both a window onto the underlying technology and the agents, as well as a view of how the business process is performing.

Further details of the ADEPT demonstrator system are available in Reference 7.

### Current and Future Directions

An APMS must address a range of issues, particularly:

- defining and agreeing the new procedures for all involved;
- defining and agreeing the resourcing levels for the new process;
- monitoring and measuring the performance of the business process, to determine its objective value, and to ensure that other processes are not compromised; and
- assessing and dealing with the constraints imposed by existing IT systems.

The first of these issues corresponds to the creation of new services, and in the APMS approach this is a key action that is still (and probably must be) performed manually. However, though a manual activity, it requires to be supported by appropriate tools and a methodology that will help business-process engineers to identify and understand the necessary and sufficient conditions for new-service

creation. Project IBS (Intelligent Business Systems) has started investigating these issues related to the development of a business monitoring and engineering environment.

The second requirement relates to service provisioning, and part of this activity (as demonstrated by ADEPT) can now be brought on-line and made part of the infrastructure itself. Using advanced resource-scheduling techniques<sup>8</sup> and emerging intelligent software<sup>9</sup>, many of the lower-level decisions concerning the mapping of resources to business functions can be made within the process management system itself. This allows an organisation to respond much more quickly to changing resource requirements.

Of course, both of the above demand the continual monitoring of the process enactment to ensure that objectives are always being met. Any organisation must monitor the services that it provides and realise when changes are necessary. New services can be integrated into the

business infrastructure in an incremental way, with much of the management of the service being performed by the components of the infrastructure itself.

With regard to the problem of dealing with existing systems, the frustration in trying to integrate with an existing IT infrastructure can lead to a tendency to try to build a way out. This tends to result in a 'new service, new system' approach. However, unless the extensibility and future-proofing of the system is considered, what gets built today becomes tomorrow's legacy system.

Perhaps a better way would be to extract the constraining process requirements from the middleware, and keep this process information at a level where it can be examined and modified more easily. Thus the IT infrastructure is geared to providing access to information and resources, but not to operating (or constraining) the process. One has only to witness the popularity of intranets<sup>10</sup> to realise that the days are numbered for restrictive IT systems. But one also has to have suffered the Internet to realise that customers cannot be kept waiting on the telephone while the IT system recovers information.

## Conclusions

The APMS approach views the business process fundamentally as being component-led; the process is created and enacted in a decentralised way. However, greater decentralisation requires more effective monitoring of performance, so that the end-to-end impact of the process is understood properly.

The partial automation of the service-provision cycle can support the early communication of changes, whereby software agents negotiate recovery through naturally occurring organisational interfaces. Stakeholders within the process are readily identifiable and can be contacted immediately. This can provide the interfaces through which effective benchmarking can be achieved.

The current trend to decentralise the business of large organisations is not just a fashion. Delegation and empowerment are the key concepts, and an infrastructure is required that is flexible enough to support this customer-focused, service-oriented approach. The APMS infrastructure has the potential to be a great 'enabler'.

## References

- 1 FINEMAN, LAURIE. Process Re-engineering: Measures and Analysis in BT. *Br. Telecommun. Eng.*, April 1996, **15**, p. 4.
- 2 SKEVINGTON, P.; VIDELO, I.; WITTEGREFFE, J.; PUTLAND, P.; SLOAN, D.; CRESSWELL, D.; and SMITH, A. TRIADS: Planning Changes to BT's Operational Support Systems Infrastructure. *ibid.*, April 1996, **15**, p. 13.
- 3 SMITH, R. Software Agent Technology. Proceedings of the 1st International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology, London, UK, 1996, pp. 557-551.
- 4 NWANA, H. Software Agents: An Overview. *Knowledge Engineering Review*, 1996, **11**(3).
- 5 ARMSON, R.; and PATON, R. Organisations: Cases, Issues, Concepts, 2nd edition. Paul Chapman Publishing Ltd, 1994.
- 6 ALTY, J. L.; GRIFFITHS, D. G.; JENNINGS, N. R.; MAMDANI, E.; STRUTHERS, A.; and WIEGAND, M. E. ADEPT-Advanced Decision Environment for Process Tasks: Overview and Architecture. Proceedings of the BCS Expert Systems 94 Conference, Cambridge, UK, 1994, pp. 359-371.
- 7 O'BRIEN, P. D.; and WIEGAND, M. E. Agents of Change in Business Process Management. *BT Technol. J.*, Intelligent Systems special issue, Oct. 1996, **14**(4), p. 133.
- 8 Advanced Information Processing Techniques for Resource Scheduling and Planning. *ibid.* (special issue), Jan. 1995, **13**(1).
- 9 Intelligent Software Systems. *BT Technol. J.* (Special Issue), Oct. 1996, **14**(4).
- 10 CALLAGHAN, J.; and FLAVIN, P. Intranets. Corporate Nirvana—The End of the Traditional Organisation. *Br. Telecommun. Eng.*, Oct. 1996, **15**, p. 224.

## Bibliography

JENNINGS, N. R.; FARATIN, P.; JOHNSON, M. J.; O'BRIEN, P. D.; and WIEGAND, M. E. Using Intelligent Agents to Manage Business Processes. Proceedings of the 1st International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology, London, UK, 1996 pp. 345-360.

KADOR, J. The Ultimate Middleware. *Byte*, April 1996, pp.79-83.

KOBIELUS, J. Workflow Standards Developers Need New Rules in an Age of Electronic Commerce. *Network World*, Sept. 1995.

MEDINA-MORA, R.; WINOGRAD, T.; FLORES, R.; FLORES, F. The ActionWorkflow Management Technology. *The Information Society*, 1992, **9**, pp. 391-404.

TRAMMEL, K. Workflow Without Fear. *Byte*, April 1996.

WIEGAND, M. E.; and O'BRIEN, P. D. ADEPT: An Application Viewpoint. Proceedings of Intelligent Systems Integration Programme Symposium, Ambleside, UK, 1996.

WIEGAND, M. E. Building Practical Agent-based Systems (tutorial notes). 1st International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology, London, UK, 1996, pp. 345-360.

WORKFLOW MANAGEMENT COALITION.  
Workflow Management Coalition  
Terminology and Glossary. Doc. No.  
WFMC-TC-1011, Issue 2.0, June, 1996.

## Biographies



**Paul O'Brien**  
BT Networks and  
Systems

Since joining BT, Paul O'Brien has worked on numerous projects concerned with the application of intelligent systems technology to telecommunications management. He has led a number of collaborative projects including the European RACE project DESSERT, which developed intelligent decision-support systems for service and network management. Presently he is leading the ADEPT project, which is applying intelligent agent technology to business process management. Paul graduated with a B.A.(Hons.) in Philosophy and Computing from Hatfield Polytechnic, and an M.Sc. in Intelligent Knowledge Based Systems from Essex University. He recently completed an M.Sc.(Econs.) in Information Systems Management at the London School of Economics. Paul is a member of the IEEE.



**Mark Wiegand**  
BT Networks and  
Systems

Mark Wiegand graduated with a B.Sc. in Mathematics with Computer Science from the University of Sussex in 1982. He developed supervisory control systems for manufacturing and process industries before moving to Edinburgh, where he obtained an M.Sc. in Knowledge Based Systems at Heriot-Watt University. At Heriot-Watt, he conducted research on the application of artificial intelligence techniques to industrial automation,

and lectured on advanced programming. He received his Ph.D. for work in qualitative reasoning and simulation in 1991. He joined BT in 1992 and worked initially on decision support for service management and configuration management of ATM networks. He currently leads projects on the use of agents for business process management and systems integration. He is a member of the Institution of Electrical Engineers and a Chartered Engineer. He is also an Associate Fellow of the Institute of Mathematics and its Applications.




**Brian Odgers**  
BT Networks and  
Systems

Brian Odgers graduated from Queens University, Belfast in 1991 with a B.Sc. in Mathematics and Computer Science. He gained an M.Sc. in Computational Science in 1992. He completed his Ph.D. in Theoretical Computational Physics at Queens, which considered computer simulations of electron collisions with atomic hydrogen. He has several publications in the *Journal of Physics*. Since joining BT in October 1995, he has been working in the Intelligent Systems Research group considering cooperative agents over a distributed environment. Within the ADEPT project his work has concerned the resourcing and scheduling of business processes and the handling of exceptions which arise when part of the business process fails. He hopes to extend his knowledge of this field and gain a wide range of skills within artificial intelligence.



**Christos Voudouris**  
BT Networks and  
Systems

Christos Voudouris received a Diploma Degree in Electrical and Computer Engineering from the National Technical University of Athens in 1992. In 1993, he gained an M.Sc. in Intelligent Knowledge Based Systems with distinction at the University of Essex. From 1994-95, he was a research assistant in the Department of Computer Science at the same university, researching neural networks and optimisation. He joined BT in October 1995 and has worked on projects involving scheduling, visualisation and agents. He is studying for a Ph.D. in Computer Science at the University of Essex. His research interests are heuristic methods for optimisation, fuzzy logic and neural networks.



**Donald Judge**  
BT Networks and  
Systems

Donald Judge joined BT in 1973 and worked on the production of electrical and electronic telecommunications equipment. In 1980, he moved on to the hardware/software design and development of embedded process controllers. In 1987, he gained a BTEC/HTEC in Electronics and Communications Engineering. In 1990, he joined the Human Factors Division where he worked on the usability of broadband multimedia systems. In 1992, he joined the Applications Research Section working on distributed cooperating intelligent agents for service management systems. Since 1993, he has been working on the ADEPT project developing a business process management system combining intelligent agent and distributed object management technologies. He has been responsible for the development of PC-based agent components.

*Stephen Morgan and Graham Barratt*

# Christmas in BT's Network Management Centre

*BT's network management centre in Oswestry has managed the traffic on the national network since 1987. In 1989, the international network management centre, previously located in London, moved to Oswestry to join the national team in a new purpose-built network management centre. The benefits of collocating national and international traffic management are especially appreciated during peak traffic periods, such as Christmas, when the teams work in close cooperation to provide the best possible service for customers. This article summarises the planning and operational activity involved in Christmas.*

## Introduction

Christmas comes but once a year. 'Thank goodness!' I can hear some of you say! We all know how children look forward to Christmas; as we grow older, of course, some of the magic disappears. Nevertheless, most of us can look forward to enjoying the festive season and a few well-earned days away from the office. Spare a thought then for the many people who must work during the festive period: doctors, nurses, police and fire brigade to name but a few. We rely on many service industries to maintain 365 day, 24 hour operations and the telecommunications industry is no exception. However, the major difference between BT and the other services is that Christmas is, without doubt, the busiest day of the year for the telephone network. For the people in BT's network management centre (NMC) in Oswestry, this means working round the clock to

ensure that the network can cope with the huge number of calls that will be generated. To achieve this, the NMC team continuously monitor network performance on huge wall display screens and implement network controls to optimise network performance.

## Planning

Planning for Christmas begins in January, when traffic statistics, call accounting data, and customer perception from the previous Christmas are critically reviewed. This is followed by a schedule of planning meetings throughout the year, bringing together NMC staff, national and international network planners, switch databuild and support staff as well as overseas telecommunications operators. The work of the various planning teams culminates in the preparation of detailed plans for managing the Christmas traffic. These

*BT's network management centre, Oswestry*



BT CORPORATE PICTURE LIBRARY: A BT PHOTOGRAPH




plans include temporary augmentation of national and international circuit quantities and reconfiguration of the network to ensure that best use is made of the available capacity. In other words, the network is 'fine-tuned' to cope with the predicted Christmas calling pattern. In the NMC, operational plans are drawn up, defining the traffic profiles to be expected and the timing and implementation of network management controls. *Expansive controls* provide temporary alternative routings to circumvent congestion while *restrictive controls* protect vulnerable parts of the network from overload. Internationally, considerable liaison with overseas telecommunications operators takes place to ensure the most effective use of the global network.


### Christmas Arrives

With planning completed, additional circuits and special network configurations in place, it is up to the network management team in Oswestry to implement the planned actions on the day and to respond to any unforeseen network events that may occur: Australia and New Zealand start getting busy on the morning of Christmas Eve and extra capacity is provided by routeing calls via the USA. At 11:00 and 13:00 UK time it is midnight in New Zealand and Australia, respectively, and the beginning of Christmas Day in the Antipodes generates a surge of traffic to the UK. During the afternoon and evening of Christmas Eve, calls to the USA, Canada, Hong Kong and Italy keep the network managers busy. However, this is just a foretaste of what is to follow on Christmas Day.

Christmas morning in the NMC starts fairly quietly, with the night rota team going home to be welcomed by their excited children and wondering whether they will get any sleep at all that day; while the day rota team settle down for what is inevitably going to be the hardest day's work of the year. Signs of what is to come appear as early as 07:00, as the calling



Traffic between digital main switching units (DMSUs) is higher than normal on Christmas morning because the average call holding time is much greater. The network is dimensioned for peak weekday business traffic and this changed profile results in potential congestion on routes which can normally handle all the traffic offered.



During Christmas Day the calling rate is lower than on a normal business day but the longer call-holding times associated with social calls fills the DMSU routes, resulting in congestion.

rate to Australia and New Zealand starts to climb. After most of the population of Britain have finished opening presents and eating their Christmas breakfast, they begin making their first telephone calls to their favourite aunt, uncle, etc. and, by 09:00, the network is already 'humming', with some local exchange routes to digital main switching units (DMSUs) going into congestion. The calling profile is very different from a normal business day because the average holding time of each call is much longer; typically, between 15 and 20 minutes, instead of the usual 5-8 minutes. Between 09:30 and 12:00, almost every route in the national network is fully occupied. The South West of England is particularly congested at this time, while central London areas have redundant capacity. This situation is exploited to good effect by the network managers, who re-route calls to and from the West Country via available capacity in the London DMSUs.

**New Routeing Mechanism**

For Christmas 1995, a new automatic routeing mechanism, called *dynamic alternative routeing* (DAR), was effectively deployed for the first time, to alleviate congestion in the West Country. Now DAR has been implemented nationally and all inland calls are given the option of several route choices. In addition, DAR is used on the national access routes to the international gateway switches. The principle of operation is explained in the panel opposite.

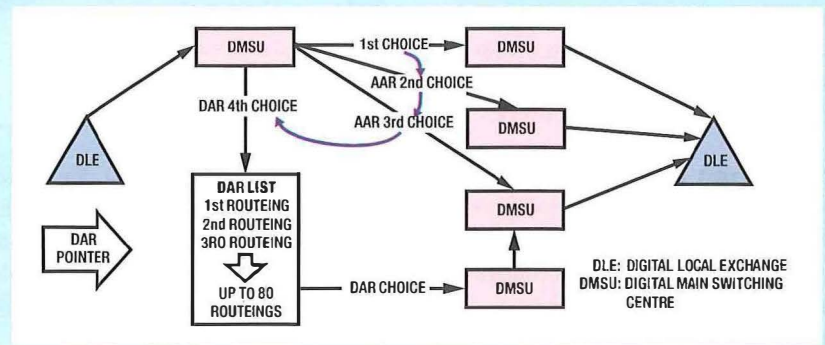
**International Relations**

Many of us have relations abroad, and Christmas is the traditional time for saying hello and catching up with family news. This creates a huge surge of calls both outgoing and incoming to the UK, and it is the task of the international network managers to ensure that as many as possible get through. The international network is volatile and depends on the

**Principles of dynamic alternative routeing (DAR)**

DAR is a form of network management that operates automatically when congestion is encountered. The DAR functionality, which is written into the exchange software, selects the most appropriate paths through the network. It has been developed as a flexible routeing strategy that can handle complex traffic patterns that occur with an increasingly volatile network environment, and yet is simple and cheap to implement and avoids the need for elaborate planning and dimensioning.


When all the fixed routeings for a call have been tried using primary and automatic alternative routeing (AAR) choices, then DAR may be used. There are up to 10 different lists available at every DMSU, each one containing up to 80 routeings. DAR will select a particular routeing by means of a pointer that moves cyclically around the list. If the call does not encounter congestion, the pointer in the DAR list stops, and that exchange choice remains selected. However, if no free circuits are available on this routeing then the pointer moves on to the next routeing in the list.



correct operation of overseas networks and the submarine cable and satellite systems linking them. Besides calls originating and terminating in the UK, many countries send traffic via BT for onward routeing to other countries. This may be because the

originating country does not have direct circuits to the destination country or because it requires temporary additional capacity to handle the Christmas peak. Managing this traffic is yet another job for busy network managers.

*Christmas Day international traffic profiles*



**International direct dialled (IDD) traffic**

The vast majority of international Christmas calls are now dialled directly and these are routed via BT's seven international gateway switches. The most popular countries called from the UK are Australia, Canada, New Zealand, South Africa and the USA. When BT's direct route capacity to these countries is fully utilised, network management controls are implemented to re-route traffic via transit carriers to take advantage of the differences in international time zones. This is done by using permanent alternative routes (PARs) and temporary alternative routes (TARs). For example, on Christmas Day, from 07:00 until 14:00, traffic to Australia and New Zealand predominates. North

America is 5-8 hours behind UK time and, consequently, its network is lightly loaded. Calls bound for Australia and New Zealand can be overflowed via the North American and Canadian carriers during this period, effectively expanding BT's network capacity. Lunchtime comes none too soon and the traffic drops off enough to allow the Network Managers to take their lunch at their consoles. Usually, Christmas lunch is prepared the day before, then served up hot at 13:00. This respite lasts only until 14:00, when the network comes alive again. International destinations change to reflect the different time zones throughout the world, and it is now the turn of Canada and USA. Eire remains very busy and this trend continues throughout the afternoon and

evening. Dinner time comes and goes, with a helping of Christmas cake, and it is almost time for the network managers who started at 08:00 to go home and spend what is left of Christmas day with their families and friends. In their place a fresh team of network managers arrives to work through the night. At 20:00 traffic starts increasing again and network controls need to be adjusted to reflect this. By 23:30 the traffic has reduced and it is time to 'tidy up'. Information that has been written hurriedly on white boards can now be transferred to the network management information system (NMIS) log, and most network controls are no longer necessary and can be disabled.

**Permanent alternative routes (PARs)**

This is the same as automatic alternative routeing (AAR) in DMSUs. It is a 2nd, 3rd, 4th, etc. routeing choice that receives calls overflowing from the previous choice. The choices that are available for a destination are called the *routeing chain*.

Overflow calls can be caused by the following:

- route congestion,
- circuits out of service, and
- *skip route* command or trunk reservation.

**Temporary alternative routes (TARs)**

As the name suggests, this is a temporary routeing choice and can be inserted anywhere in the normal routeing chain. Once this TAR has filled up the route, calls may be allowed to return to the normal routeing chain or allowed to fail. International network managers use this facility extensively during the Christmas period.

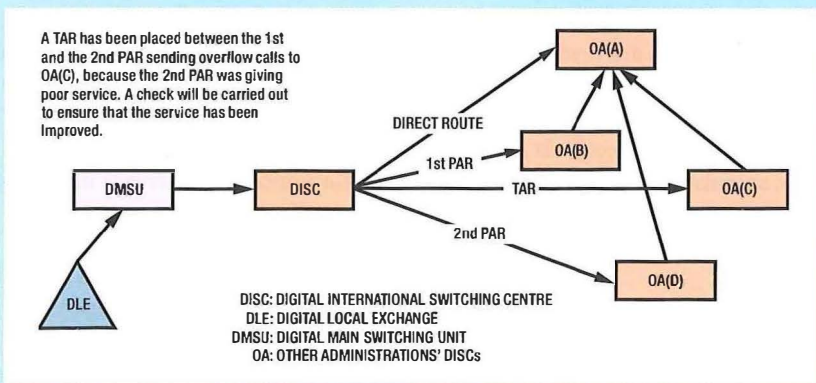
**International incoming traffic**

Incoming calls from abroad, terminating in the UK, are very profitable for BT, as BT charges the originating administration for each call for every minute's duration. Managing this traffic efficiently is therefore extremely important; in addition, it would reflect poorly on BT if calls that have been successfully transported from the other side of the world fail in our home network! As most DMSU-to-DISC routes are accessible in both directions (bothway routes), two methods can be used to improve connectivity:

- applying controls at the DISC to send overflowed homeward traffic onto a DMSU with idle circuits, and
- off-loading outgoing IDD traffic at the congested DMSU to another DMSU with idle DISC circuits.

**International switched transit traffic**

This traffic is originated in foreign telecommunications operators' networks and routed onward via BT's international gateways to its final destination. BT is one of the world's largest hubs for transit traffic. Factors influencing the volume of transit traffic include:

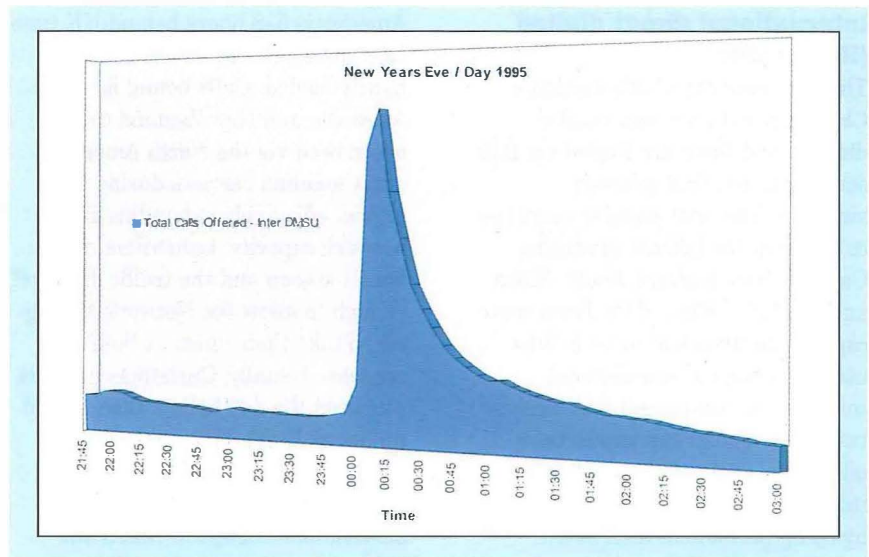


- congestion of other administrations' direct routes,
- time zone differences,
- transit accounting rates,
- perception of BT's quality of service, and
- network failure overseas.

Because of these factors, transit traffic is extremely volatile and demand cannot be forecast. The NMC has responsibility for managing this traffic in real time. This is done by identifying where idle capacity exists to the required destination and redirecting transit traffic via the appropriate BT gateways, using intergateway tie routes. The powerful circuit reservation control also enables the NMC to redirect UK originated IDD traffic to alternative gateways, if required, providing spare capacity for transit traffic.

**New Year's Eve**

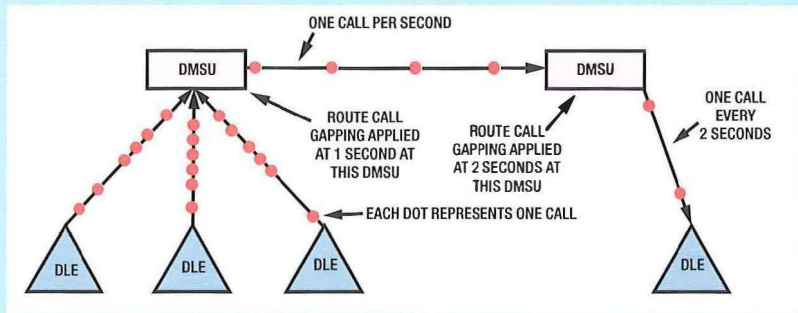
The week following Christmas is a very quiet time in the network, as many people remain on holiday. However, each year, at 23:45 on New Year's Eve, many thousands of calls are offered to the network, within a few minutes, particularly in and around Scotland and Northern Ireland. The demand is so heavy that over 200 protective controls are applied to protect the network from overload. These controls are timed to become active just before the surge of traffic hits the network; then as the traffic drops away they are automatically removed. The restrictive controls are activated at local exchange level to protect both the DMSUs and other local exchanges. On the international front; there is a surge of calls to Australia and New Zealand at this time and code gapping on the country codes 61 and 64 is applied at DMSUs to protect the international exchanges. Expansive controls are also implemented at the



*New Year's Eve / Day traffic profile*

**Protective controls—call gapping**

This control can be applied to all calls on a particular route (route call gapping) or, selectively, affecting only calls to a specified destination number or code (code gapping). It restricts the number of calls allowed to pass on to the next part of the network by prescribing a set time interval between successive calls.



international switches to maximise capacity to the popular destinations. Unlike Christmas, this peak is short lived and traffic starts rapidly decreasing at around 01:30 on New Year's Day. Checks are made to ensure that the NMC systems have correctly performed the programmed removal of protective controls and, finally, operational logs are brought up to date to bring to an end another successful festive season in the NMC.

**Conclusion**

Working over the Christmas period is hard on families as well as the

people involved. However, there is some satisfaction to be gained in the knowledge that the months of careful planning and the actions that were taken in the NMC on the day were successful in ensuring that customers have received the best possible service over the festive period. Two weeks later, when performance and financial data have been collated, a review meeting is held and the process starts all over again. Meanwhile, plans are already being put in place to manage the forthcoming Asian peak days celebrating Lunar New Year and Id Ul Fitr.

## Acknowledgements

Special thanks go to the following people for their assistance: Mike Brighty for reviewing the article and providing information on international traffic on Christmas Day, and Paul Hampson for providing the information on national traffic for a normal day, Christmas Day and New Year's Eve/Day.

## Biographies



**Stephen Morgan**  
BT Network and  
Systems

Stephen Morgan joined BT in Cardiff as a Technician 2A in 1973 after qualifying as an electrician in British Steel Corporation. He has performed a variety of duties within BT and during this time he studied computing at the University College Cardiff. In 1987 he was promoted to MPG2 and moved to Oswestry to work in the network management centre on real-time network traffic management. He had special responsibilities for phone-ins, the derived services network, Intelligent Network Phase 1 and Christmas planning. He took up his present role in 1995 and is currently involved in process and support for signalling network management and network synchronisation.



**Graham Barratt**  
BT Network and  
Systems

Graham Barratt left college in 1976 and joined BT in Telford as a TTA. After carrying out several duties within the Zone he moved to Shrewsbury in 1988 to take up a position on the System X exchange data build team. In 1990 Graham moved to Oswestry and has been a member of the international network traffic management team for five years. One of his key responsibilities is the implementation of international peak day and special event plans in the network traffic management centre.

# Learning to 'Work Smarter' in a Virtual Business

*Helping BT people to 'work smarter', the Workstyle 2000 Programme launched its new intranet web site in September 1996. An electronic library and discussion forum, it is a growing resource for BT people considering how or where to work. This article describes the development of the site's content and style.*

Figure 1—Business 'pavement café'



## Introduction

Innovation is both a problem and a solution. We have all seen how it can put businesses under new competitive pressures. Yet it also gives businesses new information and communication technology, allowing people to 'work smarter'. Some big businesses seize the technological initiative. They downsize, relocate, or try out some form of distributed working, and only then discover the know-how they need—not only need, but also have to retain, refine and supplement by continual organisational learning. They need information about work styles before opting for a radical change in working practices; for example, the type of work best suited to being done at home, how to manage a distributed team, and so on. Then there is a need to disseminate information on workplace characteristics; for example, the personal computer (PC) network facilities provided in the company's various office blocks, the choice of office environments, etc. And of course employees will need to know what information and communication technology is available to them, what it is good for, and how it can be obtained and maintained.

Figure 1 is an image of our time: the business 'pavement café': professional people working where and when it most suits them. And it is not just an artist's impression. As some BT people know for themselves, such a facility can be an everyday reality, helping them work more easily and effectively in their modern office block. People are discovering other options, too; for example, to work at customers' premises or at home, and still be able to link up with the largest computer

*'For the rest of this century, and far into the next one, the competitive battle will be won or lost by white-collar productivity.'*

Peter Drucker, *Wall Street Journal*, 1985

systems. By innovating, BT people are learning how to choose and use the latest buildings, communications and information technology; also, critically, how to organise and manage themselves in the new circumstances. BT people are seeking to develop, try out and monitor new technology and methods so as to adopt them selectively and with foresight. Success offers a 'win-win' situation with huge resulting benefits—for the participants, a higher quality of working life; for the company, competitive advantage. The company needs to keep learning from such developments. The challenge is to succeed in this, wherever its people may be.

To address such issues, BT launched a number of initiatives. One such, in 1995, was the Distributed Office Project, which recognised how an internal World Wide Web service can help. It can store not only relatively static company information but also the various lessons learned by the innovators. It can make all this freely available over the company network, and encourage users to discuss, correct or add to it. In effect, its files can be a technical prop to organisational memory, while its feedback processes contribute to organisational learning. Thus it was that the Distributed Office Project prototyped an intranet 'electronic library' on flexible working. The project's chief aim was to set up a lasting process which would help

The Workstyle 2000 (WS2000) intranet site is an information source on future working styles and everything needed to support them in BT. BT people will want to use WS2000:

- if they are managers exploring the possibilities that flexible working offers their teams;
- if they are interested in finding out more about flexible working and whether their own jobs and life-styles make it a suitable option; or
- if they themselves have information that will be valuable to BT people, and if they are looking for a way of making it known.

pass on to the whole of BT all that could be learned about the practical implementation of flexible working. BT recognised that such processes had to become generally accepted practice if it was to prosper as an effective learning organisation<sup>1</sup>. So the *Flexible Working Library* went on to be developed as the Workstyle 2000 (WS2000) intranet site, which is now available to thousands of people within the BT electronic firewall. This article describes how the site's content and style were developed.

### Who is Served by the Intranet Site?

The WS2000 web site designers aimed to serve 'all BT people' (to quote its Help page). This of course was an ideal. Apart from anything else, not everyone in the company yet had, or would need, a personal computer with an Internet connection. Yet the earliest potential users were very diverse. This was a challenge which the human-computer interface designer and the editor thought worth anticipating. Their success would minimise the risk of having to make costly page revisions later.

Users would predictably include people with widely differing equipment and network connections. They would range from mobile teleworkers using GSM phones and small laptop PCs to office-based designers with large-screen workstations on high-speed local area networks. Accordingly, the interface designer kept the main pages and their files deliberately small, so that even the remote laptop user would find them quick to download and easy to view. During development, pages were trialled with a monochrome laptop PC to ensure worst-case readability.

WS2000's basic user category is the individual. Although changing business circumstances and technologies may affect whole groups of people at a time, ultimately it is as individuals that people will most probably face choices of how and/or where to work. Hence the WS2000 section on 'Issues for Individuals', to support *informed* choices between alternative workstyles.

Managers are another category of user having different (if overlapping) perspectives; for example, as the source rather than the subject of changes in working practices. They may want help in planning the

***'By the end of September [1996], around 30 000 people in Networks and Systems are expected to be able to access BT Intranet, a pioneering service offering people all the benefits of the Internet in a secure and predictable environment. Intranet is to revolutionise the way we work by providing individuals with easy access to precisely the information they need for the task at hand - whenever they need it and wherever they are. It reduces costs, increases productivity and encourages smarter working.'***

From 'Three steps to an Intranet heaven', *On Demand*, Issue 5, August 1996

introduction of such changes; hence pages addressing 'Issues for Managers'. Incidentally, here are the pages on health and safety, as affected by flexible working. They are everybody's concern, of course, but it is the managers who have the employer's ultimate duty of care to inspect and authorise individuals' working arrangements. Predictably, as working arrangements become more flexible, managers will welcome being able to get the very latest guidance on-line.

Technical and organisational developers are a third user category, who will need something like WS2000 as a tool of their trade. They will be particularly concerned with the case studies reported, and the associated discussions.

### What the Intranet Site Should Contain

WS2000 is designed to be BT people's first choice for information on alternative working styles. And while some people will be generally browsing, others will be looking to answer particular questions. The site caters for both approaches. In the main contents page, simple section headings have been chosen, based on three propositions.

- 1 *In any work process there are only three fundamental components; the place, the technology and the human behaviour—hence the site's topic headings 'WorkPlaces', 'WorkSystems' and 'WorkStyles'. The WS2000 designers believe that these are stable categories and that populating them with reference material effectively creates primary courseware<sup>2</sup> in what amounts to an on-line tool for organisational learning.*
- 2 *People are interested to learn by example from others' comparable experiences—hence the fourth major category, 'Case Studies' in WS2000. This is the learning system's secondary courseware<sup>2</sup>.*

## *As WS2000 matures, users will be able to draw on an increasingly rich web of reference, case study and discussion material.*

3 *Users of any multi-user service will always identify further requirements after its launch* – so on-line feedback and discussion groups are provided for, to tell the WS2000 editors where and when users need other information.

As WS2000 matures, users will be able to draw on an increasingly rich web of reference, case study and discussion material. This combination of contents will offer truly practical insights about working smarter. To close the loop, the editors will review the discussions and feed new insights back into the web site, creating the equivalent of *tertiary courseware*<sup>2</sup>. All users may thus gain from the new understandings arising from the reflection and discussion of their more active and pioneering colleagues. Undoubtedly, a network of such intranet sites will improve BT's so-called *organisational memory* and better illuminate the way forward.

### **Style: Or Why not to be 'Cool'**

Readers who are already surfing the net will know the lure of 'cool sites'. If you have all the hardware and software they require, they can have more entertainment value than an arcade game. However, for an intranet site to be accessible and seem businesslike to many diverse users, the designer may have to adopt more of a 'de-tuned' approach. Certainly, it is critical for any web site that its users can easily find

*'Graphic/typographic designers have had to jettison the basic assumption underpinning their training, that they have complete control over the placement and form of text and graphics elements. One never knows how big the user's windows will be, whether they have windows or frames capability, or a particular plug-in, or what colours will be available.'*<sup>4</sup>

their way around in it. Easy navigation is achieved by careful design, both of the individual pages and the site as a whole. The development of the WS2000 style meant putting to one side the very latest features supported by the latest World Wide Web browsers. The following two accounts illustrate this.

### **The image map experiment**

Broadly speaking, the clearer users find a web site's structure, the more often they will find what they are looking for. This affects how they rate a web site's usefulness. Consequently, the first experiment was to offer users an image map (Figure 2) which (1) showed the site's information structure graphically on the home page and (2) allowed users to access any part of the site with just one mouse click. (And of course, graphics were 'cool'!)

The image map did provide a top-down perspective on the site. However, as a navigation aid for a site under development, it proved to be ineffective, since many of its entries were still only place markers. Place markers may be useful indications of intent, where information is yet to be

available but meanwhile clicking on such would-be hotspots just elicited an error message. Furthermore, the image map's file size meant that it took longer to download than the text equivalent (so it would have been doubly impractical to implement it on each page).

### **The frames experiment**

Many 'cool' sites use frames. It seemed old-fashioned not to follow suit. Frames were introduced with Netscape 2.0 and have been the most contentious of the new features offered by the browser – so much so that the most significant difference with the later version of Netscape has been the introduction of greater user control over the frames.

Two (later, three) frames were used on each page. To the left, a narrow vertical 'toolbar' style frame contained labelled icons (Figure 3). These were linked to give immediate access to any of the main areas of WS2000. At the top of each page was a header bar. The third frame was the main body of the browser window, displaying the particular information the user was trying to access. Whenever the user clicked on a link,

Figure 2 – Prototype image map of contents page

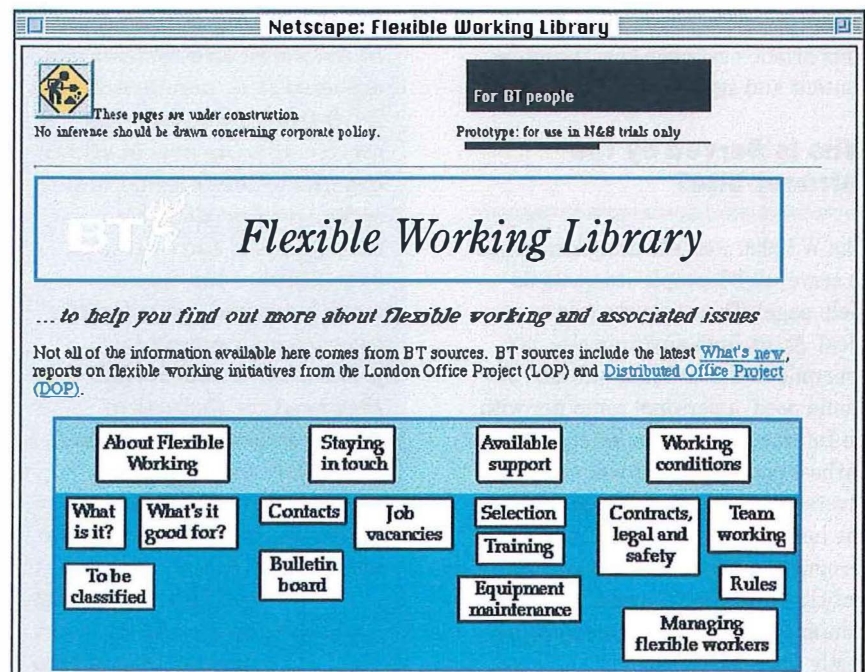




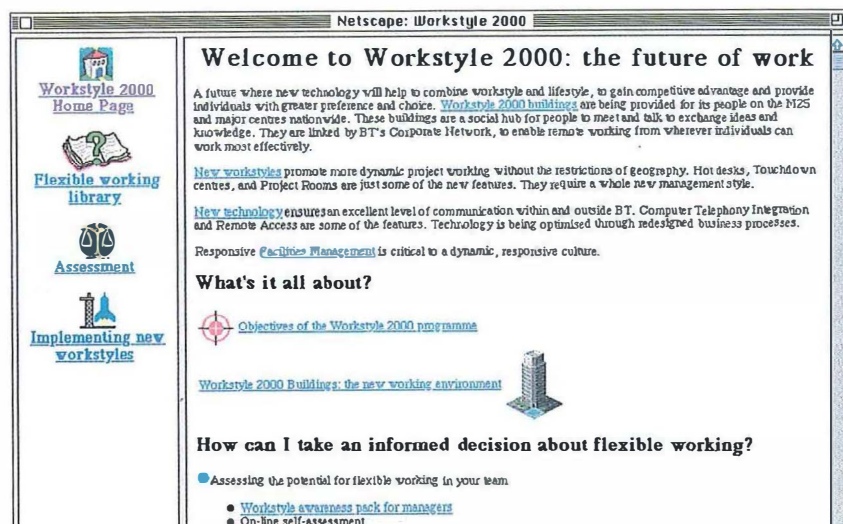
Figure 3—Screen dump of prototype frames page

only the required new information for this frame had to be downloaded; a welcome efficiency for users with slow links.

Providing a common frame could foster a common perception of the site. With little development effort it supplied a common style with the bonus that users could get about the site with the least waiting for pages to download. As mentioned earlier, this could be important for flexible workers who had relatively slow (for example, mobile) communications links. A further bonus was that where non-BT pages were accessed, the 'frames' browser alerted the user by generating a new window.

At this time, a move was made towards a more fully defined style for the pages. In addition to the disclaimer at the top, each page had a footer incorporated as standard (not implemented as a separate frame). The footer named the page owner (responsible for maintaining the accuracy of the page's information) and invited the reader to e-mail that person, using a convenient 'mail-to' hyperlink, with any corrections or suggestions.

Unfortunately, interactions with frames ran counter to the (then prevailing) Netscape 1 users' understanding of how the browser worked. Users of the prototype had problems with navigating back and forward between pages; the previously familiar 'Back' and 'Forward' buttons no longer operated in the same way when frames were used. As luck would have it, at just this time these same difficulties were highlighted in a seminar given by internet design guru Jakob Nielsen, Chief Engineer at Sun Microsystems. Furthermore, it became clear in discussions with the Royal National Institute for the Blind that the frames interface would additionally disadvantage visually impaired users. Conventional text-to-speech equipment was unable to interpret information held within frames. So the frame-capable version of WS2000 was dropped.



### Present styling and navigation

Convention would have dictated a site composed of pages *separately* owned by various information providers. With their contact details at the foot of each page, these page owners would thus become directly accessible to users with suggestions or queries. However, to achieve styling consistency as soon as possible, the design strategy favoured *centrally* owned pages. Although this meant taking on an extra editorial burden, the decision produced pages that, from the start, gave users a completely consistent visual distinction between the site's validated core and its more peripheral information. Since WS2000 amounts to a cross between a within-company advisory service and an employees' bulletin board, this visual distinction is an essential feature, clearly indicating the status of what is on screen.

To make a distinctively BT intranet site meant conferring with the company's Corporate Relations specialists, and developing pages which drew upon their existing style template. The result is a page layout that uses white space generously, in the best conventional typographical tradition. All the core pages are topped by the same white-on-black banner bearing the site name in a corporate script together with the BT logo. In the page body, fonts are used consistently for the page titles and subject headings.

The information is structured in a shallow hierarchy, minimising the number of mouse clicks needed to access the desired information. Jakob

Nielsen has since described this as the *inverted pyramid* approach. The aim is that all the site's major categories of information must be visible at the top level. (In practical terms, *top level* means the *page that users are likeliest to bookmark*; that is, the main contents page). Clicking on a selected contents page entry then displays a finer level of detail, usually a topic contents page. Clicking within that should then take the user to the required page of detailed information.

Paging between sections is made easy with a button bar on each core page. The button bar contains hyperlinks to take users to the main contents, to an alphabetical index, to a suggestions dialogue and to Help. Just as in a book, the structured contents and alphabetical index are two separate listings of the information available within the web site. The 'Help' and 'Suggestions' functions are rather more unusual, and are discussed below.

### Help with 'real world' work-related problems

We chose the label 'Help+' for the help button because the help pages offer more than just hints on using the WS2000 web site. They point to sources of help for those with real work-related problems. A whole range of problems, named and arranged in a structured index, link users with the contact details of suitable units inside the business or with outside agencies. Wherever possible, the contact details are those which are independently and routinely kept up to date by the company's directory-maintaining unit.

## A responsive managed resource

It is tempting to claim that the responsiveness designed into WS2000 was a virtue. In truth, it is a simple necessity born of two features of the working practices domain that are worth some discussion: its scope and its subtleties.

### Scope problem

For the most part, of course, we all like to complete what we set out to do. However, the scope of the working practices domain is broad and the WS2000 developers and editors are few; so providing a complete service on the subject is not so easy. Moreover, since the organisation that WS2000 serves will continue to evolve, completion may always be unattainable. Certainly at the launch of WS2000 the main contents page offered a seemingly complete four-part framework: 'WorkStyles', 'WorkPlaces', 'WorkSystems' and 'WorkStudies'. But even then, user feedback and suggestions were needed; hence the 'Suggestions' button on each page. There were questions unanswered; such as, could readers easily understand the conditions needed to make alternative workstyles viable? Could they find enough information about BT buildings actually on-line yet? Were they taken by surprise to find furniture classed as part of 'WorkSystems' along with IT and telecommunications technology? To which other intranet sites did they want links? All the time, the editors need the users' help to acquire new material and spot anything out of date or misleading that needs to be replaced or deleted. Users willing to make suggestions remain critical to the vitality and completeness of WS2000.

### Subtleties problem

The subtleties of 'working smarter' require more of WS2000 than good housekeeping. After all, where WS2000 seeks to inform BT people about flexible working, it is actually

talking about *generic* business solutions. Yet the efficacy of a generic solution depends on how people apply it to a particular situation. And they must understand that situation well to succeed. Here is a critical requirement. Gaining that understanding requires practical appraisal, not just theory. This has been identified elsewhere. For example, in a recent management text, John Seely Brown and Paul Duguid wrote that

*'... education, training and technology design generally focus on abstract representations to the detriment, if not exclusion, of actual practice. We, by contrast, suggest that practice is central to understanding work. Abstractions detached from practice distort or obscure intricacies of that practice. Without a clear understanding of these intricacies and the role they play, the practice itself cannot be well understood, engendered (through training), or enhanced (through innovation).'*<sup>3</sup>

Brown and Duguid went on to cite a graphic case in a service organisation. Its technicians sustained their ability to diagnose complex equipment problems mainly by swapping stories, building 'shared accounts'. By contrast, their formal training was based on simplistic abstractions which they found unhelpful in the face of the sophisticated yet unpredictable machines they had to work with. Brown and Duguid drew an inference which transfers directly to the flexible working domain:

*'Learning is fostered by fostering access to and membership of the target community-of-practice, not by explicating abstractions of individual practice. Thus central to the [learning-in-working] process are the recognition and legitimization of community practices.'*<sup>3</sup>

So it is that WS2000 collects stories and is a forum for discussion as well as a library of pilot study findings. As the site matures, it will

thus give users an enriched and more realistic picture of alternative 'smarter' ways of working, tuned through practice to actual and particular circumstances.

## Future Prospects

Interpersonal networking in its various forms has always been important for sharing problems and getting things done on time and to quality. This will surely benefit from intranet sites like WS2000 starting to maintain contact lists of experts. Of course the managerial challenge may then be to meet a huge and previously unknown demand for expert advice; but such an outcome should be a healthy lesson for any organisation that is interested in unlocking its potential.

When it was just a prototype, WS2000 was little more than a collection of reference material. Now it is in use, it is beginning to look a little more like the open-learning resource that was intended. However, as the site grows, people coming to it for the first time may be less content with the teach-yourself approach. Indeed, as user numbers climb, there may be a sustainable demand for on-line tutorial support, and this may herald some form of integration with the company's development and training services and other software resources. Such a development would be no surprise. It has already been noted elsewhere that the World Wide Web 'is increasingly being integrated with the tools that we already use.'<sup>4</sup>

Naturally, the BT intranet continues to grow, offering BT people an increasingly informative web of connections. More links will undoubtedly follow. Depending on usage and user feedback, WS2000 will probably reach maturity in two or three years. As part of the BT intranet, WS2000 will surely help significant numbers of BT people learn about and discuss flexible working, thence to progress with greater self-awareness and confidence to the working styles of the new millennium.

## References

- 1 SENGE, P. The Fifth Discipline.
- 2 FOWLER, CHRIS; MAYES, TERRY; and BOWLES, BERNARD. Education for Changing Times. *Br: Telecommun. Eng.*, April 1996, 15, p. 32.
- 3 BROWN, J. S., and DUGUID, P. Organisational learning and communities-of-practice: toward a unified view of working, learning and innovation. Chapter 9 in HARIDIMOS TSOUKAS, *New thinking in organisational behaviour: from social engineering to reflective action*. Oxford, Butterworth-Heinemann, 1994.
- 4 SHUM, S. B. The missing link: hypermedia usability research and the Web. *Interfaces*, No. 32, Spring 1996, British HCI Group, Report on the British HCI Group symposium. <http://kmi.open.ac.uk/~simonb/missing-link/>

### Accessing WS2000

World Wide Web users outside BT can access BT's public pages on teleworking at <http://www.labs.bt.com/innovate/telework/index.htm>; but WS2000 itself is a service for BT people only.

BT people familiar with the World Wide Web can find WS2000 via two high level pages: the BoAT home page (click on 'Programmes') and the N&S Information Map (click on 'Employee Communications').

BT people who are not yet connected to the World Wide Web themselves might consider finding a colleague who is, and browsing the site that way in the first instance.

### More links?

WS2000 users should use the site's <Suggestion> dialogue. Others who wish to contribute or to have their own World Wide Web pages made accessible from WS2000 are invited to e-mail Workstyle 2000's Steve McPherson ([mcphersw@boat.bt.com](mailto:mcphersw@boat.bt.com)).

## Biographies



**Roger Martin**  
BT Networks and  
Systems

Roger Martin is with the Flexible Working Team, BT Human Factors. He is the current leader of his unit's Organisations Learning Domain. He has degrees in psychology and many years' experience in large organisations working with computers and people. In the late 1980s, before joining BT, Roger managed a human factors consultancy team in large-scale government computerisation projects. In the early 1990s, he led a BT team which prototyped and field tested a new graphic user interface for the BT operator assistance service. He set up the boundary-crossing professional network HOFSIG (Human & Organisational Factors Special Interest Group) in 1994. He then led the Human Factors input to the Distributed Office Project (1995/96) and its vision of flexible working in the year 2000. For this, he had the idea of a managed intranet site on Flexible Working, the subject of this article. Scenario planning (helping to anticipate alternative futures) is a complementary activity. The most recent deliverable that Roger has co-authored is the BT Future Worlds intranet site.



**Peter Hornsby**  
Student

Peter Hornsby is a student who worked in the BT Human Factors Unit, 1995/96. This was his 'year out' from a first degree in Information Technology and Human Factors at Loughborough University. Peter worked on flexible working, and looked mainly at organisational learning and the human factors of World Wide Web design.

# How 'Smart' Can We Get?

## Technology and Business Opportunities for Smart Cards In The Age Of Electronic Commerce

*The last two years have seen a rapid increase in the world-wide growth of the smart card industry. This article examines the main technical and business forces driving this growth, and shows how BT and the telecommunications sector are playing a leading role in the evolution of the industry.*

### Introduction

It may have been true at one time that smart cards were only important for technology buffs and inventors. In fact, smart card technology has been around a long time without causing much of an impact in many markets—nowadays, though, we can hardly open a newspaper or magazine without the announcement of a new smart card scheme either being launched or planned somewhere in the world.

It seems that smart cards have finally come of age. Be it electronic cash, access to the information superhighway and multimedia services, cellphones on the global system mobile (GSM), or the world of electronic commerce in general, the smart card is seen as a central and indispensable component in delivering these services to customers. Why this should be the case, and the role BT is playing in this rapidly expanding market place, is the subject of this article.

### A Brief Overview of The Technology

Readers will be familiar with the traditional definition of smart cards—'a silicon chip embedded in a credit card size piece of plastic'. The International Standards Organisation (ISO) use the term *integrated circuit (IC) card* collectively to describe the full range of devices containing a silicon chip within an ISO identification card. To qualify as being really 'smart', the silicon chip must be in the form of a microproces-

sor, giving the card the processing power and characteristics of an 8-bit microcomputer.

The smart card family of products comes in two different technology guises: contact and contactless. Contact cards are easily identified from the familiar gold or silver 'contact' plate located on the front of the card (Figure 1). Power is supplied to the chip, and the card activated, by direct contact between




Figure 1—The gold contact plate

card and IC card reader (which might be free standing or embedded in a terminal). Contactless cards, as the name implies, do not rely for power/activation on direct physical contact with the terminal. Communication between card and terminal is achieved by the use of a high-frequency radio antenna embedded in the card.

Contact cards are the most commonly used type today and are the main focus of attention in this article. However, readers should be aware that contactless cards are growing in importance, particularly as a ticketing medium in urban transport systems around the world.





Figure 3—Smart card memory areas

### Standards

The ISO has mandated a series of international standards for smart cards covering, in particular, the physical dimensions (ISO 7816-1) of the card, location of the contacts (ISO 7816-2), transmission protocols (ISO 7816-3), and command and file structures (ISO 7816-4).

An ordinary smart card is the same size as any common plastic credit or debit card (Figure 2). The six contact interfaces are shown in Figure 1. The important ones to note are:

- power supply—used for supply voltage;
- reset input—signal line used to initiate state of the chip after it is powered up;
- clock input—clock signal to drive the IC logic at, nominally, 3.57 MHz;
- serial data input/output—signal line for sending commands and data between card and terminal; and
- ground—0 volts connection.

### Silicon real estate and different card types

Miniaturisation is, not surprisingly, a key variable in the economics of smart card manufacture. In general, the less silicon space needed to deliver the required specification, the more economical the cost. This space is, therefore, frequently

Figure 2—The dimensions of a smart card

referred to as *real estate*, with the obvious suggestion, (a) that there is a limited amount of it, (b) there are different types of it, and (c) occupation of the space carries a cost. The die size is restricted to a nominal 25 square millimetres due to considerations of the mechanical strength of the die.

In smart card terms this real estate is taken up by memory, which comes in three main types (Figure 3). Each of these is used for a different purpose, and the combination determines the type of application possible, and the cost. These memory types are as follows:

- **ROM**—nonvolatile read-only memory is used to hold the card's operating system and other data that is 'plumbed in' at the manufacturing stage and cannot be subsequently altered. This is the lowest-cost memory type because it occupies, relatively, the least amount of real estate on the silicon substrate. Application developers will, therefore, try to program as much as possible of the fixed data requirement into ROM.
  - **PROM**—nonvolatile programmable read-only memory comes in two different types: EPROM (erasable PROM) and EEPROM (electronically erasable PROM). EEPROM is the most commonly used type of memory in this category and is relatively expensive in terms of real estate. The advantage is that the memory area is not fixed and can be erased during the life cycle of the card and rewritten many times over.
  - **RAM**—volatile random access memory is used to perform calculations and act as a temporary store. In terms of real estate, this is the most expensive type of memory.
- Simple memory cards are generally used in older-style prepaid-disposable phonecard applications (not used in BT). They are also used in applications where security requirements are low and the need to protect information on the card is not an issue.
  - Memory cards with hard-wired logic are used in modern prepaid-disposable phonecards with RAM, ROM and EEPROM (typically, small memory). More generally, this type of card is card issuers' choice where there is a requirement for data to be processed in a preset fashion or higher levels of security, enabled by the processing of hard-wired cryptographic algorithms, are needed, and the secure personal identification of the card user is not a priority.
  - Microprocessor cards that combine ROM, EEPROM, RAM and a CPU are used in applications like electronic cash and GSM, both of which are discussed later in this article. The requirement is for (a) high, issuer-definable levels of security, especially as regards access rights to data stored on the card, (b) dynamic processing of data on the card, in particular to create cryptographic messages like digital signatures, (c) very-secure user authentication on-line and off-line, and (d) in some cases, a multi-application capability.

Modern microprocessor-based cards can contain up to 16 kilobytes of EEPROM, 24 kilobytes of ROM and 512 bytes of RAM. In line with the wider advances in microprocessor technology, this capability is expected to double within the next two years.

Most IC manufacturers offer a range of smart cards with, for example, more RAM and/or a hardware cryptography engine, usually at the expense of less EEPROM.

This capability is translated into the following different card types:

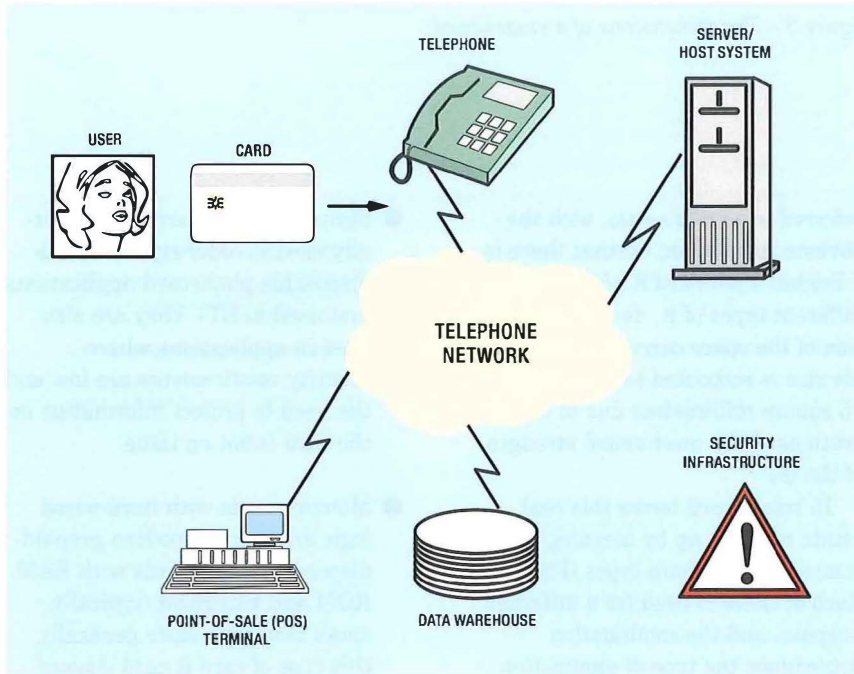


Figure 4—The discrete parts of a smart card system

The principal areas where smart cards have successfully been deployed as business solutions are:

- secure, portable cashless-payment systems by banks/telecommunications companies;
- access control and identification medium for electronic commerce, digital mobile communications and multimedia; and
- multi-application card service platforms for home banking/teleticketing.

### A secure cashless payment solution

Consumer and business confidence in payment systems depends ultimately upon the real security of the system. This is particularly true of the new electronic cash systems like Mondex, where real monetary value is being transmitted over telephone and wireless networks. (BT's role in the evolution of Mondex is explored in detail below.) Smart cards have become the chosen solution of most major payment systems organisations—VISA and Mastercard have announced the migration of their global card businesses onto a smart card platform. Both these organisations, as well as Mondex and Amex, have chosen smart technology to launch their electronic purse schemes. BT has chosen smart technology as its new platform for delivering the BT Phonecard service, which was launched to much acclaim nationally in May 1996.

The security provided by the smart card rests on a number of principal differentiators from the magnetic stripe card alternatives traditionally used by the payment systems organisations. These differentiating features include:

- tamper-resistant hardware,
- capability to perform cryptographic security functions,

This brief summary of the different card types and the memory permutations on the card illustrates clearly some of the principal factors that have driven the growth of the smart card business. The capability to store, access and manipulate or process data on the card **securely** in an economic way is what distinguishes the smart card from other forms of card technology available today.

### Cards as components of a wider system

Interesting as the technology of the cards might be, discussions on the development and use of the technology often ignore the fact that the cards are but one component of a much wider system infrastructure required to deliver service to customers (Figure 4). Later, some examples are explored of how complex this system can be, but it is sufficient to point out now the most common elements of existing smart card implementations around the world, where development activity on each of the components is required to enable the system to deliver an effective service:

- cards;
- terminals (for example, telephones, point-of-sale devices, vending machines, parking meters, and PCs);
- networks;

- server/host systems (for example, to operate the customer application and/or to manage the population of terminals and security modules);
- data warehouses (for example, in mass transit systems for tracking transactions and customer profiling); and
- security—primarily to guarantee end-to-end security of the system.

This overview of the technology has necessarily only skimmed the surface of what is becoming a global multi-billion dollar industry. But why the fuss? What exactly can our friendly silicon chip, nestled in the card, actually do for the millions of customers who currently own them, and the millions more who will do so in the future?

### Smart Card Technology as the Chosen Application Solution in the 1990s

The reason for the astounding success of smart card technology around the world is that it manages to offer market place solutions to major operators across a wide industry spectrum, from banks and financial institutions, to communications companies, to transport operators and, more recently, to governments and public utilities.

Figure 5—Public-key cryptography and digital signatures

- secure portable storage medium to carry personal data, and
- uneconomic to counterfeit.

### Access and identification in the world of electronic commerce and multimedia services

The growth of electronic communications and on-line multimedia services is creating the demand for delivery of, and access to, information 'any time, any place, any how'. Many of these services (video-on-demand, satellite TV, Internet, home banking, GSM) are also subscription based, and rely on the user establishing his/her usage rights before being logged on to the service. Content with high added value means that a premium will be placed on the security of the access and authentication mechanism. Smart card technology offers an integrated solution to this service development problem, principally on account of the secure authentication process that the card can deliver. Placing biometric templates onto smart cards will strengthen this proposition further.

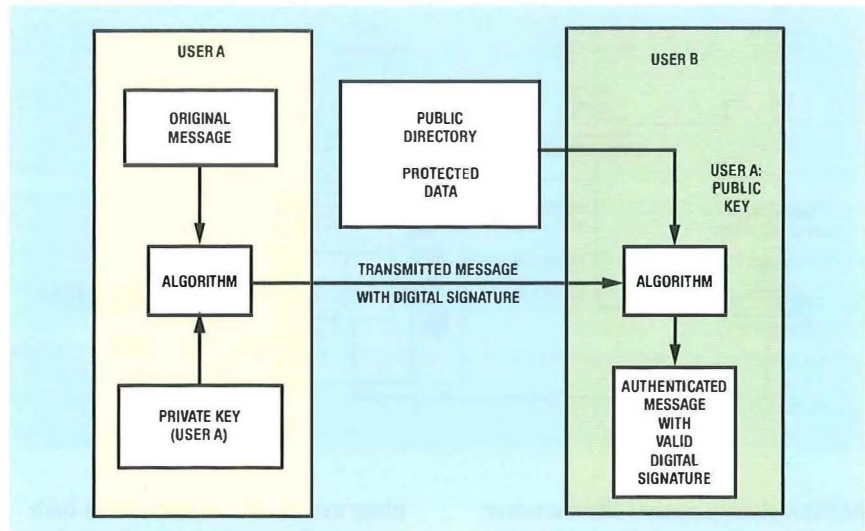
#### Smart cards and public key cryptography

Many of the new electronic commerce services will rely on public-key cryptography to deliver digital-signature and non-repudiation services, including secure electronic messaging. The principles of public key cryptography are important to grasp in the context of current discussions on the use of smart cards.

Under such a system, each user is allocated a public key, which is published in a directory (much like a telephone directory) and a private key, which must be held securely and known only to the user/owner.

#### Confidentiality

A message (or any other piece of data) locked with a private key can be unlocked only with the associated



public key, and vice versa. If user A wishes to send a message to user B, A can encrypt the message using B's public key. Only B can then unlock this message, using his/her matching private key.

#### Data integrity and authentication

Digital signatures are the logical inverse of the above process. If user A wishes to prove his or her identity to user B, A signs the message with his or her private key. (In practice, A would summarise the message into a message digest using a hashing algorithm, then sign the hash of the message.) B can then confirm that the message was indeed signed with the correct private key, by validating the signature using the associated public key, which he or she can recover from a public key directory. If the validated signature is correct, B can then accept that A did indeed sign the message and B can tell if the message has been altered by a third party as the signature would no longer match the message. Figure 5 summarises the process graphically.

In a system such as the one described above, the process of securely generating, distributing, and storing the public and private keys is known as *key management*. The principal question for users of such systems is where the individual's private key is to be stored. The key must be secure from hacking and tampering and, if it is to be truly useful, must be portable; that is, tied to the individual and not to a terminal. The solution to this problem is likely to be that private

keys will be stored on a smart card—in some cases where message integrity requirements are very high, the card will also generate the user's private keys. Alternative solutions, primarily software centred (for example, storing keys on the hard disk of a PC), are considered to be far less robust than the smart card option.

However, this is not all—with usage rights to many new services attached to the card/individual, not the terminal, users have **mobile and portable** access to services. They can then transact business anytime or anywhere at their convenience, certainly a powerful customer proposition whatever business you happen to be in! Users will also have a **mobile and portable database**, which can carry anything from a stored telephone number directory (GSM), a transaction log (Mondex), or favourite web sites (a future service). Access and service provision can be targeted at individuals within a household, rather than to the household as a unit, with billing and payment similarly enabled.

#### Multi-application cards, customer choice and Java

Most people would agree that they carry too many plastic cards. They would also probably agree that they don't want all their requirements bundled onto one card either. A rationalisation of the card portfolio will result in individuals carrying, perhaps, three or four smart cards in their wallets. But what applications and services should the service provider package onto the card? The




Figure 6—The BT Payphone system

answer, clearly, is that the customer should be allowed to choose how he/she configures the services on the card, and he/she should be allowed to change this configuration to meet changing needs and circumstances.

Traditional smart card operating systems have tended to be proprietary systems and have not provided the flexibility to allow the dynamic downloading of new applications onto a card during the card life cycle. Lead times to develop new applications have been too long, with a correspondingly slow time to market.

An important recent industry development has been the announcement by Sun Micro Systems in October 1996 that they have developed the *Java Card API Specification*, which, according to the press release, will 'bring the power of Java to hundreds of millions of cards world wide'. Such applications can, of course, also be downloaded to card terminals to change important parameters. This initiative is aimed at providing the flexibility, platform independence and swifter development times demanded by a market

place used to the rapid product turn around of the software industry. It is not the only such initiative in the market place, but it clearly illustrates once again how the industry is sensing that the market place for smart cards will evolve in the future.

#### *Some volume forecasts for smart card sales*

Translating these business opportunities into volume and revenue forecasts is no easier for smart cards than for any other service offering in a market place where the rate of change is so fast. However, industry analysts have achieved a fair degree of consensus over the mix of potential applications, which ones are going to be the most successful, and the fact that the market is definitely going to grow. Table 1 summarises the shape of the market in volume terms, comparing the situation now with a projection for the year 2001.

#### **BT at the Leading Edge**

BT is playing a leading role in the development of infrastructure and

applications for smart card based services in the UK. Before discussing some of the BT smart card service offerings in more detail, it would be useful to look at the breadth of BT's group wide activity in card based services. The range of activities spans a complete spectrum, as befits an organisation like BT seeking to address the market needs of a very large and diverse customer base. This base is served by:

- prepaid and account-based calling cards (BT Phonocard and BT Chargecard);
- cards covering fixed and mobile networks;
- electronic cash and credit and debit card systems (Mondex, Credit Call); and
- a range of access terminals and devices (payphones, screen phones, cellphones).

The card-services platform built to accommodate this range of cards will be extremely important to the on-going success of the business. The cards essentially provide the secure access and payment mechanism to BT's network service offerings in the public (away from base), home and business environments. They also serve as a powerful BT branded tool for marketing services to customers and profiling usage patterns.

The following section examines how BT has implemented some of these services using the smart card technologies discussed in this article.

#### **BT Payphones**

60 000 BT public payphones now accept the new BT smart Phonocard, which replaced the old optical technology with which the service had been launched in 1982. This new smart card platform is an important strategic asset for BT because it places a national network of smart card reading terminals in public-access locations.

Table 1 The Smart Card Market in Terms of Volume

Application	1996 Volumes	2001 Volumes
<b>Memory Cards</b>		
Phonocard; Vending; Parking	458 million	1.3 billion
<b>Microprocessor Cards</b>		
Finance	45 million	351 million
Health Care	59 million	243 million
Pay TV	26 million	81 million
Transport	7 million	189 million
Global System Mobile (GSM)	30 million	81 million
ID Cards	20 million	215 million
Other	20 million	200 million
<b>Totals</b>	<b>665 million</b>	<b>2.7 billion</b>



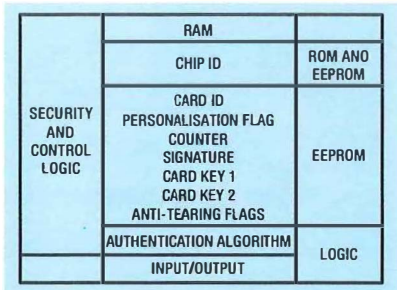


Figure 7—Memory map of BT's eurochip card

The new Phonecard system (Figure 6) integrates payphones, cards, security modules and a new management system.

### Payphone terminals

Four types of smart-card payphone were deployed on launch of the system in March 1996.

- From Landis and Gyr: the new payphone 2000 which accepts smart cards, magnetic-stripe cards, and coins.
- From GPT: two payphones, including a card-only version and a coin and card version/
- From Schlumberger: a new card-only payphone.
- From Inventec: a new low-cost card-only payphone.

### Smart cards

Initially the cards will be supplied by GPT and Gemplus and will use a Siemens 4438 chip. Payphones have chosen the *eurochip* technology for their card. In addition to the cryptographic-card-authentication mechanism, the card also has two secret keys, allowing BT, if it wishes, to offer secure segregated space on the card to a potential third party service provider. A memory map of the eurochip card is shown in Figure 7.

### Security modules

The use of the eurochip card requires a security module (SM) which performs the following functions:

- authentication of eurochip cards;
- secure collection of value from cards; and

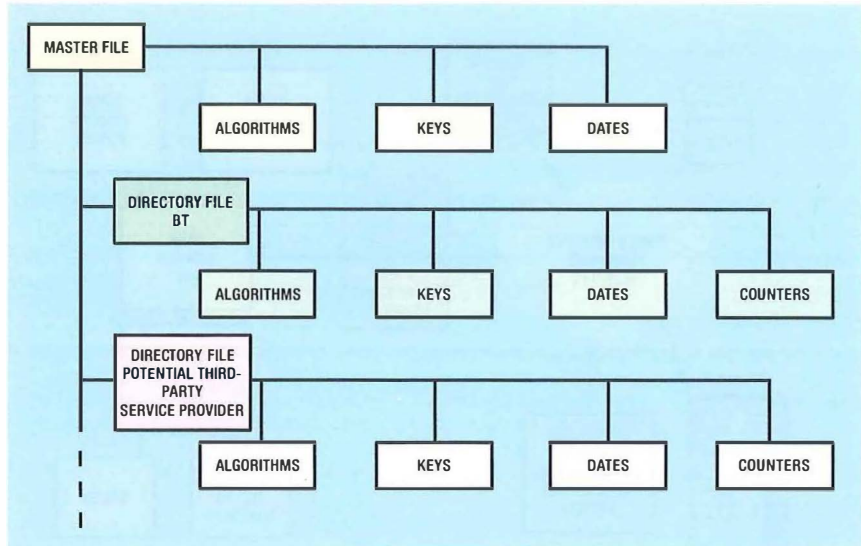


Figure 8—File structure of a BT security module

- secure communications with the management system, including accumulated counter readings as part of a secured audit trail.

The file structure of a security module is shown in Figure 8.

### Management system

The integrated payphone management system (IPOMS) has been enhanced to give it the functions required to manage the eurochip cards and the SMs. These new functions include:

- management of encryption keys;
- collection of accumulated counters;
- collection of statistics on card usage, including card ID, card type etc.
- fraud management
- provision of marketing information
- control of the validity of user cards by enabling new authentication keys and disabling old ones.

The building of this new smart card platform is an important milestone in the history of BT Payphones. Combined with the major payphone smart card activities on Mondex (see below) the once simple payphone has graduated to becoming a financial transaction and payment terminal for a whole range

of personal and business applications.

### Mondex and electronic purses

#### Electronic cash comes of age

Mondex is an electronic cash scheme based on smart card technology and was launched in the UK in Swindon in July 1995 jointly by National Westminster Bank, Midland Bank and BT. Electronic cash is, as the name implies, an electronic substitute for notes and coins. Mondex is by no means the only example of such a scheme—indeed there is a worldwide trend, being led by banks and telecommunications companies, towards the adoption of electronic cash/purse schemes. A recent estimate was that there are some 35 trials or national implementations across five continents. In the UK, VISA International announced in October 1996 their plans to launch a trial of VISACash in Leeds during 1997. American Express has also recently announced its arrival on the scene by taking a licence to use the European Proton scheme, pioneered by the banking sector in Belgium.

The main business drivers for this trend include:

- move away from branch banking to home/direct banking;
- very high overhead costs, to the national economies, of handling notes and coins;

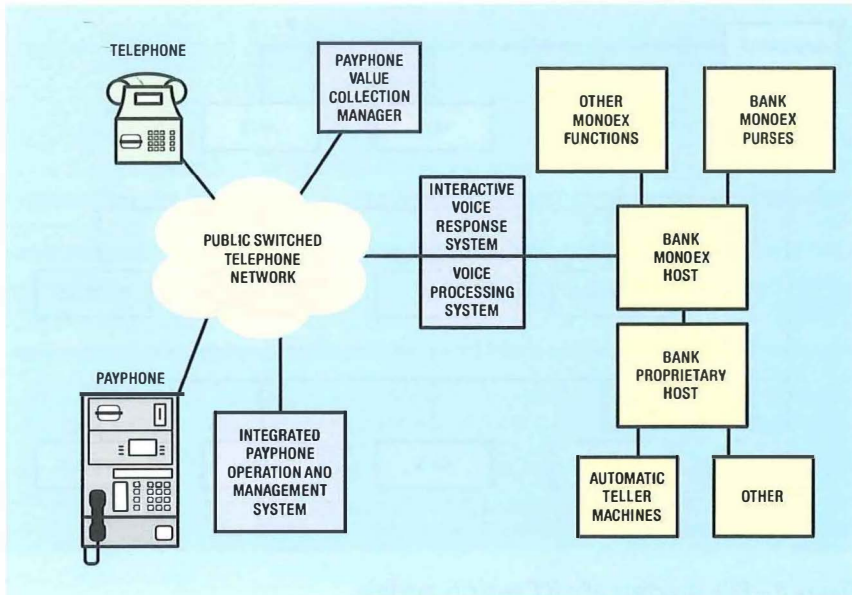


Figure 9—Mondex systems architecture

- growth of on-line commerce and electronic trading;
- the need to increase security of unattended retail outlets (pay-phones, vending machines etc.);
- need to offer merchants a card payment instrument for very low transaction values;
- need to offer payment at point of sale (pay as you go) for home based teleshopping/telebanking and multimedia services; and
- customer convenience factors.

### **Mondex as global payment scheme**

It had always been part of the aspirations of the Mondex founders to become a global multi-currency payment scheme. In November 1996, this vision took a step nearer to being realised when Mastercard International announced that it was planning to take a 51 per cent stake in the equity of Mondex International. It remains to be seen what the commercial implications of this arrangement will be, but what is certain is that Mondex now has a worldwide distribution outlet for its cards of some 22 000 financial institutions and, potentially, 13 million accepting outlets.

### **BT and Mondex**

For BT, Mondex offered a model for the future of electronic cash that

fitted BT's own vision for adding value to its basic telephony offering. The ability to transfer electronic cash securely over a telephone network was a compelling new proposition for bank and BT customers. It offered the prospect of being able to download and upload cash cheaply across the PSTN from the comfort of a payphone or the comfort of your own home 24 hours a day, 365 days a year. In particular, customers can:

- reload cards with value direct from a bank account at a payphone/homephone;
- deposit value into their bank account from a payphone/homephone;
- check their bank balance remotely;
- transfer money across the PSTN from one telephone device to another;
- read the balance on their card, lock/unlock the card at a Mondex enabled phone; and
- pay for their calls at Mondex-compatible BT payphones.

The new revenue opportunity for BT lies, of course, in the use of the core telephone network for a new range of financial services transactions, enabled by the new telephone devices especially developed for Mondex.



Figure 10—A Mondex public payphone

As part of its contribution to the launch of Mondex in the UK, BT has done some pioneering work in the field of secure value transfer over public communications networks. The telephony architecture for the project is shown in Figure 9. The principal components of this system in the context of value transfer are:

### **Public and retailer payphone**

These payphones were specially developed for the launch of Mondex (Figure 10). These two types of payphone are very similar in terms of the Mondex functionality they offer. (Note: the screen and soft button configuration are modelled on existing designs for bank automatic teller machines (ATMs)). The main difference is the more rugged construction of the public payphone. The display screen is analogue display services interface (ADSI) protocol compatible (Bellcore standard) and can support ADSI type services via menu selection. The payphones can be used in normal voice mode and can be switched into data mode (V.22bis) when transfer of value is required. The payphone contains a Mondex chip/purse of its own, in addition to being able to accept a user Mondex card. For the Mondex launch, BT has deployed 150 public and approximately 50 retailer payphones.

### **Payphone 'value collection' manager (PVCM)**

This system is responsible for collecting Mondex value from the



Figure 11 – A Mondex Homephone

internal card/purse within the payphone. This value is accumulated from calls made at the payphone and paid for by the user's Mondex card. The system also manages the clearing of error logs within the purse. Calls to the PVCN are initiated by the payphone.

### Homephone

Homephone is a new domestic telephone device that accepts user Mondex cards and features a display screen driven by menu selections. Like the payphones, the Homephone has an ADSI compatible display screen and can be used in both normal voice mode and be switched automatically to data mode for the value transfer function (Figure 11). For the Swindon launch, BT has deployed approximately 600 Homephones in residential premises, and 500 in the back offices of banks and retailer premises (so retailers can download their daily Mondex takings direct to their bank accounts at the end of each day's trading—or at any time during the day, if they wish).

### Public switched telephone network

#### Integrated payphone operation and management system

This system is the standard remote management system employed by BT for managing self-reporting payphones. The system controls internal-payphone software parameters and gathers management-information statistics. The system has been suitably upgraded with new interfaces to accommodate the reporting and management information requirements of the Mondex payphones.

#### Voice processing system (VPS)

This is the interactive gateway for telephones into the bank Mondex host system. The VPS is ADSI compatible and controls the display screen on telephones to provide the menus and voice prompts for telephone-to-bank host dialogue. The VPS also contains the V.22bis/V.42 protocol capability which is used by telephones for Mondex value transfer.

#### Bank Mondex host system (BMH)

This system is the intelligent controller associated with Mondex value transfer in the banking environment. When value transfer is required either to or from a Mondex chip/purse within telephones, the BMH manages the transaction via the VPS and controls the debit/credit of the associated bank account via the bank proprietary host system (BPH).

#### Bank Mondex purses

There are a number of Mondex purses (Mondex cards) used by the banking systems to deal with user transactions from telephones, ATMs, POS terminals etc. These purses are repositories for Mondex value held within the banks.

#### Bank proprietary host system

These are the various existing bank systems associated with the holding of customer accounts.

Mondex represents an important milestone in the evolution of the modern information age. Here is a living example of how the convergence of telecommunications, finance and microprocessor technology have contributed to improving the lives of ordinary citizens. By giving people the means to transfer and spend cash securely over a public telecommunications network, the scheme will enable people to conduct their business in a manner that suits them, again, 'any time, any place'.

Figure 12 illustrates this proposition. I live in London, my daughter lives in Scotland. I wish to send her

some money, so I call her on my Mondex Homephone and she answers with her similar device. We both insert our Mondex cards into our respective telephones and I transfer £50 across the PSTN direct to her card in her Mondex telephone. She can now spend the money in any way she pleases at any Mondex accepting outlet. The value transfer process took no more than 90 seconds.

What are the options in this case? One possibility would be to write a cheque, post it first class, and wait for the cheque to arrive. My daughter would then have to take the cheque to the bank and, hopefully, it will clear in three days, after which she can spend the money.

The development of a new payphone and a screen-based home telephone will lay the foundation for the future evolution of these smart card based services well into the future, particularly as the model for such systems can easily be translated into any environment where secure download/upload is required—one need only mention in this context the growing demand for electronic ticketing in the airline and entertainment industries.

Last but not least, in a discussion of the potential application of a smart card based payment system

Figure 12 – Using Mondex to transfer money across the country




Figure 14—An example of a smart card in a multi-application situation

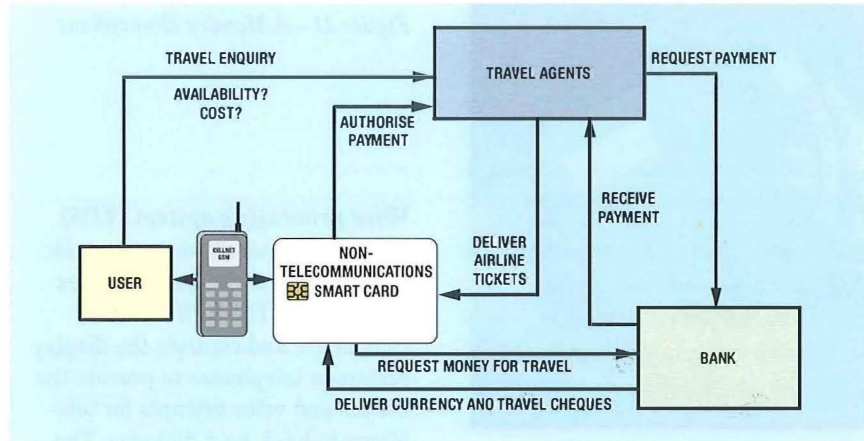

like Mondex, is its importance as an enabler for *micro-payments* and pay-per-view services on the Internet. It has been argued that the huge potential market for commerce (including pay-per-view content) on the Internet will remain untapped until a low-cost payment system is developed to address this issue. Smart card solutions from a number of potential sources, including Mondex, are currently considered among those most likely to gain widest acceptance from users and traders/merchants.

### GSM

BT, through its 60 per cent stake in Cellnet, is at the forefront of developments and customer service in the mobile telephone market. Mobile telephony is a major growth business in Europe and there are now an estimated 31 million users of mobile telephones across the continent. Some 16 million of these users are GSM subscribers. The Cellnet GSM digital service, launched in 1993, now has 650 000 customers.

The GSM service itself was developed against a heavy background of fraud on analogue systems in the USA and Europe. Customers had expressed concern that analogue services are subject to eavesdropping by inexpensive commercially available equipment. The smart card based system was, therefore, developed as a direct solution to the problem of fraud (Figure 13). The digital technology of GSM made possible active authentication between network and customer

Figure 13—Smart card for security and fraud prevention



equipment. The technology also allows encrypted speech.

The main purpose of the smart card in the GSM telephone is to authenticate the user to the network. It is the only access and billing mechanism for the network and, therefore, the key to the customer revenue stream. However, the system does separate the access/billing process from ownership of the handset, allowing users to hire handsets at will, especially when abroad.

A more in-depth consideration of the evolution of the GSM system from a smart card perspective could be the subject of a separate article, but Figure 14 hints at a possible service development opportunity in the area of multi-application cards, remote value transfer and teleticketing. In this case the services are delivered to a cell-phone over a wireless network. Here, the smart card is being used to perform a number of functions:

- to authenticate the card and user;
- to be a secure store of value for downloaded currency;
- to be a secure device for up-loading value to travel centre/bank; and
- to be a secure store for downloaded ticket.

Once again, such a service illustrates a very powerful combination of customer service functions being enabled by the use of chip technology. In this instance, the multi-function card does not offer telecommunications facilities, but one can equally imagine a situation where the card is programmed with a range of different user-selected services, which might well include a telephony element.

### A BT perspective on the future

The advent of new smart card operating systems and an environment, where dynamic downloading and deleting of applications is made possible, opens up yet more possibilities for the service provider. As was mentioned earlier in this article, the potential now exists for users to configure services on a card to suit his or her particular requirements in particular situations.

Figure 15 illustrates a possible portfolio of cards chosen by an individual from a menu of services that a service provider could offer. The user can package them in a way that suits him or her, and one hopes that, when this vision is fully realised, he or she can re-package them at will.

A further important development, that BT will encourage, is the growth

Figure 15—A possible portfolio of smart cards



of a smart card terminal infrastructure extending from the street into the home and office environments. While this article has described a service vision for the home/office context, it is true to say that the vast majority of current 'slots' capable of accepting smart cards are to be found in the public domain. Recent announcements from the PC hardware and software industries have recognised this as a major issue—clearly the next few years will see the integration of smart card reading capabilities into new PCs. The task for BT is to supplement this with telephony and other home terminal developments, and then integrate the service offerings across the range of devices.

## Conclusion

---

This article has skimmed the surface of a global business opportunity facilitated by the exploitation of a particular technology that will eventually touch all our lives. Space has prevented the discussion in more detail of a number of areas including the growth potential for contactless technology, the impact of biometrics, and the opportunity opened up by the interest of governments in smart card technologies. However, the article has outlined the main contours of the development path for smart cards and the way in which they will improve customer service and convenience for the ordinary citizen in the age of the information superhighway.

## Acknowledgements

---

The author would like to acknowledge the following colleagues for the help they have provided in preparing this article—Mike Meyerstein and Tim Hart of BT Laboratories, Steve Brown of BT Payphones, Andrew Bolton of Cellnet, and Michelle Mooney of BT Card Services

## Biography

---



**Mike Arnavutian**  
BT Personal  
Communications,  
Card Services

Mike Arnavutian is the Business Development Manager in BT Card Services, with particular responsibility for evaluating group-wide business opportunities in the area of electronic commerce and smart cards. Mike joined BT in 1979 and, after a career in what was then Group Procurement Services, spent five years with BT Payphones. Mike was part of the team that developed the strategy to move the BT Phonecard business on to a smart card platform, and was also the project manager for the payphone implementation of the Mondex project.

# Organisational Learning

## Early days in a change process

*This article reflects on the issue of organisational learning within BT. It describes how one division is managing the shift to an organisation style, built around reflection, innovation and adaptation. The primary thrust of the process is to move away from transforming the organisation, to helping the individual to change, and to recognise that organisations can learn only by helping people to learn.*

### Introduction

In May 1996, BT's Strategy and Business Management (S&BM) division decided that developing its capability to learn would be one of the main priorities. However, this presented a conundrum. There is no agreed model for organisational learning, relatively little is understood about enabling such a change, and it can take a long time to achieve such a deep structural transformation. This article briefly introduces the process used, outlines some of the key learning points, and concludes that, while the journey is difficult, it is a worthwhile one. The article suggests that organisations can learn only by helping individuals to learn<sup>1</sup>.

Peter Senge, one of the leading authors on this subject, recounts a story that illustrates many of the problems faced. He talks about an 18-year-old freshman at an American university who was studying calculus. She was totally lost. It was such a different way of thinking. She had been a good student in high school, but when she went to university and had to take calculus, all of a sudden she started to worry. Within a couple of months, she began to grasp the subject, and by the end of her first year, she understood the basics of calculus. By the end of her five-year engineering degree programme, it had become part of who she was as a professional. It had been internalised.

She made the following point: 'If calculus were invented today none of our organisations could learn it!' We would send everyone off, she says, to the three-day training programme 'Calculus In Three Days'. Then we

would send them back into the organisation and, after three, four or five months, depending on how patient our executives were, we would check to see what they had accomplished and decide if it was useful. Of course, after three, four or five months they would have accomplished nothing, and we would conclude that 'we tried calculus, now let's try something different.'

When you think about it this way, organisations are often set up to learn only what is trivial—what can be learned in one-, two-, or three-day training programmes. This offers a dilemma: how to create fundamental new capabilities of thinking and interacting that will take months or years to develop, not hours and days. Moreover, the real trick is to create a situation where it is not one or two people who are developing learning capability, but hundreds or thousands of people who are acquiring such skills.

While not professing in any way to have the answer to this conundrum, S&BM is undertaking an initiative that might offer some of the potential answers. S&BM has recently embarked on a change process that will significantly enhance its learning capability. The initiative is based on four main propositions (illustrated in Figure 1):

- Competitive advantage is facilitated by developing the organisation's capability to learn through reflection, innovation and adaptation.
- Cultural change should be focused primarily on the individual, as opposed to the organisation.

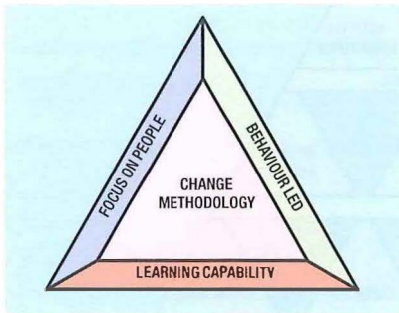


Figure 1 – Propositional themes

- Long-lasting change is delivered through leadership grounded in demonstrable behavioural change.
- The change methodology must positively reinforce the outcome that is sought by the change process.

The division is still in the early days of change, but already there are clear examples of a positive response. Self-managed teams are instigating change across the organisation. Managers are making positive and public steps to change their behaviour. Most significantly, a senior management team has committed 'personally' to deliver a training programme to people in the division.

### Background

The Strategy and Business Management division is headed by Andy Green and resides within the BT Group organisation. Its role is to work with the senior management of BT:

- to develop corporate strategy, strategic plans, quality policy and the business information strategy;
- to translate these strategies and plans into effective action;
- take a leading role in telephony product management and related activities for the business; and
- work closely with Regulatory Affairs Department in the management of key regulatory projects.

In April 1996, the S&BM Quality Council (QC) decided to place greater focus on three key areas:

- improving how people are led and motivated,
- creating a team of people who could build and inspire success across BT, and
- delivering the objectives set out in the organisational learning arm of BT's corporate scorecard.

Key to all of these objectives is the notion that people are the division's primary asset<sup>2</sup>. People must therefore be treated in a way that recognises their worth, and helps them gain a personal sense of their importance to the business and their contribution to the process of value creation.

In support of this objective, the QC decided to rethink how 'they' led people, and took a conscious decision to modify their own personal attitudes and behaviours. The increased focus on people issues will be embedded into the division's objective system, and is a key part of the director's 'vital few' objectives for the year.

The S&BM Quality Council offered a challenge to the division of achieving a 'world-class standard of people management by the end of 1997'. Demonstration of this will be through the emergence of an organisation that is more adaptive, reflective and open in style. Such an organisation will encourage the freedom to express personal views, make mistakes and stimulate a level of disagreement and debate that will contribute towards BT's development. In summary, it will actively embrace learning as a core competency at individual, team and system level.

### Focus on People

The ethos of valuing people is the bedrock of the change. It sets out the underlying philosophy that supports everything that takes place within

the division. This is not simply a downward paternalistic process, where managers just 'state' how much they value their team. It is a demonstrable indication of the 'way things are done around here'. It must be regarded as the primary axiom that defines how business is conducted, with individuals, teams, suppliers, colleagues or customers. The intention is:

- to create a culture that values people for their individual capabilities and competencies, as opposed to the job title of the function that they are 'currently' undertaking;
- to suggest that the experiences and feelings of an individual are as important as the tasks, activities, objectives and milestones; and
- to stress the importance of the 'supposedly' softer issues, such as trust, reflection and feedback, which are, in actuality, critical contributors to organisational effectiveness.

It is important to emphasise that the statement 'valuing people' is not the name of an initiative. It is very much a verb as opposed to a noun, setting out a statement of how we will strive to behave.

Examples of the theory in action<sup>3</sup> can be seen with the Open View teams. These are self-managed focus groups that meet to provide senior managers with open and honest feedback on issues that are of concern across the organisation. Once key issues have been identified, then action learning teams<sup>4</sup> are formed to take corrective action.

Already, one team has undertaken a review of the personal development process within the division. Another team is looking at the management of reward and recognition in the division. For example, they are attempting to understand the dynamics and balance that affect the criteria for recognition.

Figure 2—Learning model

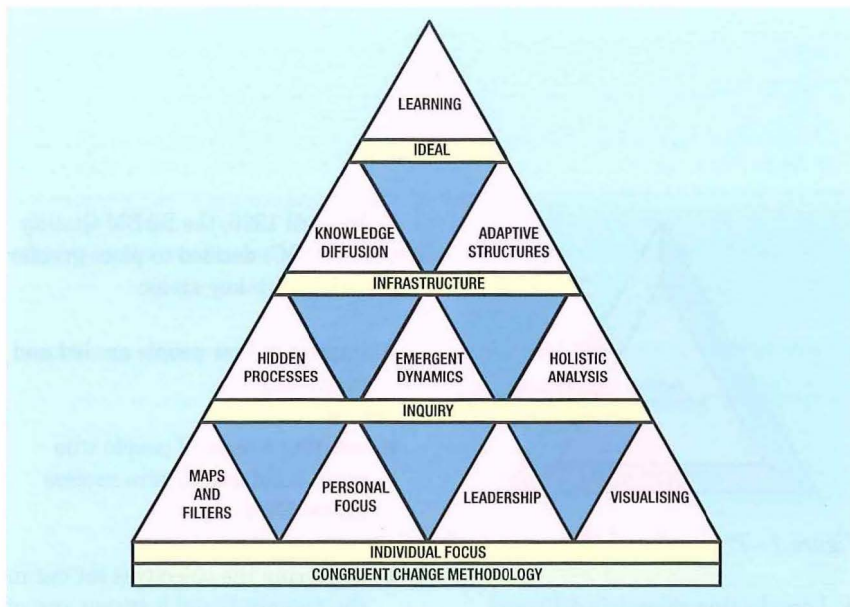
Should it be individual- or team-based? Do formal interactions with a senior manager really help an individual to 'feel' as if he or she is being recognised? Is it possible for recognition to be a destructive process, possibly by seeming to ignore those people who are not recognised? Finally, another team is looking at the issue of empowerment, and how to ensure that messages of empowerment are reinforced by the appropriate behavioural changes.

The message is that the teams do not have (or need) any formal or delegated authority to consider the problem. They have effectively created a sense of self-empowerment, breaking the traditional mould that change is led from the top. The indications are that placing the focus on people in the change process, and giving them freedom to effect changes in their own environment can create a momentum all of its own.

### Developing the Capability to Learn

The desire to enhance the organisation's capability to learn has become increasingly pronounced during the 1990s. The ability of an organisation to respond quickly to changing market conditions is advocated as being key to creating competitive advantage in the market place. Interestingly, the ability to learn and create organisational knowledge is gaining credence as a form of capital<sup>5</sup> within the business. One can almost place a value on a business ability to rapidly implement a new production process, provide a new service, or identify a particular outcome valued by the market.

One example of this might be Intel Corporation's ability to improve the speed of production runs, and thus deliver new products to the market quicker than its competitors. Other examples of companies that have adopted a learning approach include Xerox, Ford, Royal Dutch/Shell and Federal Express. They have publicly stated their intention to embark on



the journey towards becoming learning organisations. All of their strategies, processes, and initiatives are planned with the learning objective in mind. BT considers that developing such capability is one of the key success factors for the coming years. Its inclusion as part of the scorecard is a clear indication of the organisation's intent to take this seriously.

In developing the capability to learn, the first step is to define what organisational learning means, and how S&BM can interpret it locally. This is currently being defined, and the first step is the development of the model in Figure 2. This model sets out the 10 key aspects of organisational learning that are significant for S&BM. There is nothing original in the design because it simply offers a broad synthesis of ideas taken from many of the leading paradigms that are being applied across industry and academia. The crucial stage is where individuals or teams start to identify and apply those attributes and behaviours that are most pertinent for their particular context.

The intention is to avoid the trap of mandating a change in behaviour and attitudes. Instead, activities are being created so that people can voluntarily come together to stimulate and create learning at team level. By using such experiences a culture is being built based upon reflection, innovation and adaptation.

To date, some eight initiatives are being supported across the organisation (see Figure 3). These are not

being positioned as 'the answer'. They are simply examples of some of the activities that are underway across the division to engage people in change activities. A principle of the change process is 'it does not matter what you do, just so long as you do something'. If an individual chooses simply to develop the capability and desire to say 'good morning' to all of his or her team every day, that is fine. Success breeds success, and the demonstration of internalised change will help people to appreciate that a long-lasting transformation is underway and visible to everyone.

The single key message with all the development initiatives is that there is no miracle cure<sup>6</sup>. No one action will be practical or appropriate for all people. The offer made to the people across the division is to 'challenge the givens'—to find something that can foster personal development and learning, while looking to cause some form of improvement at an organisational level.

One of the ways in which this is being fostered is through the use of the Pilat or 360° feedback process. One of the first actions that the S&BM director took in launching the change was to go through his 360° feedback with his directors at a team meeting. They discussed his management style, and identified ways in which changes could be made to enhance the capability of the management team. This proved to be a highly-charged and emotional session, but has facilitated amazing results. Additionally, the director's willingness to start this



### 1. Action learning

The senior management group have formed an action learning team, where they both manage the change process on behalf of the QC, and use action learning principles to develop themselves on a personal level. There are plans to initiate more teams to look at areas such as recognition, managerial styles and personal development.

### 2. 360° feedback process

People are being encouraged to use feedback as a positive management tool—to share the results of their PILAT survey or other forms of 360° feedback with their team and ask for feedback in areas where personal improvement can take place.

### 3. Business excellence model

Many teams are now actively making use of the European Quality Award or Business Excellence model as a structured methodology for reinforcing the quality culture. The model encourages the use of dialogue<sup>1</sup>, feedback and team learning as management tools through its use of scoring workshops.

### 4. Open View teams

These are longitudinal focus groups that meet regularly to provide feedback to the senior managers on issues or concerns. This gives them access to a dynamic feedback system that can both offer up the current issues from across the organisation, and give a month-by-month review of the improvement activities.

### 5. Relationship management

One of the key aspects of the management system is the ongoing relationship between people, the manager and team member, individual and customer, et al. The division will be actively supporting the adoption of the relationship paradigm that is based upon the BT Situational Leadership model.

### 6. Influencing skills

The appropriate opportunities must be provided for people to develop the necessary skills and competencies that will help them to take a more persuasive and influencing style. The aim is to help people to understand how to influence across the business, without relying on the traditional approach of stimulating change through formal authority or power.

### 8. Covey leadership paradigm

The Covey leadership model is an excellent paradigm for helping people to understand that change starts from within, and not with other people. The 'Seven Habits...' model is being used as a stimulant for individual change, without positioning it as 'the miracle answer' to problems.

Figure 3—Divisional initiatives

process has sent a clear signal to the division that reflection and adaptation are critical issues for all managers.

## Leading Through Behavioural Change

The members of the QC were aware of the need to lead the change, not by mandate, but by setting personal

examples. They were determined to avoid this initial barrier that often surfaces at an early stage of any change process.

There are, typically, three levels of management commitment to an organisational development programme. The lowest level is 'giving permission' for change to occur. The second level is 'giving support and

encouragement' for the change effort. The highest level involves 'participation' in the change programme, including active display of a change in personal behaviour. The S&BM senior team sought to tackle this issue head on by working at the highest level from the outset. From the beginning, the focus has been on helping the senior team to identify areas where they can make a personal commitment to change, and can openly share this with their team.

A good example of this is the *gumball initiative*. Each of the QC members had a gumball machine that was full with bubblegum. Every time they demonstrated (publicly) a new positive behaviour, they wrote the behaviour on a piece of paper and put it in the machine. At the same time they rewarded themselves by taking a bubblegum. Over time, each director (and his or her team) could see the bubblegum disappearing, and the machine filling up with the pieces of paper.

While this does not suggest that the directors underwent a radical overnight transformation in how they behaved, it did make a clear positive statement that they were willing to change their personal behaviour, and to do it in such a way that reinforces the ethos of valuing people.

## Change Methodology

**'Most, if not all projects involve the need for people to do something different. Successful change requires the understanding of the soft issues. Indeed, the experience in BT is very much that the soft stuff is the hard stuff.'**<sup>9</sup>

This is the view offered by Lowry Stange in the April 1996 issue of the *Journal*, when discussing the difficulties faced with implementing business process re-engineering<sup>9</sup>. As we learn to understand and tame the management of change, it behoves us to develop the capability to understand and manage the 'soft stuff'. However, the view offered in this article is that possibly we have to

Figure 4—Change model

learn how to manage the 'soft stuff' without using the 'hard stuff'. We must appreciate how to facilitate and manage change without using formal and rigid structural processes—to navigate the waters of change with the people, and not to throw it at them<sup>10</sup>.

From the outset, S&BM has attempted to avoid many of the pitfalls that can arise when attempting to shift an organisation's style and behaviour. One important approach has been to focus on building a change ethos, or common philosophy within the senior management team, before undertaking any other actions—to consider transformation at an emotional level, before designing the outcome or the change methodology.

The change model in Figure 4 summarises the principles at the heart of the change process. This does not seek to be a prescriptive divisional model for change, but is a representation of the approach that the senior managers are attempting to follow.

One example of this model in action is the *relationship initiative*. Feedback from the last morale survey suggested that improvements are possible by focusing on the performance review process. In July 1996, the senior managers agreed to use the Situational Leadership model<sup>11</sup> as a basic paradigm for managing all work-based relationships. However, rather than simply employing trainers to deliver the programme, they decided to train people in the division themselves. They received coaching from the training department in December 1996 in the use and application of the model. Starting January 1997, they will begin the delivery of a programme across the division. This is a positive example of inspiring leadership through role modelling and personal change, as opposed to the traditional approach of mandated instruction. The goal is to facilitate improved management of organisational relationships, by using a relationship management tool. The

### **1. No delegation**

Behavioural change cannot be delegated. Public display of changes in the personal behaviour of senior managers is the first step towards significant and long-lasting cultural change.

### **2. Talk, don't count**

The only way that we will achieve true understanding of people's issues is through a balance of discourse and objective measurement. Qualitative analysis should always take precedence, and be 'supported' by quantitative assessment.

### **3. Behaviour, not words**

Most important is the development of an internalised belief in the need for change, and 'demonstration' of this through modification in personal behaviour. It is not enough for people to just state their belief in the change; some modification in behaviour, however small, is necessary to stress truly its importance.

### **4. Voluntary approach**

The adoption of a new way of working cannot easily be mandated. It must be grounded in the process of personal selection and choice. So, wherever possible, activities should be voluntary. If this is not practical, the rule must be to seek to convince people of the benefits and wisdom of transforming how they work, before forcing any change to take place.<sup>7</sup>

### **5. Principle, not project**

The approach is built more around a philosophy or set of values than a set of techniques. Any tool or technique that adds value is welcome irrespective of its source.

### **6. Context is key**

It is difficult, if not impossible, simply to import models and paradigms from other corporate cultures. A more realistic approach is to nurture and build a paradigm that is appropriate for the situation in which it is to operate. Trying to force real people and organisations into preset theories is not beneficial in the long run. It is better to try to understand the realities and nuances of the local situation, and then adapt any theories accordingly.

### **7. No miracles**

There are no miracle cures or panaceas for all evils. Typically, there are always a selection of management theories and views that tend to be formed and circulated over time. The role of the manager is to select those which are appropriate for his or her team.<sup>8</sup>

### **8. Invest wisely**

There is often little correlation between the amount of money invested in a change process, and the actual outcomes. The investment of money should be on the basis that it purchases skills, knowledge or experience that is not readily available within the division. If one views the definition of organisation development as 'maximising the use of "all" the organisations' resources', then any organisation development investment should be considered as closely as any other investment programme.

### **9. The energy is in the relationship**

We cannot hope to influence any situation without respect for the complex network of people who contribute to the organisation. It is crucial to recognise that the value of the business actually lies in the effectiveness of the relationships, and not just the people.

### **10. Focus on the purpose, not the change**

The desired outcome is not the completion of a project. It is clear cultural and economic improvement across the organisation.

# *Developing an open and adaptive system by a closed system methodology is like trying to mix oil and water: there will always be some form of conflict or separation.*

focus is not to say 'you must improve how you deal with others', but to say 'we will all learn to work better'. This is a true example of 'top led' as opposed to 'top down'.<sup>12</sup>

The message is that in attempting to develop a new culture, one has to use a change methodology that is complementary to the desired change. Developing an open and adaptive system by a closed system methodology is like trying to mix oil and water: there will always be some form of conflict or separation.

## Lessons Learnt

### Positive feedback

The experience of building change around a less formal and structured model makes the process somewhat fraught and painful at times. It means that people must be willing to relinquish control and free others to make and take decisions at a much lower level within the organisation. The positive side is that once people have taken up the cudgel, they fight with it with such a passion and force that little support is necessary to progress and support the change.

### The learning is in the journey

This aspect is key to the whole process, because the focus must be on valuing the journey as a learning experience in itself. It is not simply a project to 'install or implement' the new approach. Improving the capability to learn is grounded in the process of giving people the opportunity to experiment. The change journey is the prime opportunity for this to take place.

### The power of dialogue

The benefit gained from allowing people the time and space to simply 'talk' is quite staggering. Creating an environment where people can meet to reflect and talk about 'how life is' is already realising benefits. The idea of having meetings without fixed agendas, rules or a chairperson, really allows people to surface their beliefs, thoughts, and ideas. Most important

is the idea of a meeting of equals, where rank and power do not influence either the process or the outcome. It is through this type of meeting that people can really experience the power and energy of reflective and adaptive thinking. The challenge is to turn this approach into a personal and commercial advantage.

### Measuring the change

Measuring the progress of the change can be difficult, because so many of the issues under consideration are relatively soft in form. The short-term solution has been to develop a matrix (Figure 5) that identifies the key areas for measurement in relationship to the areas where the change is taking place. In many instances, it might be appropriate to focus on organisational hard issues, simply to gain a degree of high-level momentum. In other areas, if the shift is underway, it might make more sense to focus tightly on the soft/individual level, and to consider a degree of introspective reflection to understand what movement has transpired.

### Short-term versus long-term returns

The adoption of a change process focused on long-term return might be in conflict with the nature of short-termism contained within most organisations. The rational/formal approach to managing change can, by its very nature, deliver short-term outcomes. The question is 'Are they robust and embedded in the deep structure of the organisation?' The emergent model allows for the change to integrate with the existing norms

of the organisation, but that in itself supposes that the organisation will not change. How can a long-term change paradigm exist in an organisational structure that has only a limited life?

The one way to avoid this is by creating an organisation that is able to thrive on constant flux and complexity, where the notion of 'resistance to change' becomes obsolete since change is the natural way of life for the organisation.

### Reflection is difficult

Taking time out to reflect and think about past actions often seems unnatural and smacks of time-wasting. The question is how to reverse this view, and stimulate a more reflective style within the organisation. As a process, it is slowly increasing, but it is often a catch-22 situation. People recognise the benefits only when they have experienced the process.

## Conclusions

S&BM has started a culture-change process that will deliver long-term organisational effectiveness. Placing focus on the individual, behaviour, learning and change will help create a more adaptable, flexible and innovative organisation.

The approach adopted is not definitive, and it is not 'the' right way, but it is starting to work. One certainty is that this model may be different in one month, will probably be modified in six months, and will be totally different in a year's time. This is of little concern, because the objective is not to deliver the change programme, but to improve the capability of the organisation to learn at individual, team and system level. Most importantly, the aim is to develop an organisation that, as its primary ethos, consciously recognises that organisations only learn through individuals who learn. In doing so, the value of people and their contribution to the business will be consciously and continually valued.

Figure 5—Measurement matrix

	HARD	SOFT
ORGANISATIONAL	CARE EQA CUSTOMER REVIEW ISO 9001 INVESTORS IN PEOPLE	OPEN VIEW TEAMS  PEOPLE DEVELOPMENT FORUMS  ACTION LEARNING TEAMS
INDIVIDUAL	360 DEGREE REVIEWS PERFORMANCE REVIEW YEARLY APPRAISAL ISO 9001	PDPs CELEBRATION OF MISTAKES RECOGNITION SCHEME 'LEARNING TO LEARN'

## Acknowledgements

The author thanks Paul Oliver for his support and ideas offered in developing many of the models used in the paper. He also thanks the people who helped to formulate both the structure and content of the article, including Lizzie Beesley, Paul Burns, Linda Cope, Rod Gray, Dan Moorhead, Pat Spink and Sara Thomas.

Finally, this change would not have been possible without the tolerance and support offered by Andy Green and Ian Morfett and their senior management teams and, most importantly, all the people involved in the S&BM initiatives.

## References

- 1 SENGE, PETER M. *The Fifth Discipline*. Century Business, 1990.
- 2 GREEN, ANDY. S&BM People Power Statement issued in Aug. 1996.
- 3 ARGYRIS AND SCHON. *Organisation Learning 2, Theory, Method and Practice*. Addison Wesley, 1996.
- 4 REVENS, REG. *Action Learning. The ABC of Action Learning*, London, Blond and Briggs, 1983.
- 5 EDMONSON AND MOINGEON. *Organisational Learning and Competitive Advantage*. Sage, 1996.
- 6 WILLIAMSON, DEREK. Presentation at S&BM State of the Nation, Oct. 1996.
- 7 COVEY, STEPHEN, R. *The Seven Habits of Highly Effective People*. Simon and Schuster, p. 103.
- 8 SHAPIRO, EILEEN, C. Fad Surfing in the Board Room. Capstone, 1996.
- 9 STANAGE, LOWRY. Making Breakthroughs. *Br. Telecommun. Eng.*, April 1996, 15, p. 3.
- 10 LITWN, *et al.* *Mobilizing the Organisation*. Prentice Hall, 1996.
- 11 HERSEY, and BLANCHARD. *Situational Leadership, Management of Organisational Behaviour*. Englewood Cliffs, Prentice Hall, 1982.
- 12 LITWN, *et al.* *Mobilizing the Organisation*. Prentice Hall, 1996.

## Biography



**Mick Cope**  
BT Strategy and  
Business Management

Mick Cope completed his MBA in 1994, and is currently undertaking a Doctorate in Business Administration at Henley Management College. He is also a member of the Academy of Human Resource Development. He joined BT in 1974, and rapidly established himself as a particularly uninspiring apprentice, with the words still ringing in his ears 'you'll never become a Technical Officer, lad'. After spending a number of years on the CSS project, he moved into the organisation development field, where he supported the introduction of Total Quality into Northern London District. Following this he managed the project office for ISO9000 implementation in Worldwide Networks. He subsequently managed the BT pricing process and worked as a global pricing analyst. He is currently organisational development manager for the Strategy and Business Management division.

# Workplace Assessment for National Vocational Qualifications

*Gradually, people in the United Kingdom are becoming aware of National Vocational Qualifications (NVQs). Previous articles in the Journal have discussed the design of NVQs and how they differ from existing qualifications. This article looks into the assessment process in the workplace and the collection of suitable evidence to prove competence.*

## Introduction

A previous article<sup>1</sup> in this *Journal* described the background and development of NVQs produced by the Telecommunications Vocational Standards Council (TVSC) for the telecommunications industry, and gave an introduction to the assessment process. This article provides more detail on the practicalities of assessment from the candidate's point of view. It also discusses the benefits of the relationship between the candidate and assessor, where these people are also the employee and his or her supervisor. It also includes guidance for candidates when compiling a record of evidence of their competency.

## Assessing Performance

### Assessment of NVQ candidates

As the previous article stated, assessment of candidates involves the three activities: generating, collecting and judging evidence. These activities are shared between the candidate and assessor. Figure 1 shows a model where:

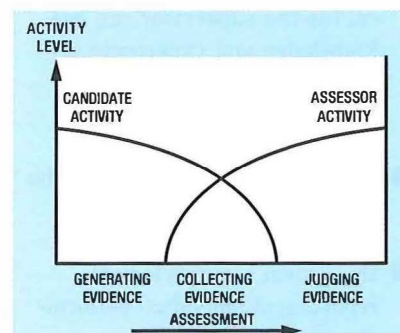
- the candidate generates evidence through normal work activities;
- the candidate and the assessor collect and record this evidence; and
- the assessor judges the evidence against the national standard.

The evidence generated, collected and judged is used to determine whether the candidate has met the national standard for occupational competence in the workplace. If this competence is to be demonstrated in the workplace, it should not seem unreasonable to replace *candidate* with *employee*, and *assessor* with *employer* or *supervisor*.

## Appraisal or Assessment of Individual Performance

The majority of appraisal systems are based on job descriptions, written in suitably vague terms, and use subjective measurements of staff performance. Some systems have a levelling process which attempts to reduce the degree of individual subjectivity. However, these staff appraisal systems cannot produce reports on staff performance which are based on objective measurements against national standards. Standards should be known and agreed to by both the appraised and the appraiser. The

Figure 1 – Candidate and assessor activities during assessment



occupational standards that form NVQs can be used as the standards for the appraisal or individual occupational performance.

The building block of all NVQs is the *element of competence*. A statement of occupational competence can be written for any role using these elements. Each element describes an activity that is performed in the execution of an occupational role. Occupational roles can therefore be described by grouping the appropriate elements of competence. These elements, with their respective performance and knowledge criteria, are the standards against which individual performance can be measured. A benefit of this approach is that as an occupational role changes or develops, then its statement of competence can be reviewed and revised. A new statement of occupational competence can then be written, using the national standards.

### Building on the Candidate and Assessor Partnership

Where the NVQ candidate and assessor are employee and supervisor, this partnership has real benefits for the employing organisation. The following list indicates just some of the possible benefits:

- the employee and supervisor use the same standards;
- they have shared expectations;
- they use a common vocabulary;
- the assessment is more economical (as the supervisor can use knowledge and experience of past performance by the employee);
- there is greater objectivity in the measurement of performance;
- the system requires formal reporting of individual performance and development;

- the employee has more precise details of what is meant by competent performance;
- the supervisor develops better appraising skills; and
- individual motivation is increased.

Organisations that encourage employees, at whatever level, to discuss competence will reap the rewards of making their employees more aware of competence. Each employee will appreciate more how they are individually contributing to the overall increased effectiveness and efficiency of their organisation. Several NVQ candidates have commented how undergoing NVQ assessment has made them aware of being professional in their job.

Assessment of NVQ candidates by their supervisor can also be integrated with other quality checking systems, where they exist. Although an organisation's staff performance standards and the NVQ standards may be written using different words, they could be describing the same characteristics of competent performance. Any organisation implementing or considering implementing NVQs would be advised to identify where there is correlation

between their existing performance standards and the appropriate NVQ standards. This could lead to more economic assessment of NVQ candidates in the workplace.

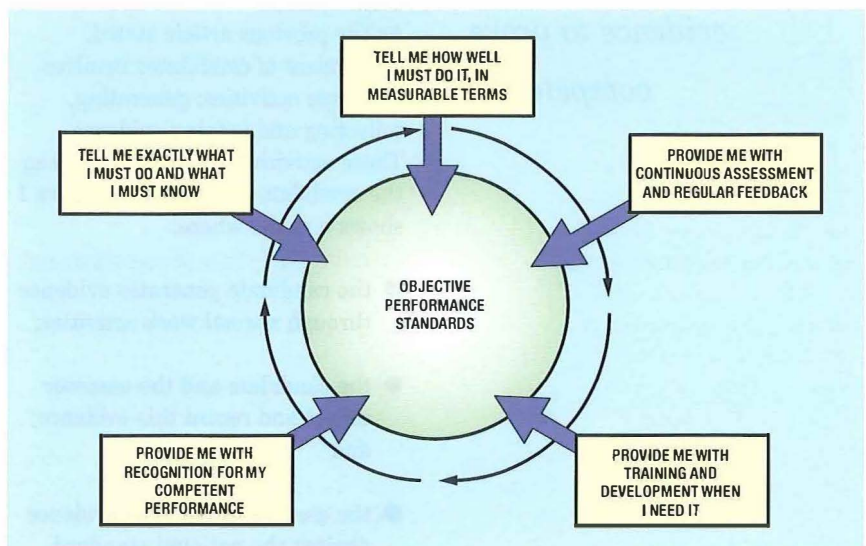
### Increasing Motivation

It is widely accepted that increased motivation is an important ingredient for the success of organisations. It is the people within an organisation who really determine the success of that organisation. Occupational standards can form a focus for increasing motivation. Figure 2 shows the standards as being at the centre of the motivation cycle.

### Collecting and Recording Evidence of Competence

Unlike traditional examination-based qualifications where the evidence performance is recorded on students' answer papers, the evidence used in NVQ assessment is often the observation of the candidate activity by the assessor. It is not sufficient to record simply the details of what was done, but also details of how well it was done and the underpinning knowledge. This is often where difficulties arise in the assessment of NVQ candidates. It is also where integrity and credibility

Figure 2 – The motivation cycle



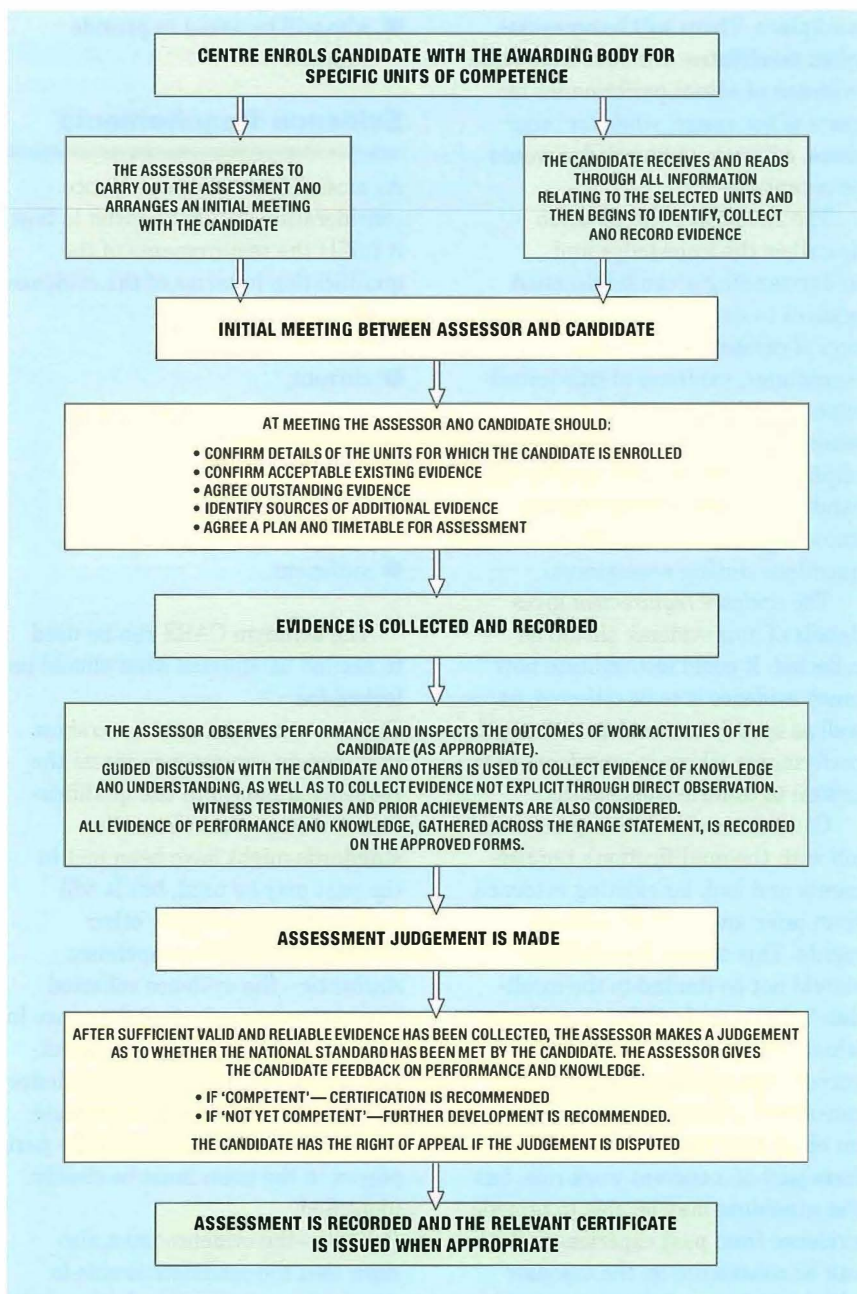


Figure 3—Outline of the assessment process

of NVQs could be questioned. The quality of NVQ assessment, and thereby the quality of the award, rests on the competence and professionalism of the assessors. They have a significant responsibility for the consistent interpretation and implementation of the occupational standards. They also have the responsibility for ensuring that their assessment decisions can be audited by the verification process. This process relies on several factors,

including the way in which evidence is recorded.

### Outline of the Assessment Process

The assessment of all NVQ candidates has to be through approved centres. These centres have demonstrated to the NVQ awarding body that they can meet certain criteria relating to management and administration systems, sufficiency of

resources, and quality assurance arrangements. Approved centres are able to obtain advice and guidance from the appropriate awarding body on how to operate as an assessment centre.

The following should help candidates who are undergoing NVQ assessment, or who are considering applying to be assessed. However, as each assessment for an NVQ is individual, the following information must be used as guidance only. A candidate's particular assessment may follow a slightly different process. Typically it should follow this general path:

- candidate registration or enrolment;
- candidate and assessor meet to discuss and plan assessment;
- evidence of competency is collected and recorded;
- evidence is judged against the national standards;
- assessor provides feedback to the candidate; and
- outcome of assessment is confirmed and recorded.

If the assessment outcome is confirmation of the candidate's competence, then a certificate is issued by the awarding body for the units or full award achieved. If, on the other hand, competence cannot be confirmed, the assessor should discuss ways in which skills and knowledge, needed to become competent, could be developed. The candidate may also wish to discuss further development with the person in the organisation who has responsibility for training and development. Figure 3 gives an outline of this process.

### Planning Assessment

It is the candidate's responsibility to prove competence to the assessor. It

is not the assessor's job to find evidence for the candidate. However, the assessor will explain what evidence is required and acceptable. At the initial assessment planning meeting, the assessor and candidate will agree what evidence is already available, what evidence is outstanding and how it will be collected.

This last aspect of assessment could require the generation of evidence through workplace activities. The candidate will therefore need to identify ways in which this could be achieved. To assist this, candidates should make certain they are familiar with all the qualification's requirements in terms of:

- units;
- elements;
- performance criteria;
- range statements;
- knowledge specification; and
- evidence requirement.

*Units* (of competence) reflect a significant proportion of a candidate's work role, or sub-area of competence (to use NVQ terminology). Each unit is made up of a number of *elements* (of competence), each of which describes a subdivision of its unit and identifies an activity that the candidate should be able to carry out at work. *Performance criteria* set the performance standard against which candidates are measured. These criteria enable assessors, and others, to distinguish between acceptable and unacceptable performance. Each criterion defines a characteristic of competent performance and comprises an observable outcome, with an evaluative statement against which actual performance is measured.

The *range statement* describes the contexts and applications in which competence would be expected in the

workplace. There will be occasions when candidates will have to provide evidence of actual performance for parts of the range; while for other parts, evidence of knowledge would be acceptable.

The *knowledge specification* describes the knowledge and understanding a candidate must possess to support their demonstration of competent performance. Sometimes, evidence of this knowledge is explicit in performance, while at other times it may only be implied. Assessors will confirm that candidates possess the required knowledge, usually by asking questions during assessment.

The *evidence requirement* gives details of how evidence should be collected. It could also indicate how much evidence is to be collected, as well as specify any evidence of actual performance which is considered to be critical to confirm competence.

Candidates should compare their job with the qualification's requirements and look for existing evidence from prior and current achievements. This search for evidence should not be limited to the candidate's current job. Other experiences which may provide evidence of current competence should also be considered. An activity required by an element of competence may not form part of a current work role, but the candidate may be able to provide evidence from past experience which can be considered by the assessor when making an assessment decision.

After this initial inspection of evidence currently available, the candidate should prepare a plan for generating and collecting the outstanding evidence. The assessor will need to agree this plan, which should include details under the three headings of:

- what action needs to be taken to generate and collect evidence;
- what the timescales are for this action;

- who will be asked to provide support?

## Evidence Requirements

As a candidate collects evidence, consideration should be given to how it fulfils the requirements of the qualification in terms of the evidence being:

- current;
- authentic;
- reliable; and
- sufficient.

The acronym CARS can be used to remind candidates what should be looked for.

*Current*—the evidence has to show that current competency meets the required standard in the qualification. Evidence that shows the standards might have been met in the past may be used, but it will have to be supported by other evidence of current competency.

*Authentic*—the evidence collected must be as a result of performance in the workplace. It must be attributable to the candidate. If the evidence was generated while the candidate was working in a team, then the part played in the team must be clearly identified.

*Reliable*—the evidence must also show that the candidate is able to repeat the competent performance. It must not be just a lucky chance that the standard was met on a particular occasion. The assessor will need to have confidence that the performance demonstrated during assessment is an indication of normal performance.

*Sufficient*—The assessor will need to consider a balanced mix of quality evidence, before making a judgement. The candidate does not need to produce a mountain of evidence and hope the assessor will be impressed by quantity. The assessor will give guidance on the sufficiency of evidence.



Evidence can be thought of as being like a jigsaw, whose pieces when put together give a clear picture of a candidate's competency. It is only after the picture is complete that the assessor is able to make a fair judgement.

## Compiling a Record of Evidence

Many NVQ awarding bodies provide forms on which to record evidence, that should be completed throughout the assessment process and kept safe. Some awarding bodies also allow candidates to generate their own record format which suits their individual approach to assessment. The TVSC, as an awarding body, provides masters of forms to be copied and completed by candidates and assessors, but also allows other formats to be used. The principal criterion is that all forms used must provide a record of evidence and assessment which can be audited.

The guidance given to candidates by the TVSC suggests the following steps are taken when compiling their evidence:

- complete an *Evidence and Assessment Record* title sheet for each of the units for which you are to be assessed;
- write a *Personal Profile* giving details of your experience and the contexts in which you have generated evidence appropriate to this unit;
- follow this, if applicable, with a list of witnesses who have or will provide evidence on your behalf;
- complete *Record of Evidence* sheets ensuring that all relevant details are complete. (Your assessor will probably also provide details to be entered on this form); and
- compile an *Evidence Record Index* to enable the source of evidence to be located.

The order of these steps has been designed to lead candidates from providing information with which they are familiar (that is, about themselves) to providing information which requires more understanding of NVQs and their assessment. As NVQ candidates become more experienced in recording evidence of their competency, they will become more familiar with recording this evidence as it is generated.

The Evidence and Assessment Record should be compiled unit by unit, although a candidate and his or her assessor may wish to produce a record of all the units assessed in one document, if this is more appropriate.

## Ownership of the Record of Evidence

The TVSC believes the Record of Evidence belongs to the candidate, who may wish to produce it to supplement his/her NVQ certificate; for example, when seeking a new employment position and it will become increasingly necessary as candidates gain NVQs which comprise mandatory and optional units. Candidates will be able to provide details of the specific evidence they produced when they were assessed as competent and awarded the NVQ certificate.

However, although candidates will retain their record of evidence and assessment, these documents must be made available on request to the internal verifier, the external verifier or other officer of the awarding body. It is, therefore, an important document and candidates are strongly advised to keep it safe.

## Conclusion

This article builds on the information given in a previous article in the *Journal*<sup>1</sup>. It is recommended that both articles are read to obtain a fuller understanding of Telecommunications NVQs and how candidates are assessed in the workplace.

The information in this article suggests that assessment of NVQ candidates could be integrated with an organisation's existing performance quality systems and staff appraisal procedures. The objective measurement of an individual's competency can be made against industry-wide standards. The existence and use of these national standards can encourage a common understanding of what is required and expected for competent performance in employment.

The NVQ system encourages candidates, that is, employees, to be more proactive in considering how well they are doing in their work role. The system also encourages employees to record evidence of how well they are doing. A candidate's record of evidence and assessment is a useful addition to his or her NVQ certificate. The candidate can proclaim 'I am competent, and I hold a NVQ certificate to show it', and can add 'and my record of evidence shows why I am competent'. This record can also be used during appraisal of the employee, as it is a permanent record of the range of work activities which have been performed, and how well they have been achieved. It can also be used to record identified training and development needs.

Although each assessment for NVQs is individual to the candidate being assessed, it is accepted that some guidance is required 'to get started'. This article has provided an outline of what candidates should be doing during their assessment. Some larger organisations, such as BT, may wish to provide more detailed guidance to assist their employees. This is especially helpful where there are significant numbers of employees performing similar work roles.

In order to support the implementation of the telecommunications NVQs available from the TVSC, it is intended to provide samples of actual evidence supplied by candidates. In this way, it is hoped to remove any

mystery surrounding the NVQ assessment process. These examples will be published as part of *British Telecommunications Engineering*, starting later this year.


## References

---

- 1 MARTIN, A. National Vocational Qualifications for the Telecommunications Industry, *Br: Telecommun. Eng.*, July 1996, 15(2), p.186

## Biography

---



**Alan Martin**  
Head of Standards  
and Awards  
Telecommunications  
Vocational Standards  
Council

Alan Martin has the responsibility for the development of occupational standards and vocational qualifications for the telecommunications industry. He began his career in London with the Post Office, working on exchange installation. A move into the Research Department at Dollis Hill, and then Martlesham Heath, required a move from Strowger systems to electronic-based equipment. He spent three years working with the South African Post Office Research Laboratories, and five years in Zimbabwe working at Zimbabwe Posts and Telecommunications Corporation Training College and Phillips Electrical. His interest and experience in vocational training began in Zimbabwe and have included positions as training manager and further education college lecturer. He has been with the TVSC since the Council was established as the lead body for telecommunications. The TVSC may be contacted at:

Blackfriars House  
399 South Row  
Central Milton Keynes  
MK14 5DJ  
Telephone: (01908) 240120  
E-mail: TVSC@wildnet.co.uk

**Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his regular column in the Journal by discussing the advances in office equipment leading to the concept of the office you wear.**

Only five years ago I used to carry a large briefcase full of paper, a pocket calculator, a Dictaphone plus a few other sundry items. This was extremely heavy, and cumbersome, but necessary to manage my department of 160 people. Today my briefcase is a fraction of the size, contains virtually no paper and is twice as heavy. The reason? Batteries and connectors! Batteries to power my laptop (including a five hour long-life pack), my mobile telephone, Dictaphone, pager, and calculator plus a selection of cords, connectors, chargers, screwdrivers, crocodile clips, and all the other sundries necessary to make the mobile, electronic office work. While this is a miracle of technology, it is also essential to help me run my now larger, 650 strong, department. This mobile office is extremely compact and powerful, but is gradually lengthening my arms as I lug it from one location to another in cars, trains and aircraft.


A critical look at the individual functionality soon reveals that a considerable degree of integration is both feasible and necessary. For example, £100 now buys a digital wrist watch with an integrated paging device and display; £1000 buys a complete cellular telephone that can be worn on your wrist; laptop computers can be condensed into an organiser or pocket-sized device with a considerable saving in weight. What then is the chance for complete integration of all these units into one entity, in the not too distant future?

Just over 250 years ago, travellers in the horse-drawn carriage would take with them a time piece almost as large as my present briefcase. Naturally enough it was called a carriage clock! After due evolution of

the technology it became a pocket watch about the same size of a small dinner plate. This soon progressed to become a fob watch and ultimately a wrist watch, and today we regard it as jewellery or body furniture. Might we, therefore, expect a similar progression for my briefcase full of separate bundles of electronics?

The principal limitation to realising information technology that we wear, and perhaps information technology that wears us, is a requirement for batteries to power the mobile communications, data processing, storage, and, most critically of all, displays. Today the smallest cellular radio is composed of three integrated circuits, though it is feasible to reduce this to one chip requiring a single watt of power. Similarly, a powerful personal computer can be reduced to two integrated circuits requiring only 2 W, plus a liquid-crystal display consuming 3 W. The main difficulty that we have encountered is the necessity for back lighting to make the display easy to read. Typically, newsprint has a contrast ratio of 200, a standard domestic television set is around 50, and a black-on-grey liquid-crystal display without back lighting can be as low as five. This is low enough to be both painful and detrimental to our eye sight.

A further impediment to the office you wear is the need for a keyboard. Anyone using a laptop computer for long periods will have suffered the *praying mantis* syndrome of rounded shoulders and clawed fingers. The electronic organiser is even worse and reduces us to one finger and one key at a time. Perhaps the next vital step is to replace the keyboard with voice recognition and speech synthesis technology. Today, voice synthesis is just about adequate for text-to-speech conversion while speech to text still leaves a lot to be desired, and will probably require a further five years of development before it can hope to replace the keyboard. Unfortunately, speech technology is more power hungry in terms of processing power, and could add a further 5 W to the power budget. Hopefully, during the



next five years, integrated circuit technology will realise ultra-low power demand, and massive processing power without a significant increase in weight.

An intermediate solution that might be tenable is a minimal keypad or stylus for immediate access, and a fold-away keyboard that could be plugged in when necessary, with a head-mounted screen giving high-definition video directly into the eye of the user. All of these components are available now so such a device is feasible, and could even be practical. However, the critical problem remains—power storage, and its requirement for large and cumbersome batteries. The solution may reside with the wearer. A fully-grown person sitting doing nothing radiates approximately 60 W from his or her torso and head areas; when animated, this can exceed 100 W, and is a potential source of power to drive the office you wear. It appears that all of the solutions needed to realise the wearable office are to hand! Why then is no one marketing such a device?

As a general rule, the human race makes progress by incremental change, with slight extensions of existing paradigms that do not upset or compromise existing working practises or protocols and social sensibilities. The migration from a book full of paper to a book full of electronics (the laptop) or from the diary to the personal organiser are obvious examples. More recently the emergence of an electronic purse or wallet with electronic functionality and plug-in cards that mimic the

earlier coins and paper notes is another. More radical changes that have involved the eradication of coins and money and seen the migration to the plastic card have taken 20 years to become established. But perhaps computer power and total communication that you wear will not take quite so long. It is likely, however, to

introduce some very different and interesting modes of operation.


If I were talking to you face-to-face, and someone interrupted the conversation, we would think it most rude. However, if during our conversation, they were to call by telephone I would feel totally relaxed about momentarily ignoring you, answer-

ing the telephone, and starting a new conversation then coming back to you. This has become acceptable. How will you respond when, in mid-conversation, I gaze into a head-mounted screen, adjust controls on my wrist, and start a new conversation with some ethereal being or machine?

## book reviews

### TCP/IP Tutorial and Technical Overview

by Eamon Murphy, Steve Hayes and Matthias Enders



This is the fifth edition of this IBM Redbook. According to IBM, 'redbooks are superb technical publications written by and for technical and marketing professionals. ... Redbooks provide positioning and value guidance, installation and implementation experiences, sample code, typical scenarios, and step-by-step 'how-to' guidelines.'

The more cynical may latch on to the marketing input and see them as IBM glossies. Although there is an element of this, this book is of greater interest than that.

It goes without saying that the holy grail of open systems is interoperability. It is also a fact of life that, for many people, one of the big issues is how to interconnect IBM's offerings with those of other manufacturers. Therefore, to many people the IBM-centric view which this book offers will be entirely appropriate to the problems they are facing. For the more general reader seeking an overview of the TCP/IP protocol set,

this book has to compete with several impressive rivals.

The book starts with a good historical overview of the Internet and its predecessors such as ARPANET. The meat of the book, as one would expect deals with architecture and protocols. The basics of the TCP/IP architecture are followed by in-depth descriptions of the Internet protocol (IP), Internet control message protocol (ICMP), Internet group management protocol (IGMP), address resolution protocol (ARP), reverse address resolution protocol (RARP), user datagram protocol (UDP) and transmission control protocol (TCP). There is also a brief discussion on TCP/IP versus open systems integration (OSI) and IP: The Next Generation. Routing protocols are accorded the importance they deserve in their own section.

Key applications such as Telnet, file transfer protocol (FTP), domain name system (DNS), simple mail transfer protocol (SMTP) and multi-purpose Internet mail extensions (MIME), also get an in-depth treatment with examples from the IBM product offerings. This could either be a life saver or a mild irritant depending on how big your site's investment is in IBM kit.

The fifth section of the book deals with the connectivity options available to the various IBM TCP/IP product offerings. Again, the relevance of this to an individual will depend on whether the book is being read as a subject-area overview or as a precursor to solving an IBM-related interconnectivity problem.

The scanty coverage given at the end of the book to Internet access and navigation and distributed computing

environment (DCE) is not so much a criticism of this work, but more an indicator of the speed of change occurring in the area.

Overall, this is a well-written book of particular value to people with a requirement to interconnect IBM proprietary offerings and open systems from other vendors. It is of interest to the general reader seeking education on TCP/IP, but in that market place it is merely one of many books on the subject and some readers may want a more independent stance.

Published by Prentice Hall


ISBN 0134608585

£26.95 xxiii + 580 pp.

Reviewed by Robert Temple

### Instant Java

by John A. Pew



This book is the entry-level book in a series of four books from SunSoft. Its aim is to enable non-programmers to use Java applets, instantly.

There are a wide variety of applets described in the book, with ready-to-run examples and source code. Each applet and its options are described


in detail, with helpful pictures showing the effect on graphics or text.

Unfortunately, that's all the book does. If all you ever want to get out of Java is to be able to do clever things with rotating pictures and the like, then this may be the book for you; however, it left me feeling dissatisfied.

The books in the series are all packaged with the same CD, containing the software and code for all four books. You may find one of the more advanced books will teach you to do something useful with the language—and if you ever need to rotate flashing text in an offset frame with an embossed shadow, you will still have the applets available.

*Published by Prentice Hall  
ISBN 0135658217  
£24.95. xi + 340 pp. + CD-ROM  
Reviewed by Nicky Moxey*

**Java by Example**  
by Jerry R. Jackson and  
Allan L. McClellan



This book is the second in a series of four books from SunSoft: *Instant Java*, *Java by Example*, *Just Java* and *Core Java*.

The book is well structured, with clear, thorough examples logically arranged; after working through it, I have managed to write my first Java application! The book is not exactly a light read but it is laid out with lots of logical breakpoints, with chapter summaries.

Each of the books in the series comes with the same CD, a copy of Café Lite and an enhanced version of Winedit. Having read this book, I am


now busy going through the source code to try and consolidate all that information! All of the examples I've tried so far have worked.

The book also draws parallels between C and C++ coding, and talks you through incorporating non-Java code into applications. They warn that this can be complex; although I haven't tried it, the instructions looked clear.

For me, going through this book was a worth-while exercise. I now have a much clearer picture of the way Java is organised, with enough source code to take apart to keep me happy for some time. I know the difference between applets and javascript, and when to use each; I even discovered some useful ways of using applets! I now also feel that it will be worth me visiting some of the Java tutorials available on the net—they were meaningless before.

*Published by Prentice Hall  
ISBN 0135657636  
£27.95 xxv. + 345 pp. + CD-ROM  
Reviewed by Nicky Moxey*

**HTML for fun and profit**  
by Mary E. S. Morris



HTML for fun and profit is about writing hypertext markup language (HTML) pages for the World Wide Web (WWW). HTML is the language of the WWW. It is a derivative of the standard generalised markup language (SGML), a document formatting language that could be consistently implemented on different computer platforms. HTML uses hypertext linking conventions and a reduced set of formatting codes.

The book is based on a hands-on tutorial style and comes with a CD-

ROM containing shareware. The reader is expected to use this in order to follow the examples and understand the descriptions. It makes the content difficult to understand, for a complete novice, if one is not sitting in front of a computer and following the examples. However, for someone familiar with writing HTML documents, the book provides a good collection of information in one place.

The book is structured into 14 chapters, each ending with a summary of the chapter which includes a list of the HTML tags introduced with a brief description. After providing a condensed introduction to the Internet and the WWW, an overview of the software on the CD-ROM is given. This is followed by an explanation of how to use these for setting up servers on various platforms. Interestingly, the Windows 95 platform is not covered.

Chapters 2–5 describe how different types of formatting can be implemented using HTML tags. The user is gently taken from using simple tags, to linking documents, creating tables and incorporating multimedia into web creations.

In chapter 6, the common gateway interface (CGI) is introduced. CGI enables developers to tie scripts to web pages which can be used to receive information from clients, making the web pages more interactive. Chapters 7–9 build on this by showing how to create forms and process data from forms.

One of the main selling points of using the WWW for information distribution is that information can be navigated in a way which is non-linear. To make this effective, the documents to be put on the WWW need to be designed so that they are easy to navigate and have a consistent look and feel. Chapter 10 outlines guidelines to make web pages more appealing to all audiences. Writing HTML documents can be a vexed task if one has to do this in text mode. Fortunately a number of authoring tools are available to help with this. Chapter 11 explains the classes of available tools.

Chapter 12 considers practical issues in putting data out for public

## book reviews

---

consumption. This is followed by a chapter specifically devoted to using the Netscape Navigator. This is understandable as Navigator is perhaps the most popular browser. Netscape Navigator implements a superset of HTML 2.0 and some HTML 3.0 specifications, and, in addition, some of the proprietary tags used by Navigator may be included in HTML 3.0 specification. The chapter

on 'future directions' highlights the trends and directions of the world wide web, covering virtual reality markup language (VRML) and security issues in brief. The appendixes cover the use of Mosaic on different platforms, a complete list of HTML tags and lists pointers to additional information resources.

The book is best suited to people starting to write documents for

publishing on the web and are prepared to load the examples included on the CD-ROM. If you are already familiar with HTML, reading the book may refresh your knowledge and perhaps act as reference material.

*Published by Prentice Hall*

*ISBN 0-13-242488-6*

*£31.95. xxviii + 306 pp. + CD-ROM*

*Reviewed by Dipak Rajani*

## BT news

---

### BT To Merge With MCI

BT and MCI have announced that they have entered into a merger agreement to create a premier global company well positioned to take advantage of the rapidly evolving telecommunications industry.

The merger combines the substantial financial resources and global position of BT with the growth momentum and market expertise of MCI, known for its success in the competitive US long distance market.

The new company will be called *Concert plc*. The company, incorporated in the UK, will have headquarters in London and Washington DC.

Concert will operate under the BT and MCI brand names in the UK and United States respectively.

The merger will create a world-leading communications power house with annual revenues of over £25 billion, annual cash flow from operations of about £7.5 billion, and 43 million business and residential customers in 70 countries.

Concert will provide its customers with the full range of communications services, including local, long distance, international, mobile, multi-media, Internet-based services and systems integration for business customers.

Sir Iain Vallance, Chairman of BT, said: 'Concert will be exceptionally well placed to play a leading role in the major growth areas of the changing global communications market place.

'The complementary strengths and skills of BT and MCI will enable

Concert to take full advantage of the great opportunities provided by the forthcoming liberalisation of telecommunications markets in the US and Europe.

'We believe this merger will provide major benefits for the shareholders, customers and employees of both BT and MCI.'

Bert Roberts, Chairman and Chief Executive Officer of MCI, said: 'This merger creates the premier telecommunications company of the new millennium. Financial muscle, global customers and brands and customer-driven innovation will provide a strong competitive advantage as we open up communications markets, both domestically and around the world. Concert's scale will allow it to pursue major opportunities in new markets while maintaining the financial stability that comes from strong core businesses in the developed markets of the United States and the UK.'

### MCI US business

The merger with MCI will give BT's shareholders a full and direct exposure to the United States, the world's largest and most dynamic market place where 40 per cent of the world's largest multi-national companies are headquartered.

MCI's core business is the provision of domestic and international long-distance telecommunications in the United States, accounting for more than 90 per cent of MCI's operating revenues and operating income during each of the last three financial years.

The business has been growing rapidly and over the last four financial years revenues grew by an average of 12.6 per cent and earnings per share by 11.6 per cent per annum.

The passing of the United States Telecommunications Act 1996 is expected to open the United States market to full competition. In particular, it should enable MCI to compete in the local market from which it has been largely excluded to date. This market is immense, with total revenues in 1995 of nearly \$100 billion, and profitable (earnings before interest, tax, depreciation and amortisation in 1995 of \$43 billion).

By combining with BT, MCI will gain access to BT's technical expertise in the provision of local market products and services. In addition, the financial strength of the combined group will enable an aggressive move into this market at the earliest available opportunity.

With customer surveys supporting the desire for one-stop shopping, MCI believes that customers are very likely to award MCI their local telecommunications business in conjunction with embedded MCI services.

In its core business, MCI derives two thirds of its revenue from business customers. Increasingly, these customers have come to rely on MCI for sophisticated, value-added services. MCI believes that the combination of MCI's strong customer relationships and its robust portfolio of integrated services will be difficult for the Regional Bell Operating Companies to replicate.

## BT Tests Digital Audio Broadcasting in London

BT Broadcast Services has joined forces with a number of independent radio broadcasters, including GWR, Classic FM and London News Radio, to carry out a series of digital audio broadcast (DAB) experimental transmissions in the London area.

The transmissions will trial in-car data services, such as traffic and weather information, that can be carried within the DAB signal. The DAB data stream could offer a variety of interactive services including paying for a particular music track, stock market information or video pictures.

A DAB multiplexer, which will be sited at BT Tower, combines data services with high-quality audio feeds from the studios of each broadcaster. The combined signals, forming a DAB data stream, will be delivered over BT circuits to the BBC Crystal Palace transmitter for onward broadcast.

'DAB allows broadcasters to transmit both CD-quality audio and additional data at the same time. The radio industry is eager to evaluate the potential afforded by DAB's ability to deliver value-added data services,' explained John Swingewood, BT Broadcast Services general manager.

In the Birmingham area during 1996, BT provided circuits and transmitter access in a DAB trial with GWR and Classic FM to test coverage, quality of reception and carry out market research to gauge listeners' feedback.

## BT To Join Bharti Cellular Consortium

BT announced that it intends to acquire a significant stake in Bharti Cellular Ltd. (BCL), the largest mobile operator in India.

The BCL consortium presently consists of Bharti Group; STET of Italy; GMC—a subsidiary of CGE Group of France; Emtel—the cellular operator in Mauritius; and MSI UK.

Necessary applications have been made to the Indian Department of Telecommunications (DOT) for BT to

take over GMC, with its 22.5 per cent stake in BCL.

BCL operates one of the two cellular services in the Delhi license area comprising Delhi, Faridabad, Noida, Gurgaon and Ghaziabad. With more than 65 000 subscribers it has the largest subscriber base of all the Indian cellular operators. With this alliance, the Bharti Group will be able to draw on the strengths of two major European operators, BT and STET.

Sir Iain Vallance, Chairman of BT, said: 'Close cultural and historical links between India and the UK make our relationship with India special for BT. India is a key market for BT as it is one of the world's largest emerging markets for telecommunication services.

'BT believes that wireless has an increasingly important part to play in the telecommunications market of the future as fixed and wireless converge.

'We have mobile interests in the UK, France, Spain and Japan, and have applied with VIAG for the fourth mobile licence in Germany.'

Sunil Mittal, Chairman and Group Managing Director of Bharti Enterprises, said: 'This alliance would not only begin a new chapter in Indo-UK cooperation, but also reinforce global interest in the Indian telecommunications industry.

'It will have far reaching benefits, not merely for BCL but for the Indian telecommunications industry as a whole.'

## UK's First Managed ATM Network Contract

Halifax Building Society has awarded BT a new contract to provide and maintain the UK's first managed automatic teller machine (ATM) network solution. The managed network contract, valued at approximately £130 million over five years, comprises an ATM core supplied over BT's recently announced CellStream service, with specially-provided ATM nodes delivering the network to the Halifax's UK offices and branches.

BT has designed a managed ATM network, capable of transporting voice, data, images and high-quality video, to

give Halifax an enabling infrastructure to deliver new business services and applications to its branches and estate agencies. BT will manage the physical network, which includes multiple Cisco routers and StrataCom BPX switches, up to the router itself, leaving the Halifax free to focus upon delivering new, customer-oriented multimedia applications.

Halifax will benefit from the creation of a virtual office via an intranet, enabling dispersed cross-functional groups to work together across existing business boundaries to deliver new services to customers.

In the future, the managed ATM network will also provide Halifax customers with new alternative routes to access products and services by using new multimedia technologies such as the Internet, intelligent scripting, interactive voice response, interactive video and, ultimately, electronic kiosks and unmanned branches.

BT's CellStream network will be introduced in Halifax's corporate centre, located in West Yorkshire. The network will be rolled out to the Society's 1000 branches across the country over the next six months, with the 650 estate agency offices to be completed later in 1997.

John Miller, Business Strategy and Operations Director for Halifax, explained: 'Technology will play an increasingly important role in high street banking as we move into the next century, and the Halifax intends to lead the way. BT's ATM technology provides us with a managed infrastructure to meet not only our immediate data requirements, but also a flexible, strategic platform to deploy our new multimedia applications.'

## BT, MCI and Microsoft to Develop Global Intranet Services

BT, MCI and Microsoft have expanded their alliances to develop a new range of global intranet services. The new services will be marketed worldwide by BT and MCI and offered by Concert, the existing BT and MCI joint venture global communications company.

The portfolio of intranet services (private Internets) which are based on open Internet standards, will combine the global networking expertise of BT, MCI and Concert with the leading network and desktop applications offered by Microsoft. It will include Concert managed-data networking services provided by BT and MCI.

Concert InternetPlus Service, the world's first global Internet backbone

will support intranet transport for the new services. Along with Concert Frame Relay Service and other BT and MCI communications services, the new intranet range will offer customers the service-level assurances required to support vital intranet communications.

The intranet services will also offer the full range of Microsoft's messaging and intranet platforms, including the recently launched Microsoft Commer-

cial Internet System, Microsoft Exchange Server and the complementary intranet server products in the WindowsNT-based BackOffice family. In addition, they will utilise the Microsoft Internet Explorer 3.0 web browser and exploit the full capabilities of ActiveX.

BT, MCI and Microsoft are combining these capabilities in an expansion of their existing alliances to address the rapidly growing intranet market.

## industry news

---

### Nuremberg Interactive-Video Pilot Project

---

Deutsche Telekom has launched a pilot project for interactive video services in Nuremberg. For the first time in Germany, individual programmes with moving pictures will be transmitted via the telephone line. 100 subscribers living within the city limits will be able to select films offered via video-on-demand, obtain information-on-demand or engage in home shopping. The home shopping application was developed in collaboration with the mail-order company, Quelle. There are plans to steadily increase the products offered.

The system that was put together for the Nuremberg pilot project is made up of a server, set-top box and transmission system. Deutsche Telekom was responsible for adapting the individual components. The complex software was adapted and expanded by the Dresden-based subsidiary of Deutsche Telekom, Multimedia Software GmbH.

The transmission technique used in Nuremberg is asymmetrical digital subscriber line (ADSL) technology. With this technology, it is possible to transmit data streams asymmetrially and place calls via the telephone line simultaneously.

The Nuremberg pilot project is the second pilot project after Berlin that is currently in operation for interactive video services. The demonstration project in Berlin, in operation since February 1995, has continuously been expanded.

The system will be further tested in a pilot project that is scheduled to get under way in Cologne/Bonn early in 1997. There, the transmission platform will be a fibre-optic cable network.

### MCI One Makes International Calling Easier

---

MCI has announced a new international calling option for MCI One. MCI One customers who select this option pay one low rate per country when calling from home, when calling from anywhere in the United States with the MCI Card, and when accepting an international collect call using MCI's access numbers.

With MCI One's international plan, customers pay the same low rate whether calling Poland from home, accepting a collect call from relatives and friends in Poland using MCI's access number, or when using an MCI Card in the United States to make a call to Poland while away from home. MCI offers one low rate per country for all of these types of calls.

'We listen to what our customers want and continuously improve to meet our customers' needs,' said Lesley Weller, director, MCI International Marketing. 'MCI One gives our customers the freedom to choose how they want to keep in touch with their relatives and friends without having to worry about what time of day they call and what rate they will receive.'

Other telecommunications companies typically charge custom-

ers different rates when placing calls away from home or for international collect calls received. Now with MCI One, customers can call Poland from anywhere in the United States—such as the office or a friend's house—or receive an international collect call and still pay the same low rate.

'When using MCI One, MCI customers and their families gain control and freedom,' added Weller. 'They don't have to worry about cost or when and where they are placing an international call. Their family members and friends get the same benefit when calling collect using MCI access numbers—all for one low rate per country.'

### MCI To Transmit At 40 Gigabit/s

---

MCI is offering customers the fastest network in the world, with the deployment of advanced technology that enables it to transmit data and voice traffic at speeds of 40 Gbit/s, or  $4 \times 10^{10}$  bits of information per second. The technology also improves the performance of the network, with a bit-error rate that is 1000 times better than previously.

The 40 Gbit/s traffic will be carried initially on a 275-mile route between Chicago and St. Louis over MCI's existing single-mode fibre facilities, using technology developed with Hitachi Telecom and Pirelli.

MCI began deployment in January 1997, nearly two years ahead of schedule, having announced just a



year ago the deployment of 10 Gbit/s technology.

'Today, MCI is building the network for the next century—a network capable of supporting the future generations of high-bandwidth products and services that our customers will demand,' said Fred Briggs, MCI's Chief Engineering Officer. 'There were a lot of sceptics who said you could not run 40 Gbit/s on existing fibre. MCI has proved it could be done.'

MCI eventually plans to use the high-speed, high-capacity technology throughout its network, giving customers swifter access to a wide range of services, including interactive multimedia, teleconferencing and medical imaging. In teleconferencing, 40 Gbit/s speeds allow for the simultaneous transmission of 6400 channels of premium-quality video.

Briggs noted that the technology adds speeds and capacity to MCI's network, and does so very efficiently and at a significant cost savings.

'From the growth of the Internet to developing multimedia applications, MCI is seeing tremendous demand for bandwidth,' Briggs said. 'We are using state-of-the-art technologies to expand the capacity of our network to meet the needs of customers today, and in the future in a very cost-effective manner.'

MCI's 40 Gbit/s route contains two full-duplex channels of 10 Gbit/s terminal transmission equipment, bi-directional line amplifiers at four intermediate sites and OC-192 regenerator equipment in the middle. Two wavelengths of OC-192 travel in each direction on a single fibre, through a wave-division multiplexing (WDM) system, enabling telecommunications companies to increase their network bandwidth without the cost of laying new fibre.

### **Mondex USA**

A group of seven major United States organisations have announced that they are establishing Mondex USA Services Limited Liability Company (Mondex USA) to commercially develop and implement the Mondex electronic-cash system in

the United States. As a separate company from Mondex International Limited, Mondex USA will be owned by AT&T (through a wholly-owned subsidiary of AT&T Universal Card Services), Chase Manhattan, Dean Witter Discover (NOVUS), First Chicago NBD, MasterCard, Michigan National Bank and Wells Fargo. The Office of the Comptroller of the Currency has approved investment by each of the national banks in Mondex, which is necessary to implement Mondex in the United States.

'We are delighted to have such prestigious companies as owners of Mondex USA,' said Janet Hartung Crane, president and CEO of Mondex USA. 'Each owner brings a unique set of skills, expertise and geographic reach to Mondex. The power of this group will propel Mondex as the preeminent electronic cash payment system for the physical and virtual world in the United States.'

'With Mondex, consumers will be able to store electronic cash onto a smart card embedded with a micro-processor chip. They can then use the card to make small dollar purchases at stores and over the Internet. Consumers can load their card with cash from their own home, via telephones and personal computers, and from public locations like stores via cashless ATMs. Merchants will reduce their cash-handling costs and be able to make electronic cash deposits from their own store.'

Mondex is the only system which enables the direct transfer of electronic cash from one card to another. The cash value is immediately accessible without centrally routing and collecting each point-of-sale transaction. This feature makes it the most cost-effective electronic cash solution for both retailers and financial institutions.

Mondex will allow Internet commerce to expand rapidly by introducing a cash option for customers purchasing small-dollar services, such as information and entertainment. In the future, the Mondex platform will offer loyalty programs and personal identification applications added to the chip to increase the

card's utility for customers and merchants. The Mondex system has been piloted internationally in Swindon, United Kingdom, New Zealand, Hong Kong and Canada. In the United States, Wells Fargo has been testing the system among employees and merchants in San Francisco since August 1995 and AT&T has begun a pilot with employees using Mondex cards at the Jacksonville campus of AT&T Universal Card Services.

Mondex USA is an open system and its owners will be actively recruiting other financial organisations as licensees. Licensees will enjoy the full commercial opportunities of Mondex by offering Mondex cards to their customers and merchants. In addition, Mondex USA will work closely with technology and equipment companies on development efforts to further enhance Mondex's value for consumers and merchants.

### **Cable Deal For Schools**

Don Cruickshank, Director General of Telecommunications, welcomed the cable industry's new deal for schools, and announced that there was no bar to BT competing in similar vein. He said 'The cable deal for schools meets many of the needs identified by Oftel's Education Task Force. It is important for two reasons: first, it represents a positive step to improve schools' access to the developing range of interactive services, including appropriate use of the Internet. Second, it is an excellent illustration of collaboration on a strategic issue to strengthen telecommunications competition in the UK.'

'The Task Force has been looking at how the telecommunications industry can best play its part in improving the opportunity for schools, colleges and public access points to exploit the benefits of on-line technology. Usage of the Internet and on-line technology by schools, colleges and libraries has been limited because of concern about running up large and unpredictable call bills and the need for cheaper telecommunications links offering more capacity than ordinary telephone lines. The cable industry's

deal meets schools' immediate needs for affordable and predictable charges providing unlimited and faster access to on-line services. An important feature of this deal is that the whole cable industry is involved so that any school which has been passed or will be passed by the cable industry's network can benefit. The cable industry estimates that around 17 000 UK schools will be able to benefit from this deal by the turn of the century when the build in current cable franchise areas is complete.

'I have been challenged many times with the question: Would you allow BT to offer lower and non-usage based charges? I can say that the answer is unequivocally yes. I have written to BT setting out the rules I will apply if BT wishes to reduce its prices to schools in response to this cable deal. These rules give BT considerable scope to reduce its charges to schools providing: schools pay the same price regardless of location; the prices cover BT's incremental costs, on average, of offering the service to schools wherever they are in the UK; and any independent service provider can get the tariff on the same terms as BT's service provider arm so that schools have a choice of the service provider they use. This would be a competitive response which will give wide-ranging benefits for the schools as customers and for the community at large.

'I will be having further discussion with the rest of the industry about the precise details. The Task Force has identified a number of other challenges for the telecommunications industry. These include: the need to ensure that there are no 'have-nots' and that all schools can get access to a baseline of ISDN or its equivalent; and the need of further education colleges and public libraries for similar deals.

'The Task Force is well aware that access to networks and the cost of that access are but a small part of the obstacles faced by teachers in using communications technology to the full. But multifaceted problems need multifaceted solutions, and I

am very pleased to see the telecommunications industry tackling creatively and, I believe, effectively one significant part of the problem.'

### Oftel Determines Costs For Portability

Don Cruickshank, Director General of Telecommunications, announced that he has determined, for the first time, BT's costs and charges for geographic number portability.

Number portability is the ability of customers to retain their telephone numbers when changing operators. Some local cable companies and other licensed operators already offer this service to telephony customers moving from BT's network or returning to it. Oftel considers the availability of number portability vital to the promotion of competition in telecommunications.

As a result of the reference to the Monopolies and Mergers Commission (MMC) made by Oftel in 1995, BT's licence was amended to facilitate the introduction of number portability. The new licence provisions allow the Director General to determine BT's reasonable costs in providing portability and the charges BT can make to other operators in order to recover these costs.

Copies of the determination are available from Oftel's Library.

### Misuse Of Call-Queuing Systems

'Consumers are rightly concerned about the call-queuing systems used by a number of firms,' says Moira Black, Chairman of the English Advisory Committee on Telecommunications (ENACT), 'but adoption of good practice could soon put things right.'

Automatic call-queuing apparatus attached to a company switchboard accepts an incoming call and holds it until someone is available to answer it. This can be for a considerable period of time. However, the caller incurs call charges from the moment the call is accepted. 'Unfortunately,' says Moira,

'many consumers are still unaware of that. Although we publicised the problem several years ago, there is a steady stream of complaints from people who could not wait indefinitely and hung up, only to find themselves charged for a call which they regarded as unanswered. A warning in the system's voice message should be the norm.'

In most cases, voice messages do explain to the customer what is happening and are repeated at regular intervals. However, there are instances where acceptance of the call is hidden: no message is given and a ringing tone is produced by the system which is indistinguishable from the normal one. 'And that,' says Moira, 'is simply unacceptable. So is the failure of some systems to deal with calls in strict rotation. Waiting callers should never be leapfrogged by later arrivals in the queue.'

ENACT's advice to consumers is that, if they feel a call-queuing system operates in a misleading manner or that the time taken to answer their call is unacceptable, they should write to the managing director of the company concerned.

Peter Calver, Chairman of the Business Advisory Committee on telecommunications, agrees with Moira that improvement is needed. 'Call-queuing enables firms to service their customers more efficiently—the best systems even tell callers how many people are ahead of them in the queue. If all you get is an engaged tone, it tells you nothing and you may have to redial over and over again.

'Unfortunately, some businesses are using less-advanced call-queuing systems, which while reasonably effective, can cause problems at times of peak demand. Irrespective of the limitations of the technology, our aim must always be to make clear to the caller what is happening and to deal with waiting calls as quickly as possible. If we do not, then customers will start hanging up whenever they encounter a queuing system, and both businesses and their customers will lose out.'

# BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by  
The Institution of British Telecommunications  
Engineers, Post Point G012, 8-10 Gresham Street,  
London, EC2V 7AG.

Formerly *The Post Office Electrical Engineers'  
Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any  
statements made nor the opinions expressed in  
any of the articles or correspondence in *British  
Telecommunications Engineering* or the *Structured  
Information Programme* unless any such state-  
ment is made specifically by the Board.

© 1997: The Institution of British Telecommuni-  
cations Engineers.

*British Telecommunications Engineering* and the  
*Structured Information Programme* are printed in  
Great Britain by Headley Brothers Ltd.

## Subscriptions and Back Numbers

Annual subscriptions (including postage)—com-  
panies, universities, libraries and other bodies  
£50.00 UK, £55.00 overseas; private individuals  
(paid by personal cheque) £30.00 UK, £35.00 over-  
seas.

Single copies and back numbers—companies etc.  
£12.50 including postage UK, £13.75 including post-  
age overseas; private individuals (paid by personal  
cheque) £7.50 including postage UK, £8.75 includ-  
ing postage overseas.

Overseas customers should pay by a sterling  
cheque drawn on a London bank, a sterling  
Eurocheque or a sterling travellers cheque pay-  
able in London.

Price to BT employees: call +44 171 356 8050 for  
details.

Orders, by post only, should be addressed to *British  
Telecommunications Engineering Journal* (Sales),  
PP: G012, 8-10 Gresham Street, London EC2V 7AG.

Remittances for all items should be made payable  
to 'IBTE' and should be crossed '& Co'.

## Advertisements

All enquiries relating to advertisement space res-  
ervations should be addressed to The Advertise-  
ment Manager, *British Telecommunications  
Engineering Journal*, PP: G012, 8-10 Gresham  
Street, London EC2V 7AG. (Telephone: +44 171  
356 8050. Fax: +44 171 356 7942.)

## Communications

All communications should be addressed to the  
Editorial Office, *British Telecommunications Engi-  
neering Journal*, PP: G012, 8-10 Gresham Street,  
London EC2V 7AG. (Telephone: +44 171 356 8050.  
Fax: +44 171 356 7942.)

## Copyright

The entire contents of *British Telecommunications  
Engineering* and the *Structured Information Pro-  
gramme* are covered by general copyright and the  
Board of Editors' permission is necessary for  
reprinting extracts.

Authorisation to photocopy items for internal or  
personal use, or the internal or personal use of  
specific clients, is granted by The Institution of  
British Telecommunications Engineers for users  
registered with the Copyright Clearance Centre's  
(CCC's) Transactional Reporting Service, provided  
that the base fee of \$5.00 per copy is paid directly  
to CCC, 27 Congress Street, Salem, MA 01970,  
USA. For those organisations that have been  
granted a photocopy license by CCC, a separate  
system of payment has been arranged. Payment is  
additionally required for copying of articles pub-  
lished prior to 1978.

0262-401X/97 \$5.00 + .00



# THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

## IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone / Centre	Contact	Telephone Number
London	Terry McCullough	(01707) 601151
Martlesham Heath	Peter Shaw (in Ipswich)	(01473) 224488
Midlands	Mark Broughton	(0116) 253 4476
North East	Geoff Jenkinson	(0113) 237 8163
North West	David White	(0161) 600 2468
Northern Home Counties	Terry Birdseye	(01702) 372557
Northern Ireland	Allan McDougall	(01232) 215301
Paris	Garry Platt	+33 1 46 67 27 76
Scotland	Raymond MacMillan	(01463) 665515
South West	Peter Shaw (in Bristol)	(0117) 972 3216
Southern Home Counties	Nick Smith	(01227) 474823
Wales	Ferg Brown	(01952) 249886

**IBTE International HelpLine: +44 171 356 8008**

**IBTE On-Line (BT intranet): [http://mhwww.boat.bt.co.uk/ibte\\_nat/ibtehome.htm](http://mhwww.boat.bt.co.uk/ibte_nat/ibtehome.htm)**

## BRITISH TELECOMMUNICATIONS ENGINEERING

### Journal Subscription Order (Not BT Employees)

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH  
THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager  
Post Point G012  
8-10 Gresham Street  
London EC2V 7AG  
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£50.00 UK; £55.00 Overseas**  
Private individuals (to be paid for by personal cheque): **£30.00 UK; £35.00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose  
a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', should be enclosed with the order.  
Overseas customers should pay a sterling cheque drawn on a London bank, a sterling  
Eurocheque or a sterling travellers cheque payable in London. Alternatively, customers  
may wish to use a local subscription agent.)

Title (Mr/Mrs/Ms etc) \_\_\_\_\_ Initials \_\_\_\_\_ Surname \_\_\_\_\_

Mailing Address \_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

Postcode (Zip Code) \_\_\_\_\_ Country \_\_\_\_\_

Signature \_\_\_\_\_

Date \_\_\_\_\_

Please state the issue you wish your subscription to commence: \_\_\_\_\_  
(April, July, October or January)

VAT Registration Number (zero rate): 243 5917 54

*Photocopies of this form are acceptable.*

## Contents

Editorial: Welcome to Our New Readers	289
Integrating the Internet with the Voice Telephone Network	290
Callscape	293
Computer Telephony Integration for the Small Business	
Work Manager	300
A Service System for BT's Major Customers	309
Superhighways...Supergridlock?	315
Advocating a Managed Optical Platform	322
Using Software Agents for Business Process Management	326
Christmas in BT's Network Management Centre	334
Learning to 'Work Smarter' in a Virtual Business	340
How 'Smart' Can We Get?	346
Technology and Business Opportunities for Smart Cards in the Age of Electronic Commerce	
Organisational Learning	356
Early days in a change process	
Workplace Assessment for National Vocational Qualifications	363



**BT**

The New BT Phonecard

with this you can use one of these

£10

INSERT CARD THIS WAY UP

**MONDEX**

**MONDEX ELECTRONIC CASH**

OTHER