Customer requirements on teleservice management

J. Hall, I. Schieferdecker, M. Tschichholz

GMD-FOKUS

Hardenbergplatz 2, 10623 Berlin, Germany

Phone: +49-30-25499200, E-mail: hall@ina.tschichholz@fokus.gmd.de

Abstract

This paper examines some of the issues arising from customer requirements concerning the management of end-to-end services and guaranteed end-to-end quality of service. The implications of supporting the desired management capabilities both horizontally (inter-domain cooperative management) and vertically (from the service to the network elements) are discussed using as examples work currently being undertaken in two complementary projects.

Keywords

TMN, user requirements, inter-domain management, quality of service (QoS)

1 CONTEXT

The developments in the telecommunications world that are leading to an integrated broadband environment will result in an open service market where a variety of advanced multimedia services will be on offer in a competitive arena. These developments have been initiated by two main thrusts - liberalisation and advances in technology. Liberalisation implies greater consideration of customers' needs as teleservice providers will only be successful if their services are viable in the market place. The ability to meet customers' requirements will therefore play an increasingly important role not only in the national arena but also internationally as improved fast speed communication links promote the emergence of a global market place. Liberalisation in offering services is being accompanied by an evolution, if not revolution, in networking and information technology. High speed integrated broadband communications (IBC) over ATM, together with faster LANs, can support the transmission of multimedia streams, including audio, video, and text, over the same digital infrastructure. Highly sophisticated multimedia services will be able to provide support for cooperative working in a variety of areas and as users become familiar with the availability and flexibility of such services they will make greater demands on the services being provided.

The trend is therefore towards a service-driven market where services are offered on a competitive basis in response to specific customer needs. This increase in services being offered is creating new challenges for management, both in managing advanced multimedia services with different characteristics and quality of service requirements as well as in meeting customer demands for more control over the services that are offered.

These issues are being investigated in the European research and development programme RACE, which is promoting research into advanced technologies enabling the next generation of services to be created and openly available in an IBC environment. This paper discusses work from two RACE projects investigating how customer requirements affect, and are being supported by, management. Section two discusses customer and end-user requirements vis-à-vis the management of the telecommunication services they purchase and use. Sections three and four introduce two examples of how these requirements can be met by management, first for managing an organisation's virtual private network (VPN) as part of its corporate telecommunications network (CTN), and then for ensuring end-to-end quality of service in a multimedia collaboration service. Conclusions are presented in section five.

2 REQUIREMENTS

Corporate customers in an integrated teleservice environment are expected in the first instance to be organisations operating in a distributed, increasingly global, market where communications and the distributed handling of information are essential to success in their core business. Such customers are becoming more demanding and sophisticated; they must see the benefit from subscribing to a new service or to new features in an existing service, and this must be at a price that they are prepared to pay. Services are judged not only according to cost, but also on quality of service, which is defined as user satisfaction with service performance as it is perceived at the user interface, including service availability, reliability, and flexibility. Customers expect high levels of connectivity, bandwidth on demand, convenience, and teleservices tailored to their specific requirements, and they will select the services that most closely meet their requirements.

The impact of corporate customer and end-user requirements on the management of teleservices is being investigated here with respect to two groups of teleservice which have been selected as representative of the type of service that customers will be purchasing: VPN data services, and multimedia teleservices. VPN data services offer a more efficient and flexible alternative to leased lines for organisations wishing to connect geographically distributed sites. End-to-end multimedia teleservices will be offered by value-added service providers in a variety of areas, with multimedia collaboration services and multimedia mail among those currently being developed. In many cases multimedia services will be used in connection with a VPN service providing the underlying data communication service, and possibly offered together with a VPN service by the same service provider as part of a one-stop-shopping package.

In a competitive environment, customer requirements concerning management of the services purchased will determine what is offered. Customers using VPNs instead of leased lines for connecting their various sites will regard the VPN together with their local networks as forming one CTN (ETSI, 1993). Customers will therefore request a certain quality of service and control over their network and may wish automated access to the service provider's management system via standardised management interfaces that can interoperate with their own management systems to provide end-to-end CTN management. The customer management

services required include requesting end-to-end bandwidth dynamically according to the application being used over the VPN; modifying on-line the configuration of their VPN, for example by reallocating addresses; changing user profiles and access controls; obtaining information about the status and usage of the VPN; receiving notifications about faults or alarm thresholds being exceeded; or modifying QoS choices dynamically. The customer management system and the service provider management system have to cooperate to support such management functionality. For example, if the customer management system needs information about its VPN or wants to change the VPN configuration parameters, the service provider must offer the required functionality to the customer over a well-defined interface, and must also interact with the network management systems across which the VPN is offered in order to pass down the modifications requested by the customer to the network service operators.

As users become more familiar with multimedia teleservices, quality of service must be approached from the users' point of view and not from the network-oriented view as has traditionally been the case (Seitz, et. al., 1994). QoS attributes are user-specific and should be determined and controllable by users before and during service use respectively. Customers will not be interested in how a service is provided and over which networks it runs, but they will be interested in comparing services in terms of QoS user-oriented attributes that the service provider can support to suit their particular needs. The choice and trade-offs between the various options available must be comprehensible and meaningful to both customer and end user. As bandwidth alone will no longer be the determining factor, customers will wish to make their own decisions regarding quality of service and price, determining the kind of quality they wish to receive for the price that they are prepared to pay. When using a multimedia teleservice, users will want to specify their performance requirements and obtain assurance that these requirements can be met. If they cannot be met users may prefer to modify their original requirements. If during the course of using a service the QoS guarantee fails, the user should be informed and given a choice of alternatives, including terminating the service, or dropping one of the media streams (Ferrari, 1990).

The high performance necessary for real-time multimedia services encompasses not only speed, but also reliability and availability, integrity, operability, delay, and accuracy, including synchronisation accuracy between media streams. Different kinds of media stream require different levels of latency, bandwidth, jitter, and error resilience. Advanced multimedia services are time-critical and need management support for ensuring agreed qualities of service. In the case of multimedia collaboration services, management has to promote end-to-end service guarantees for the different time scales of a service. The overall quality of service depends on the management of both the service and all the networks on which it runs from end to end, including the customer's own networks. The capabilities of each end system also need to be considered. The end-to-end management therefore has to include management capabilities for each layer of the communication stack, i.e., for the multimedia service itself, for the end system, and for the various networks used for interconnecting the conference partners. Consequently, the end-to-end management has to adopt a vertical management approach. Specifically, it needs more elaborate network management interfaces which will be the basis for building the service management.

The above requirements from customers and end users highlight specific management challenges in an advanced service environment. In addition, teleservices such as VPN and multimedia collaboration are end-to-end services crossing several networks and administrative domains. Management of such services therefore involves both security and inter-domain management issues. Customers will only use the teleservices being offered if these services can

support customer security requirements concerning confidentiality and integrity of the data being handled both by the teleservices themselves as well as by the management services offered to the customer (O'Connell and Donnelly, 1994). Inter-domain management issues concern particularly the question of how several autonomous domains, including the customer premises network, can cooperate to provide end-to-end QoS management, and how network operators can make available the required functionality to service providers over the network/service management boundary. It also requires an understanding of the intra-domain enterprise network capabilities in order to integrate them with the end-to-end inter-domain capabilities (Tschichholz, et. al., 1995).

3 INTER-DOMAIN VPN SERVICE MANAGEMENT FOR CORPORATE NETWORKS

The RACE project PREPARE (PREPilot in Advanced REsource management) is investigating the issues arising from managing end-to-end IBC services spanning heterogeneous networks belonging to public and private organisations. IBC VPN services providing corporate networking facilities between customer premises networks (CPNs) have been selected as examples of end-to-end services requiring inter-domain cooperative management. The CPNs are part of the CTN and are therefore also involved in providing end-to-end services to users. PREPARE has been developing support for cooperative end-to-end management, focusing on performance and configuration management, including management support for guaranteeing quality of service. Customer management services are a specific subject of investigation in the project, particularly the interactions between the customer's management system and the service provider's management system. Advanced management services are being produced for the customer management system that enable it to obtain information about its virtual networking facilities and to make configuration and other changes to these facilities.

PREPARE has adopted the Telecommunications Management Network (TMN) framework for designing and implementing its inter-domain management system, including the relevant management services and management information model. However, in the many cases where standards did not exist, PREPARE took a pragmatic approach about where to place the management functionality. An inter-domain management architecture based on the TMN functional architecture (CCITT, 1992) was defined for the PREPARE testbed*. It consists of four subnetwork management domains and one or more end-to-end service management domains, or TMNs, which incorporate the customer management systems and the service provider management system and which interface to each subnetwork TMN (see Figure 1). The end-to-end services are managed by cooperating service operations systems that interact over the x reference point. The service provider is responsible for providing on-line service management capabilities to the customer at the x reference point (via the X interface at the service level) to allow, for example, the customer to reconfigure the VPN or to change QoS parameters. The customer management system is not concerned with the underlying networks and their management. It sees only the end-to-end service with the management services

* The PREPARE testbed consisted in the first phase of an ATM WAN and DQDB MAN providing the public network which connects customer premises networks comprising token ring LANs and ATM MUXs. This allows for both connectionless and isochronous VPN services. In the second phase (1995) ATM LANs are being connected via the European pilot ATM network.

offered to the customer being mapped to management actions across the service/network management boundary (Schneider and Donnelly, 1993).

Figure 1 The PREPARE Management Architecture.

In order to manage an end-to-end service there must be an integrated view of management between the CPN management system and the service provider's management system at the service level. This implies not only standardised interfaces but also a shared understanding of the management information being manipulated. Inter-domain and intra-domain management information models have been specified according to (ISO, 1992) using where possible existing managed object definitions from standards bodies and other organisations, such as the Network Management Forum. Internal intra-domain subnetwork information models have been defined for the subnetworks that enable the service-level functionality available to the customer management system to be supported within the individual TMNs of the subnetworks involved and made available over an open X interface from each TMN. The inter-domain information model designed for end-to-end management of the VPN services across the testbed ensures that all TMNs involved in managing the end-to-end service share the same understanding of the management information and of the management functionality in the form of the operations that can be performed on the managed objects, i.e., shared management knowledge (PREPARE, 1994). As the service-level information must be available on an end-to-end basis, a globally consistent view is maintained by an Inter-Domain Management Information service which enables the management information to be uniquely identified throughout the testbed and retrieved in a uniform manner wherever it is located (Tschichholz and Donnelly, 1993).

Some examples of the types of managed object are given as an illustration. At the service level, information about the end-to-end service is supported by the operations systems (OSFs) within the end-to-end service TMN, i.e., the service provider's domain and the CPN domains at the customer's locations (see Figure 2). In the CPN OSF, *cpn* gives the service provider details of the customer network that are relevant to operating the service at that site. In the same OSF are *endUser*, *te* (terminal equipment), *vdl* (virtual direct line), and *userStream*, representing a communication path between two or more end points which has specific QoS requirements and which can be booked in advance with a certain priority level. In the service provider's VPN OSF are held details about customers and their VPNs, for example, *pVLL* (public virtual leased line) which is a logical representation of public network resources available between customer access points in the same VPN and which is specified using a profile containing information about bandwidth, start time, end time, and quality of service, *userStream* which represents a communication path between two or more TEs, and *cap* (customer access point) (PREPARE, 1994). Managed object specifications have also been produced for each subnetwork in order to provide the functionality required to manage the end-to-end service within the subnetwork, such as monitoring the QoS parameters of a connection.

Figure 2 The End-to-End VPN Service Containment Tree.

Customers can be provided with a set of management services allowing them to manage the VPN service according to their particular needs. Management service functions supporting both end-to-end management across the testbed as well as allowing value-added service providers access to the bearer services of the subnetworks were specified. The service functions granting the customer organisation management capabilities concerning its VPN service and made available by the service provider at the X interface were classified into four groups. The first group is concerned with customer administration, providing a means of managing customer service information, such as the customer profile, list of permitted end users, customer access ports. A customer management system can retrieve information and also modify parts of the customer data using the services in this group. The second group consists of the end-user access management service functions which enable the customer to retrieve information about end users and to modify this information, for example, to add and remove end users both of the VPN service itself and of the VPN management service. The traffic and switching management

service functions in the third group enable connections between service end users to be created, modified and deleted. Information on the bandwidth for a connection can be retrieved and dynamically modified by the customer management system. The fourth group of service functions are concerned with service performance and quality of service, enabling service performance information to be retrieved and analysed and end-to-end connections to be tested. Notifications when the performance of the VPN goes beyond specified threshold values can also be transmitted to the customer management system. Management service functions have been defined for each subnetwork, so ensuring that customer requests can either be met at the service level or can be mapped to the network management level, for example, by allocating virtual paths through the ATM network (PREPARE, 1993).

Scenarios for two VPN services were designed in order to demonstrate the various activities that can take place during the provision and operation of the VPN services and which focus on different aspects of inter-domain end-to-end service management. They were specified in terms of the CMIS operations that the relevant management service functions are mapped down to. The scenarios developed for the connectionless bearer service show, for example, how information relating to the customer's VPN can be retrieved, how the customer can modify information, such as that relating to the customer's access bandwidth, in order to improve the quality of service, or add and remove allowed VPN connections between customer access points. The scenarios specified for the connection-oriented isochronous bearer service in conjunction with a multimedia teleconferencing application running over it enable the initial teleconference to be set up and bandwidth to be renegotiated during the session, and also provide a situation where QoS degradation during the conference has to be handled. For example, when reacting to the QoS degradation, the customer management system can invoke service functions such as *getPVLLStatus*, and can request changes to the quality of service with a *modifySession* request to the service provider's management system (PREPARE, 1993).

The scenarios were realised for the PREPARE demonstrator, which enabled the PREPARE TMN architecture to be validated and which showed how a self-contained design for a specific testbed could be used to investigate the requirements of customers and end users for active participation in end-to-end service management. Many of the issues are not adequately covered in the standards work, if at all, and so the work in PREPARE can, by testing many of the ideas in a real demonstrator, show how customer requirements on network management need to be met by value-added service providers and network operators and how the relevant management architecture and information model should be designed to support this from the service level down to the network element level.

The PREPARE work has demonstrated the need for an inter-domain architecture with clearly defined interfaces between different domains in order to provide cooperative end-to-end service management. An end-to-end service management information model allows the management operations available across the X interface to be clearly specified and made available to external TMNs. Current work is extending this architecture in order to investigate, among other things, customer management services for multimedia teleservices in an environment composed of multiple value added service providers and network operators.

4 QoS SUPPORT FOR MULTIMEDIA TELESERVICES

The RACE project TOMQAT (Total Management of End-to-End Service Quality for Multimedia Applications in IBC Networks) is investigating support for end-to-end QoS management for

multimedia services in an interconnected broadband environment and, in particular, how the customer perception of quality of service can be translated into network performance parameters and mapped to specific requirements on the network provider. It is designing and implementing QoS management, including management services based on appropriate information models and performance models. The project is using the TMN framework for its management architecture and is testing its approach over an ATM testbed* using the multimedia application Joint Viewing and Teleoperation System which has high QoS requirements as it combines video conferencing with (multimedia) application sharing between several participants in a wide area environment [Dermmler, et. al., 1993].

Figure 3 The Quality of Service Management Configuration.

The QoS requirements from customers and users need to be supported by a management system that can monitor the performance of the teleservice being used and make the appropriate modifications to ensure that the performance can be maintained. QoS parameters are subjective and so it is necessary to look at the QoS requested by users and map this to network performance parameters which are objective and can be measured. Performance measurements will provide the management system with information about the status of the network. If performance is not sufficient to support the required QoS, corrective action can be taken. There currently exists a lack of guidelines for end-to-end management of the QoS perceived by users, including the management of all factors affecting this quality. Total quality of service management is concerned not only with the network element and network TMN layers but also with the service layer and necessitates both static planning of network and component resources as well as the dynamic real-time management of QoS attributes of running multimedia services.

* The testbed is based on BALI (Berlin LAN Interconnection Net), an ATM infrastructure connecting ATM LANs in Berlin that is connected to the German and European pilot ATM networks.

Total management of end-to-end QoS means that the end-to-end system performance is relevant and that the global, end-to-end system has to be considered when making local choices about the best communication system to use. QoS management at the end system therefore includes the end-user communication stack as well as the cumulative effect of the performance of each subnetwork constituting the whole end-to-end network. Management is used to tune the network performance by observing network element performance, with performance models aiding the decision-making process. This configuration is shown in Figure 3.

In general, a multimedia application is realised by a number of different service elements supporting, for instance, the audio and video communication between end users, the coordination of joint processing on shared documents or the session management to establish/renew or release connections. Some of these service elements carry functionalities of local influence (for example, the control of local peripheral devices such as cameras, monitors, microphones), whereas others provide functionalities of global significance (such as joint editing). It is not sufficient to rely exclusively on the information and management capabilities of the used service alone but on the whole end-to-end network. More information about the data flow within the network is needed so that suggestions can be made, based on long-term end-to-end observations, for new network parameters/thresholds to be used within the network control strategies. It should also be possible to tune the network using global information, as opposed to local information available to the network. This implies a more elaborate management interface between network operator and service provider to allow for network monitoring down to network element monitoring and network tuning.

A management architecture has been designed to meet the management requirements for total quality of service (see Figure 4). It has adopted the domain approach as a means of structuring the environment and has designated QoS managers with two kinds of roles: a Service QoS Manager (SQM) for each domain and a local End-System QoS Manager (ESQM) for each end system. A master SQM will be provided by the value-added service provider of the multimedia collaboration service (TOMQAT, 1994).

The ESQMs realise QoS management functionalities in the user's local environment and rely on network element (NE) functions offered in this environment. Basic network element functions in an ATM network are, for instance:

- *ATM switch*: set-up thresholds for traffic management strategies, change routing tables.
- *Measurement equipment and mediation devices*: collect data, alarm on performance degradation.
- *Multimedia application*: change compression algorithm for audio or video, adjust frame size of a video at the source, adjust playout buffer.

An ESQM will support QoS management functions such as: QoS negotiation; monitoring of NEs to collect performance-oriented data; QoS analysis and evaluation to derive network performance parameters and to forecast future or to determine current QoS bottlenecks; QoS decision-making based on appropriate performance models to decide on better/optimal parameter settings for NEs; and local QoS configuration by tuning NEs with the above identified parameters to remove performance bottlenecks. The SQMs realise QoS management functionalities that are responsible for coordinating the multimedia application within a specific domain, and for maintaining required global end-to-end user QoS characteristics between different domains. Typical functions offered are: QoS negotiation coordination; global monitoring of network performance and evaluation of end-to-end QoS; global decisions on NE parameters settings in the whole network using end-to-end performance models; and global

network tuning in order to guarantee end-to-end QoS. The QoS management strategy refers to the way of finding better/optimal parameter settings for the required QoS. After predicting/observing a QoS degradation the manager will either tune the network or renegotiate QoS requirements with the end user.

Figure 4 The Quality of Service Management Architecture.

One example - QoS influenced routing - is considered here. During connection set-up QoS requirements are specified by the end user. In the case of the video stream, its frame refresh rate and frame size are given. Due to costly communication resources, the multimedia collaboration service uses a variable bit rate (VBR) service with sustainable cell rate (SCR) and peak cell rate (PCR) suitable for video transmission. A virtual channel with corresponding SCR and PCR will be established. The QoS requirements for the multimedia collaboration service are forwarded to the ESQM responsible for maintaining the contracted QoS.

The throughput and cell error rate of a specific virtual channel is measured using ATM measurement equipment which is connected to a physical link of the ATM network. Due to overload and congestion in the ATM network, the throughput at some communication link may drop below the specified SCR, which implies an unacceptable QoS degradation for the end user – the video might simply stop. This causes the generation of a notification which will be forwarded to the ESQM via the operations system of the measurement equipment. In order to re-establish the contracted QoS the communication link can be tuned. The ESQM must have knowledge about alternative links and their mean load observed in the past. Based on this knowledge, the ESQM can choose a new route for the virtual channel that is used for transmitting the video stream. It will modify the routing tables of some switches to avoid the congested link. If it is not possible to achieve the required bandwidth through alternative links,

renegotiation of end-user QoS requirements will be necessary. The user will be informed that only smaller video frames or a lower frame rate can be supported. The user can decide whether to reduce the video in size or the frame rate. Of course, another possibility for the user will be simply to terminate the connection.

TOMQAT is producing a series of similar scenarios such as QoS based call admission control, playout delay adaptation, and end-to-end application control, showing that QoS can be maintained by using the QoS management service but will degrade crucially otherwise. Special emphasis is being placed on the mapping from end users' QoS requirements down to the network element performance. Based on this, appropriate monitoring functionalities to observe the local and global network state as well as network tuning strategies at each vertical layer of the communication stack will be derived. As TOMQAT aims at developing a running QoS management environment, deeper insights into the layered management architecture are expected. As far as the management interfaces of public network providers are concerned, it has emerged that more management capabilities are needed at the boundary between network and service provider in order to offer QoS management to the end users.

5 CONCLUSIONS

The work discussed here is investigating how customer and user requirements can be met by management in an advanced service environment. The result of the work so far shows that any management capability offered to customers and end users must be underpinned by an architecture and corresponding information model that structures the essential management functionality and responsibility at all levels of this environment.

The development of management services for the customer management system in a CTN has emphasised the significance of a standardised open interface over which management functionality can be offered by the service provider so that it can be used with the customer's existing management system. The work has shown that such functionality, although offered at the service level, must be mapped to operations offered by each of the underlying networks supporting the end-to-end service. A modelling approach is required that can structure the management functionality in an architectural framework, specify the management information model that is needed in each TMN, and designate the operations available at each interface in the architecture. Without this underlying support, customer management requirements cannot be met. End-to-end management of teleservices enables end-to-end services to be managed efficiently and their quality of service to be supported. The work in designing QoS management for such teleservices has shown that the high requirements from demanding multimedia applications necessitate a well thought-out performance management and quality control system in order to guarantee end-to-end QoS to the end user.

In a heterogeneous environment cooperative management between the customer, service provider and network operator entails both an end-to-end inter-domain (horizontal) approach providing cooperation between customers and service providers, as well as a vertical layered approach ensuring management capabilities at each communication level in order to guarantee end-to-end QoS. The boundary between service and network management is proving to be crucial as it is here that the interaction between network operator and service provider takes place. The conclusion from the work discussed in this paper is that the current functionality offered by network management is not sufficient to support end-to-end service management for quality of service requirements and that work on the end-to-end management support and its

interoperability with intra-domain management is essential before customer and end-user requirements can be satisfied. The two projects discussed here are contributing to this work.

ACKNOWLEDGEMENTS

The authors wish to thank their colleagues at GMD-FOKUS for many fruitful discussions and the partners of the PREPARE and TOMQAT consortia for their contributions to the ideas presented here. This work was partially supported by the Commission of the European Communities (CEC) under projects R2004 PREPARE and R2116 TOMQAT of the RACE II programme. This paper does not necessarily reflect the views of the PREPARE and TOMQAT consortia.

REFERENCES

- CCITT (1992) Principles for a Telecommunications Management Network, CCITT Recommendation M.3010, ITU, Geneva.
- Dermler, G., et. al. (1993) Constructing a Distributed Multimedia Joint Viewing and Tele-Operation Service for Heterogeneous Workstation Environments, in *Proceedings Fourth IEEE Workshop on Future Trends of Distributed Computing, Lisbon*.
- ETSI (1993) Strategic Review Committee on Corporate Telecommunications Networks. Report to the Technical Assembly, SRC5 Final Report.
- Ferrari, D. (1990) Client Requirements for Real-Time Communication Services, *IEEE Communications Magazine*, **28**, 65–72.
- ISO (1992) Information Technology – Open Systems Interconnection – Structure of Management Information, Part 4: Guidelines for the Definition of Managed Objects, ISO/IEC International Standard 10165-4.
- O'Connell, S. and Donnelly, W. (1994), Security Requirements of the TMN X-Interface within End-to-End Service Management of Virtual Private Networks, in *Proceedings of the RACE International Conference on Intelligence in Broadband Services and Networks, Aachen, September 1994*, 207–217.
- PREPARE (1993) CNM and PNM Specification Based on MIS, PREPARE deliverable R2004/LME/WP6/DS/I7011/b1.
- PREPARE (1994) Final TMN Information Model Specification, PREPARE deliverable R2004/BRI/WP2/DS/P/017/b1.
- Schneider, J.M. and Donnelly, W. (1993) An Open Architecture for Inter-Domain Communications Management in the PREPARE Testbed, in *Proceedings of the 2nd International Conference on Broadband Islands, Athens, June 1993*, 77–88.
- Seitz, N.B., et. al. (1994), User-Oriented Measures of Telecommunication Quality, *IEEE Communications Magazine*, **32**, 56–66.
- TOMQAT (1994) Architecture of the TOMQAT System and Definition of the Net Infrastructure, TOMQAT Deliverable R2116/TUB/WP2/DS/P/006/b1.
- Tschichholz, M. and Donnelly, W. (1993) The PREPARE Management Information Service, in *Proceedings of the RACE International Conference on Intelligence in Broadband Services and Networks, Paris, November 1993*, IV/3/1–12.
- Tschichholz, M. et. al. (1995) Information Aspects and Future Directions in an Integrated Telecommunications and Enterprise Management Environment, to be published in *Journal of Network and Systems Management*.

BIOGRAPHIES

Jane Hall is a senior scientist at the Research Institute for Open Communication Systems (FOKUS) in the Management in open Systems (MinoS) group at GMD Berlin. She has been working in several European projects (COST11, ESPRIT, RACE) in the area of network and service management. Her current research interests are quality of service management and management of teleworking environments.

Ina Schieferdecker graduated in mathematics at the Humboldt University in Berlin in 1990 and received her PhD in computer science from the Technical University of Berlin in 1994. She is currently working at the Research Institute for Open Communication Systems (FOKUS) at GMD Berlin in the Performance Laboratory. Her current research interests are formal specification, verification and performance evaluation of communication networks and management systems.

Michael Tschichholz received his diploma in computer science from the Technical University of Berlin in 1982. He has been working in the area of open communication systems (E-mail, Directory, OSI Management, TMN) since 1980, and is actively contributing to international standardisation work. He is working at GMD-FOKUS and is the head of the Management in open Systems (MinoS) group. He has participated in several national and international management related projects. His current research interests are related to multi-domain management based on TMN and ODP.