An Inter-domain Virtual Private Network Management Service

David Lewis
University College London
Computer Science Department, Gower Street, London, WC1E 6BT, UK
tel:+44 171 391 1327
fax: +44 171 387 1397
email: dlewis@cs.ucl.ac.uk

Lennart H. Bjerring
L.M. Ericsson A/S
Sluseholmen 8, DK-1790 Copenhagen, Denmark
tel:+44 3388 3057
fax: +44 3388 3129
email: lmdlhb@lmd.ericsson.se

Ingi H. Thorarensen
L.M. Ericsson A/S
Sluseholmen 8, DK-1790 Copenhagen, Denmark
tel:+44 3388 3312
fax: +44 3388 3129
email: lmdiht@lmd.ericsson.se

ABSTRACT The evolution of an open market for telecommunication services provides a wide range of opportunities for the provision of value added services by providers other than public network operators. One service already available today is that of a Virtual Private Network that provide dispersed corporate sites with wide area data and telephony capabilities using public network services. The introduction of B-ISDN allows for the provision of VPN services in a much more integrated fashion than currently. However in the short to medium term any effective VPN service must be able to deal with the technological and organisational heterogeneities that are present when providing a high level of service functionality over an arbitrary number of administrative domains.

This paper presents the design of a management service for a VPN that addresses some of these multi-domain and network heterogeneity issues. It outlines how a design based on the ITU-T's TMN recommendations was developed and how working prototypes were developed over a series of demonstration testbeds based on real broadband networks. (This work was partially funded by the Commission of the European Union as part of the PREPARE project under contract number RACE 2004)

Introduction

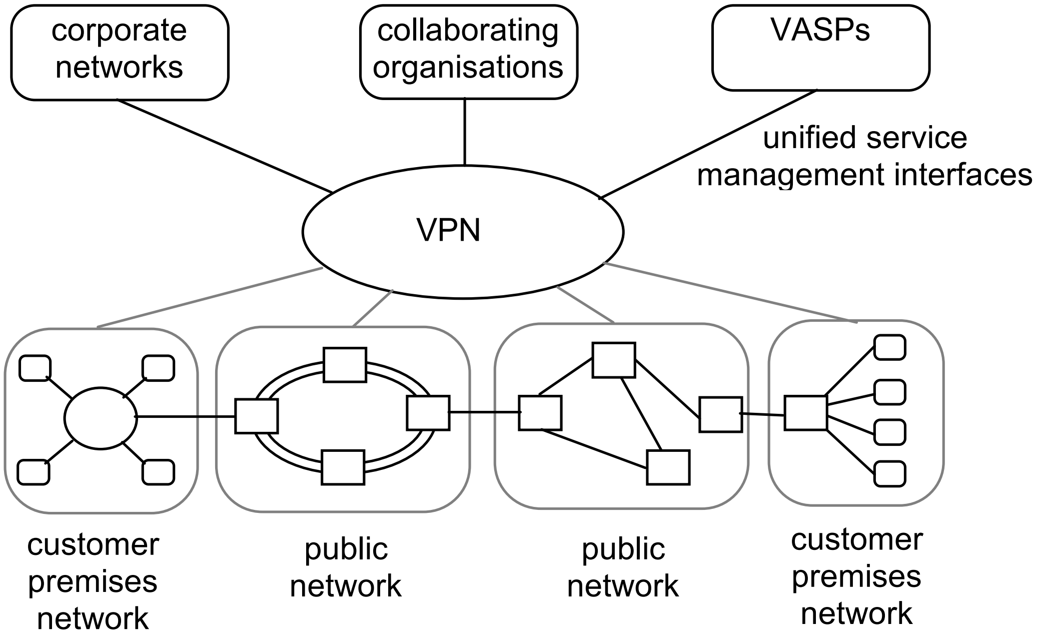
corporate networks
collaborating organisations
VASPs
unified service management interfaces
VPN
customer premises network
public network
public network
customer premises network

The general aim of this work was to design, implement and demonstrate a management service that would address the requirements of a VPN service operating in a future, liberalised telecommunications services market. A key aspect of this work was therefore the clear identification of the various organisational stakeholders that would be involved in the provision and consumption of VPN services, and the relationships between them.

Though it currently provides little in the way of inter-domain service layer interfaces the ITU-T M.3000 series of recommendations (TMN) was adopted as the basis for the specification work. This provided a clear framework within which to develop management services and provided standard interfaces for network and network element management that could be integrated into the service. In addition several implementation platforms were available. A TMN architecture was therefore developed for the VPN management service in which each administrative domain (mapped from stakeholders) would operate its own TMN with operation systems functions (OSFs) at the service layer and also, if required, at the network and network element layers. Communication of VPN service related information would only be performed over interfaces between the service layer OSFs via the TMN x reference point. The relationships between stakeholders provided the basis for defining these inter-domain interfaces.

The target scenarios for the VPN covered support for the communications requirements between sites of a large multinational corporation; between internationally distributed, collaborating companies and for Value Added Service Providers (VASP), e.g., multimedia conferencing or multimedia mail providers. The service provides a unified interface, i.e., one stop shopping, to the customer for several management functions, e.g., configuration, fault, performance and accounting management. The VPN management service operates end to end, i.e., between terminal equipment running distributed applications communicating over the VPN. This implies that management of the customer network is also covered by the VPN service.

The VPN management service aims to be essentially technology independent, in that its basic design does not assume some specific underlying network technology. Where there is a requirement for the VPN service to communicate technology specific information between domains this is done through well defined specilisations of the model for specific technologies. This allows the VPN service to be instantiated over different types of network technologies.

Stakeholders and Relationships

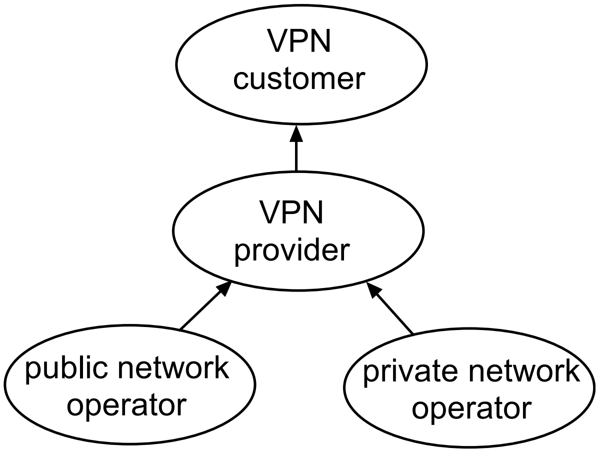

The analysis of requirements for the VPN management service was based on the identification of the stakeholders involved in the provision and consumption of the service. Primarily this is the *VPN provider*, which offers the service in an open services market, and the *VPN customer* which subscribes to the services and pays the provider for its use. In addition however the following stakeholders, which are related to the primary ones by the provision of the underlying network services, are also identified:

- *Public Network Operator*- This is the operator of a network which offers its services on the open service market and which are used by the VPN provider in the provision of end to end network resource management services in the public domain to the VPN customer.
- *Private Network Operator*- This is the operator of a network which offers operational management services to the VPN provider in much the same way as the Public Network Operator does, but without a public service obligation. This has implication on how some management activities, e.g. fault resolution, are reported and dealt with by stakeholders.

For the primary relationship between the VPN provider and the VPN customer, configuration, fault and accounting management services are provided. Configuration management services include the identification of VPN end points, the reservation of network resources between these end points and the allocation of network resources which constitute end to end communication paths. Fault management services involve the identification of network faults that effect the VPN service, the extent of that effect and the responsible administrative domain. Accounting management involves the charging for usage of network resources and management services through the assembly and delivery of itemised bills. The private and public network operators provide network related configuration, fault and accounting services related to their specific network domains to the VPN provider.

By performing the requirements analysis at this level of atomicity it is possible, from a mapping from management services to stakeholder relationships, to determine which management services are required by an administrative domain playing two or more stakeholder roles- a situation likely to occur in practice.

VPN Resource Model

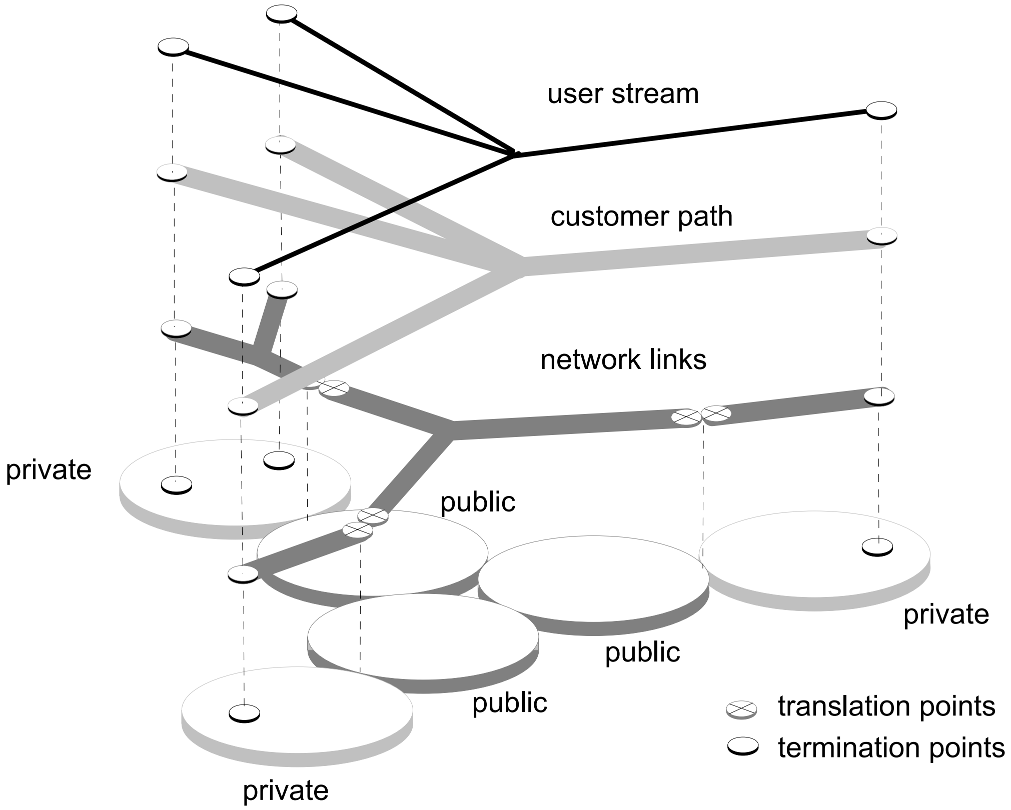

This figure shows the core structural concepts identified in the information viewpoint of the VPN design and how they map onto the different types of network domains involved.

This information viewpoint needs to provide a basis for the required configuration, fault and accounting management on an per domain basis. When discussing domains, however, the primary distinction drawn in the view seen by the customer is between the domains that the VPN provider takes responsibility for on behalf of the customer (i.e., the public network operator domains) and the ones which the VPN provider monitors but does not take responsibility for in case of a failure (i.e., the private network operator domains).

The basic end to end concepts are *customer paths* and *user streams*. Customer paths define the end points which may potentially communicate with each other as well as the total amount of network resources available, i.e. maximum quality of service, that is available between these end points. The customer path is therefore a mean of reserving network resources for a customer. A user stream is a unit of resource allocation. It defines a group of end points between which communication has been enabled with a specific quality of service. The aggregation of network resources allocated to user streams available in any section of a VPN must not exceed the total resources reserved by a customer path in that section.

To provide a domain specific view the resource reservation topology of the customer path is broken down into *network links* that show resource reservation topologies for either individual private domains or for collections of adjacent public network domains. Network faults can then be mapped to changes in the operational states of network links and thus the stakeholder responsible for dealing with the fault is easily identified.

TMN Architecture

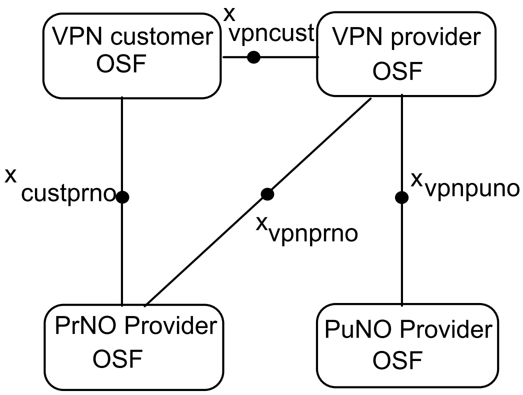

The figure above shows the functional TMN architecture for the VPN management service for a stakeholder model where the customer of the VPN uses the private networks to provide some value added service to its own customers. This architecture identifies several different inter-domain TMN reference points; the $_{vpncust}$ reference point between the VPN customer OSF and the VPN provider OSF; the $x_{vpnpuno}$ reference point between the VPN provider OSF and the public network OSF; the $x_{vpnprno}$ reference point between the VPN provider OSF and the private network OSF and the $x_{vpncust}$ interface between the VPN customer OSF (acting in the private network customer role) and the private network OSF.

The $x_{vpncust}$ reference point provides the customer with access to end-to-end VPN resource reservation, resource allocation, fault and accounting management services from the VPN provider and is expressed purely in terms of VPN concepts described previously. The $x_{vpnpuno}$ reference point allows the VPN provider to access configuration, fault and accounting services from the public network operator. These services can take two forms. They can either be expressed purely in terms specific to the network technology used in the network domain being represented, e.g., an ATM virtual path service, or they can be expressed in terms of VPN concepts in a similar way to those available over the $x_{vpncust}$ reference point. This latter case allows us, therefore, to effectively have multiple VPN providers since the $x_{vpnpuno}$ reference point could provide access to a further layer of VPN service provision.

Service over the $x_{vpnprno}$ reference point can also be expressed in technology specific or VPN terms. However the wide range of different private network types renders the technology specific case largely impractical and they have therefore been defined as VPN based services within our implementation. These are resource reservation, resource allocation and fault management services for the domain concerned. The $x_{vpncust}$ reference point, for similar reasons, also provides a VPN based view of the private network state to the VPN customer. In this way the VPN customer and provider share a clear view of whether VPN faults originate in the private or public networks.

Technology Specialisation

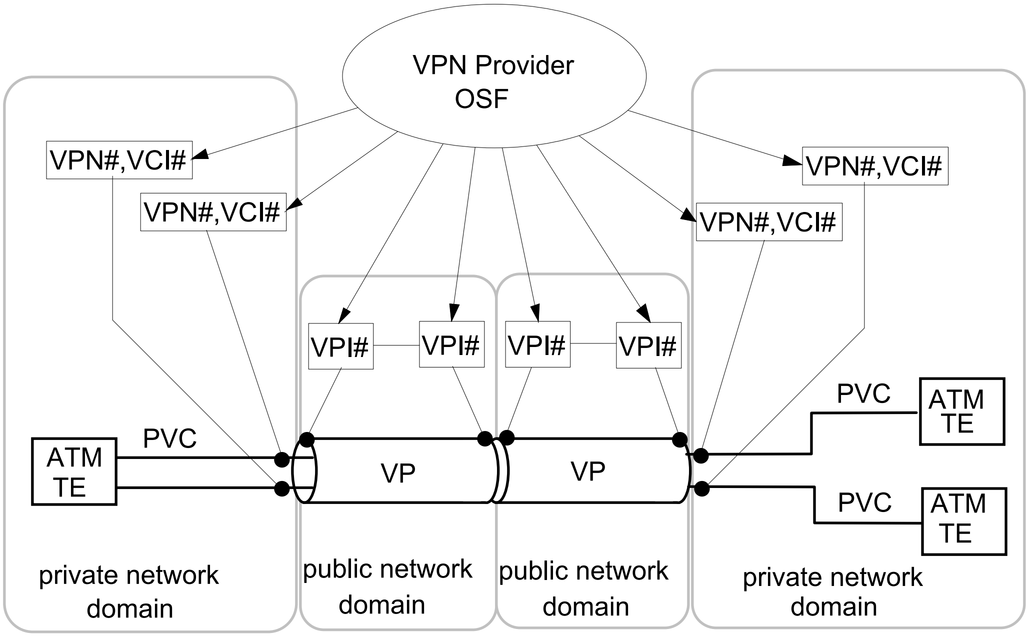

The VPN concepts described on the previous pages allow us to build an abstract model of network resource management that provides an end to end view as well as a view detailing the breakdown of end to end entities into entities relevant to specific domains. However to apply this to a real situation involving different network domains, the model has to be refined to incorporate the network technology specific information that need to be exchanged at network domain boundaries in order that the physical networks can interwork. The coordination of the exchange of information required for end to end connections is performed centrally by the VPN provider thereby protecting the customers and network operators from having to deal with the full array of different network technologies.

The diagram above shows an example where two public networks offering an ATM virtual path (VP) service are used to interconnect terminal equipment (TE) in private networks that provide for ATM Permanent Virtual Circuit (PVC) set up. In this example the VPN provider must set up an end to end connection in response to a request for point to multipoint user stream from the VPN customer. This therefore requires the VPN provider OSF to coordinate the VPI value at the boundaries between the public networks for VPs it requests across each of these network domains. The VPN provider OSF is also required to coordinate VPI and VCI values for the individual connections between TEs within the private networks.

To exchange this information over the x reference points identified in the VPN TMN architecture, the general VPN concepts are refined to include connection description information at translation points for each user stream. These can be specialised to accommodate the parameters needed for specific network technologies.

Information Model

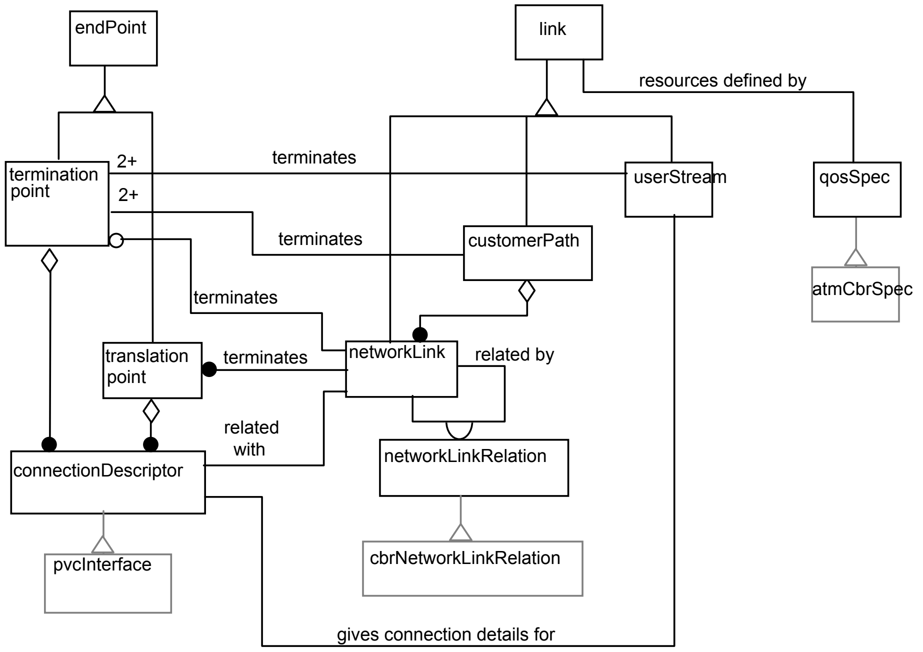

The figure above shows the overall information model for the VPN service expressed in OMT notation. The core concepts of customer path, network links, user streams, termination points and translation points have all been mapped straight onto this information model.
Additional objects have then been introduced in order to deal with the details of grouping related network links at translation point (networkLinkRelation), managing individual connections between termination points within a user stream at the boundaries of networks (connectionDescriptors) and managing quality of service specifications for customerPaths, networkLinks and userStreams (qosSpec). These additional objects therefore contain the network related information needed to map the high level VPN abstractions onto network level abstractions. These objects are therefore inherited from when supporting a specific network technology.

Information objects for dealing with the ATM public and private networks example from the previous page and the management of constant bit rate quality of service over them are inherited from generic objects and shown shaded above.

Individual TMN interfaces based on these objects are defined in more detail by the detailed GDMO definition of these objects and the subset of them available at a specific interface. Additionally information flows detailing the sequence of CMIP primitives that flow over an interface for a particular management scenario, e.g., resource reservation set-up, are required to define the dynamic behaviour of the interface.

Implementation

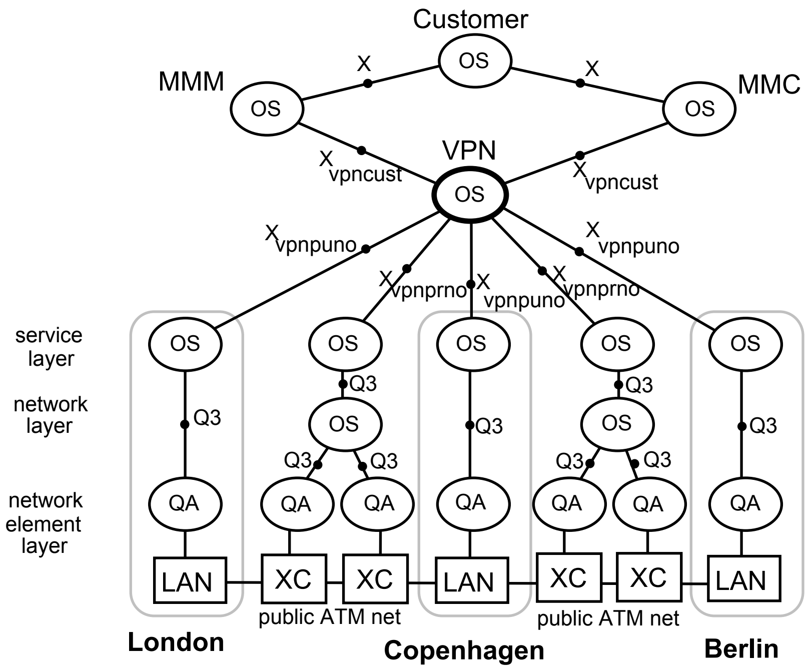

The first implementation of the VPN management service was completed and publicly demonstrated in December 1994 and was assembled in Copenhagen, Denmark. The physical network consisted of an ATM WAN (the Danish national ATM pilot BATMAN), and a DQDB MAN acting as the public networks and ATM cross connects (XC) and Token Rings acting as private networks. This implementation was designed to support scenarios where a single VPN customer owned the private networks and connected them via the heterogeneous public networks using the VPN management service. Configuration, performance and fault management scenarios were performed.

The second implementation, pictured above, was designed to support a more complex scenario where the VPN customers were a multimedia conferencing (MMC) provider and a multimedia mail (MMM) provider. These customers used the VPN provider's services as part of the service they offered to their own customer, which in turn owns the private networks. The public networks were implemented as three ATM cross connects from DSC and one from IBM, situated at sites in London, Berlin and Copenhagen. These were connected by Virtual Paths over the European PNO ATM Pilot. The private networks were ATM LANs based on switches from Fore Systems.

The various service and network OSs for both testbeds were implemented on a variety of platforms including OSIMIS (a research platform developed by the RACE project ICM and the ESPRIT project MIDAS), Ericsson's TMOS platform, Hewlett-Packard's OpenView and IBM's TMN/6000. Network element management applications were implemented ar Q-adaptors (QA) using TMN/6000, or Q3ADE from UHC, or used existing SNMP MIBs through the OSIMIS SNMP-CMIS proxy agent. A public demonstration in November 1995 showed configuration, fault and accounting management functions being exercised using this architecture.

Assessment and Conclusion

Our experience has shown that it is possible to achieve flexible and powerful management of customer services on the basis of a very simple conceptual model. Based on the TMN and OSI Management principles the VPN management interface information model consists of only a very few generic MO classes which can be easily specialised to accommodate technological specificities of individual networks. By this we have achieved our design goals of coming up with a truly generic concept which allows the implementation of broadband data VPNs of a variety of network technologies with a minimum of adaptations.

This approach was greatly eased by the use of the OSI Management object oriented information modelling concepts. From a methodological point of view it proved to be not only very useful but also necessary to combine a top-down approach with a bottom-up approach. The first approach analysed the multi-stakeholder characteristics of an open service market and provided a very broad view of the functional and non-functional requirements imposed on the VPN model by a multiplicity of stakeholders with diverging objectives. The second approach revealed a lot of low-level technical problems and limitations of technologies which needed to be taken into account in order to achieve a working system. The iterations between the two approaches in all stages of the development assisted us in arriving finally at a robust concept which has been proven to work in real implementations.

The use of TMN as the architectural framework provided a stable ground for the development. However, the lack of service layer management information models a problem of growing importance. As such the VPN management concept presented here can be regarded as evidence that TMN is a feasible technique, however a major standardisation effort is required to produce the amount of service layer management specifications needed in order to facilitate an open service market for management services.

The VPN management service described here has been developed through analysis of open service market needs and validated through implementation over real broadband testbeds. The VPN service is now entering its third phase of implementation. The EU ACTS project PROSPECT investigates new applications of advanced IBC teleservices based on the VPN service and these will be fielded in a Pan-European trial involving groups of end users using tele-education applications.

The VPN design has been documented fully in the Network Managment Forum Ensemble format and has been submitted to the NMF for comments and as input to the formation of a broadband managment group within the Forum. It is our hope that by publishing this Ensemble we will see both new implementations based on the core concept, as well as more confidence in TMN as an appropriate implementation option for advanced inter-domain management systems. The VPN model and implementation thus provides important indications and input to standards bodies working on the development of service layer information modelling principles and standardisation of the TMN X interface.

References:

“Methodologies for Multidomain Management Service Design”, D.Lewis, L.H.Bjerring, J. Hall, T.Tiropanis, Proceeding of the IS&N 95 conference Crete, Greece, 1995

“End-to-end communications management using the TMN X interface”, J.Schneider, W.Donnelly, Journal of Network and Systems Management, Vol.3, No.1, June 1995

RACE Common Functional Specification H412, Management of IBC VPN services in the IBC environment, 1995