Different Flavours of VPN: Technology and Applications

Victor Olifer

With contributions from
Duncan Rogerson, Steve Williams,
Rina Samani, David Salmon,
Chris Cooper, Andrew Cormack

Table of Contents:

1. Introduction

Virtual Private Networks, or VPN,[1] provide a customised enterprise IP network service among several sites, belonging either to the same organisation or to collaborating organisations, over an IP network such as the Internet or JANET. The generic term 'VPN' currently covers different kinds of services which can benefit different kinds of applications through improved security and/or performance in the way they transport traffic.

This document briefs network managers on JANET's current position as regards VPN, the different flavours of VPN available, and the current position of VPN on other networks around the globe. The Appendix is a technical supplement that provides background information about VPN-enabling technologies.

2. VPN on JANET

JANET provides its users with a *basic transmission service,*[2] which is a regular IP best-effort service. Every packet is treated alike, with the same chance of being delayed or dropped if network congestion occurs. However, network applications available to JANET users may benefit from enhanced transport services. For example, multimedia applications like IP videoconferencing and VoIP[3] may benefit from enhanced network performance (for example lower latency and delay variation parameters than are available from IP best-effort) that IP QoS[4] could provide. Multicast transport could also save bandwidth on low-speed links.

Currently JANET supports several prototype or experimental enhanced transport services, including multicast, IPv6 and QoS. Another candidate IP technology for consideration is VPN.

The commonest form of VPN in use within the JANET community is a user-provisioned encrypted VPN, meaning that the VPN is provisioned by the computer service staff of a JANET-connected organisation. This kind of VPN provides secure access to the networked resources of an organisation (a university, a college etc.) for its remote users. There is not currently a central VPN service (i.e. a service centrally managed by the JANET NOC[5] or RNOs[6]) on JANET.

We would like to investigate:

- what requirements, if any, does the community have for centralised JANET VPN services, and what type of service could be implemented?
- possible scenarios of VPN use within the JANET community
- would the balance between the benefits of centralised VPN services and the cost of their deployment/maintenance justify deploying them as JANET production services?

The first step in this investigation was the VPN survey which UKERNA conducted in March-April 2006. The results can be found at **http://www.ja.net/development/vpn/VPNsurveyresults.html**. They showed that the JANET community has a significant interest in VPN services in general but no strong demand for centrally managed VPN services. Taking this survey result into account, centrally managed VPN services will not be deployed across JANET in the foreseeable future. However, if demands for such services arise within the JANET community then this area may be explored further.

1 For background reading see *Virtual Private Networks*, Second Edition by Charlie Scott, Paul Wolf and Mike Erwin. 1999, O'Reilly & Associates.
2 JANET SLA, http://www.ja.net/services/publications/policy/sla/operational-production-services.html#A21
3 Voice over IP.
4 Quality of Service.
5 Network Operations Centre.
6 Regional Network Operators.

3. VPN Definitions and Understandings

General Definition

Material in this section is based on material from *Computer Networks: Principles, Technologies and Protocols for Network Design* by Natalia Olifer and Victor Olifer (pub. John Wiley & Sons 2005).

The term 'VPN' has no standard interpretation. Different networking specialists and different organisations may understand it in different ways.

Historically, the term was first introduced by telephone companies. The main feature of a telephone VPN is that it can provide users from an organisation which uses a public provider's telephone service instead of its own private PBX[7] with something very close to PBX functionality (commonly known as Centrex, and popular in North America). For example, they can dial using convenient private (usually short) numbers; certain phones can be fully or partly isolated from the public telephone network; and users can use PBX-style telephone services like call forwarding, call rerouting, voice mail etc.

For data networks the term came to be used later, at first mainly for services which provide more security than the standard Internet service due to user data encryption. However, there are also services that do not encrypt user data but create logical channels for users within public data networks and provide controllable connectivity between VPN users and with the outside world. (Some of these elements were available to the data networking community in the closed user group of X.25 and the filtering capabilities of SMDS[8] in previous incarnations of the JANET network.)

One of the possible broad definitions for VPN could be:

> 'a network (or service) that reproduces (emulates) the properties of an *actual private network* using a shared public networking infrastructure.'

This definition could be applied both to telephone and data networks. The remainder of this document will focus only on data networks; the example of telephone VPNs was simply used as an analogy.

Emulated Features of a Private Network

So, what does it mean to say that a data network (or a packet-switched network) is private? It can be considered truly private only when the body using it owns all the elements (and hence has full control) of all the network infrastructure – cables, channel-building equipment, switches, routers and other communications equipment. However, a network is often considered private even though an organisation leases rather than owns all the channels that connect its sites. This is because the technical effects of traffic transmission are the same whether physical channels are owned or leased, as these channels always have a known and fixed bandwidth. By contrast, when an organisation uses a public data network to connect its sites, traffic goes through shared public channels and receives an unknown share of the channels' bandwidth.

As well as having a known channel bandwidth, a private network is distinguished from a public network by its isolation from any other network – the private channels only connect the sites of one organisation.

An actual private network can provide the following benefits for its users:

A. ***Improved security.*** Lack of connections to the external world considerably reduces the possibility of an attack on the network from the outside, as only certain users are physically connected to it. It also reduces the probability of eavesdropping on the traffic.[9]

7 Private Business Exchange.
8 Switched Multimegabit Data Service.
9 We cannot exclude traffic tapping for private networks completely because, where a leased line is used, traffic

B. ***Predictable performance***. Ownership of the communication links guarantees the bandwidth between the user sites and can make network performance more predictable.

C. ***Independent choice of network transport technologies*** for user site networks. The possibilities are limited only by the choice of a vendor or manufacturer, and an organisation-owner can use Ethernet, Frame Relay, IP, IPX or any other networking transport technology for connecting its sites.

D. ***Independent IP address space***. In private networks it is possible to choose any address. For example, almost all VPN services support the use of private IP addresses such as 10.0.0.1 or 192.168.0.3, which could not be routed over the public network.[10]

These features will be useful for certain users, though the relative importance of each can vary. The vulnerability and poor performance predictability of the Internet or public IP networks make ***'improved security'*** and ***'predictable performance'*** the most important features of a private network. Recently, ***'independent choice of network transport technologies'*** and ***'independent IP address space'*** seem to have become less important: the former because of the domination of a single technology (Ethernet at Layer 2 and IP at Layer 3) and the latter because IPv6 is expected to eliminate IPv4's current deficit of public addresses. However, another reason for having independent address space is security, as an organisation's address range can be used for access restriction within an organisation's sites.

On the other hand, a private data network is very expensive as it uses its own (or leased) channels with dedicated bandwidth (TDM[11] or optical) to interconnect LANs at different sites. A VPN data service tries to improve standard data transmission by providing some (though usually not all) of the features of a private network using a shared packet-switched infrastructure, such as JANET, commercial provider networks or the Internet as a whole.

> The aim of VPN of any kind is to provide communication between all a network's sites in a way which emulates as closely as possible their being connected by dedicated physical channels.

Different VPN services

There are many kinds of VPNs and understanding tends to vary for each type. We will attempt to classify them on the basis of three factors:

1. **Which features of a private network does a VPN service emulate, and to what extent?** For example, some VPNs support a very high level of data privacy with no performance improvements, whereas others support performance improvements but have a rather basic level of data privacy. The VPN survey results showed how useful different VPN features are considered by users. The priority list looks like this:
   - site protected from unauthorised access
   - strong confidentiality based on data encryption
   - traffic protected from non-VPN users, with the possibility of encrypting it
   - improved performance (low latency, low loss)
   - improved bandwidth guarantees
   - independent addressing
   - non-standard connectivity between sites (e.g. multicast through unicast-only network).

might be eavesdropped upon by unscrupulous employees of a leased line provider. Tapping can still take place even when channels are the property of a corporation (i.e. they own the physical cables), e.g. by detecting the low-power electromagnetic radiation which exists near cables, even optical ones.

10 The addresses that can be used in this way are defined in RFC1918.

11 Time Division Multiplexing.

It was quite expected that security features would be at the top of the list (occupying the first three positions); however, the relatively high placing of improved performance and guaranteed bandwidth shows there is potential for VPNs that support QoS.

2. **Whether the VPN is provisioned by a customer (a JANET-connected organisation) or by a provider (a JANET network operator).** In this document we will focus on provider-provisioned VPNs, since our aim is to explore the suitability of a JANET centralised production service managed by JANET network operators.
3. **Location of VPN equipment:**
   - network-based VPNs are built on equipment that is located within a provider network network-based VPN (i.e. where VPN equipment are located within a provider network) are provisioned by provider
   - customer-based VPNs use equipment located within a customer's network or the customer's computer.

Usually, network-based VPNs (i.e. where VPN equipment are located within a provider network) are provider-provisioned and customer-based VPNs are customer-provisioned. There are some exceptions where, for example, a provider can manage customer-based VPN equipment; but in practice such situations are relatively uncommon. Providers quite often manage users' access equipment (for example, JANET NOC manages some RNO router interfaces) but this is not the case for VPN as it would require organisations to reveal their security policy details, which most organisations prefer to avoid.

As well as the user-oriented features (i.e. features important for VPN users) described above, VPNs also have provider-oriented features. The most important of these are:

- scalability, i.e. the ability to support a large number of VPNs and sites within each VPN
- manageability, i.e. a low level of effort should be required to configure and support VPN. Provider-provisioned VPNs consume extra resources in network equipment and add to the complexity of the provider's configuration, which potentially threatens the manageability of the VPN
- ability to work in a multi-domain environment. This is very important for the JANET community as the JANET backbone, Regional Networks and JANET-connected organisations' networks are managed independently and form a distinct three-layer structure. In most applications, the VPNs will need to cross two or three of these layers.

Currently, there is no single technology that can provide a VPN service with all the desired features of a private network as described above. There are several technologies which could currently be used to create VPN-like services which will differ in terms of the user-oriented and provider-oriented features. Generally, such technologies, e.g. encrypting, tunnelling, QoS or MPLS,[12] have not been designed especially to provide a VPN service; each has its own functionality and might be used as a building block to create different services. It is a challenging task to combine several such underlying technologies and create a specific VPN service that will maximise user benefits and minimise provisioning and maintenance overheads.

Existing VPN Types

There are currently three broad classes of VPN services:

- encrypted VPNs
- tunnel-based VPNs
- optical private networks.

12 Multiprotocol Label Switching.

Encrypted VPNs

This type of VPN encrypts user data so that potential eavesdroppers cannot understand the content even if it is intercepted. Secure data exchange through a public network provided by encrypted VPNs usually complements an organisation's firewall service that protects the data inside its network.

Today, practically all organisations with home-working employees use encrypted VPNs to give those employees secure remote access. Encrypted VPN services are also used for connecting distributed offices through the Internet.

Within JANET-connected organisations this kind of VPN is provisioned by computer service departments, as they are generally based on dial-up style technology (even over broadband) where a user initiates a client on a PC and connects to a server. It is unlikely that such arrangements could be provider-provisioned.

Encryption can be added to the types of VPN listed below: however, the group using the VPN needs to consider carefully whether having a third party provide security is advisable. In addition, the provider will need to consider its legal position in the case of the encryption failing and confidential data being exposed.

According to the VPN survey, encrypted VPNs are currently the most popular kind of VPN in use within the JANET community: 85% of respondents use them, either along with other kinds of VPN (56%) or as the only kind of VPN (44%).

Tunnel-based VPNs

There are several VPN services that fall into this category, based on the different networking technologies available. The common features which they all share include the provider transmitting traffic between the VPN sites using tunnels, a.k.a. logical channels, within a provider network. As a result such VPNs may provide:

- *improved security as a result of apparent traffic separation*, since using logical channels separates the VPN traffic from other Internet traffic. However, the traffic will still be sharing the same physical network. Some tunnelling techniques can provide a very high degree of separation, in effect providing **traffic isolation**. Traffic isolation means security improvements in two ways:
  - *data security improvement* as user data is isolated 'on-the-fly' from other users' data within the ISP's network
  - *site security improvement* as the provider has control over the connectivity between each of the user sites; therefore an intruder connected to the public domain of the Internet will not be able to direct traffic towards the VPN sites and attack them (assuming that the VPN does not also provide Internet connectivity).
- *potentially improved performance*. VPN in itself does not necessarily improve network performance unless appropriate QoS methods are implemented in parallel, but VPN within a public network can simplify QoS implementation as it provides increased knowledge of and hence control over individual traffic flows between VPN sites (as opposed to the unpredictable connectivity of a common IP network).

Different tunnelling technologies could be used for this kind of VPN: the more sophisticated the technology, the greater the VPN functionality it can support. However, at the same time it must be noted that sophisticated tunnelling technologies require more complex configuration and management. The ideal requirement would be to find a technology which can provide the desired functionality with minimum overheads. The most popular tunnelling technologies used for VPN build are L2TP[13] and MPLS. Another has emerged very recently: PBT[14] from Nortel. Currently it is a proprietary technology but Nortel is taking steps to standardise it.

13 Layer 2 Tunnelling Protocol.
14 Provider Backbone Transport.

L2TP (used by about 24% of the respondents) is simpler to implement and support, as the technology encapsulates user frames or packets into standard IP packets which can be transferred transparently by any standard IP network. L2TP transparency allows organisations and end users to self-provision L2TP through any provider IP core network.

MPLS (used by 7% of respondents) is not transparent as it requires MPLS support within a provider network. This may be a problem for non-MPLS enabled providers (e.g. for JANET backbone which currently does not support MPLS); however, MPLS has more potential for improving performance than pure IP networks. This is because of the native MPLS capability to control flows within a network and hence provide strict admission control to elevated QoS resources.

Tunnel-based VPNs could be provider-provisioned (within a provider network) and hence could potentially be implemented as a JANET VPN service. It is equally possible for a customer to provision VPN tunnelled over IP inside their own networks; however, any QoS requirements in the wide area network will, of course, have to be handled by the provider and need coordination between the user and provider.

Optical Private Networks

Optical private networks are a fast progressing area which use achievements in high-speed networking based on SDH[15] and DWDM[16] technologies. Not so long ago, top-speed optical channels with a bandwidth of 2.5Gbit/s and 10Gbit/s were only available for carriers and large ISPs, but now they tend to be accessible to enterprise users (organisations large enough to need them and with the ability to fund them).

Modern SDH/DWDM optical networks provide users with fixed bandwidth channels, which are similar in many ways to the much slower copper leased line services used for building private networks in the past.

SDH/DWDM based services are generally not viewed as true VPNs because they do not use a shared packet-switched infrastructure. However, they are sometimes called VPNs for several reasons.

- This kind of service became available to enterprise customers due to price reductions and wide implementation by telcos and ISPs. They have therefore shifted from telco-only services to commodity services for mass users.
- Though this kind of service does not use a shared packet-switched infrastructure it does use a shared circuit-switched infrastructure, and if we do not restrict ourselves to considering packet networks only then we can extend the definition for VPN to include both kinds of networks.
- This kind of service has become much more dynamic as providers and sometimes even customers themselves can configure the necessary connections between the customer sites on demand (for example, using UCLP[17] software developed by CANARIE[18]). This makes private optical networks quite similar to traditional self-provisioned encrypted VPNs.
- The definition of VPN is quite broad, including not only shared packet-switched but also circuit-switched infrastructures (i.e. SDH/DWDM).

Optical private networks are truly private and have all the desirable features of a private network. However, they also have some limitations: they are still relatively expensive and not as widespread as IP or Ethernet services.

About 7% of the VPN survey respondents currently use optical private network services.

Descriptions of VPN-enabling technologies can be found in Appendix 1: VPN-Enabling Technologies.

15 Synchronous Digital Hierarchy.
16 Dense Wavelength Division Multiplexing.
17 User Controlled Lightpaths.
18 Canada's non-profit advanced Internet development organisation.

Examples of Centrally Provided VPN Services

RedIRIS

RedIRIS (the Spanish NREN – **http://www.rediris.es**) started work on VPNs in 2004 and has been providing a point-to-point Layer 2 VPN service since the end of that year. The most remarkable project making use of this service was controlling reception of several HDTV20 video sessions via UCLP. This project is supported by CANARIE, in which i2CAT (one of the centres within RedIRIS) is involved. In particular there were two demonstrations for the project during 2005 for which a Layer 2 VPN had to be configured across GÉANT between RedIRIS, CESCA and CANARIE. More information can be found at:

- UCLP Demonstration at APAN:
  http://www.canarie.ca/canet4/uclp/apan/demo.html#Tab2
- UCLP Demonstration at Viola Workshop 2005:
  http://www.canarie.ca/canet4/uclp/viola2005/demo.html
- CESCA
  http://www.cesca.es

With respect to providing a VPLS,21 RedIRIS has been working in both the intra- and interdomain environment, but in both cases for testing purposes only. VPLS is not yet a production service in RedIRIS, although RedIRIS is evaluating the benefits of this technology in order to inform its customers. If it is found that VPLS is useful for research organisations, RedIRIS will put it in production in the near future.

HUNGARNET

HUNGARNET (the Hungarian Academic and Research Network Association) is using Layer 3 MPLS VPNs for several projects.

1. **Providing Layer 3 VPN for the ClusterGrid infrastructure**.[19] ClusterGrid has been using Layer 3 VPN for more than four years. It was deployed so that a completely virtual infrastructure could be put on top of the existing routing infrastructure. Sometimes it also goes through the firewalls of partner organisations.

   Some problems of the existing VPN infrastructure were:

   - MPLS capable equipment was not available on some sites
   - separate VLANs therefore had to be used on backbone devices for GRID VLANs
   - unfortunately there was no IPv6 capability in the equipment in use for Layer 3 MPLS VPN.

   These three reasons led HUNGARNET to jump into testing Layer 2 VPNs, especially VPLS. The outcome of the VPLS test was not very satisfactory as it required special linecards and equipment. There was little compatibility among the vendors. Therefore HUNGARNET decided to use Layer 2 VPN only in point-to-point environments until the technology evolves.

2. **Financial system for museums**. This project to provide a separate infrastructure for a financial system for museums was initiated by the Ministry of Culture. A separate VPN was set up for this purpose. Only a dedicated system could be attached to the financial system VPN.

3. **HUNGARNET Directory service management**. The management VLAN of the Directory servers was put into VPN so that only dedicated systems can access it.[20]

19 More information can be found at http://www.clustergrid.niif.hu.
20 More information can be found at http://www.directory.iif.hu/ (unfortunately in Hungarian only).

4. **HUNGARNET VoIP service management**. The management VLAN of the VoIP call managers was put into VPN so that only dedicated systems could access it.
5. **Dedicated e-learning systems**. Several e-learning systems that had more than one site wanted to have virtual interconnection but without Internet access. Only dedicated systems could access the e-learning systems.

BT Infonet VPN Services

BT Infonet offers different kinds of VPN services:

- **An ATM or Frame Relay VPN** – traditional services which work within BT Infonet's privately owned The World Network. (This is the network which belonged to Infonet before it was bought by BT in 2005.)
- **Private Internet**[21] is a service which combines security and performance on the basis of MPLS VPN within The World Network. BT Infonet Private Internet provides connection speeds ranging from 64kbits/ to 45Mbit/s (higher in some locations).
- **IP VPN Secure**[22] looks similar to Private Internet as it also provides security and improved performance for users connected to The World Network. However, it is an advanced version of VPN as it supports five classes of services, whereas Private Internet does not support such a differentiation.
- **IP VPN OffNet service** – an encrypted centrally managed VPN service that lets users be connected to different ISPs. BT Infonet manages VPN gateways at client premises to build VPN tunnels across the public Internet.
- **MobileXpress**[23] services are used for IP VPNs for travelling users or for small offices that only need a dial-up connection to the network

4. Possible Areas of VPN Use Within the JANET Community

Generally, VPN is only worth consideration when collaboration between sites or remote users and sites is long-term relative to the time needed to provision the service. If VPN use is short-term then the overheads associated with service provisioning may not be justified. For example, if the provisioning of a VPN service takes, say, two days, it is not worth using such a service for a 10 minute VoIP phone connection. However, it might be effective to connect several sites for a research trial lasting some six months. The emergence of dynamically provisioned VPN with a lead time measuring in seconds, which is currently still at the research stage, may change this situation.

Something else to take into account when considering potential areas for VPN use is of course application requirements. There should be a need for some VPN functionality like strong data confidentiality, site protection from external deliberate or accidental harmful activity, or guaranteed bandwidth.

In the context of VPN, the term 'site' means a physically or logically separated part of a campus network. For example, it could be a subnet which has no physical connection to other organisation subnets and uses only the VPN tunnel for external communications (a remote user's computer falls into this category). A more widespread example is a VLAN logically separated from other subnets and nodes by the respective configuration of LAN switches.

Within the education and research community we can suggest several areas of specific research and education activity where VPN services, possibly of different types, might be beneficial for users. Some such areas are given very high-level descriptions below. This is

21 http://www.bt.infonet.com/services/internet/private_internet.asp
22 http://www.bt.infonet.com/services/intranet/ip_vpn_secure.asp
23 http://www.bt.infonet.com/services/access/MobileXpress.asp

a very preliminary list of possible VPN use and further detailed discussions, investigations and (if necessary) trials are needed to finalise the set of areas and requirements for VPN services (if any) for each one.

- **Traditional intra-organisation multi-site network-based applications** like e-mailing, database access and web-surfing. Usually, such applications require improved security to protect an organisation's networked resources from unauthorised access. At the same time, applications in this area do not have any special requirements for improved performance as they are not delay sensitive (also known as elastic applications). Respondents to the VPN survey indicated that these applications are the most popular among VPN users (79% of VPN users use e-mail, 74% use web services and 68% use database access).
- **e-Learning applications** were identified as the VPN applications in use by 35% of the VPN survey respondents, which shows their importance as a driving force of VPN deployment. The e-Learning area generally includes very diverse applications, from video clips and other materials which can be downloaded in advance to real-time teaching which might use high quality video formats like HDTV (which results in high bandwidth demands) and interactive communications. Hence, while some e-Learning applications can easily be served by standard IP services, others might benefit from the improved security and performance of VPN services.

  As an example, one can imagine several e-Learning studios in different colleges which are used by course attendees twice a week over half a year. Each studio is equipped with video facilities which are used by a lecturer and students to communicate during course delivery. The known pattern of stable long-term connectivity between sites, the advanced security requirements (including protection from external attacks to provide reliable connectivity and protect learning content from unauthorised access) and the advanced requirements for guaranteed bandwidth and low loss/delays to provide high quality are all strong reasons for these studios to use VPN.

  Of course, there are many details which should be taken into account, discussed and investigated before a decision can be made about using a VPN service with particular functionality for e-Learning applications.
- **e-Science/Research collaboration**. As with e-Learning, the e-Science area includes a wide spectrum of applications, some of which can benefit from VPN services.

  At one end of the spectrum are the most demanding networking applications in terms of performance and bandwidth parameters. Examples of such applications are astrophysics collaboration and high energy physics collaboration.

  Such applications require real-time data processing and hence require a very low level of delays and jitter. Loss of synchronisation (even in the milliseconds range) between data source and data processing centres might devalue an entire experiment.

  The bandwidth demands of such applications can also go beyond the capabilities of modern packet switched networks, or more precisely, beyond the limits within which this kind of network remains effective. When, for example, 1Gbit/s needs to be allocated to a few virtual connections between two VPN sites, a packet-switched network with 10Gbit/s core links will lose its advantages as a shared environment. This will happen because shared packet-switched networks were designed to work effectively when every user flow consumes only a small percentage of the link's bandwidth. A user could try to consume as much as 10% of total core link bandwidth, which might monopolise a network and prevent other users from receiving a proper service.

  Generally, such requirements are very difficult for standard IP networks to satisfy so such extreme projects tend to use private optical networks based on SDH, DWDM or even on dark fibre. In fact, several such applications already make use of the UKLight bandwidth channel network that is operated in parallel with the JANET IP network. As was mentioned before, such a service might be seen as a kind of VPN.

At the other end of the spectrum of e-Science applications are the applications that can easily be served by the standard best-effort IP service. The examples include medium size ftp downloads of non-real time data, wikis, and other means of online collaboration.

Between the two extremes we can expect to find some e-Science applications which have medium demands in terms of performance and security. On the one hand, such demands might be too high to be satisfied by a standard IP best effort service; on the other they may not be too high to be satisfied by some kind of packet-switched VPN with a low level of delays/loss and strong protection of traffic. We can imagine several sites collaborating in a relatively long-term project which need to exchange data with delays less than 200 ms, guaranteed bandwidth up to 10Mbit/s and strong requirements for stable, non-interrupted communications during experiments. Of course, such an e-Science application would benefit if its sites were connected by SDH channels. However, it might well be that such requirements can be met by a packet-switched VPN service with improved performance and security functionality as well. As packet-switched networks are generally cheaper and better known for end users, it is worth investigating in more detail what e-Science applications can benefit from packet-switched VPNs and what are their requirements.

- **Art collaboration** (so called Cyber Arts or Humans Interacting with Virtual Realities) which needs real-time interactive data exchange. The use of HDTV and other high-quality video standards makes such applications quite demanding in terms of bandwidth, whereas the real-time nature of collaboration requires improved performance in terms of delays and jitter. The security offered by VPN might be very useful to protect art studios from outside intervention (accidental or malicious) during an art performance.

5. Conclusion

The use of VPN technologies within the JANET community has been considered and the requirements for centralised provider-provisioned VPN services have been assessed. The assessment was carried out in March-April 2006, when the VPN survey was conducted. 121 of the responses received showed strong interest in VPN services within the JANET community: over 80% of respondents indicated that they use some form of VPN service; about 55% of the rest plan to use VPN in the near future. However, the respondents' answers also indicated that the community has no strong demand for centrally managed VPN services. Taking into account this result of the survey, centrally managed VPN services will not be deployed across JANET in the foreseeable future. However, if demands for such services arise within the JANET community then this area may be explored further.

Appendix 1: VPN-Enabling Technologies

Table of Contents

A1. Introduction

Earlier in this document it was stated:

> **'Currently, there is no single technology that can provide a VPN service with all the desired features of a private network ...** There are several technologies which could currently be used to create VPN-like services which will differ in terms of the user-oriented and provider-oriented features. Generally, such technologies, e.g. encrypting, tunnelling, QoS or MPLS, have not been designed especially to provide a VPN service; each has its own functionality and might be used as a building block to create different services. It is a challenging task to combine several such underlying technologies and create a specific VPN service that will maximise user benefits and minimise provisioning and maintenance overheads.'

This supplement describes particular aspects of each technology:

- technique, i.e. how it works
- ability to support emulated user-oriented features of a real private network. These features are considered in detail in the main document but for ease of reference are:
  A. Strong security
  B. Predictable (or improved) performance
  C. Independent choice of network transport technologies
  D. Independent address space
- topologies supported
- provider-oriented features:
  - scalability
  - manageability.

This supplement does not include a description of QoS, the technology which improves network performance in terms of packet latency and loss. This is described in documents published by the JANET QoS Development Project.[24] However, QoS is mentioned in the descriptions of the packet-switching technologies considered here: some of the VPN-

24 http://www.ja.net/development/qos

enabling technologies are described as QoS supportive or QoS neutral. (QoS supportive technologies can simplify QoS deployment or strengthen QoS guarantees, while QoS neutral technologies have no additional functionality which QoS can exploit and benefit from; their traffic looks like standard IP traffic.) Of course, both QoS supportive and QoS neutral technologies can benefit from QoS if it is deployed across a network, but QoS supportive technology tends to guarantee a higher level of QoS and QoS deployment tends to be simpler. There are some VPN-enabling technologies which have built-in QoS functionality, e.g. ATM and some versions of Frame Relay, but they are becoming rarer and are not in widespread use within JANET.

A2. Encrypted Tunnelling (IPSec/SSL)

The overwhelming majority of current VPN implementations are encrypted VPNs, for security purposes. If a networking professional is asked about VPN in general and the type of VPN is not specified, encrypted VPN is the first association which comes to mind.

Technique

Encrypted VPNs use a **secure channel** (or tunnel, or association) for data transmission between VPN sites. This means:

- authenticating the two end points of a secure channel so that only authorised users have access to an organisation's network, and only the organisation's authentication server can request the user's secret credentials
- encrypting a user's packet/frame (or packet/frame-passenger, as it is sometimes known) and encapsulating it into another packet/frame (the delivery packet/frame) which seems 'normal' to the ISP's network equipment. Therefore encrypted traffic is absolutely transparent to a provider network and is served the same way as any other traffic.

IPSec[25] and SSL[26] are the most popular protocols used nowadays for establishing secure channels. PPTP[27] is another example, although less popular, probably because it remains the Microsoft® proprietary protocol whereas the first two are IETF standards. All these technologies encapsulate secured data into IP packets, which is unsurprising given the domination of the Internet and IP.

The secure channels of encrypted VPNs are usually complemented by firewall services and the security features within computer operating systems. Secure channels protect an organisation's data while it is being transported through public networks, whereas firewalls and operating systems protect an organisation's data (and other networking resources like computers, routers and switches) from external attacks. Secure channels also provide some additional protection against external attacks as they do not accept encrypted traffic from non-authenticated users.

Emulated User-Oriented Features

The main goal of an encrypted VPN service is to ensure **secure end-to-end data transmission** through a public packet-switched network: that is, in the list of private network features given above, they aim to emulate **security** (feature A in the list above). Encrypted VPNs provide data integrity, authenticity and confidentiality during data transmission between VPN sites.

Encapsulation of user packets can emulate another feature of a private network, namely **independent addressing system** (feature D), as the address of the encapsulated user packet

25 IP Security, described in RFC 2401.
26 Secure Sockets Layer, described in RFC 3546. (This RFC actually describes the Transport Layer Security protocol, which is the successor of SSL and the Internet standard; however, SSL is used for both these protocols as a brand.)
27 Point-to-Point Tunnelling Protocol.

cannot be used for transportation through public networks. A delivery packet address is used for this purpose.

Providing **independent choice of network transport technologies** (feature C) is generally not possible in a private network with current encrypted VPN services. In fact IPSec channels only accept IP packets from an organisation's sites; SSL is also implemented only for IP networks. PPTP is more flexible: on the one hand it accepts only PPP[28] frames but on the other these PPP frames can carry practically any traffic – IP, IPX, frame relay etc.

Encrypted VPN tunnelling is QoS neutral because (as emphasised above) encrypted VPN traffic resembles a sequence of standard IP packets to provider equipment, and hence does not give any additional support for a QoS implementation.

Topologies

Secure channels/tunnels are usually point-to-point channels. Very often they create a hub-and-spoke topology, with the VPN gateway on a main enterprise LAN and VPN clients on the remote computers of employees working from home or travelling. One-to-many (multicast) topology is not presently used in practice; however, some work in this area has been established within the IETF.

Provider-Oriented Features

Encrypted VPNs implemented today are mostly customer-provisioned and customer based (i.e. all the VPN-specific software/hardware is only located on the customer site) and therefore could not be part of any centrally managed JANET service. End users (more precisely, their IT Support departments) define an appropriate security policy and implement it by configuring VPN gateways and software clients on users' computers.

As encrypted VPNs use normal IP packets for transferring user packets/frames through a provider network, the service is very convenient for end users and ISPs. End user sites can be connected to different ISPs providing nothing more than a regular Internet access service. (This is the up-side of the 'normality' of delivery packets; the down-side is their inability to provide improved performance.) ISPs in their turn also do not need to provide any service other than regular best-effort, any-to-any transmission to support encrypted VPNs.

Despite the fact that most encrypted VPNs are customer-provisioned, encrypted VPN could in principle be provider-provisioned. The provider can remotely provision and administer VPN gateways and clients located at customer premises (note that such a VPN will still be customer-based). In this case the customer has to formulate their security policy and inform the provider about it. The provider may then use an management system to support the customer's VPN devices. One of the most powerful and scalable specialised security management systems is Provider-1 from Checkpoint Software technologies. IP VPN OffNet from BT Infonet is an example of an encrypted VPN service where VPN devices on customer premises are managed by a provider.

Scalability of encrypted VPNs depends on their topology.

- Mesh topology requires roughly N^2 secure channels for N sites which means quite poor scalability. IPSec is generally considered less scalable than SSL as IPSec often requires complicated configuration involving key distribution.
- Hub-and-spoke topology has better scalability as it requires only roughly N secure channels for N sites.
- When encrypted VPN is self-provisioned, scalability is often not a problem as the number of sites and users is not as big as in the case of a provider-provisioned VPN where the provider supports hundreds of organisations.

28 Point-to-Point Protocol.

Manageability of encrypted VPNs is quite poor because of the complexity of distributing, configuring and storing authentication and encryption information: user IDs, passwords, digital certificates, secure keys etc.

A3. GRE/L2TP Based VPNs

This kind of VPN aims to transfer any kind of traffic (IP and non-IP) through an IP network by tunnelling.

Technique

Packets from the user's site are encapsulated into normal IP packets and then transferred through a provider network to other site(s) on the VPN, creating a transport tunnel. GRE[29] and L2TP[30] are the standard mechanisms for establishing transport tunnels through IP networks, and are supported by all major vendors of network equipment. IP packets with encapsulated GRE or L2TP data are treated as regular IP packets by the provider network, going through it in the normal way along its regular routing path. (Both mechanisms actually allow any routable protocol to be used as the delivery protocol but IP is the only practical option for this role.) Any passenger protocol can be used for GRE; PPP is used for L2TP. The latest version of L2TP, version 3, can be used with any Layer 2 protocols; however, as PPP can carry almost any other protocol data, L2TP older than version 3 can still be used for transferring most protocols through a provider network.

Both customer equipment and provider-edge equipment can do this sort of encapsulation, so this kind of VPN could be user-provisioned or provider-provisioned. The latter means that GRE/L2TP VPNs could be deployed as a possible service on JANET.

There are no mandatory encryption or authentication procedures for GRE or L2TP tunnels; L2TP specifies an optional authentication mechanism which works while a tunnel is being established.

Emulated User-Oriented Features

Being able to encapsulate any type of packets into IP packets means that the technology emulates **independent choice of network technologies** (feature C). This could be re-formulated more generally as 'transferring non-standard traffic', meaning that the transport protocol of the user site is not supported by the provider network. However, as all provider networks today support IP we can say that the first definition is sufficiently general.

Examples of such a service might be transferring IPX traffic between sites using native Novell protocols, or SNA[31] traffic of old mainframes that do not support IP. Another (and more up-to-date) example is transferring IPv6 traffic through an IPv4 network, which is quite a common task for providers who do not support native IPv6 transport.

The use of encapsulation also emulates **independent address space** (feature D)**.**

Security (feature A) is provided because of tunnelling. Tunnelling improves security by means of apparent traffic separation, as using logical channels separates the VPN traffic from other Internet traffic. However, the traffic will still be sharing the same physical network. Some tunnelling techniques can provide a very high degree of separation, in effect providing **traffic isolation**. Traffic isolation means security improvements in two aspects:

- data security is improved as user data is isolated 'on-the-fly' from data of other users within the ISP's network

29 Generic Routing Encapsulation: RFC 2784.
30 RFC 3931.
31 System Network Architecture.

- security of sites is improved as the provider has control over the connectivity between each of the user sites; therefore, an intruder connected to the public domain of the Internet will not be able to direct traffic towards the VPN sites and attack them (assuming that the VPN does not also provide Internet connectivity).

However, as authentication and encryption are optional for GRE/L2TP, we can expect that this kind of VPN will have only moderate security.

GRE/L2TP VPN security can be strengthened by using encryption/authentication techniques on top of GRE/L2TP tunnels. However, this would not be a provider-provisioned VPN service and therefore cannot be supported by JANET.

GRE/L2TP VPNs are QoS neutral because for a provider network their traffic looks standard.

Topologies

GRE/L2TP tunnels are point-to-point, so full mesh or hub-and spoke topologies are possible.

L2TPv3 is generally capable of handling IP multicast; however the efficiency of transport will depend upon the exact topology deployed. For example, in some cases where multiple L2TP tunnels traverse the same physical link, each tunnel may carry a copy of the multicast data stream. The IETF is currently working on some improvements of L2TP support for multicast.

Provider-Oriented Features

This kind of VPN is quite easy to implement as it only needs additional configuration (and processor power) for customer or provider-edge routers.

Scalability of GRE/L2TP based VPNs is as poor as the scalability of encrypted VPNs, as tunnels must be established through a provider network between all sites belonging to the same VPN (mesh VPN topology with N^2 tunnels) or between a central site and all the others (hub-and-spoke topology).

Manageability of GRE/L2TP VPNs seems to be better than for encrypted VPNs as it is not necessary to maintain authentication and encryption information for VPN sites.

A4. Policy-Based VPNs

This kind of VPN uses policies and access lists to create special routes for VPN site traffic through a provider network. There is very little information available about experiences of policy-based VPN deployment, so any estimation of their features can only be approximate.

Technique

Policy-based VPNs provide **traffic separation** inside a provider network. This restricts normal IP connectivity between customers' sites so that only sites belonging to a particular VPN can communicate with each other. With normal IP connectivity (the datagram style of communication), no preliminary procedure for establishing a session is needed and any Internet-connected computer can communicate with any other. This connection-less feature provides a very simple and effective way of communicating on the global scale (i.e. the Internet scale), but at the same time it creates a very good opportunity to attack an organisation's resources from any point of the Internet. Normally, IP routers forward any packet dedicated to a particular network if the router has an entry for that network in its routing table.

Policy-based VPNs do not use tunnelling of any kind, in contrast to encrypted and GRE/L2TP based VPNs. All user packets go through a provider network without modification or encapsulation. Instead, policy and access lists are created in the routers of the provider network (the technique is vendor-dependent) which change the normal routing of user packets. This technique alters the connectivity between sites (allowing communication only between sites belonging to the same VPN) and can alter the routes through a provider network, providing a sort of traffic engineering.

> When creating a policy-based VPN service, it is necessary to specify a set of rules that only allow packets from one VPN site to be forwarded to another site in the same VPN, and block traffic from other VPN sites and from the public Internet.

Most router vendors' equipment can support policy-based routing; however, the respective configuration commands used for different vendor routers will tend to be completely proprietary.

Emulated User-Oriented Features

Policy-based VPNs provide **security** (feature A referred to in the Introduction), as they restrict connectivity between VPN sites and the rest of the world. Traffic separation makes users' sites more secure, and encryption of traffic on top of that is also possible.

Policy-based VPNs do not support **independence of address space** (feature D) in itself as they do not change the IP addresses of incoming user IP packets.

Policy-based VPNs do not support **independent choice of network transport** technologies (feature C) for an organisation's site network as they only accept IP packets from VPN sites.

In regard to **improved performance** (feature B)**,** policy-based VPNs might be QoS supportive. There are two possibilities for implementing policy-based VPN:

- policy-based routing is configured only on provider-edge routers so that VPN packets go along the normal paths within a provider network. This kind of VPN is QoS neutral
- policy-based routing is configured on all provider routers including core ones. This feature might be used to control a VPN traffic path through a network, thus making it simpler to control privileged bandwidth consumption by elevated QoS classes. In this case policy-based VPN might be considered as QoS supportive.

Topologies

Any topology can be supported as this kind of VPN uses IP routing, not tunnelling. Multicast could also be supported.

Provider-Oriented Features

Scalability of policy-based VPNs depends on two factors: the number of VPN sites and the complexity of routing rules. The first factor makes policy-based VPNs more scaleable than encrypted and GRE/L2TP VPNs, as a number of configuration procedures are proportional to N when we have N VPN sites (and not to N^2 as for GRE/L2TP VPNs). The second factor makes the scalability quite poor because the routing rules could be very complex. This also means that **manageability** of policy-based VPNs tends to be very poor.

A5. MPLS VPN

MPLS VPNs are another example of using traffic separation to provide VPN functionality.

Technique

Traffic separation is made easier if a provider-network supports some kind of virtual circuit technique, for example ATM,[32] Frame Relay or MPLS. ATM and especially frame relay based VPNs were very popular in the 1990s but their implementation now is almost unknown, so we will not consider these technologies in this document.

Generally, a virtual circuit is a stable path through a network that passes particular network nodes. Virtual circuits provide a far greater degree of control for a provider over traffic paths, so this feature could be used for:

- traffic engineering, i.e. optimal use of all network resources due to a rational choice of traffic paths
- guaranteed QoS due to the possibility of providing proper utilisation of network resources for different traffic flows
- VPN functionality due to separation of traffic flows from different users and their sites.

MPLS is a relatively new technology which combines IP networks with the virtual circuit technique, drawing on the advantages of both. It is now becoming quite popular among providers; for example, all Tier 1 providers currently have MPLS-enabled backbones.

MPLS could be used to support different network applications. Most popular at the moment are:

- VPN – now the most popular MPLS-based service; all Tier 1 providers offer the service to their commercial customers
- traffic engineering (internal improvement of a provider's network resources utilisation) with QoS support (mostly as an add-on service for VPN users).

An MPLS-enabled network consists of IP routers which can establish virtual circuits through a network and forward incoming traffic either on the basis of IP addresses (acting as a normal IP network) or on the basis of MPLS labels. This means that ISPs should not need to buy any additional equipment to deploy MPLS-based services on their existing network devices; to make their networks MPLS-enabled they just need some additional configuration of routers, since most backbone routers can already support MPLS.

MPLS-forwarding of some traffic can be combined with regular IP-based forwarding of other traffic, which makes IP/MPLS networks a universal transport system and lets new services be deployed smoothly.

MPLS virtual circuits are called Label Switched Paths but in essence they are just another kind of virtual circuit. They make a particular kind of tunnel through a network, and therefore have features in common with the tunnel-based VPN services previously considered such as GRE/L2TP VPNs.

MPLS VPNs can be second layer (Layer 2) or third layer (Layer 3). For MPLS VPN Layer 2, a provider network acts as a big LAN switch, supporting VLANs for their customers. MPLS VPN Layer 3 acts like a normal routed network, in effect creating a dedicated network for every user. An MPLS VPN Layer 3 service can provide additional IP functionality for its customers as it accepts data in the form of IP packets, rather than LAN frames like an MPLS VPN Layer 2 service.

32 Asynchronous Transfer Mode.

Emulated User-Oriented Features

Improved security (feature A) is provided by means of traffic isolation. This feature of tunnelling was described above in the section dedicated to L2TP.

Independent address space (feature D) is supported by using labels and a special form of IPv4-VPN address that is different to public IP addresses.

Independent choice of network transport technologies (feature C) for an organisation's site networks is provided. In the commercial world some customers still use frame relay services and use an MPLS VPN service to transfer frame relay traffic through ISP backbones. Any other kind of customer traffic can also be transported by an MPLS-enabled ISP backbone.

At the same time the control over traffic paths makes MPLS VPNs QoS supportive. MPLS label-switched paths provide a good basis for guaranteed QoS as they could be established taking reserved bandwidth into account, and be re-established quickly in the event of network faults.

Topologies

In principle there could be any logical topology between sites; however multicast is not currently supported in practice. There are several activities within IETF aiming to develop multicast support for MPLS.

Provider-Oriented Features

Scalability of MPLS VPNs is better than other tunnel-based VPNs like GRE/L2TP, because they exploit the hierarchical layered design of MPLS. The number of tunnels through a provider network does not depend on the number of sites and is proportional only to the number of provider edge routers. This is because MPLS VPNs use second layer labels for traffic separation which do not need separate tunnels for customers' sites.

Manageability of MPLS VPNs needs to be explored. On the one hand they may be more manageable than GRE/L2TP VPNs because they need fewer tunnels. (Another positive factor is the existence of specialised management systems aiming to automate MPLS VPN configuration.) On the other hand there are more configuration operations than with GRE or L2TP tunnelling.

A6. Optical Private Networks (SDH/DWDM)

The perceived ambiguity of SDH and DWDM networks having a dual private-public nature was explained above under 'Optical Private Networks' on page 7. Being building blocks for truly private networks, these circuit-switched technologies are a good reference point for comparing the features of different kinds of VPNs and their possible applications. However, as they are not true VPNs we will not discuss them here in the same format as previously.

The access speed of modern SDH/DWDM private networks is in the range of 155Mbit/s (SDH access channels) to 10/40Gbit/s (DWDM access channels). The access speed granularity of SDH/DWDM services is quite poor as it depends on the speed hierarchy of SDH and DWDM technologies (155Mbit/s – 622Mbit/s – 2.5Gbit/s – 10Gbit/s – 40Gbit/s, with an additional 1Gbit/s for Gigabit Ethernet, which is not a standard SDH/DWDM hierarchy speed but is supported by many SDH/DWDM vendors due to the domination of Ethernet).

User provisioned optical networks are an emerging and very promising area. There are several international research projects exploring the area, including VIOLA,[33] GÉANT2

33 Vertically Integrated Optical Testbed for Large Applications project: http://www.viola-testbed.de/

JRA3,[34] HOPI[35] and MUPBED.[36] Probably the most popular tool for user provisioning of optical paths is UCLP, developed and funded as an initiative of CANARIE and Cisco® Canada.

The circuit switched nature of private optical networks has its down-side – a user request to establish a connection can be blocked due to lack of network capacity. This is a well-known drawback of telephone networks and Erlang produced several formulae to evaluate the probability of such blocking events occurring.[37] The risk of request blocking is the price to be paid for the excellent quality of circuit-switched services based on dedicated bandwidth for every connection. At the early stages of optical services for end-users (both provider and user provisioned), when the number of potential users is small, such a probability tends to be low if not negligible. That is why this problem is rarely mentioned at the moment; however, the more such a service becomes popular, the more likely request blocking is to occur.

34 See http://wiki.perfsonar.net/jra1-wiki/index.php/JRA1-JRA3
35 Hybrid Optical Packet Infrastructure project: http://networks.internet2.edu/hopi/
36 Multi-Partner European Testbeds for Research Networking: http://www.ist-mupbed.org/
37 see http://www.erlang.com

Appendix 2: Table of VPN-Enabling Technologies

Technology	Provisioned by	Security	Private user address space	Independent transport technology for end user sites	Improved performance[41] (guaranteed bandwidth, limits for delays and loss)	Multi-domain capability	Multicast support	Scalability and configuration complexity / automation software
Encrypted VPN (IPSec or SSL based secure channels) Two versions: - remote access, when a user accesses the central site of an organisation - site-to-site, when sites of an organisation connected to each other create a mesh.	Mostly: IT Support staff of an organisation. Quite rare: by a provider who manages VPN gateways on the central site and VPN software clients on remote PCs.	Confidentiality – excellent as all fields of packet can be encrypted. Protection of end site – good as only VPN end-point need have an Internet-routable address.	Yes	In practice, no, as the two currently popular technologies – IPSec and SSL – are IP-oriented.	QoS neutral	Yes, as security tunnels are transparent for providers.	In practice probably not, but there is an ongoing activity in IETF.	Scalability is an issue in the case of site-to-site as number of tunnels is proportional to N^2 of sites if a full mesh is configured. Remote access version is much more scalable as it uses the hub-and-spoke topology.
GRE or L2TP tunnels-based VPN	Users and providers	Confidentiality – encryption can be added on top of tunnels. Protection of end site – moderate: only the paired endpoint can insert traffic into the tunnel.	Yes	Yes, packets/ frames of different technologies – 'passengers' could be encapsulated in IP.	QoS neutral	Yes, as tunnels are transparent for intermediate providers.	Yes	Scalability is an issue as number of tunnels is proportional to N^2 of sites if a full mesh is configured.

Policy-based routing VPN (no tunnelling)	Providers	Confidentiality – no. Protection of end site – possible if policy only permits traffic from other trusted sites.	No	No	Possible stronger basis for QoS as control over flows is tighter than for a plain IP network.	Yes	Yes	Poor scalability as specific routes for VPN flows need to be configured manually at every router.
MPLS VPN (layer 2 or 3)	Providers	Confidentiality – moderate: traffic is protected by separation due to using different LSPs; independent tests have showed good degree of protection. Protection of end site – possible if the site only accepts MPLS tagged traffic (as for (3)).	Yes	Yes	Advanced support for guaranteed QoS if provider supports MPLS-based traffic engineering.	Not widespread, but some providers have started to do this.	Not in practice yet; however there is an IETF development activity.	Quite good scalability as it is possible to use hierarchy of MPLS VPN providers. Provisioning automation software exists from Cisco® (probably from Juniper as well); no information at present about its usefulness/ convenience.
SDH and DWDM private optical networks	Providers and customers (through software like UCLP from CANARIE)	Confidentiality – good (network provider may be able to read traffic unless additional encryption is used). Protection of end site – excellent.	Yes	Yes, including non-standard physical coding of optical signals.	Yes: excellent performance is provided natively as it is circuit-switched technology which guarantees bandwidth per user.	Unknown	No, it uses point-to-point circuits.	Poor, as it needs manual configuration of every point-to-point connection. Another negative factor is a shortage of wavelength in most DWDM backbones (no more than 40 generally).

41. None of the technologies listed in the table improves performance natively. However this can be done by deploying QoS across networks. For example, if IP QoS is implemented over the whole path of the L2TP tunnel it can give some guarantee of performance/bandwidth to tunnelled traffic. However, some of the VPN-enabling technologies can be called QoS supportive as they have functionality which can simplify QoS deployment or strengthen QoS guarantees. Other VPN-enabling technologies are QoS neutral as they have no additional functionality which QoS can exploit and benefit from; their traffic looks like standard IP traffic. Of course, both QoS supportive and QoS neutral VPN-enabling technologies can benefit from QoS if it is deployed across a network but in the case of QoS supportive technology the level of QoS guarantees tends to be higher and QoS deployment tends to be simpler. There are some VPN-enabling technologies which have built-in QoS functionality, e.g. ATM and some versions of Frame Relay, but they are rather in decline and not in widespread use within JANET.

Tell us what you think

We welcome your comments on all aspects of this document and on any other of our publications.

Please direct feedback to the JANET Service Desk, at the address below, or e-mail us directly at:

documentation@janet.ac.uk

The JNT Association manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's academic and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Service Desk
The JNT Association
Lumen House
Harwell International Business Centre
DIdcot
Oxon OX11 0SG

Tel: 0870 850 2212
Fax: 0870 850 2213
E-mail: service@janet.ac.uk

Availability:

Further copies of this document may be obtained from the JANET Service Desk at the above address.

This document is also available electronically from: **http://www.ja.net/development/vpn**

ST / VPN / DOC / 004