

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 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 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. (We cannot exclude traffic tapping for private networks completely because, where a leased line is used, traffic 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.

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

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 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).

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, 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.

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 and MPLS. Another has emerged very recently: PBT from Nortel. Currently it is a proprietary technology but Nortel is taking steps to standardise it.

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 and DWDM 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 software developed by CANARIE). 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.

Examples of Centrally Provided VPN Services

RedIRIS

RedIRIS (the Spanish NREN – <http://www.rediris.es> ^[1]) 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, CESCO and CANARIE. More information can be found at:

- UCLP Demonstration at APAN: <http://www.canarie.ca/canet4/uclp/apan/demo.html#Tab2> ^[2]
- UCLP Demonstration at Viola Workshop 2005: <http://www.canarie.ca/canet4/uclp/viola2005/demo.html> ^[3]
- CESCO: <http://www.cesco.es> ^[4]

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. 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.

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** 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 64kbit/s to 45Mbit/s (higher in some locations).
- **IP VPN Secure** 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** services are used for IP VPNs for travelling users or for small offices that only need a dial-up connection to the network

Source URL: <https://community.jisc.ac.uk/library/advisory-services/vpn-definitions-and-understandings>

Links

[1] <http://www.rediris.es/>

[2] <http://www.canarie.ca/canet4/uclp/apan/demo.html#Tab2>

[3] <http://www.canarie.ca/canet4/uclp/viola2005/demo.html>

[4] <http://www.cesca.es/>