Appendix 1: VPN-enabling technologies

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. However, QoS is mentioned in the descriptions of the packet-switching technologies considered here: some of the VPN-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 and SSL are the most popular protocols used nowadays for establishing secure channels. PPTP 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 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 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 N2 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.

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 and L2TP 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 reformulated 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 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
- 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 N2 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 N2 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, Frame Relay or MPLS. ATM and especially frame relay based VPNs were very popular in the 190s 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.

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 155 Mbit/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 (15 Mbit/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, GÉANT2 Vertically Integrated Optical Testbed for Large Applications project (http://www.viola-testbed.de/ [1]) JRA3, HOPI and MUPBED. 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. 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.

Links
[1] http://www.viola-testbed.de/