

VPN Technologies: Definitions and Requirements

VPN Consortium, July 2008

1. Introduction

This white paper describes the major technologies for virtual private networks (VPNs) used today on the Internet. The VPN market has changed significantly in the past ten years as the Internet has grown and as vastly more companies have come to rely on the Internet for communications.

The landscape of VPN products and services offered by a wide variety of vendors continues to evolve. This has caused companies whose networks need protection to become confused about what is and is not a VPN, and the features of the different VPN systems that are being offered to them. The descriptions and definitions in this white paper should help to reduce the confusion for VPN customers, as well as to aid VPN vendors in describing their offerings in a useful fashion.

2. VPN Terminology

A **virtual private network** (VPN) is a private data network that makes use of the public telecommunication infrastructure, maintaining privacy through the use of a tunneling protocol and security procedures. A virtual private network can be contrasted with a system of owned or leased lines that can only be used by one company. The main purpose of a VPN is to give the company the same capabilities as private leased lines at much lower cost by using the shared public infrastructure. Phone companies have provided private shared resources for voice messages for over a decade. A virtual private network makes it possible to have the same protected sharing of public resources for data. Companies today are looking at using a private virtual network for both extranets and wide-area intranets.

This document describes three important VPN technologies: trusted VPNs, secure VPNs, and hybrid VPNs. It is important to note that secure VPNs and trusted VPNs are not technically related, and can co-exist in a single service package.

Before the Internet became nearly-universal, a virtual private network consisted of one or more circuits leased from a communications provider. Each leased circuit acted like a single wire in a network that was controlled by customer. The communications vendor would sometimes also help manage the customer's network, but the basic idea was that a customer could use these leased circuits in the same way that they used physical cables in their local network.

The privacy afforded by these legacy VPNs was only that the communications provider assured the customer that no one else would use the same circuit. This allowed customers to have their own IP addressing and their own security policies. A leased circuit ran through one or more communications switches, any of which could be compromised by someone wanting to observe the network traffic. The VPN customer trusted the VPN provider to maintain the integrity of the circuits and to use the best available business practices to avoid snooping of the network traffic. Thus, these are called **trusted VPNs**.

As the Internet became more popular as a corporate communications medium, security became much more of a pressing issue for both customers and providers. Seeing that trusted VPNs offered no real security, vendors started to create protocols that would allow traffic to be encrypted at the edge of one network or at the originating computer, moved over the Internet like any other data, and then decrypted when it reached the corporate network or a receiving computer. This encrypted traffic acts like it is in a tunnel between the two networks: even if an attacker can see the traffic, they cannot read it, and they cannot change the traffic without the changes being seen by the receiving party and therefore rejected. Networks that are constructed using encryption are called **secure VPNs**.

More recently, service providers have begun to offer a new type of trusted VPNs, this time using the Internet instead

of the raw telephone system as the substrate for communications. These new trusted VPNs still do not offer security, but they give customers a way to easily create network segments for wide area networks (WANs). In addition, trusted VPN segments can be controlled from a single place, and often come with guaranteed quality-of-service (QoS) from the provider.

A secure VPN can be run as part of a trusted VPN, creating a third type of VPN that is very new on the market: **hybrid VPNs**. The secure parts of a hybrid VPN might be controlled by the customer (such as by using secure VPN equipment on their sites) or by the same provider that provides the trusted part of the hybrid VPN. Sometimes an entire hybrid VPN is secured with the secure VPN, but more commonly, only a part of a hybrid VPN is secure.

3. Usage scenarios for secure VPNs and trusted VPNs

The main reason that companies use **secure VPNs** is so that they can transmit sensitive information over the Internet without needing to worry about who might see it. Everything that goes over a secure VPN is encrypted to such a level that even if someone captured a copy of the traffic, they could not read the traffic even if they used hundreds of millions of dollars worth of computers. Further, using a secure VPN allows the company to know that an attacker cannot alter the contents of their transmissions, such as by changing the value of financial transactions. Secure VPNs are particularly valuable for remote access where a user is connected to the Internet at a location not controlled by the network administrator, such as from a hotel room, airport kiosk, or home.

Companies who use **trusted VPNs** do so because they want to know that their data is moving over a set of paths that has specified properties and is controlled by one ISP or a trusted confederation of ISPs. This allows the customer to use their own private IP addressing schemes, and possibly to handle their own routing. The customer trusts that the paths will be maintained according to an agreement, and that people whom the customer does not trust (such as an attacker) cannot either change the paths of any part of the VPN or insert traffic on the VPN. Note that it is usually impossible for a customer to know the paths used by trusted VPNs, or even to validate that a trusted VPN is in place; they must trust their provider completely.

It is clear that secure VPNs and trusted VPNs have very different properties. Secure VPNs provide security but no assurance of paths. Trusted VPNs provide assurance of properties of paths such as QoS, but no security from snooping or alternation. Because of these strengths and weaknesses, **hybrid VPNs** have started to appear, although the list of scenarios where they are desired is still evolving. A typical situation for hybrid VPN deployment is when a company already has a trusted VPN in place and some parts of the company also need security over part of the VPN. Fortunately, none of the common trusted VPN technologies prevent the creation of hybrid VPNs, and some manufacturers are creating systems that explicitly support the creation of hybrid VPN services.

4. Requirements for VPNs

There is one very important requirement that is common to secure VPNs, trusted VPNs, and hybrid VPNs: **the VPN administrator must know the extent of the VPN**. Regardless of the type of VPN in use, a VPN is meant to have capabilities that the "regular" network does not. Thus, the VPN administrator must be able to know at all times what data will and will not be in the VPN.

Each of the four types of VPNs have their own additional requirements.

4.1 Secure VPN requirements

All traffic on the secure VPN must be encrypted and authenticated. Many of the protocols that are used to create secure VPNs allow the creation of VPNs that have authentication but no encryption. Although such a network is more secure than a network with no authentication, it is not a VPN because there is no privacy.

The security properties of the VPN must be agreed to by all parties in the VPN. Secure VPNs have one or more tunnels, and each tunnel has two endpoints. The administrators of the two endpoints of each tunnel must be able to

agree on the security properties of the tunnel.

No one outside the VPN can affect the security properties of the VPN. It must be impossible for an attacker to change the security properties of any part of a VPN, such as to weaken the encryption or to affect which encryption keys are used.

4.2 Trusted VPN requirements

No one other than the trusted VPN provider can affect the creation or modification of a path in the VPN. The entire value of the trusted VPN is that the customer can trust that the provider to provision and control the VPN. Therefore, no one outside the realm of trust can change any part of the VPN. Note that some VPNs span more than one provider; in this case, the customer is trusting the group of providers as if they were a single provider.

No one other than the trusted VPN provider can change data, inject data, or delete data on a path in the VPN. A trusted VPN is more than just a set of paths: it is also the data that flows along those paths. Although the paths are typically shared among many customers of a provider, the path itself must be specific to the VPN and no one other than trusted provider can affect the data on that path. Such a change by an outside party would affect the characteristics of the path itself, such as the amount of traffic measured on the path.

The routing and addressing used in a trusted VPN must be established before the VPN is created. The customer must know what is expected of the customer, and what is expected of the service provider, so that they can plan for maintaining the network that they are purchasing.

4.3 Hybrid VPN requirements

The address boundaries of the secure VPN within the trusted VPN must be extremely clear. In a hybrid VPN, the secure VPN may be a subset of the trusted VPN, such as if one department in a corporation runs its own secure VPN over the corporate trusted VPN. For any given pair of address in a hybrid VPN, the VPN administrator must be able to definitively say whether or not traffic between those two addresses is part of the secure VPN.

5. Technologies Supported by VPNC

The following technologies support the requirements from the previous section. VPNC supports these technologies when they are implemented by users themselves and when they are implemented in provider-provisioned VPNs.

5.1 Secure VPN technologies

- **IPsec with encryption** in either tunnel and transport modes. The security associations can be set up either manually or using IKE with either certificates or preshared secrets. IPsec is described in many RFCs, including 2401, 2406, 2407, 2408, and 2409 (for IKEv1), and 4301, 4303, 4306, 4307, and 4308 (for IKEv2).
- **IPsec inside of L2TP** (as described in RFC 3193) has significant deployment for client-server remote access secure VPNs.
- **SSL 3.0 or TLS with encryption.** TLS is described in RFC 4346. An excellent book on SSL 3.0 and TLS is "SSL and TLS: Designing and Building Secure Systems" by Eric Rescorla (ISBN 0201615983).

These technologies (other than SSL 3.0) are standardized in the IETF, and each has many vendors who have shown their products to interoperate well in the field.

5.2 Trusted VPN technologies

Modern service providers offer many different types of trusted VPNs. These can generally be separated into "layer 2" and "layer 3" VPNs.

Technologies for trusted layer 2 VPNs include:

- **ATM circuits**
- **Frame relay circuits**
- **Transport of layer 2 frames over MPLS**, as described in draft-ietf-l2vpn-vpls-bgp and other related Internet Drafts.

Technologies for trusted layer 3 VPNs include:

- **MPLS with constrained distribution of routing information through BGP**, as described in RFC 4364 and other related Internet Drafts.

It is widely assumed that both will become standards in the future. Also, the service provider industry has not embraced one of these technologies much more strongly than the other.

5.3 Hybrid VPN technologies

- **Any supported secure VPN technologies running over any supported trusted VPN technology.**

It is important to note that a hybrid VPN is only secure in the parts that are based on secure VPNs. That is, adding a secure VPN to a trusted VPN does not increase the security for the entire trusted VPN, only to the part that was directly secured. The secure VPN acquires the advantages of the trusted VPN, such as having known QoS features.

6. About VPNC

The VPN Consortium (VPNC) is the international trade association for manufacturers in the VPN market. The primary purposes of the VPNC are:

- Promote the products of its members to the press and to potential customers
- Increase interoperability between members by showing where the products interoperate
- Serve as the forum for the VPN manufacturers throughout the world
- Help the press and potential customers understand VPN technologies and standards
- Provide publicity and support for interoperability testing events

It should be noted that VPNC does not create standards; instead, it strongly supports current and future IETF standards.

More information on the VPN Consortium can be found at [the VPNC web site](http://www.vpnc.org).