The Internet Protocol Journal

June 1998 Volume 1, Number 1

*A Quarterly Technical Publication for Internet and Intranet Professionals*

## In This Issue

From the Editor ..... 1

What Is a VPN?–Part I ..... 2

SSL: Foundation for Web Security ..... 20

Call for Papers ..... 30

Book Reviews ..... 31

Fragments ..... 35



## From The Editor

Welcome to the first edition of *The Internet Protocol Journal* (IPJ). This publication is designed to bring you in-depth technical articles on current and emerging Internet and intranet technologies. We will publish technology tutorials, as well as case studies on all aspects of internetworking.

Our first article is a detailed look at *Virtual Private Networks* (VPNs). Many organizations are turning to VPNs as a cost-effective way to implement enterprise networking, but the industry has not yet settled for a single approach, nor even a single definition of the VPN concept. The article by Paul Ferguson and Geoff Huston is in two parts. Part II will follow in our second issue, due out in September.

When the Internet Protocol suite (TCP/IP) was first designed, security was not a major consideration. Indeed, the primary goal in the early days of networking was sharing of information among academics and researchers. Today, TCP/IP is being used for mission-critical applications and for the emerging area of electronic commerce. As a result, security mechanisms are being added at all levels of the protocol stack. In this issue, we take a closer look at the *Secure Sockets Layer* (SSL), which is used for Web transactions. William Stallings explains how SSL works and how it is becoming the standard for Web security.

If you want to learn about computer networks, many options are available, including conferences, journals, standards documents, Web sites, glossaries and, of course, books. Our *Fragments* page gives you some pointers for further reading, and every issue will include at least one book review.

A detailed description of the scope of this journal can be found on page 30 in our *Call for Papers*. We want your input in this new publication. Please send comments, suggestions or questions to `ipj@cisco.com`. You may also use this address to request a complimentary copy of the next issue of IPJ. If you would like to write an article, send me e-mail and I will send you author guidelines.

*—Ole J. Jacobsen, Editor and Publisher*
`ole@cisco.com`

# What Is a VPN? — Part I

*by Paul Ferguson, Cisco Systems*
*and Geoff Huston, Telstra*

The term "VPN," or *Virtual Private Network*, has become almost as recklessly used in the networking industry as has "QoS" (Quality of Service) to describe a broad set of problems and "solutions," when the objectives themselves have not been properly articulated. This confusion has resulted in a situation where the popular trade press, industry pundits, and vendors and consumers of networking technologies alike generally use the term VPN as an offhand reference for a set of different technologies. This article provides a common-sense definition of a VPN, and an overview of different approaches to building one.

> "The wonderful thing about virtual private networks is that its myriad definitions give every company a fair chance to claim that its existing product is actually a VPN. But no matter what definition you choose, the networking buzz-phrase doesn't make sense. The idea is to create a private network via tunneling and/or encryption over the public Internet. Sure, it's a lot cheaper than using your own frame relay connections, but it works about as well as sticking cotton in your ears in Times Square and pretending nobody else is around."[1]

## A Common-Sense Definition

As *Wired Magazine* notes in the quotation, the myriad definitions of a VPN are less than helpful in this environment. Accordingly, it makes sense to begin this examination of VPNs to see if it is possible to provide a common-sense definition of a VPN. Perhaps the simplest method of attempting to arrive at a definition for VPNs is to look at each word in the acronym individually, and then tie each of them together in a simple, common-sense, and meaningful fashion.

Let's start by examining the word "network." This term is perhaps the least difficult one for us to define and understand, because the commonly accepted definition is fairly uncontroversial and generally accepted throughout the industry. A network consists of any number of devices that can communicate through some arbitrary method. Devices of this nature include computers, printers, routers, and so forth, and they may reside in geographically diverse locations. They may communicate in numerous ways because the electronic signaling specifications, and data-link, transport, and application-layer protocols are countless. For the purposes of simplicity, let's say that a "network" is a collection of devices that can communicate in some fashion, and can successfully transmit and receive data among themselves.

The term "private" is fairly straightforward, and is intricately related to the concept of "virtualization" insofar as VPNs are concerned, as we'll discuss in a moment. In the simplest of definitions, "private" means communications between two (or more) devices is, in some

fashion, secret—that the devices that are not participating in the "private" nature of communications are not privy to the communicated content, and that they are indeed completely unaware of the private relationship altogether. Accordingly, data privacy and security (data integrity) are also important aspects of a VPN that need to be considered when implementing any particular VPN.

Another means of expressing this definition of "private" is through its antonym, "public." A "public" facility is one that is openly accessible, and is managed within the terms and constraints of a common public resource, often via a public administrative entity. By contrast, a private facility is one where access is restricted to a defined set of entities, and third parties cannot gain access. Typically, the private resource is managed by the entities who have exclusive right of access. Examples of this type of private network can be found in any organizational network that is not connected to the Internet, or to any other external organizational network, for that matter. These networks are private because there is no external connectivity, and thus no external network communications.

Another important aspect of privacy in a VPN is through its technical definition. For example, privacy in an addressing and routing system means that the addressing used within a VPN community of interest is separate and discrete from that of the underlying shared network, and from that of other VPN communities. The same holds true for the routing system used within the VPN and that of the underlying shared network. The routing and addressing scheme within a VPN should, in general, be self-contained, but this scenario degenerates into a philosophical discussion of the context of the term "VPN." Also, it is worthwhile to examine the differences between the "peer" and "overlay" models of constructing VPNs—both of which are discussed in more detail later under the heading "Network-Layer VPNs."

"Virtual" is a concept that is slightly more complicated. *The New Hacker's Dictionary* (formerly known as the Jargon File)[2] defines virtual as:

> virtual /adj./ [via the technical term "virtual memory," prob. from the term "virtual image" in optics] 1. Common alternative to {logical}; often used to refer to the artificial objects (like addressable virtual memory larger than physical memory) simulated by a computer system as a convenient way to manage access to shared resources. 2. Simulated; performing the functions of something that isn't really there. An imaginative child's doll may be a virtual playmate. Oppose {real}.

Insofar as VPNs are concerned, the second definition is perhaps the most appropriate comparison for virtual networks. The "virtualization" aspect is one that is similar to what we briefly described previously as private, but the scenario is slightly modified—the private communication is now conducted across a network infrastructure that

is shared by more than a single organization. Thus, the private resource is actually constructed by using the foundation of a logical partitioning of some underlying common, shared resource rather than by using a foundation of discrete and dedicated physical circuits and communications services. Accordingly, the private network has no corresponding private physical communications system. Instead, the private network is a virtual creation that has no physical counterpart.

The virtual communications between two (or more) devices is because the devices that are not participating in the virtual communications are not privy to the content of the data, and they are also altogether unaware of the private relationships between the virtual peers. The shared network infrastructure could, for example, be the global Internet and the number of organizations or other users not participating in the virtual network may literally number into the thousands or even millions.

A VPN can also said to be a discrete network[3]:

> (discrete \dis*crete"\, a. [L. discretus, p.p. of discernere. See Discreet.] 1. Separate; distinct; disjunct).

The discrete nature of VPNs allows both privacy and virtualization. Although VPNs are not completely separate, intrinsically, the distinction is that they operate in a discrete fashion across a shared infrastructure, providing exclusive communications environments that do not share any points of interconnection.

The combination of these terms produces VPN—a private network, where the privacy is introduced by some method of virtualization. A VPN could be built between two end systems or between two organizations, between several end systems within a single organization or between multiple organizations across the global Internet, between individual applications, or any combination.

It should be noted that there is really no such thing as a nonvirtual network, if the underlying common public transmission systems and other similar public infrastructure components are considered to be the base level of carriage of the network. What separates a VPN from a truly private network is whether the data transits a shared versus a nonshared infrastructure. For instance, an organization could lease private line circuits from various telecommunications providers and build a private network on the base of these private circuit leases, but the circuit-switched network owned and operated by the telecommunications companies are actually circuits connected to their *Digital Access and Crossconnect Systems* (DACSs) network and subsequently their fiber-optics infrastructure. This infrastructure is shared by any number of organizations through the use of multiplexing technologies. Unless an organization is actually deploying private fiber and layered transmission systems, any network is layered with "virtualized" connectivity services in this fashion.

A VPN doesn't necessarily mean communications isolation, but rather the controlled segmentation of communications for communities of interest across a shared infrastructure.

The common and somewhat formal characterization of the VPN, and perhaps the most straightforward and strict definition, follows:

> A VPN is a communications environment in which access is controlled to permit peer connections only within a defined community of interest, and is constructed though some form of partitioning of a common underlying communications medium, where this underlying communications medium provides services to the network on a nonexclusive basis.

A simpler, more approximate, and much less formal description follows:

> A VPN is private network constructed within a public network infrastructure, such as the global Internet.

It should also be noted that although VPNs may be constructed to address any number of specific business needs or technical requirements, a comprehensive VPN solution provides support for dial-in access, support for multiple remote sites connected by leased lines (or other dedicated means), the ability of the VPN service provider (SP) to "host" various services for the VPN customers (for example, Web hosting), and the ability to support not just intra-, but also inter-VPN connectivity, including connectivity to the global Internet.

### VPN Motivations

There are several motivations for building VPNs, but a common thread is that they all share the requirement to "virtualize" some portion of an organization's communications—in other words, make some portion (or perhaps all) the communications essentially "invisible" to external observers, while taking advantage of the efficiencies of a common communications infrastructure.

The base motivation for VPNs lies in the economics of communications. Communications systems today typically exhibit the characteristic of a high fixed-cost component, and smaller variable-cost components that vary with the transport capacity, or bandwidth, of the system. Within this economic environment, it is generally financially attractive to bundle numerous discrete communications services onto a common, high-capacity communications platform, allowing the high fixed-cost components associated with the platform to be amortized over a larger number of clients. Accordingly, a collection of virtual networks implemented on a single common physical communications plant is cheaper to operate than the equivalent collection of smaller, physically discrete communications plants, each servicing a single network client.

DOCKET ALARM

# Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

## Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

## Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

## Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

## API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

**LAW FIRMS**
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

**FINANCIAL INSTITUTIONS**
Litigation and bankruptcy checks for companies and debtors.

**E-DISCOVERY AND LEGAL VENDORS**
Sync your system to PACER to automate legal marketing.