Adding the “N” to Virtual Private Networking

The term virtual private networking (VPN) has had several incarnations in the telecom and networking industries. It has been applied to carrier-based, corporate voice networks and to frame relay and ATM networks that use a shared network to deliver private services to customers. For the purposes of this white paper, VPN is the secure transport of private traffic over a public IP network. This function is supplied by combining security (encryption, authentication and access control) with tunneling protocols to create authenticated, encrypted tunnels through a public shared IP network. By interconnecting many of these private tunnels, logical private networks are created.

For businesses, virtual private networking is a new method of building corporate communication networks. VPNs enable enterprises to reduce their dependence on expensive, leased-line networks and troublesome remote-access solutions by establishing VPN connections across shared IP networks such as the Internet. Most analysts agree that VPNs will experience explosive growth in the coming years for the simple reason that they save money. By leveraging Internet economies of scale, VPNs offer the promise of reducing the cost of operating the critical corporate IT infrastructure. Improved flexibility, security and global reach are additional benefits provided by this innovative method for setting up private networks.

However, despite the fact that the value of VPN technology has been accepted for several years, VPNs have not received the widespread market penetration that analysts anticipated. If there is such broad acceptance that the positive return on investment (ROI) for VPNs is accurate, why have network managers been slow to adopt the technology on a larger scale? The answer is that, although the first generation of VPNs offered network managers the ability to save money by replacing expensive private line facilities and cumbersome remote modem banks, the cost savings were limited. First-generation VPN products did not include many of the networking and automation features upon which network managers have grown to rely.

Many early products established point-to-point tunnels between two sites by encrypting the data, authenticating the end points, and authorizing the user through access controls. Those tunnels offered the advantage of preventing data from being viewed or altered by external sources, but also the disadvantage of requiring personnel resources to manage all creation, establishment and changes to tunnels.

The network is a combination of point-to-point tunnels. VPN tunnels can be provisioned between a software client and a central office concentrator to support remote access users and between two hardware devices for site-to-site connectivity. Initial VPN products provided the ability to create manually configured, static, point-to-point tunnels between two locations. By establishing and maintaining static routing tables in the hardware devices, the tunnels could be used to develop a static network (see Figure 1). However, manually configured static tunnels are difficult to administer, preventing meshed connectivity required for larger scale deployments.

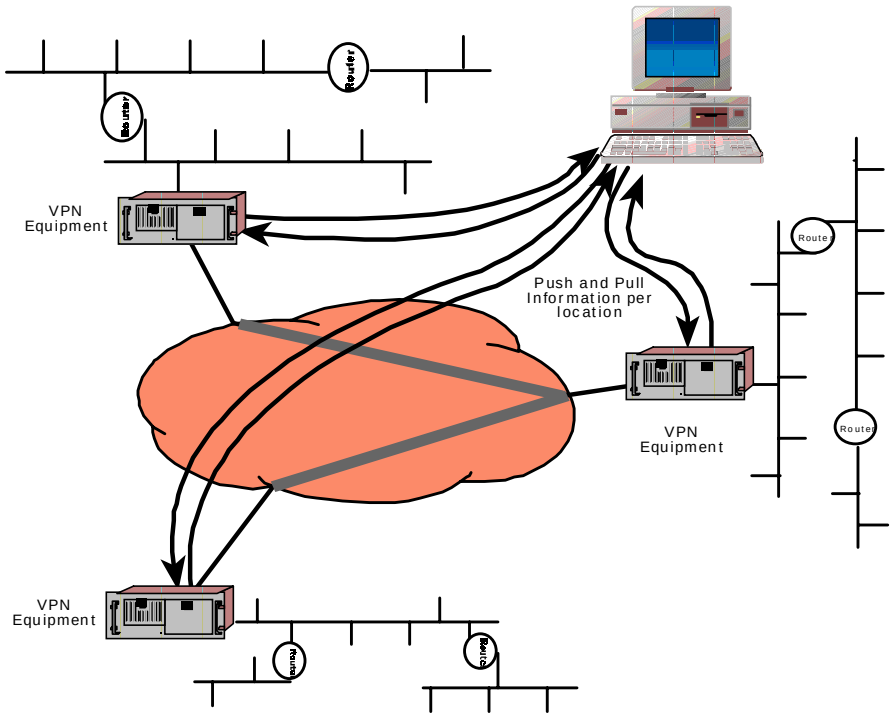

Fig 1. Static VPN Architecture

Many early VPN implementations were also architected in a hub-and-spokes configuration to handle the topology and management tasks that point-to-point tunnels require. Each tunnel needs to be managed individually and meshed configuration topologies become virtually impossible to maintain and control. Today, network managers are accustomed to automated topology configurations and fully meshed connectivity. Reverting to secured point-to-point tunnels, where connection configurations are controlled by a third party, is akin to returning to the days of manual telephone switchboards, which required human intervention to establish and maintain connections. In a way, first-generation VPN products were a step forward for Internet security, but a step backward for networking and ease of use.

New DHCP and BootP relay features were also not available in early VPN offerings. The result was that network managers who chose to implement first-generation VPN devices were not only required to learn a new set of security technologies, but also had to revert their network-provisioning model to more manual processes. Since the trend in networking is automation, network managers were reticent to take this step back to manual processes.

Automated VPNs

Equating tunnels with VPNs is a very constrained and limiting use of the technology. Network managers now rely on automated topology configuration and network discovery. These two features are critical for VPN devices to provide networks that are easy to deploy and maintain. Now, second-generation products are coming to market and alleviating these limitations by providing policy-driven automation of the topology configuration and security functions. They also embed automated networking features, such as dynamic routing protocols, into the products, so that VPNs can be used as networks with fully meshed connectivity where connection set-up and ongoing topology changes are dynamically managed.

To provide network managers with the high ROI they expect, VPN products must be easy to install, operate and maintain. VPNs must not burden network managers with additional operations and maintenance functions that generate hidden costs. If operational costs are too high, the savings from the

use of a shared infrastructure to provision the corporate backbone will be offset, and the promise of VPNs will not be realized.

Centralizing the provisioning information for large-scale networks enables automated topology configuration; network and security parameters need to be configured only once rather than multiple times to accommodate each location. This facilitates the addition and removal of devices and locations by automating the establishment of tunnels through the centralized provisioning server. Running dynamic routing protocols, such as RIP, OSPF and IGRP, across the tunnels provides automatic network discovery. This simplifies ongoing network operations and automates the constant moves, adds and changes in corporate networks, eliminating the need to maintain static routing tables in each of the network devices. Corporate network routers automatically discover the presence of new routers and select the new best routes for traffic (see Figure 2). These dynamic routing protocols offer the ability to create meshed networks that are easy to maintain. By meshing the network, redundant paths are created automatically and enabled upon link or device failure. Some dynamic routing protocols even offer the ability to balance the load between paths, essentially providing network managers with additional bandwidth. Network managers have grown accustomed to these features and expect them in networking devices.

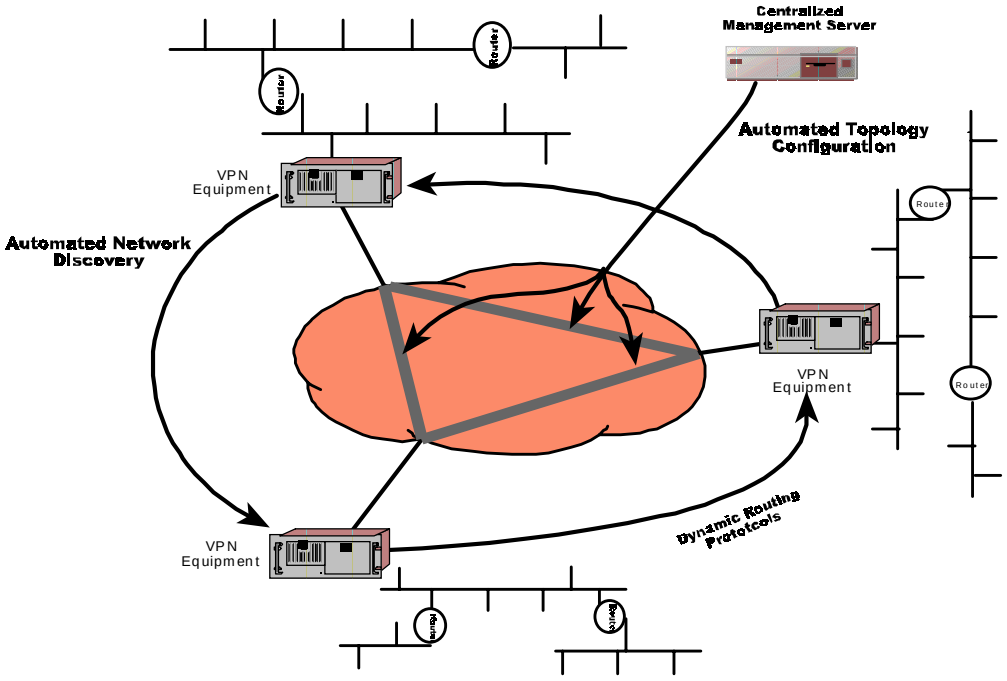

Fig 2. Networked VPN Architecture

VPN products now offer the networking features that are embedded in other internetworking devices. Paramount among these are integrated dynamic routing protocols. By embedding dynamic routing in VPN devices, fully secure and meshed networks can be developed. Network topology information is exchanged through tunnels, allowing VPN devices to automatically learn the topology of the network on opposite ends of a tunnel. In addition, these new products automate the provisioning of security parameters through policy-driven provisioning systems based on public key infrastructure and directory-driven configurations. These advances make VPNs easier and less costly to deploy and manage.

By developing meshed VPNs, network managers enjoy many benefits, including embedded networking and automation. Such features will propel this emerging technology into the mainstream of enterprise network provisioning.