Global Information Assurance Certification Paper

Implementing Virtual Private Networks – Observations from the field.
GSEC Practical Assignment v1.2d
Andrew Egorov
29th April, 2001

Introduction

Private leased lines have traditionally formed the core of the majority of the corporate worlds' WAN infrastructure. The reliability and point to point security provided by private leased lines serves its purpose well when used to connect a small number of remote sites to a central administration site, typically in a star network topology.

However, the burgeoning ecommerce industry requires that effective communications exist between many different branch or remote sites within a corporation, as well as between business partners. Meshed network topologies are the only logical way to provide the many-to-many connectivity that ecommerce depends on.

Unfortunately, the point to point nature of private leased lines mean that they do not scale well for use in meshed network topologies. The number of private leased lines that would be required to implement such a topology rise exponentially with the number of sites, resulting in increased cost and network complexity.

The introduction of frame relay has enabled corporations to continue using their existing star topology networks cost effectively. However, like private leased lines, if meshed networks are required, the costs of deploying and managing the associated PVCs, not to mention the complexity of designing such a network infrastructure, mean that frame relay suffers from the same scalability problems.

The dramatic growth of the IP-based public network, the Internet, has inadvertently provided a solution to the meshed network dilemma facing corporations. The very nature of the Internet, with its pervasive spread across national and international boundaries, both political and geographical, provides a common intricate backbone linking together the majority of corporations. Being a public network, however, means that the security and reliability offered by private leased lines no longer exists. However, it wasn't long before it was found that the potential cost-savings realised by utilising the cheap connectivity offered by this infrastructure, would far outweigh the disadvantages, provided a secure means of transmitting data over the internet could be found.

Enter the Virtual Private Network or VPN. VPNs are so called because they "appear" to provide the reliable and secure services of a private leased line, even though these services are being delivered by the technically unreliable and insecure public Internet.

The requirements for ecommerce and more recently, telecommuting, has resulted in an ever-increasing demand for cost effective and secure meshed network connectivity. By using the Internet and VPNs, this can be achieved relatively cheaply and easily. VPN use has been on the increase for several years, with most of the big

telecommunications companies now providing IP VPN services.

Whilst it cannot be denied that VPNs are a cost effective way to achieve business-to-business, remote site and mobile worker connectivity, it has become apparent that they are not as simple to implement as all the hype would have us believe.

This paper will present the observations of the author, who was involved in a large international VPN implementation. It will focus on various aspects of the VPN implementation in which problems were encountered, and with benefit of hindsight, will discuss ways in which these problems could have been avoided.

It is hoped this paper will provide practical insight and guidance for those contemplating the implementation of similar VPN infrastructures in the future.

The VPN Implementation

The VPN implementation project was initiated by an international client who expressed a wish to trial a pilot VPN across several branch offices in various countries.

The client had two specific requirements that it wanted the VPN implementation to fulfil.

The primary requirement was that the VPN provide remote dial-in access via the Internet for all its mobile sales staff. These sales staff would be constantly on the move not only within their home country, but abroad in other countries as well. They should be able to access data from their base corporate office, at any time, from anywhere.

The secondary requirement was that the VPN provide Intranet / Extranet site-to-site connectivity via the Internet for specified sites in remote locations that could not be serviced by traditional leased lines.

Remote VPN Dial-in Access

VPN client software

Perhaps the biggest problem faced when implementing the remote VPN dial-in access component of the solution, was the large diversity and number of desktops that had to be accounted for. This was the focus of an article by Salvatore Salamone[1] that found…"Companies moving to VPNs to cut networking costs have found that while there may be toll savings on paper, expenses related to deploying and maintaining VPNs can nullify the benefits".

We too, found that the effort required to audit and determine what desktops existed, then testing the VPN client on all the various desktop configuration permutations was drastically underestimated. Even when this was completed, problems arose when the VPN client installation and use policy conflicted with the security policy of the client.

What the VPN appliance vendor failed to inform us of, was that on NT v4.0

workstations, the VPN client had to be installed, AND used by a Local Administrator of the desktop. This was in direct violation of the clients', and I suspect most other companies' security policy, which specified that normal users of the desktop were NOT given local administrator rights to their desktop.

This problem was eventually fixed by the vendor with the next version release of the VPN client software.

Of course the other problem relating to the VPN client software was the logistics involved in distributing the VPN client software to the 5000+ users. The client overcame this only after much planning and expense. Again, this aspect of any VPN implementation should not be underestimated, as failing to do so severely undermines the cost benefits that would have otherwise been realised.

Maximum MTU size

While not as an important factor to consider as the above, several problems arose as a result of incorrect MTU size being set on the VPN client.

The default MTU size of 1500 bytes resulted in adverse response times being experienced by remote users. The 41-byte encryption header added to the packet by the VPN appliance caused excessive fragmentation of the packet as it traversed the Internet. The majority of traffic on the Internet consists of packets smaller than 576 bytes in size. Smaller packets traverse the Internet much faster, as most routers can handle smaller packets without the need for fragmentation.

After much testing of different packet sizes, it was found that as soon as the maximum MTU size on the client was set to 576 bytes, the response time was immediately improved.

Universal dialling

One of the requirements put forward by the client was that mobile sales staff be able to establish a connection to their base corporate office from anywhere at anytime. Whilst connectivity at any time of the day wasn't a problem, trying to provide a single dialup connection on their laptop from which they could establish an internet connection from anywhere in the world proved to be almost impossible.

Not a single carrier had a PoP in all the countries that the clients' mobile sales staff could potentially travel to. The only option available to us at the time was to select a small number of carriers that between them could provide the PoP localities required. Bespoke software was then written that provided a common interface that accessed the required PoP depending on which country the mobile sales staff member was visiting.

Outsourcing peripheral VPN components

A problem that will continue to be encountered, at least into the immediate foreseeable future, is the lack of a single vendor suite of VPN components.

In our VPN implementation, we had inbound authentication being provided by RSA SecureID tokens, intrusion detection by ISS RealSecure, outbound authentication by CiscoSecure ACS TACACS+, VPN termination by Cisco Altigas, and the list goes on. The list of interoperability problems encountered is a long one, and many times throughout the testing features had to be disabled or severely limited to allow the solution as a whole to continue functioning. The issuing of digital certificates from one vendor's digital certificate server, for example, would not work with the other vendor's VPN appliance, forcing the use of pre-shared keys.

The lack of interoperability will continue to be a disadvantage when implementing a feature rich VPN implementation, however, as the VPN industry is still in its infancy, it won't be long before significant improvements are made to rectify this. Indeed, in the 8 months it took for this particular implementation to be completed, numerous code and firmware revisions were released by several vendors, which drastically improved the operation of their respective components.

Intranet / Extranet Site-to-Site VPNs

The second requirement requested by the client, was that VPNs be used to provide Internet / Extranet site-to-site connectivity via the Internet for specified sites in remote locations that could not normally be serviced by traditional private leased lines.

Multiple Carriers

Again, the main problems, apart from the interoperability issues discusses previously, were experienced by the fact that not a single carrier had PoPs in all the locations required.

In some of the remote locations, the only PoPs that were locally available were 3rd party partner ISPs. Usually these ISPs were of poor quality and did not have the support infrastructure and technical expertise to ensure that the service was of a satisfactory standard and that faults be resolved expeditiously. As a result, SLAs failed to be met during critical outages, and change requests sometimes took days to be completed.

Although unlikely at the moment, carriers providing IP VPN services should attempt to guarantee that they administer all the PoPs that any potential VPN dial-up client may need.

The unreliable Internet

Perhaps the most underestimated and forgotten aspect of VPN solutions in general is that all this connectivity is traversing an inherently unreliable and non-centrally administered network. This fact made itself blindingly obvious over a period of several months, when it was discovered that asymmetric routing was occurring between two sites.

The main site in question was a major hub site in Europe, whilst the other, a remote site based in Asia Minor. The problem began manifesting itself when users at the remote site noticed erratic performance problems whilst accessing systems at the main hub site.

Extensive troubleshooting involving packet traces from various sources revealed that the source path (from the remote site to the main hub site) differed from the return path (from the main hub site to the remote site). It was also discovered that packets returning to the remote site were being filtered out by an ISP in the United Arab Emirates region. To make matters worse, it was thought that the carrier involved had partnerships with all the ISPs on the route, when in fact, because of the asymmetric routing, quite a few additional ISPs were involved.

Over the course of the next few months, considerable effort and expense was wasted in tracking down the administrators of the various AS areas involved (AS areas are autonomous systems used in BGP) because several different ISPs were involved in the different paths. Eventually it was discovered that IPSEC traffic was being filtered out and that extensive fragmentation was occurring.

Conclusion

When faced with the prospect of implementing a VPN solution, particularly an international solution, for a client, never underestimate the planning and commitment that needs to be made for the project to succeed.

Recommendations that should be noted before any VPN implementation can be summarised as follows:

1. Ensure that all desktops requiring VPN client installation are identified and tested for interoperability with the VPN client. If possible, try to standardise on a particular build specifically designed for VPN use.

2. Ensure that a full analysis of the current network topology and traffic composition is performed and a performance baseline obtained. This will enable any problems with the introduction of VPN (IPSEC) packets to be more easily resolved.

3. Ensure that if requiring PoPs in various locations that the carrier used has ISPs in all locations, has full administrative control of same and a guaranteed SLA. If possible, also ensure that the carrier can guarantee that they control or own all the networks from source to destination.

4. Ensure that all efforts are made to minimise the number of different components and vendors when designing a VPN solution. This will ease interoperability and support problems.

References

1. ***Keep an Eye Out for the Hidden Costs***
   Internet Week, Salvatore Salamone, March 29, 1999
   http://www.internetwk.com/VPN/supplement329-2.htm

Other Sources

Deployment Hurdles: Look Before You Leap
Internet Week, Salvatore Salamone, March 29, 1999
http://www.internetwk.com/VPN/supplement329-5.htm

Virtual Private Networks
Ennovate Networks
http://www.ennovatenetworks.com/technology/apps/vpnsol/overview.htm

The Vaunted VPN
Network World, Tim Greene, 27th September, 1999
http://www.nwfusion.com/buzz99/buzzvpn.html

VPN: Light at the end of the tunnel
Federal Computer Week, Tom Yager, 3rd July, 2000
http://www.fcw.com/fcw/articles/2000/0703/tec-vpn-07-03-00.asp

Virtual Private Network Overview
http://www.tda.ecrc.ctc.com/kbase/virtual/VPN.html