Providing Secure Access to Information Using the Internet

John Holmwood, *TransCanada*
Kevin Reichert, *Alterna*
Blaine Feniak, *TransCanada*

Introduction

Historically, TransCanada PipeLines Limited, like most companies, has had a security policy that required its computing infrastructure (networks and computing systems) to be isolated from public networks, including the Internet. The existing infrastructure took advantage of this isolation when it was designed. New business opportunities require TransCanada to open up its computing infrastructure to public networks, particularly the Internet. The security policies have been revised to ensure that, while everyone who has a need to access TransCanada information can do so, individuals who do not have a legitimate reason to see TransCanada information cannot do so.

Security Requirements

There are a few main security requirements for connecting TransCanada's computing infrastructure to public networks. First, the connection must not provide an unintended access capability into our infrastructure. Secondly, we must ensure that no one can access TransCanada's information passing through the public network, without permission. Finally, it must be possible to ensure that anyone, who accesses TransCanada's general computing infrastructure from outside that infrastructure, whether using public or private networks, is properly authenticated. This paper describes the analysis TransCanada undertook to determine its strategy and architecture for providing secure access to company information over the Internet.

What is a VPN

A Virtual Private Network (VPN) is composed of two fundamental components:

- A tunnel –a network of virtual circuits for carrying private traffic [1]over the Internet that is created by encapsulating the data in special IP packets. The virtual in VPN.

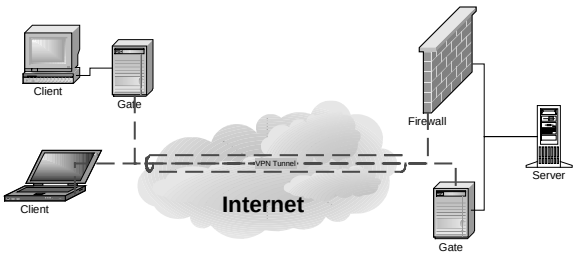

Figure 1 Possible Virtual Private Network Configurations

Figure 1 above depicts possible VPN configurations. A tunnel is set up from either an individual workstation, or from a gate on a subnet at a remote site to a gate or firewall at the local site. There are a number of protocols that can be used to encapsulate the data in special IP packets, thereby creating the tunnel. See the section on protocols below.

- Security Services – the authentication, access control, confidentiality and integrity services. The private in VPN. These services are provided by cryptographic procedures such as encryption.

How to choose the right VPN

There are a very large number of products on the market that purport to be VPNs. For example, at the 1999 Interop+Networld conference in Las Vegas there were over 50 vendors who had VPN products based on IPSec. Therefore, it is very important to understand the features that distinguish these products, and to determine which of these features are needed for a particular application. The main distinguishing characteristics are:

- Firewall or Stand-Alone device
- Security services included
- Protocols supported
- Cost

Firewall/Stand Alone

Typically a VPN requires a device at the perimeter of the company network on which the tunnel terminates. Most VPN products can be divided into either products with tunnels that terminate on firewalls, or products with tunnels that terminate on stand alone devices (e.g. routers or servers).

Firewall

The products which terminate on firewalls generally come from firewall vendors (Quel Surpris☺). They are

either integral to the firewall product, or come as add on modules.

Advantages

The main advantage offered by firewalls that act as terminators for VPN tunnels is that only one device needs to be exposed to the Internet. This lowers the exposure to attack.

Concerns

A firewall that has to support a VPN tunnel as well as do its normal job of controlling access to the network means it must be more complicated. More complicated devices are more difficult to harden properly to prevent successful attacks. The device must also have more computer power to handle the added functionality.

Router

The products that terminate their tunnel on routers or router-like devices (gates) are usually stand-alone proprietary hardware with embedded software. The device is connected in parallel to or in front of the firewall.

Advantages

The main advantages are speed and simplicity. The software is embedded in a special purpose device so relatively inexpensive hardware provides fast throughput with simple operations and maintenance. The manufacturer has hardened the device.

Concerns

The main disadvantage is that it is an additional device that is accessible from the Internet. For consulting firms that must support many VPNs the additional devices can become a large maintenance cost.

Servers

The products that terminate their tunnels on servers are generally software-only solutions that use general-purpose servers (MS Windows or Unix) rather than proprietary hardware. The server is connected in parallel to, or in front of, the firewall.

Advantages

Software only solutions are generally less expensive. General-purpose hardware is typically less expensive to operate than proprietary hardware. For consulting firms that must support many VPNs, being able to install multiple VPNs on a single server can reduce the operating and maintenance costs.

Concerns

The main disadvantage is that there is an additional device that is accessible from the Internet, which the administrator must harden from attack

Security Services

Authentication

This service is a measure of how sure you can be of the identity of the person setting up the tunnel. The New Yorker published a cartoon a few years ago that depicted two dogs looking at a desktop computer. The caption read “On the Internet no one knows you’re a dog”. There are three basic types of authentication:

- Challenge/response (something you know)
- Token (something you have)
- Biometrics (something you are)

Challenge/Response

This is the most familiar authentication method. The user provides a name and password. Over the years, Remote Access Services (RAS) have developed a number of protocols based on this method. These include: PAP, CHAP, RADIUS, DIAMETER, and EAP.

Advantages

It’s simple and inexpensive. No special hardware is required at the client or remote end.

Concerns

Its not very strong. Usernames are generally based on an algorithm such as last name/first initial. Once the algorithm is known, it is simple to guess someone’s username. There are lots of password-cracking programs available on the Internet.

Token

This authentication method is based upon a ‘token’ that has been given to the user. The person attempting to set up the tunnel must demonstrate they have the token as well as knowing a name and password. A simple example of a token is a bankcard. Some genius realized that the average person carries so many credit and bank cards that they can’t memorize all of the numbers on the cards. By requiring the number on the card to be entered before the connection is completed, it is possible to be confident that the person setting up the connection has the card in their possession.

Other common forms of tokens include public key certificates, smart cards and one-time password or code generators. Public key certificates are software that contain one or more encryption keys based on the X.509 standard. Smart cards also usually contain X.509 encryption keys, however, since they are stored on a memory chip, the user is more likely to know if the keys have been stolen. (The card is gone☺). One time passwords are usually based on the S/Key algorithm and code generators are special purpose computing devices (e.g. SecurID, SafeWord or CryptoCard) that calculate a large pseudo-random number.

Advantages

Tokens are usually considered strong authentication because you have to have something (the token) as well as know something (a password).

Concerns

Smart cards require that a special reader be attached to the client in order to read the card.

Certificates must be stored on the client computer limiting the versatility of the token.

One-time codes require an application at the protected site that can calculate the same number as the token that is handed out to the user.

A process must be developed and implemented to quickly replace lost tokens, as the person setting up the connection must have a token in their possession.

Biometrics

This authentication method is based upon things that is intrinsically unique about an individual such as voice prints, fingerprints or retinal patterns.

Advantages

There is a high probability that a person identified by biometrics is who they claim to be although there are the scenes in *Terminator II* and *Judge Dread*, where a helpless victim is dragged to an information kiosk to provide identification. The Black Hat then disposes of the victim and proceeds to extract the information being sought.

Concerns

This method of authentication requires specialized equipment to be attached to the remote end of the connection.

Access Control

This service provides the capability of restricting in some manner what the user can access once they have been authenticated. The restriction can either be done on the basis of packet filtering, which restricts either certain types of packets, or packets with specific addresses, or it can be done on the basis of policies, which define parameters such as application, individual, group and/or subnet.

Administrators may need to identify the individual user, not just the address at which the tunnel starts. They may need to route the traffic coming through the tunnel in order to maintain network security. The ability to control exactly who has access to what resource on a network is a key characteristic of the level of security that the VPN provides.

Confidentiality

This service is a measure of how certain the user can be that unauthorized personnel cannot see the material. It is provided by a combination of strong encryption algorithms (128 bit RSA) and strong authentication. It is important to understand the level of confidentiality required. Obviously, applications that support private companies bidding on government contracts in countries that are known to be corrupt require a higher level of confidentiality than applications supporting general e-mail.

Protocols used to create a VPN

VPNs are commonly created at Layer 2, the link layer, Layer 3 the network layer, or layer 5, the session layer, of the OSI model[i]. Each of these protocols brings strengths and weaknesses to the VPN solution. The VPN Overview web page [3] provides a detailed comparison between VPNs built from different protocols. I have summarized and added to that information below.

Session Layer

Advantages

Session layer VPNs can provide more detailed control of the flow of data than VPNs based on lower layer protocols. They can work with a variety of security technologies for authentication and encryption. They establish a virtual circuit between the client and host on a session by session basis allowing monitoring and access control that is based on user authentication. They proxy traffic between the source and destination, so all data packets can be transferred through a single port on a firewall. Since session layer VPNs operate above the TCP/IP stack they don't require changes to the network drivers.

Concerns

The main disadvantage is that session layer VPNs proxy all traffic. This means that they are slower than lower layer VPNs. Their more sophisticated access

control is more complicated to set-up manage and maintain than address based access control schemes.

The common layer 5 (Session) protocols are SSL, SOCKs and SSH.

- Secure Sockets Layer (SSL), which is an IETF/W3C standard designed to encrypt http traffic.
  Advantages
  The protocol is available in all modern browsers so special client software to initiate the tunnel to the proxy server is not needed.
  Concerns
  Since it uses the browser to establish the tunnel to the proxy server, it only works with web-based applications or applications that have been adapted to use a web-based interface (web based X11 or ICA).
- SOCKS is an IETF standard designed for authenticated firewall traversal. It is the protocol used in Aventail's VPN solution.
  Concerns
  It requires software on the remote computer to set-up the connection to the proxy server or special versions of software that are SOCKs-aware.
- Secure Shell (SSH2) is a draft IETF standard designed to provide secure remote shell commands. The Secure Shell program, which runs mainly on Unix computers, uses it.

Network Layer

The only Layer 3 (network) protocol used in VPNs is IPSec.

- Secure IP (IPSec) is the IETF standard for encrypting IP packets. It is included in Microsoft Windows 2000.
  Advantages
  IPSec is an extension of IP so there is the promise that it will be widely implemented in the VPN market. It supports a variety of encryption algorithms and checks the integrity of the transmitted data to ensure they have not been tampered with en route. Since it was designed to provide security between firewalls and routers it is an optimal solution for gate to gate VPNs.
  Concerns
  The main disadvantages are a result of the advantages noted above. First, because IPSec ensures the integrity of the data, it won't work from behind a Network Address Translation (NAT) firewall. It assumes that the address change means the integrity has been compromised. Because it was designed to work between gates, it does not have any routing capability built in and its authentication is based in part on the IP address so it identifies the device not the user. Since IPSec is a relatively new protocol and allows a variety of encryption algorithms to be used, most vendor products do not inter-operate. Finally, it does not support protocols other than IP.

Link Layer

Advantages

Link layer VPNs were originally designed to extend Remote Access Services over the Internet. Their number one advantage is that they come as part of Microsoft Windows so it are free and already installed. Link layer protocols can tunnel additional protocols (IPX, NETBEUI, Appletalk, etc.) not just IP. They can provide flow control thus optimizing transmission by cutting down on dropped packets.

Concerns

Link Layer VPNs are targeted at the Microsoft client space. There are very few clients for other operating systems such as Unix, Linux or Macintosh.

The common Link Layer protocols are PPTP and L2TP.

- Point to Point Tunnelling Protocol (PPTP) is an industry standard, designed to tunnel PPP traffic within IP packets. It is included with Microsoft Windows NT4 and 98.
  Concerns
  There has been a lot of discussion on the web, news groups and mailing lists regarding the quality of the encryption and authentication provided by Microsoft's VPN based on PPTP. The VPN FAQ web page [4] gives a balanced coverage of the issue. Some people do not consider the security to be adequate. Once a layer 2 tunnel is set up it is bi-directional as long as the tunnel stays up. No additional authentication is provided. Access control is based upon packet filtering.
- Layer 2 Tunnelling Protocol (L2TP) is an industry standard resulting from the marriage of PPTP and Cisco's Layer 2 Forwarding (L2F) protocol. It was designed to encapsulate IPX, AppleTalk and NetBEUI traffic within IP packets. It is included with Microsoft Windows 2000.
  Advantages
  It comes as part of Windows 2000. Microsoft has done a lot to eliminate the security concerns by supporting additional authentication facilities such as RADIUS and tokens such as one-time codes and smart cards.

Cost

The cost of a VPN product is a function of the characteristics described above. Solutions which require something (software, token, specialized reader, etc) to be distributed to each user are more expensive than those that only require services already available in the client computer.

For the same reason products that use dedicated hardware to terminate the tunnel are more expensive to purchase and maintain than products that simply add software to an existing firewall.

Effect of Windows 2000

Windows 2000 comes with a VPN client built in. RAS Services support PPTP, L2TP and L2TP over IPSec protocols and MS-CHAP, RADIUS, and smart card authentication methods. Most VPN manufacturers have revised, or are revising, their products to support this built-in client. In the future only special applications (e.g. government diplomatic services) will require specialized clients. These will typically be applications with very high authentication requirements.

TransCanada's solution

Strategy

TransCanada's VPN strategy comes in two parts. First, the main VPN will be based on the SSL protocol. Second, all non web-based traffic that passes over public networks will use IPSec.

Reasons for Choice

TransCanada wanted the VPN to run on stand-alone hardware, either a proprietary device or a Sun Unix server (company standard for servers). TransCanada did not want to add additional intense processing to a device whose original purpose was to control network access. Remember, the firewall is an intentional bottleneck on the network. This requirement eliminated all of the firewall vendors from the list.

TransCanada has chosen web-based delivery as its application display mechanism of choice. Over the next few years (target completion date - 2005), TransCanada will be converting all of its applications to a web-based user interface. By providing a VPN based on HTTPS, all applications will be available to authorized users no matter where or how they access the application. Information Technology won't have to be concerned about what client device is being used. In the spring of 1999, the only VPN based on HTTPS shipping was Sun's iPlanet, now known as iPlanet's Portal Server.

Internally, all traffic on TransCanada's networks will be IP packets. Therefore, an IPSec VPN will be used to extend TransCanada's networks through the Internet. However the requirement for IPSec compliance in a particular product does not by itself provide a means for selecting a particular VPN vendor solution. Remember the earlier statistic that more than 50+ vendors at the 1999 InterOp+Networld sported IPSec systems.

TransCanada identified a second requirement of relatively strong user authentication, which the company defined as "something you know, something you have and if you lose the thing you have, you know you've lost it." This requirement meant that either tokens or biometrics would be required to identify the person requesting a VPN tunnel be set up. Using biometrics and smart cards was deemed to be too costly, since special hardware had to be attached to each client. An X.509 certificate stored on the hard drive did not meet the requirement of knowing when it was lost, since the certificate could be copied without the owner's knowledge. This requirement did little to narrow the list as all of the IPSec VPNs at InterOp+Networld provided various mechanisms for performing user authentication. Fortunately, iPlanet's Portal server supports multiple tokens.

Both IPSec and HTTPS require Public Key Encryption to work. TransCanada had already implemented an Entrust Public Key Infrastructure (PKI), so another requirement was that the VPN had to work with the Entrust PKI. The certificates are used to identify the device that terminates the tunnel at TransCanada's network and to identify remote devices in situations where two networks are joined by a VPN, such as branch offices. In the spring of 1999 the only shipping VPN which integrated with the Entrust Certificate Authority (CA) was Timestep. IPlanet can use X.509v3 certificates. Entrust claims its CA will be X.509v3 compliant with its next version.

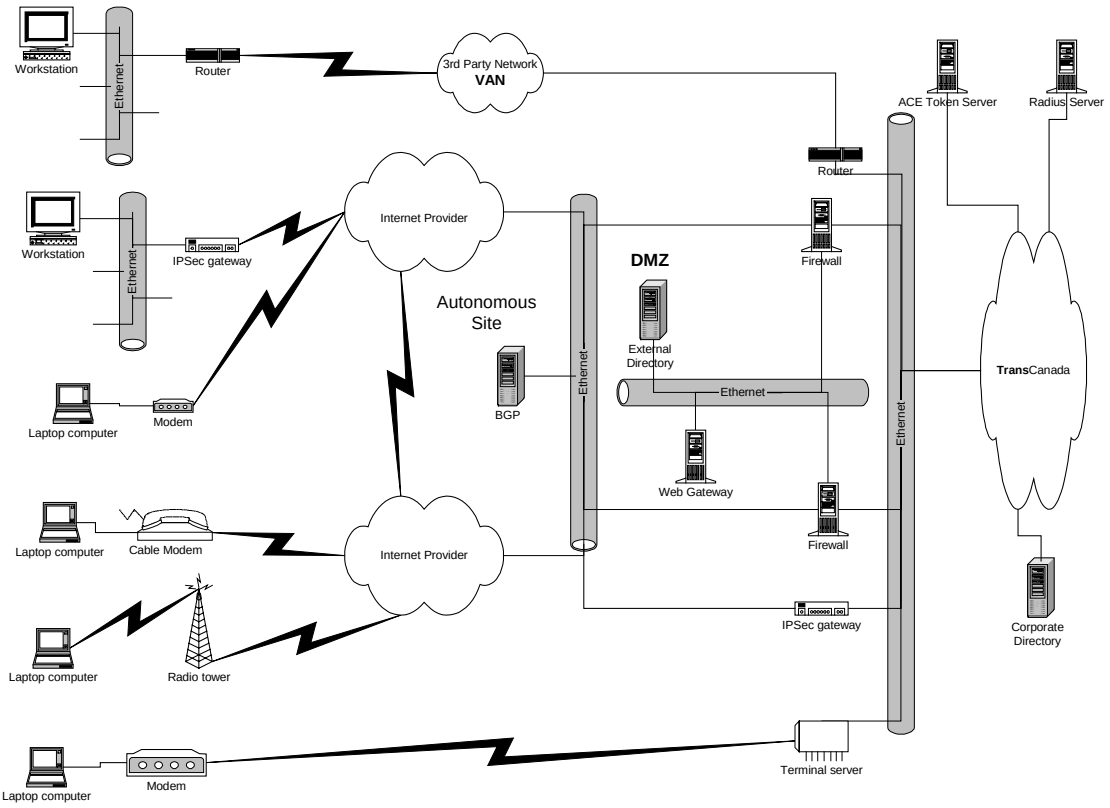

Figure 2 TransCanada External Access Architecture

Implementation

Figure 2 is a model of TransCanada's External Access Architecture. The Timestep VPN was put into production in April of 2000 after testing for six months. It uses the Corporate Directory (a Netscape Enterprise Directory) to store Entrust Public Encryption Keys. Two Shiva Access Switches provide RAS services. The switches use RADIUS to authenticate users against the Corporate Directory. This service will not be going away any time soon as our 1-800 service is cheaper than a low speed ISP connection. Finally, the iPlanet web proxy service is currently in pilot mode. It will be placed into production by the end of the year.

Thus by the end of the year TransCanada will have a general purpose web-based VPN service that can be used with any browser, on any type of computer, and from any ISP connection. TransCanada will also have a special purpose IPSec VPN. This VPN will be used to selectively replace our existing dedicated lines in both our Intranet and our Extranet.

The replacement will be done on an as-needed basis meaning that if a change to an existing dedicated line is needed for business purposes (office moves, significant changes to bandwidth requirements, etc.) an analysis will be performed to see if the VPN can provide a more cost effective solution. The IPSec VPN will also allow TransCanada to extend its network to individuals' homes to support teleworkers at LAN speeds. Finally, there will be a low speed direct dial connection for those instances where Internet access is not available.

Appendix

Requests for Comment

Topic	RFC Number
Public Key Encryption	2407, 2408, 2409, 2459, 2510, 2511, 2527, 2528, 2559, 2585, 2587, 2560, 2692, 2693
IPSec	1828, 1829, 2104, 2085, 2401, 2410, 2411, 2402, 2412, 2451, 2403, 2404, 2405, 2406
L2TP	2661
SOCKS	1928, 1929, 1961
SSH2	In Draft
SSL	2246, 2712, 2716, 2817, 2818
S/Key	2243, 2289, 2444
CHAP	1994, 1472

Bibliography

1. Kosiur, Dave, *Building and Managing Virtual Private Networks*, Wiley Computer Publishing, 1998.
2. Ball, Larry L., *Cost-Efficient Network Management*, McGraw-Hill, 1992.
3. http://www.maximalsolutions.com/current/VPNs/VPNOverview.html, Virtual Private Networks - An Overview.
4. http://kubarb.phsx.ukans.edu/~tbird/vpn.html, Virtual Private Networks Frequently Asked Questions.
5. http://www.epm.ornl.gov/~dunigan/vpn.html, Tom Dunigan's Virtual Private Networks page.
6. http://www.mnmteam.informatic.uni-muenchen.de/projects/vlan/vpn.shtml, Virtual Private Networks.

---

[i] The International Standards Organisation (ISO) developed an Open Systems Interconnect (OSI) model of network communications in the 1980s. The model segmented the communications process into 7 functional layers for which standards of operation were written[2]