SECURITY ISSUES FOR REMOTE ACCESS AND VIRTUAL PRIVATE NETWORKS INVOLVING CABLE MODEMS

Daniel Howard
Digital Furnace Corporation
Atlanta, GA

Abstract

The use of cable modems in the small office/home office (SOHO) market may appear to create special considerations when setting up remote access for employees who are telecommuting and virtual private networks (VPNs) for branch offices or small businesses with distributed employees. Shared access of the medium, multiple users and applications of home PC's, inability of telecommuters to properly administer their home networks, and the potential for multiple cable modems in a premise are all issues which impact the threats, policies, architectures, and solutions for secure networking using cable modems. In this paper, the key security issues, sample architectures for VPNs involving cable modems, cable modem security mechanisms (such as Baseline Privacy Plus), and approaches for providing secure remote access and VPNs involving cable modems will be discussed. While the suspected challenges have to do with privacy over the cable network, it will be shown that with the advent of Baseline Privacy in DOCSIS cable modems, the most likely problems have to do with developing and implementing the security policy as it relates to home networks.

Introduction

Cable modems are gaining in numbers and the delays in the deployment of alternate technologies such as xDSL will create a tremendous revenue opportunity for using cable modem networks to address the small office/home office (SOHO) market. Unfortunately, one of cable's perceived drawbacks is data integrity and security over cable modem networks. Further, since cable has the capability to address previously untapped segments of the SOHO market, the types of users, their skills, bandwidth and security needs will vary so widely that much of the challenge will be in marketing to the various customers.

It is the duty of technology developers is to ensure that technology solutions exist and are cost effective for providing whatever capabilities are required and enumerated by marketing personnel. The developers of the DOCSIS standard and cable modem vendors have embraced this philosophy and have provided an arsenal of technologies that can be used to guarantee that cable modem networks have security features that meet or exceed those of alternate technologies. But understanding how each of the available technologies relates to the customer's needs and cost constraints is a new problem for cable modem service providers.

The purpose of this paper is to explore the security technologies and issues associated with using cable modems for the SOHO market. In particular, the available security technologies will be applied to the following two applications: remote access to corporate networks and virtual private networks (VPNs). It will be shown that the capability to securely employ cable modems is readily available, but that the presence of the home in the network creates unique issues that are best addressed via development of proper security policies.

Introduction to Network Security and Security Technologies

The basic security concepts that impact cable modems for the SOHO market include security types, needs, threats, policies, and architectures. Generally, the security threats are used to develop a security policy for the individual or corporation, and the policy dictates the security architecture and the technologies, including both hardware and software, that are required to implement the policy. The policy also includes specification of the types of security to be provided, including privacy, authentication, data integrity (including protection against virus attack), accountability, and robustness. Security robustness defines conditions for non-repudiation of service and the ability to detect and characterize security attacks. Security technologies employed to implement security policy include authorization, authentication, certification, and public and private key infrastructures. Finally, the security architecture includes the functional architecture implementing security, as well as component networks and software to be addressed. For example, if remote access via cable modems is permitted as part of the general security architecture, then it must consider access network security as well as end to end security.

First, consider the security threats. A beauty shop using a cable modem for high speed internet access so that customers can access images of hairstyles may have no real threats to data security if the cable modem client computer is not connected to the computer used for record keeping, billing, and so on. At the other extreme, a large corporation in a highly competitive, international field with government contracts may have threats that include independent hackers, political extremists and groups thereof, industrial espionage and virus attacks, foreign government organizations, and former or current (but disgruntled) employees. It is interesting to note that approximately 70% of corporate security breaches have been from inside the network rather than from outside.

The threats and associated risks and consequences are used to define a security policy which enumerates the security needs of the corporation, the techniques and technologies used to meet those needs, and the process by which the security procedures are monitored, managed, and updated. Further, security policies determine who gets a given level of access, how passwords are distributed and updated, and how new and terminating employees are handled by the security system. The aspects of the security policy of greatest interest for cable modem networks are those which define the forms of remote access permitted, the types of authentication and authorization that are used, and how digital certificates, certification authorities, and public/private key infrastructures are used. Policies for traveling employees will also be relevant to the present discussion if the traveler is accessing his home or branch system via cable modem.

While ideally, cost and convenience of security policies would not enter into the policy specification, it is nonetheless a fact of life that security technologies cost both time and money to implement, and often reduce the ease of use and performance of applications that are running in secured modes. In order to make tradeoffs in cost, convenience, and security achieved, the actual technologies involved must be considered, since computing performance can depend heavily on which layer of the OSI model has the security implementation, as well as how it is implemented. For example, data link and physical layer security measures for filtering the packets attempting to pass into inside the trusted network from outside the firewall will generally produce higher speed and performance in client applications, especially if the security can be implemented in hardware. But hardware implementations can often be limited in flexibility and are not easily upgraded when new technologies are available.

Hence, in order to understand the tradeoffs in cost and performance versus security provided, we must examine the specific technical functions provided by security technologies. We begin with authentication. Authentication in its most generic sense is a means of identifying individuals and verifying their eligibility to receive specific categories of information. In data networks, authentication is the act of identifying or verifying the eligibility of a station (e.g., a cable modem), originator, or individual to access information. Authentication establishes the validity of a transmission, message, station, or originator, and a provides positive identification with a degree of certainty sufficient for permitting certain rights or privileges to the person or station positively identified.

One type of authentication results in access to network resources, and usually starts with user ID's and passwords, but can ultimately result in a chain of challenges and replies involving exchanges of encrypted data and digital signatures. The other type of authentication is source authentication, which usually involves imbedding digital signatures in documents, files, email, and other data to ensure that the items transmitted were from the stated source and have not been altered. Each type of authentication has a process associated with it that varies in complexity depending on the level of protection desired.

For access to network resources, simple user ID and password authentication, such as Telnet and FTP.TCP/IP for TCP/IP networks, have many limitations and vulnerabilities, the majority of which relate to the ability of hackers on the Internet to easily receive unencoded packets with such information. More advanced methods such as Kerberos, cookie implantation during initial registration, public key cryptography systems, the DOCSIS registration procedure for cable modems and Baseline Privacy Plus for cable modems use some form of encryption and are thus more secure. In another technique used in authentication for remote access services, the client computer hangs up after initial logon via dialup connection and is called back at a prearranged number stored in an authentication database. Hardware devices (or authentication equipment) can also be used, where secret algorithms and keys are used in the device to convert a user input into a response that is recognizable only by another machine with the same hardware device. Other implementations of network authentication include cryptographic authentication, time based authentication (access can only be granted at specific times or for specific time durations), peer-entity authentication, self-authentication, and smartcard and/or token authentication.

The heart of all advanced systems used for authentication is encryption, where the methods and algorithms used for encryption vary, as does the way in which cryptographic keys are obtained and how they are used. For example, one of the main differences between techniques such as Kerberos and public key cryptography systems is that the former relies on a trusted third party, the Kerberos server, which if compromised, opens up the entire system to attack. Key based cryptography systems, on the other hand, rely on stored keys which if compromised, only betray the user who's private key is discovered. There are two main varieties of key based encryption: secret key (or symmetric key) systems, and public key systems. In secret-key cryptography, the same key is used for both encryption and decryption. An example of this is the Data Encryption Standard or DES system. Often this system is implemented in hardware to speed up performance.

Key pair systems, usually called public key encryption, require the use of two keys, a public key for encryption and a private key for decryption. A private key can also be used to sign items for the purpose of source authentication, with the public key being used for said authentication. An individual can also encrypt their own items with their public key so that only they can open them.

The so-called Pretty Good Privacy or PGP system uses a public-key system which employs IDEA encryption and RSA encryption. The Digital Signature Algorithm (DSA) is another popular public-key technique used for signatures. There are also cryptosystems based on elliptic curves, and a key agreement protocol called Diffie-Hellman for establishing secret keys over an insecure channel.

Certificates are often used in public key systems to verify that a person and a public key are correctly associated. The most important information in a digital certificate is a public key and a name. Often a certificate also contains an expiration date, the name of the certifying authority, a serial number, and other information. It can also contain the digital signature of the certificate issuer. During authentication, a chain of certificates can be created, each one certifying the previous one until the parties involved are confident in the identity of each other.

Once the user is identified and authenticated, the network authorizes the user and grants to a user, program, or process the right of access to the network resources requested. The authorization provided can depend on the level of access requested by the user, the user type or category, time of day, network loading at the time of authorization, availability of human operators to validate the authorization, and whether or not network intrusion has been recently detected.

All of the preceding is of course specified in the security policy, in which is also specified the security architecture. A proper security architecture should not rely on any particular encryption scheme, but rather be able to insert new encryption schemes as they become available. The reason is that hackers are constantly trying to find ways to break known schemes, and when a scheme is cracked, the security policy should dictate that a new encryption scheme is immediately substituted.

In general, the security architecture specifies the components of the system, the interfaces between components and with the outside world, and desired scenarios and permitted layouts for interconnection of components. The location of trusted and untrusted networks/entities must be defined in the architecture, as well as whether access to the trusted networks/entities is via a single point or via multiple locations. In the case of large corporate networks, it is quite common to have several access points to the network, each of which is protected by a firewall. Small offices, on the other hand, may only have a single access point to the outside world, and the premise of this paper is that this access point will increasingly be provided by a cable modem.

Likewise for homes, although it must be recognized that both single and multiple access architectures are possible. Many homes currently have multiple PC's with analog phone modems. Homes with digital video and cable modems will likely have at least two cable modems: one will be in the home office cable modem and another in the digital set top box. It is nontheless desirable to have a single access point in the customer premise, since this provides the maximum control over the security in the premise. For example, a security architecture developed for homes [Calvert1] uses a single gateway access point and contains the following elements: devices, terminals, authentication mechanisms, identity validator, gateway, enforcement engine, introduction mechanism, policy database, secure channel, participant, and demarcation point or threshold (where the inside, trusted home network meets the untrusted outside network). The architecture specifies how the elements are used and interconnected to provide security for a home network with broadband access. Many of these same elements apply to the case of a SOHO premise in which workers desire secure access to and from the premise. But to characterize how security plays a role in such applications, we must first explore the network and security technologies applicable to SOHO networking involving cable modems.

Network and Security Technologies for Cable Modem Networks and VPNs

First, let us define some network architectures for SOHO applications using cable modems. The two key applications are remote access to a corporate network and a virtual private network or VPN. A VPN is a network that connects remote offices or employees to a private corporate network through a commercial internet service provider, or ISP, instead of through the more traditional private network. For the purpose of this paper, we assume that the commercial ISP is a high speed cable access provider, such as the At Home corporation or ServiceCo. VPNs can involve services such as remote access, data, fax, voice over IP (VoIP), and video conferencing.

The figure on the next page shows how cable modems can be used in remote access and VPN applications. In the most general case, the private network to which the VPN is connected can run IP or non IP packet traffic such as IPX, AppleTalk, SNA or DECnet. The most common method by which the VPN is connected to the corporate or private network through a non secure IP network is IP tunneling, where packets on the virtual network are encapsulated in IP packets for transport over the public network.

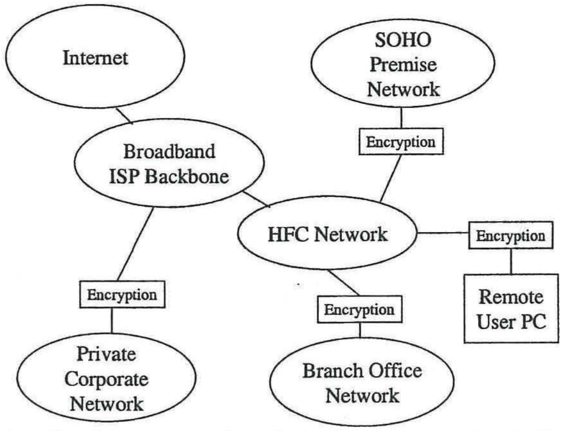

Usually the encapsulated packets are encrypted using techniques described previously. Examples of VPN tunneling approaches [Cabletron 1] include:

1. IP tunnels between a remote user and a corporate firewall with tunnel creation and deletion controlled by the user's computer and the firewall
2. IP tunnels between an Internet service provider and a corporate firewall with tunnel creation and deletion controlled by the ISP
3. IP tunnels between sites over the public Internet
4. IP tunnels over a service provider's backbone IP network that is separate from the public Internet

VPNs based on IP tunnels can be deployed and managed by either the corporate user (especially if there are no quality of service requirements), or they can be deployed and managed by the cable ISP. The latter is usually done under a service level agreement that ensures a minimum quality of service to support real time applications such as voice over IP (VoIP) and video conferencing. For option 4 above, the strength of cable ISPs is that they can offer shared bandwidth with complete control over both the IP backbone and the cable access network.

There are several different tunnel protocols available [Cabletron2], including Level 2 protocols such as the Point to Point Tunnel Protocol (PPTP), the Level 2 Tunnel Protocol (L2TP), and a level 3 tunneling protocol, IPSec, which is an effort by the IETF to add standards-based authentication and encryption to TCP/IP. IPSec is an evolving set of specifications for cryptographically-based authentication, integrity, and confidentiality services at the IP datagram layer. IPSec is specified so that many methods of encryption and key management can be supported. Although IPSec is the trend for VPN systems, the standard is still evolving and all vendor equipment employing IPSec is not yet interoperable.

Authentication in VPNs is often via estlablished methods such as Radius (Remote Authentication Dial-In User Server, of the Terminal Access Controller Access System (TACACS), although there are also proprietary methods used for authentication. A firewall based VPN can also be used, where the firewall performs address translation, user authentication and user authorization, in addition to the well known packet filtering function. While there are often significant performance degradations associated with firewalls (lack of support for some services, slow throughput due to using a host based operating system), firewall-based VPNs are probably fine for limited numbers of simultaneous remote access users or relatively small amounts of traffic passing site to site over the network.

There are also software based VPNs such as those designed for use on NT server which are flexible enough for support of both VPN traffic as well as non VPN traffic (e.g., web surfing). The faster encryption performance of hardware systems coupled with the greater protection often afforded by hardware based encryption has also led to hardware-based VPNs. These use hardware encryption modules in between the CPE computers and the access point (cable modem) and which can be built into routers or combined with software systems such as firewalls and authentication servers to provide an integrated solution for the SOHO market.

VPNs with cable modems can also be set up entirely over the cable modem network when all offices, branches, and homes involved have access to cable. In this case, IP tunneling may not be required since the cable network will soon support DOCSIS security measures such as the Baseline Privacy Interface (BPI), a method of encrypting data for transport between the cable modem (CM) and the cable modem termination system (CMTS). In BPI, data is encrypted between the CM and CMTS at the MAC sublayer (which is based on the DOCSIS standard for cable modems). While BPI used only the 6 byte MAC address to authenticate the cable modem to the CMTS (and thus has limited protection against theft of service), BPI Plus addresses that vulnerability by adding digital certificate-based CM authentication to its key exchange protocol.

In BPI, the service ID (SID) is used to identify a security association in the CM. During initialization, the CM requests an authorization key via transmission of the CMs MAC address, the CMs public key, and a list of unicast SIDs corresponding to provisioned class-of-service settings configured for BPI [Judge1] [BPI+]. BPI+ uses a variety of encryption techniques for various stages of the authentication, authorization, and encryption of actual data traffic, including RSA for encryption of authorization keys (using the CMs public key that, along with its RSA private key, is factory installed in the CM). The US Data Encryption Standard (DES) is used for for traffic encryption and for traffic key encryption, although different modes are used for each. Finally, note that while the entire packet data unit (PDU) of the cable modem packet can be encrypted, the MAC header and portions of the extended header are not encrypted.

The result is that DOCSIS cable modems now have the ability to support privacy over the access network that is equal or better than the privacy afforded by the public switched telephone network or leased lines, even though the RF downstream signals of cable modems are seen by other stations. Encryption in both the access network as well as in the IP datagrams provides a level of security that is only limited by the strength of the underlying encryption algorithms.

Security Issues for Remote Access and VPNs Over Cable Modems

A widely publicized "flaw" in the At Home network's cable modem service was the folder sharing capability of Microsoft Windows which permitted other cable modem customers to see files on an individual's internal network if he had folder sharing turned on [Pelline1]. The At Home network administrators had turned this feature off of all systems connected to the cable modem, but users turned it back on in order to set up home networks. A quick solution was found by launching a software upgrade which disabled external file sharing but not internal network file sharing.

This security event points out what is felt to be the most serious challenge to home based cable access for remote access to corporate networks and VPNs: the home environment itself. For example, there is a security issue with home networks that have multiple users on the a single system connected to the cable modem who are not employees of the company. Another issue is the use of home and office networks that employ wireless LAN systems in addition to wired LAN systems. The security of wireless LANs is generally not as extensive as that of the cable modem access network using Baseline Privacy due to the latency and performance degradations which attend more secure techniques. Hence, the use of wireless networks in the customer premise for either data or voice should be seriously considered in the security policy, since it can become the Achilles' heel of the network.

Multiple users on a home based system creates potential problems for employees who typically post their password and user ID near the computer, making it possible for others in the house to gain access to the VPN. The most prudent remote access approach would be for corporate users to use their office laptop computers at home when connecting and state in the security policy that no other home users can use the company laptop. Failing this, the security policy should dictate strict rules for how the home user accesses the corporate network and how he or she should prevent unauthorized access within their home. Over time, there will likely be configuration-oriented methods to prevent other home users from accessing corporate resources, such as logon procedures which rely on smart cards carried by the employee, voice identification, and so on. Many of these tactics are in their infancy now but are expected to be refined over time as more experience with home based high speed access is gained. It should be noted that the 'always-on' characteristic of cable modems creates unique access control problems in that home systems are generally more open to casual visitors than office based systems. Again, it is felt that as the desired capabilities are defined in detail for small office/home office systems used for remote access or VPNs, proper security policies can be developed which address the security needs and provide the required level of protection.

Conclusions

The most important conclusion is that technology exists to make remote access and VPNs using cable modems as secure as necessary for most corporate applications. It is really the behavior of humans in their homes that creates the challenges for cable modems, and indeed any broadband access technology which addresses the SOHO market.

MSOs or cable ISPs interested in addressing the SOHO market will likely have two main types of users: those seeking to set up their own VPN and remote access, and security facilities, and those who wish to rely on the cable ISP to set up, monitor, and maintain the system. For the first type, Cable ISPs should nonetheless work with the customer to assist them in developing a security policy so that when flaws in the policy lead to security breaches, the onus is on the corporate user to update their systems. For the second type, cable ISPs should develop an arsenal of solutions to offer with tradeoffs between security, cost, and performance and services provided. In this manner, the cable ISP can address the many types of remote access and VPN customers for cable modems which are likely to exist.

Further, since cable ISPs rely more heavily on the underlying encryption technologies for security over the access network, they should aggressively seek upgrades of security technologies that address known flaws or enhance the network performance of shared access security systems.

Finally, whatever solution is used, the cable ISP should include monitoring and detection of security events as key to the long-term process of maintaining the integrity of the overall cable modem network. This will assist in developing the perception that cable modem networks are as secure as alternative technologies that are similarly priced. Since VPNs in particular will likely be deployed using a variety of technologies such as xDSL, wireless, and conventional telco solutions, the cable ISP must be prepared to use a thorough understanding of the technology to market it effectively to SOHO customers.

References

[Calvert1] Calvert, K. and Aylesworth, D. "Toward a security architecture for home information infrastructure, Broadband Telecommunications Center technical report, summer 1998.

[Cabletron1] "What is a virtual private network (VPN)?" Cabletron white paper, http://www.cabletron.com/vpn/VPNwhatis.htm.

[Cabletron2] "A ten point plan for building a VPN" Cabletron white paper, http://www.cabletron.com/vpn/VPN10ptplan.htm.

[Judge1] P. Judge, "A review of the baseline privacy interface of the data over cable system interface specification (DOCSIS)" final report CS7100, http://www.cc.gatech.edu/~judge/cs7100/project03/cable.htm

[BPI+] "Baseline Privacy Interface Specification, Draft" SP-BPI-D01-990212, 2/12/99.
[Pelline1] "@Home user files may be at risk", J. Pelline, http://www.news.com/news/Item/0,4,15924,00.html.