Enterprise Wireless WAN Security

Enterprise Wireless WAN Security

Wide area wireless networks can be an enormous benefit to corporation because they have the potential to extend the reach of an enterprise application to a staggering proportion of the earth's surface. However this expanded range also increases the vulnerability of the company's devices, applications and data. In order to ensure their viability we must validate the security of this new infrastructure.

Today's legacy and emerging Wireless Wide-Area Networks, such as GSM, GPRS,EDGE, UMTS and cdma2000 already include security provisions that are enforced by the mobile terminals and the base stations. However, there are still shortcomings in the security model that can only be addressed with an end-to-end approach. This White Paper will explore the options an enterprise has at its disposal for securing remote connectivity over wireless WANs.

I would like to state up front that a great deal of this applies equally well to wired network access. My title is not meant to imply that the considerations and protocols discussed below are necessarily distinct from those of a non-wireless remote access solution. My focus in this Knowledge Brief is simply reserved to the perspective of wireless networks.

Trusting the Mobile Operator

In addition to the challenge of securing data as it rides the airwaves we also need to consider the vulnerabilities of the network between the base station and the application server. One approach could be to come to an arrangement with the mobile operator to secure this channel. This raises an important issue that is pervasive throughout any discussion of security: trust.

To what extent can an enterprise trust a carrier to provide a secure connection? Will they systematically abuse the data, filter sensitive information and sell it to competitors? Probably not. But, can you be sure that all the telecom operator's employees are trustworthy? Do you know for certain that they have adequate protection vis-à-vis hackers and industrial spies? These are legitimate questions. There may be no particular reason to distrust the mobile operator but corporations with strict security policies may be averse toward outsourcing these processes to an external entity whose operation is not completely transparent.

Consequently, we need to find connectivity options that do not rely on the security provisions of the mobile operator. A secure air interface is a nice benefit but only a small piece of the puzzle of remote connectivity.

Overview of WWAN connectivity

Fundamentally there are two different means a mobile network may offer to transfer data. It can provide a packet-data network or else it can use circuit-switched connections. A packet data network is simpler. CDPD, Mobitex and GPRS would all be examples of packet data networks. In these cases, the mobile device has an IP address and it transfers data through the mobile network, which is connected to the Internet. No special configuration is typically required at the mobile end. Its data access is transparent. If the IP address given to the device is fixed then a minimal amount of authentication is also implicit in any packets originating from it.

Data communication over primarily voice networks, such as GSM, IS-136 and IS-95, is not quite as straightforward. Typically a Point-to-Point Protocol (PPP) connection must first be

made from the device to a dial-in server. The dial-in server will assign an IP address and relay all the traffic between the device and any application servers.
This implies some configuration at the mobile end. The user must specify a phone number and then authenticate to the dial-in server using an authentication protocol such as PAP, CHAP or MS-CHAP. So the dial-in server knows who the user is but the application server does not. It cannot determine the phone number easily and the IP address is meaningless. If necessary it would then re-authenticate the user, which means additional work for the user.

It would be possible to bypass the first authentication by storing the mobile phone number on the dial-in server and then comparing the caller-id of incoming calls. However, this would provide unlimited access to the corporate network when a device was lost or stolen. Solutions to address this dilemma must combine security with ease of use, for example by using biometric authentication. (In 1999 Siemens already showed prototypes of a mobile phone that incorporated a fingertip biometric sensor. Although not available in production at this time, such combinations of technology are clearly possible, and offer considerable advantages in the battle against theft and fraud.) They must also ensure (for example by encrypting the file system) that unauthenticated users cannot access any information on the device. It is then feasible to cache some of the network credentials on the device. Nonetheless, some authentication to the network should always be based on an action or token that is separate from the device.

Circuit-switched data

The practical steps required to set up a secure WWAN connection depend on which of the two categories of network we are considering.

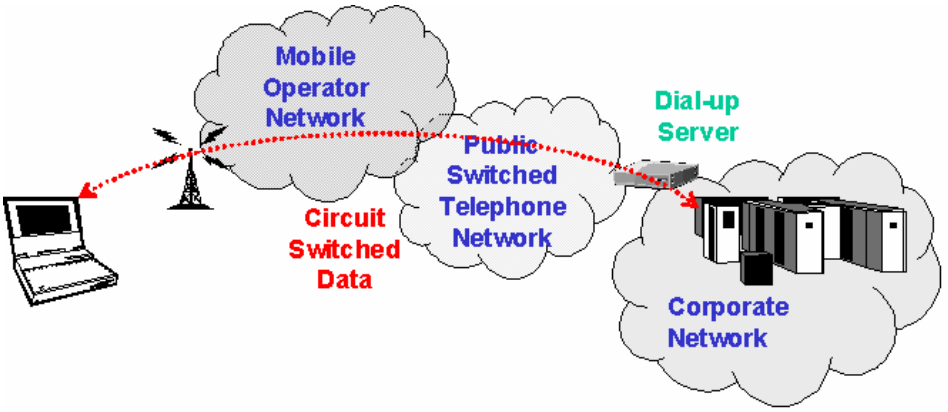

Figure 1: Dial-up to private network

Figure 1illustrates a typically circuit switched connection, whereby the user connects via the mobile operator's network into the Public Switched Telephone Network. The path is relayed on to the private dial-up server based on the phone number that was dialled. In this case we need to use a protocol that will encapsulate all the traffic between the mobile client and the corporate dial-up server. In practice, this means we need to use Serial Line Internet Protocol (SLIP) or PPP. Both are communication protocols for serial data transmission between two devices. They allow a computer connected to a server via a serial line (with a modem) to gain access to the Internet.

SLIP is not an Internet standard but is described in RFC1055. It is a simple framing scheme for putting IP packets on a serial line. SLIP's main advantage is its simplicity and consequently its implementation. Its drawback is ease of use. With SLIP, you have to know your own fixed IP address and that of the remote system you are dialing into. If IP addresses are

dynamically assigned by your service provider, your SLIP software needs to be able to pick up the IP assignments automatically or else you have to setup them up manually. You may also need to configure such details as MTU (maximum transmission unit), MRU (maximum receive unit), and the use of VJ compression headers, etc.

PPP is an Internet standard described in RFC1171. Unlike SLIP (which can only transport TCP/IP traffic), PPP is a multi-protocol transport mechanism that can accommodate various network protocols, like IP, IPX and Appletalk simultaneously. It does essentially the same thing SLIP does, but with a more complete set of features like error detection in every frame, IP address negotiation, automatic compression, login and connection configuration.

Most importantly, PPP is now supported and favoured on almost all dial-up products through the industry. It can offer a relatively secure connection without any Internet exposure and supports multiple encryptions algorithms such as RC4. PPP permits several authentication protocols including PAP, CHAP, and MS-CHAP. These are not confined to PPP, or even to CSD connections, so I have reserved the next section for them.

Authentication protocols

Authentication is at the foundation of any security scheme particularly when it involves remote access. The primary objective for an enterprise is to ensure that only legitimate users may access the resources and data on its network.

PAP - Password Authentication Protocol (RFC 1334)

PAP is the least sophisticated authentication protocol. It uses a simple, clear text authentication scheme. The authenticator requests the user's name and password, and PAP returns them in clear text (unencrypted). This authentication scheme is not secure because a third party could capture the user's name and password and use it to get subsequent access to the authenticator and all of the resources provided by the authenticator. PAP provides no protection against replay attacks or remote client impersonation once the user's password is compromised.

Because PAP uses clear-text passwords, you would use PAP in only two circumstances: when you're dialing in to a Point-to-Point Protocol (PPP) server that does not support encrypted authentication and when you're dialing into a Serial Line IP (SLIP) server. (SLIP servers understand only clear-text passwords.) Simply stated, you use PAP only when the client and server cannot negotiate a more secure form of authentication.

SPAP- Shiva Password Authentication Protocol

SPAP is Shiva's proprietary version of PAP. SPAP is more secure than PAP because SPAP uses a two-way (reversible) authentication method that encrypts passwords. Thus, SPAP offers a medium level of security for remote access. However the proprietary nature of the protocol has hindered widespread adoption. It is mainly of historical interest for legacy implementations.

CHAP - Challenge-Handshake Authentication Protocol (RFC 1994)

CHAP provides a higher level of security for remote access than PAP. CHAP is an encrypted authentication mechanism that avoids transmission of the actual password on the connection. The authenticator sends a challenge to the remote client, consisting of a session ID and an arbitrary challenge string. The client then uses the MD5 one-way hashing algorithm to return the user's name and an encryption of the challenge, session ID, and the client's password.

CHAP is an improvement over PAP because the password is not sent over the link in the clear. Instead, the password is used to create an encrypted hash from the original challenge. The server knows the client's clear text password, and can replicate the operation and subsequently compare the result to the password sent in the client's response. CHAP protects against replay attacks by using an arbitrary challenge string for each authentication attempt. Furthermore, it protects against remote client impersonation by unpredictably sending repeated challenges to the remote client throughout the duration of the connection.

CHAP uses a three-way handshake to provide encrypted authentication. The authenticator first sends out a challenge to the client. The client responds with a one-way encrypted value. The authenticator checks to see whether the value matches. If it does, the authenticator acknowledges the authentication. CHAP then periodically verifies the client's identity. It changes the challenge value every time it sends out a message, which protects against playback attacks (i.e., a hacker records the exchange and plays back the message to obtain fraudulent access).

MS-CHAPv1 - Microsoft Challenge-Handshake Authentication Protocol (RFC 2433)

MS-CHAP is the Microsoft version of CHAP, using Microsoft's version of RSA Data Security's MD4 standard. MS-CHAP uses a one-way hash function to produce a message-digest algorithm. A hash function takes a variable-size input and returns a fixed-size 128-bit string. This type of algorithm produces a secure checksum for each message, making it almost impossible to change the message if you don't know the checksum.
MS-CHAP differs from the standard CHAP as follows:

- MS-CHAP does not require the authenticator to store a clear or reversibly encrypted password.
- MS-CHAP provides an authenticator-controlled authentication retry mechanism.
- MS-CHAP provides an authenticator-controlled change password mechanism.

Microsoft CHAP is an encrypted authentication mechanism very similar to CHAP. As in CHAP, the authenticator sends a challenge, which consists of a session ID and an arbitrary challenge string, to the remote client. The remote client must return the user name and an MD4 hash of the challenge string, the session ID, and the MD4-hashed password. This design, which manipulates a hash of the MD4 hash of the password, provides an additional level of security because it allows the server to store hashed passwords instead of clear-text passwords. Microsoft CHAP also provides additional error codes, including a “password expired” code, and additional encrypted client-server messages that permit users to change their passwords. In Microsoft's implementation of Microsoft CHAP, both the Client and authenticator independently generate an initial key for subsequent data encryption by MPPE (Microsoft's Point to Point Encryption).

MS-CHAPv2 - Microsoft Challenge-Handshake Authentication Protocol Version 2.0 (RFC 2759)

Microsoft CHAP 2.0 offers improved security features over V1. These improvements include a server authentication scheme and a single change password packet. The most significant changes from MS-CHAPv1 to MS-CHAPv2 are:

- The weaker LAN Manager hash is no longer sent along with the stronger Windows NT hash. This thwarts automatic password crackers like L0phtcrack which first breaking the weaker LAN Manager hash and then use the information to break the stronger NT hash.
- An authentication scheme for the server has been introduced. This prevents malicious servers from impersonating legitimate servers.

- The change password packets from MS-CHAPv1 have been replaced by a single change password packet in MS-CHAPv2. This addresses the active attack of spoofing MS-CHAP failure packets.
- MPPE uses unique keys in each direction. This is to prevent XORing the text stream in each direction to remove the effects of the encryption.

EAP - Extensible Authentication Protocol (RFC 2284)

The Extensible Authentication Protocol (EAP) is a PPP extension that provides support for additional authentication methods within PPP. Transport Level Security (TLS) provides for mutual authentication, integrity-protected negotiation, and key exchange between two endpoints.

EAP does not select a specific authentication mechanism at Link Control Phase, but rather postpones this until the Authentication Phase. This allows the authenticator to request more information before determining the specific authentication mechanism. This also permits the use of a "back-end" server which actually implements the various mechanisms while the PPP authenticator merely passes through the authentication exchange.
The authenticator does not necessarily have to understand each request type and may be able to simply act as a pass-through agent for a "back-end" (e.g. RADIUS) server on another host. The device only need look for the success/failure code to terminate the authentication phase.

Packet Data Networks

Packet Data Networks are substantially different from Circuit Data Connections since each packet is routed separately and therefore should be authenticated and encrypted individually. Since this would complicate the life of the network layer in an excruciating way it makes sense to consolidate the authentication and encryption into a virtual connection using a tunnel, also called a virtual private network.

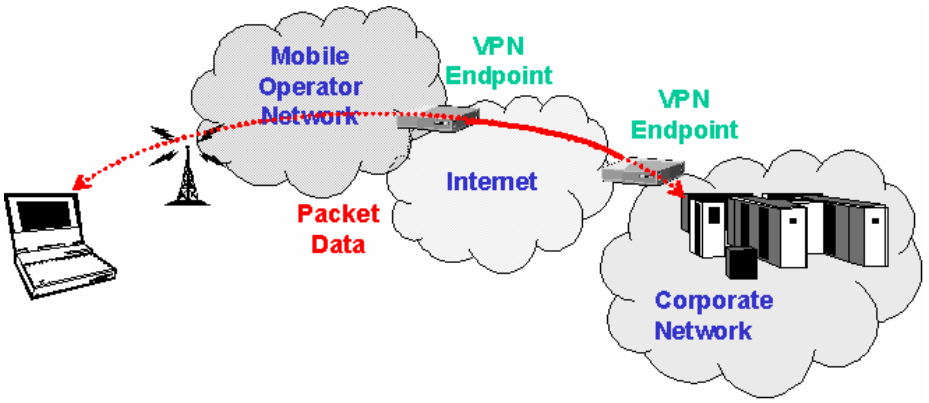

Figure 2: Dedicated VPN to Mobile Operator

There are at least two ways in which we could approach VPNs in a mobile scenario. The first would be to create a dedicated tunnel between the mobile operator and the corporate network. All traffic from the client is intercepted by the VPN on the mobile operator's network and tunneled over the public Internet into the VPN server on the corporate network and likewise in reverse corporate data passes through the tunnel on its way to the client.

The advantage of this approach is that can be completely transparent to the user. No special software is necessary on the client or the corporate servers. However, there are

also some distinct disadvantages. This type of connection typically costs a lot to set up and requires major hardware investments. It is not a viable solution except for larger corporations and even then it may not necessarily be the most cost-effective.
It is also a restrictive approach since it binds the company to a single mobile operator and implies that all users must have subscriptions with that carrier. In international scenarios this can become complex to implement consistently. There are no worldwide mobile operators so it would require a number of dedicated connections. And roaming users would be forced to use inefficient routing topologies that looped through their home country.

Additionally, we have the issue discussed at the beginning of this KB. The enterprise must trust the mobile operator. If all the traffic is routed through the mobile operator in the clear then this presents a security risk for the company. They may choose to accept the risk but at the very least they should carefully consider the implications and alternatives.

End-to-end Virtual Private Network

In addition to the dedicated VPN there is also the possibility of a client VPN which extends to the perimeter of the corporate network.

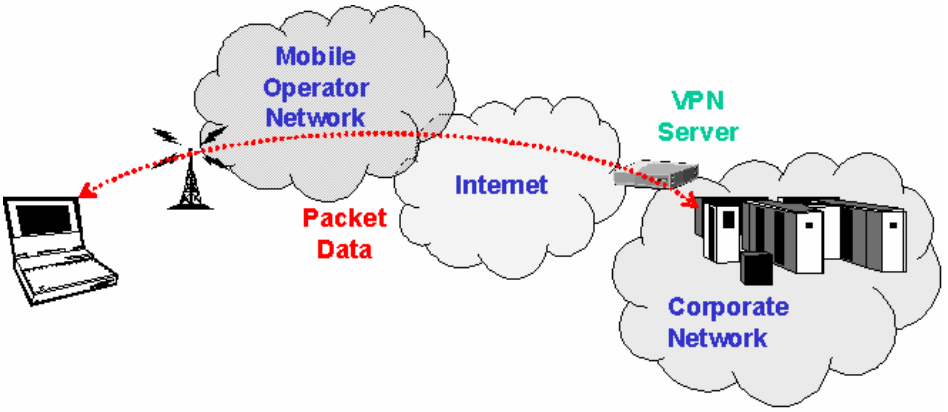

Figure 3: End-to-end VPN from client to private network

It can encapsulate all the data over the mobile operator's network in addition to the public Internet and is therefore more secure than the dedicated VPN . It is also operator agnostic since all the encryption and authentication take place on the device and the enterprise VPN server. Since this approach can use any available VPN technology it can be implemented more simply and cost-effective than the dedicated VPN described earlier.

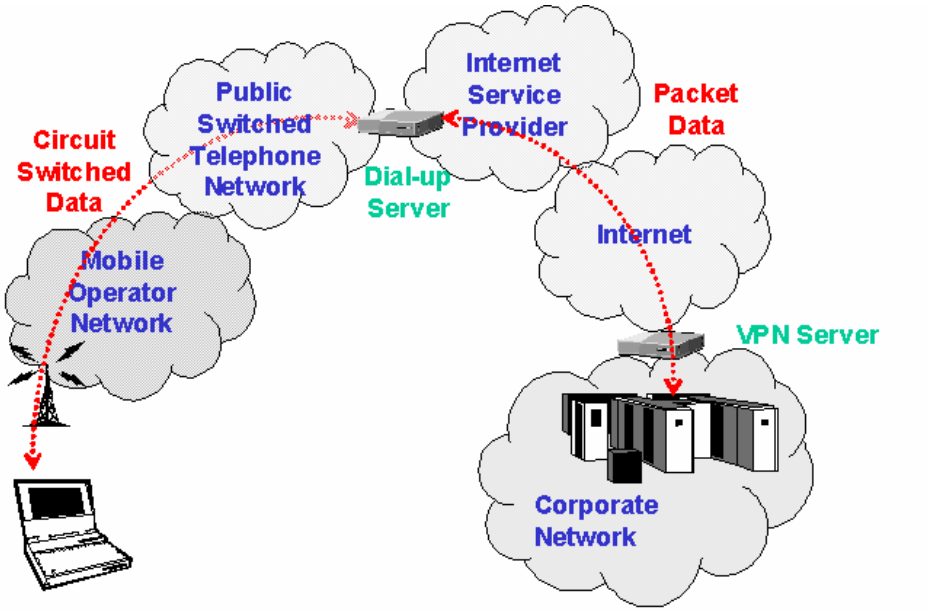

Figure 4: Dial-up via ISP

A further advantage is that this can be used as a general-purpose remote access solution for all wireless and wired networks. Not only does it span all mobile operators running packet data networks but it can also be used for circuit-switched data by allowing users to dial-into any Internet Service Provider and run a VPN over that connection.

Whether dialing up over GSM or a simple analog/ISDN line, users can connect to any ISP From there the user will have IP connectivity and can connect to any VPN server. To the VPN server the users look the same whether they are connecting via a broadband cable-modem/DSL-modem, a GPRS packet data network, or any ISP through GSM or a phone line. It is just necessary to have Internet connectivity.

The general-purpose nature of this approach makes it attractive but it is worth noting that everything has its price so there are drawbacks to consider here too. First of all, we will suffer performance degradation (typically on the order of 10%-30%) since the encryption introduces overhead into the transmission.

The second obstacle is that we must load a VPN client onto each mobile device. In addition to the administrative effort of loading the software this can be a challenge since not every VPN client is available for every platform. We must carefully select the VPN protocols and products to maximize the reach across our user base.

VPN Protocols and Alternatives

There are several VPN protocols we can consider for mobile access. The industry protocols, such as PPTP, L2TP and IPsec, are well described in many published sources, so I will not repeat the discussion here. There are, however, some problems with standard VPNs that you should be aware of in a mobile environment including their susceptibility to latency and incompatibility with network address translation (NAT).

Summary

In many ways Wireless Wide Area Networks solutions can be integrated into a general purpose remote access solution for the enterprise. Whether they are circuit-switched dial-up solutions or packet-data Internet connections wired and wireless remote access can look the same to the corporate perimeter.
They can use the same protocols and remote access products. In fact there is a compelling case for creating unified and simplified approach to remote access. However, we do need to ensure that wireless networks have unique requirements and therefore a traditional and standard solution may not necessarily be optimal.

References

IETF RFCs

RFC1055 – Serial Line Internet Protocol (SLIP)
http://www.ietf.org/rfc/rfc1055.txt
RFC1171 – Point-to-Point Protocol (PPP)
http://www.ietf.org/rfc/rfc1171.txt
RFC 1334 - Password Authentication Protocol (PAP)
http://www.ietf.org/rfc/rfc1334.txt
RFC 1994 - Challenge-Handshake Authentication Protocol (CHAP)
http://www.ietf.org/rfc/rfc1994.txt
RFC 2433 - Microsoft Challenge-Handshake Authentication Protocol (MS-CHAPv1)
http://www.ietf.org/rfc/rfc2433.txt
RFC 2759 - Microsoft Challenge-Handshake Authentication Protocol Version 2.0 (MS-CHAPv2)
http://www.ietf.org/rfc/rfc2759.txt
RFC 2284 - Extensible Authentication Protocol (EAP)
http://www.ietf.org/rfc/rfc2284.txt