



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Appropriate Use of Network Encryption Technologies

GIAC GSEC Practical Assignment v1.4, option 1

Author: Kenneth Forward
Submitted: September 20, 2002

Abstract

Not so long ago, it seemed the single most prevalent misconception of network security was that of ‘firewall as magic bullet’ – the belief that, all else aside, if you had a firewall you *somehow* had to be secure.

Happily, this misconception seems less prevalent than it once was; unfortunately, it seems to have been displaced in part by an equally misplaced trust in virtual private networks and other forms of network encryption technology.

This paper will describe virtual private networks and other network encryption technologies such as secure sockets layer – what they are, and what protections they provide. Equally important, this paper will also examine the flipside; namely, what network encryption cannot provide; how one can actually compromise overall security through poor implementation; and measures one can take to minimize these risks. The necessity of *defence in depth* will be emphasized throughout.

Intro: Cryptography and Encryption

For an entertaining introduction to the field of cryptography, the reader is referred [Schneier, Secrets, ch. 6]. For the purposes of this paper, it will suffice to briefly define encryption, distinguish between symmetric and asymmetric encryption, and to describe one-way hash functions.

Encryption refers to the process of converting intelligible, unencrypted *plaintext* to unintelligible, encrypted *ciphertext*, by application of an encryption algorithm or *cipher*. Least the secrecy of the ciphertext depend on the secrecy of the algorithm (and thereby require each set of correspondents to use a unique one), encryption algorithms usually also take as part of their input an *encryption key*. It is the uniqueness and secrecy of the latter, not that of the algorithm, which renders a given ciphertext secure.

In symmetric or *secret key* encryption, the same key used to encrypt the plaintext is used to decrypt the ciphertext. Secret key encryption is much faster than asymmetric or *public key* encryption (described below), but provides no inherent means of authentication, and suffers the chicken-and-egg dilemma of how to securely distribute the secret key without benefit of the encryption not yet conferred.

In asymmetric or *public key* encryption, there exists a pair of complementary keys: a so-called private key, which must be kept secret by its owner, and a public key, which must be made available to prospective correspondents. The public key can be used by anyone to encrypt plaintext, such that the resulting ciphertext can only be decrypted by the intended recipient retaining the corresponding private key. This ensures message confidentiality, as per symmetric key encryption, while neatly sidestepping the challenge of secure key distribution. Additionally, messages can be cryptographically signed by an author using his/her private key. This signature can be verified by the recipient by application of the author's public key, thereby authenticating the message's origin. As alluded to above however, the drawback with public key encryption is that it is substantially slower than secret key encryption.

Given these complementary strengths and weaknesses, one often finds these two forms of encryption used in tandem, even in applications – such as GnuPG or SSL - that are generally thought of as public key based. Typically, slower public key encryption is used to (optionally) authenticate the correspondents, and to securely exchange a secret key, which is then used to symmetrically encrypt and decrypt the remaining conversation.

Finally, a one-way hash function or *message digest* is a cryptographic algorithm which, given an input and optionally a key, calculates a (statistically) unique, irreversible, fixed length representation of the input. One-way hashes can be used to detect possible tampering with or corruption of a binary object, provided a hash of the object in its known good state has already been calculated. To verify an object's *integrity* in this manner, a recipient merely recomputes the object's hash upon receipt; if the latter value matches the original hash, the object is assumed unmodified.

Virtual Private Networks: Definition and Motivation

Simply defined, a VPN or *virtual private network* refers to the emulation of "...a secure, private network over a public network such as the Internet" [Scott et al, p.2].

The traditional private network referenced above is typically created using leased WAN connections and/or dial-up remote access (RAS) servers, with the former providing connectivity between entire sites, and the latter permitting travelling or otherwise remote employees to connect back to the LAN. Given they run over dedicated infrastructure, traditional private networks generally provide a high level of network security.

The drawback of the traditional private network however, is expense. Wide area network connections are costly, and scale poorly; to provide full redundancy between n sites requires $n \times (n-1) / 2$ leased lines. Long distance and toll-free charges incurred by dial-up RAS users can be similarly prohibitive.

Reducing these costs by availing of the cheap ubiquity of the public Internet has historically been the prime motivation behind VPN deployment. More recently, this adoption has been accelerated by the increasing availability of consumer broadband, which offers WAN-like bandwidth at near local dial-up prices. In many respects, it is this emerging class of VPN users which presents the greatest security concerns. But first....

How does VPN work?

Generally speaking, the VPN attempts to provide security comparable to that of the traditional private network by implementing the following [Nortel, pp. 4-5; Smith]:

Authentication - verification of client identities, both human and machine;

Integrity checking - confirmation that the data was not changed *in transit*;

Confidentiality – assurances that the data, if intercepted *in transit*, cannot be read by unauthorized third parties.

Of these three aspects, confidentiality – as provided by data encryption - is typically most associated with the virtual private network function. The symmetric encryption algorithm used will vary from VPN implementation to implementation: MPPE/RC4 for PPTP-based VPNs, DES for L2F-based ones, or any of the IPSec-supported algorithms - which include DES, 3DES, Blowfish, CAST, Twofish and AES/Rijndael [BorderWare] - for L2TP- and IPSec-based VPNs [Scott et al, p.63]. The reader wishing additional information on tunnelling protocols is referred [International Engineering Consortium]; on algorithms, [Schneier, Applied]. What is important to note for present purposes is, by rendering data unreadable, VPN encryption facilitates confidential data exchange over otherwise shared networks.

The two other, less obvious provisions of VPN are no less important than confidentiality. Just because ciphertext cannot be decrypted by an interloper does not mean it cannot be modified by one; nor is there any guarantee that the resulting changes to the decrypted plaintext would be so obvious as to raise suspicion. Hence we have the VPN requirement for data integrity checks, typically implemented using MD5 or SHA-1 message digests. By ensuring the data received is identical to the data sent, the VPN attempts to provide a level of data integrity comparable to that of the traditional private network.

Finally, just as membership in the traditional private network is limited, so too must admittance to the virtual private network be policed. And while authentication - or verification of a client's identity – is obviously used in both scenarios, the traditional private network enjoys a literal isolation which the virtual private network does not. Therein lies the extreme importance of authentication to VPN deployments, and the reason why our definition of authentication does not distinguish between human and hardware supplicants. Done properly, authentication is reciprocal: in the remote access

scenario, the VPN endpoint or gateway authenticates to the user, just as the user authenticates to it; while in the LAN-to-LAN instance, bi-directional gateway to gateway authentication occurs.

Given the importance of authentication to VPN security, the existence of legacy (i.e., non-VPN) authentication systems, and the variety of human-to-host and host-to-host authentications which can take place, it is not surprising that a multitude of VPN authentication mechanisms exist. All can be categorized as primarily shared secret or public key based however. Once again, for the purposes of this paper it is sufficient to note that, through authentication, VPNs attempt to provide the sort of access controls that are intrinsic to the traditional private network.

The Security Issues

To recap, a VPN attempts to emulate the security of a traditional private network through authentication of participants, and encryption and integrity checking of the data exchanged.

Insofar as it provides a secure communications channel, the VPN succeeds. Said channel cannot protect data before or after transmission however, nor can it guarantee security of the end nodes, or protect other networks to which they may connect. In short, the security of resources both inside and out of the VPN very much depends on the virtual private network always remaining *private*.

This statement is of course equally true of the traditional private network; when isolation, virtual or otherwise, is breached, security is lost. The added risk with a virtual private network of course, is that it does not run over dedicated infrastructure; it runs over a public network and, in remote access and business-to-business situations especially, may include remote nodes over which one has inadequate administrative control. As such, the VPN is far more likely to suffer breaches in isolation than the traditional private network.

The remainder of this paper examines some of the many ways in which these lapses can occur, and steps which can be taken to minimize these risks.

Network architecture and VPN termination

A factor over which one typically does have a large degree of control is where one locally terminates the VPN tunnel. For purposes of this discussion, it is assumed you already have a firewall; if not, in the words of Scott et al [p.6], "...don't bother with a VPN until you get one – you're already exposing yourself to considerable risk."

VPN terminates inside the firewall

This option initially appears attractive. By placing the VPN gateway behind the firewall and permitting only the necessary AH, ESP, GRE, etc, packets through on the required ports, the VPN endpoint and the internal network as whole are seemingly well protected.

The problem with this arrangement of course, is that you've not really blocked all other protocols and services, you're merely tunnelling them in such a way that you've lost the ability to filter and log them at the firewall!

Note external-client-to-internal-host and internal-client-to-external-host are merely specific if slightly more subtle cases of this 'internal termination' problem.

The classic example of the former is simple SSH port forwarding. SSH was of course among the earliest network encryption solutions, and has quite possibly prevented more passwords from passing in the clear than all other VPN solutions combined; as such it is widely and rightly venerated. That said, its use across the firewall – particularly to internal machines not directly under one's administrative control – is perhaps worthy of review¹.

As firewall administrator at an institution of research and higher education, I occasionally receive requests to support VPN connectivity of the second type above, namely internal client to external gateway or host. Typically, these requests arise because an external government, research, or certification agency requires all connections to them to be encrypted.

This motivation is of course laudable; if the remote site has situated and secured their endpoint as per the types of suggestions given in this paper, the requested arrangement is probably also a fairly secure one – for them! The problem of course is that once again communications are tunnelling across *your* perimeter, with no local logging or filtering of same. The lesson in this is perhaps two-fold:

One, never assume another party's security will also secure you. Always instead ensure any internetworking arrangements meet your local security requirements;

Two, also never assume the converse, namely, that your security measures will suffice or be compatible with those of another party. In particular, if your policies and procedures would necessarily lower the security posture of a site with which you must connect, then your policies and procedures need reworking.

VPN terminates outside the firewall

The advantage of this scenario is that you've regained full filtering and logging of your cross-perimeter traffic at the firewall; the negative of course is that your VPN gateway is

¹ Note well, if your only alternative to cleartext one time passwords across the Internet is SSH across your firewall, then you may well be right in concluding its continued use is necessary. The purpose of this or any review is simply to re-examine the basis of your policies and procedures – just as this paper plays devil's advocate with the assumption that deploying network encryption always brings an improved level of security.

now fully exposed to the Internet.

If truly faced with choosing between this and the previous option, you are in a rather unenviable position. That said, chances are you can actually avail of the following option; namely...

VPN terminates inside the DMZ

A DMZ or demilitarized zone refers to an intermediate zone between your trusted internal network and the untrusted Internet, from which one runs servers (WWW, ftp, SMTP, DNS, etc) that by necessity must be publicly visible.

The 'classic' DMZ picture is that of a network segment between two firewalls, with the outside firewall providing a modicum of protection to the public servers, and the inner one more fully shielding the internal network. More recent firewall designs typically contain three (or more) network interfaces; by applying different rulesets to each, the classic DMZ can be emulated.

Either way, VPN termination inside a DMZ provides the best of the previous two options combined: reduced VPN endpoint exposure, yet full filtering and logging of (decrypted) traffic at the inner firewall.

If you currently do not have a DMZ, have only a single firewall, and that firewall supports only two interfaces, all is not lost. By applying a set of carefully crafted ACLs or *access controls lists* to your outside border router, you can limit traffic to your VPN gateway as per the *inside* option above, decrypt at the gateway, then pass all traffic to the firewall inside; in effect, you can construct a simple DMZ for deployment of your VPN gateway.

Integrated firewall/VPN platform

Further improving on the DMZ placement above, is the idea of an integrated firewall/VPN platform – essentially, a single device offering both functions, perhaps best visualized as a DMZ-based VPN which actually resides inside the firewall.

It may be argued that an integrated firewall/VPN gateway intrinsically offers no greater security than a DMZ-based VPN; and that indeed, an integrated platform represents a single point of failure.

Complexity is the adversary of security however. By providing a common management interface, integrated logging, and guaranteed single vendor support, an integrated firewall/VPN will likely prove less complex to run than a DMZ-based alternative - which in turn will likely mean fewer security- and performance- related problems.

As to the integrated platform representing a single point of failure, it is perhaps true that a classical DMZ is more difficult to breach, especially if the two firewalls are different

models not sharing the same exploits. Insofar as the components are connected in series however (FW-VPN-FW), the classical DMZ deployment offers no physical or *denial of service* redundancy beyond that of the integrated platform.

VPN termination relative to the NIDS

Discussion of NIDS or *network-based intrusion detection systems* is beyond the scope of this paper. The author would simply note that placement concerns similar to the above also apply relative to the NIDS. If for research or other reasons you wish to maximally detect incoming attacks (assuming minimal filtering at the outside firewall), you need a classic DMZ: FW-VPN-NIDS-FW. If on the other hand you are only concerned with attacks not stopped at the perimeter, a FW-VPN-FW-NIDS, or integrated FW/VPN-to-NIDS architecture, will suffice.

Note as before, traffic allowed to tunnel directly to/from internal hosts will circumvent your network-based IDS, just as it will your firewall.

Remote Access VPNs - Client Side Issues

The previous section addressed VPN endpoint placement relative to the enterprise network perimeter (as defined by the firewall and NIDS). As such, the recommendations therein largely apply to LAN-to-LAN VPNs, and the enterprise or 'head end' of remote access VPNs.

That the security of one's VPN and indeed one's enterprise depends on the trustworthiness of both endpoints is particularly true of the remote access VPN. Unlike the intranet VPN, where both ends are under common administrative control, or even the extranet / business-to-business VPN, where one can at least hope the other endpoint is professionally administered to common criteria, the remote access node is likely - at least initially - to be less secure.

This section attempts to address these concerns.

The Ideal

Ownership: Ideally, the remote access VPN node is enterprise owned. Virtually all of the issues to follow are matters of policy, and it is difficult if not impossible to impose policy on resources one does not own. This point cannot be overemphasized. If you truly hope to maintain enterprise security, all remote access VPN nodes should be company issued, and subject to the constraints below.

Single purpose: The remote access node should only be used for enterprise purposes, never for personal or third party ones, and certainly never used by third parties.

Single homed: In order to remain secure, the remote access node should never connect to any other network. To do otherwise is to violate the virtual isolation the VPN attempts to provide.

As limiting as these requirements appear, they are in fact achievable. If your remote users are travelling sales or support personnel equipped with company laptops, who are only permitted to dial a local point of presence in order to access the company VPN, then congratulations; you are already in good stead.

For the rest of us higher education types, we have...

The Reality

The remote access node is not company owned; it is used by spouse and kids; and it is regularly – indeed, primarily – used to connect directly to the Internet, perhaps via ‘always on’ broadband.

As the enterprise security administrator, what are your options?

First and foremost, one still – or perhaps *especially* - needs to pursue policy. Make it as stringent as possible under the circumstances; try to shift the playing field back towards the ideal, and document remaining shortcomings for management sign off. Issues to address include:

Personal firewall and antiviral software: In all instances the remote access node should be equipped with personal firewall and antiviral software; in the present case, it is imperative.

Many recent remote access VPN clients include an integrated and/or centrally managed personal firewall; vendors of products of which the author is aware include BorderWare, Check Point, Cisco, NetScreen, and WatchGuard. If you are shopping for a VPN solution, or already have this capability, its implementation is highly desired. Depending on the features of the integrated client, other configuration checks may also be enforced: the remote access node may not be permitted to connect if antiviral software is not also running; if the personal firewall policy or AV signatures are outdated; if the host is currently dual homed; etc.

While these measures fall short of providing users with dedicated hardware, equipping remote access users with this level of protection can be a cost-effective compromise at roughly US\$100 per client.

Split Tunnelling: Split tunnelling refers to the practice of tunnelling only enterprise-bound traffic through the VPN; public- or Internet-bound traffic is directly routed by one’s ISP as usual. The alternative to this is to tunnel all traffic through the VPN, regardless of destination.

The advantage of the latter approach – provided your gateway is situated as per the recommendations of the previous section - is that all traffic is subject to the same firewall, AV, content filtering, etc, that would apply if the remote access client were truly internal.

The disadvantages of the ‘full tunnelling’ approach are increased latency, bandwidth consumption, and VPN processing. A GET of a public web document, for example, would entail encrypting the request, routing it to the enterprise VPN gateway, decrypting and verifying the request, and routing the resulting plaintext out to the Internet; the public website’s reply would in turn route back to the enterprise gateway, at which point it would be encrypted, routed to the client, then decrypted and verified for viewing.

The pros and cons of split versus full tunnelling must be weighed according to your requirements, both security and otherwise. Of particular consideration however are the personal firewall and AV precautions suggested above. In effect, the split tunnel ‘dual homes’ the remote access client, placing it simultaneously in the private and public networks; and our preferred model is that the client never attached to another network, let alone simultaneously. Great effort must therefore be taken to minimize split tunnelling risks, if indeed that mode of connectivity is permitted.

Secure Sockets Layer

Like VPNs, SSL or *Secure Sockets Layer* attempts to provide a secure communications channel through authentication of participants, and encryption and integrity checking of the data exchanged. Both support many of the same encryption and hash algorithms, with SSL authentication relying almost exclusively on the same public key digital certificates sometimes used in VPN authentication.

There are of course differences between the two; whereas VPNs variously operate at OSI layers 2 or 3 for instance, SSL is implemented at OSI layer 6. As such, the former can tunnel any application transparently, while support for SSL must be programmed into each application. While secure sockets layer can be used to secure virtually any network connection, its best-known application is of course https or secure web browsing.

Given their similarities, VPNs and SSL raise many of the same basic security concerns; namely, potential tunnelling of undesired protocols/content past one’s perimeter defences, and a lack of protection once the data is no longer in transit. As typically deployed in the secure web arena however, SSL usually does one worse than most VPNs; all too frequently, strong authentication is only required of the server by the client, and not vice versa.

As an example of their relative strengths, and of how any encryption technology should only be deployed in a *defence in depth* context, we consider the following scenario:

An enterprise has a web-based application it would like to make available to its travelling

and telecommuting employees. Realizing these employees will connect via shared insecure networks, the company wisely decides to provide only secure https access to this application; however, no other means of access control beyond simple username and password authentication are implemented

That telecommuters' passwords and other proprietary information will not pass as cleartext is certainly a plus in the above scenario. Not addressed however is the fact that the service itself is globally visible - and that, e.g., an *SSL-encrypted* exploit of either the URL, username, or password buffer would still amount to a successful exploit.

To better secure this application, one needs a way of denying connection attempts from unauthorized users. With a small number of remote users connecting from fixed IPs, this could perhaps be attempted via layer 3 access controls at the border router or firewall.

If the user base is large and/or connecting from dynamic or unknown IPs however, layer 3 access controls simply will not scale. Security in this instance can best be provided through a full VPN infrastructure, with the clients strongly authenticating to the enterprise VPN gateway before they are permitted access to the internal web-based application.

Wireless and WEP

The topics of wireless and WEP are germane to this paper for two reasons:

One, as a reminder that encryption itself can be broken. The emphasis of this paper is of course that network encryption is not a panacea, that it only has a place in a *defence in depth* approach to security. That said, bad encryption (or bad implementations of good encryption) do exist – which is yet another reason to not put all one's eggs in the encryption basket.

Secondly, and returning to this paper's central theme of appropriate use, VPNs are now regularly being touted as a way of securely connecting your wireless and wired networks, which means many of the forgoing considerations re gateway placement, client security, etc, may be applicable to your wireless deployment.

Of particular but perhaps not obvious concern are the risks of split tunnelling in this environment. A VPN gateway between wireless and wired networks may indeed make direct wireless attacks against the wired network difficult. Said deployment would not prevent attacks against VPN-authenticated clients however, if association could be achieved and split tunnelling were allowed [Dismukes]. Once compromised, the VPN-authenticated client could then be used to launch an attack against the wired network.

Once again, the solution to this scenario is defence in depth: disable or firewall split tunnelling, and make unauthorized association difficult by requiring user or MAC authentication, by using non-default SSIDs, by not broadcasting SSIDs, etc.

Further considerations for defence in depth

The case for *defence in depth* when deploying network encryption technologies has hopefully been made. Additional defence in depth considerations include:

Security of information when not in transit: Referring back to our discussion of integrity checking and confidentiality as implemented by VPNs and SSL, we are reminded that both provisions only apply to data *in transit*.

Insofar as portable laptops and office-to-home sneakernet have long existed, the problem of securing information residing on nodes outside the enterprise is not new. With the advent of shared household computing and ‘always on’ broadband however, the risks to this information are far greater than ever before.

Technical approaches to mitigating these risks include individual file and/or filesystem encryption, XML element encryption [Forum], and the sorts of personal firewall and antiviral measures previously discussed.

That said, greater security may lie in your answer to the question, do you really want your company’s financials sitting on the same box that somebody’s kid uses for ICQ or P2P file sharing? If the answer is ‘no’, then put that objection into writing (i.e., policy) and into practice (provide your telecommuters with company-owned machines).

patching and Administration: This could go without saying, or perhaps bears repeating: a VPN is a complex system which will require patching and administration. Once again, this is particularly true of remote access VPNs, which may involve dozens, hundreds, or even thousands of end points.

Can you secure this many remote nodes? Does your overall ‘solution’ enable you to easily push client updates? Personal firewall policies? AV signatures? What about security-related patches of non-VPN-related software (including the OS)?

The potential enormity of this task should not be confused with the (relative) ease of securing an equal number of internal clients (which are more likely to adhere to company hardware and software standards; not be dual homed; not have non-enterprise users; etc – or failing all those assumptions, they are at least behind your centrally managed firewall and NIDS).

Conclusions

Depending on implementation, a VPN may safely extend the LAN, or endanger it by

circumventing enterprise firewall, NIDS, AV, and content filtering efforts.

Implementation issues to consider include: positioning of VPN endpoints relative to the enterprise security perimeter; security of the endpoints themselves; dual homing / split tunnelling issues; and the security of data before and after transmission.

It is only through deployment with these sorts of *defence in depth* concerns in mind, that virtual private networks and other network encryption technologies can prove safe and cost-effective additions to the enterprise networking environment.

References

Borisov, Nikita, Ian Goldberg, and David Wagner. "Security of the WEP Algorithm." URL: <http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html> (10 Jul 2002).

Dismukes, Trey. "Wireless Security Blackpaper." July 2002. URL: <http://arstechnica.com/paedia/w/wireless/security-4.html> (05 Sep 2002).

Forum Systems, Inc. "SSL: Not enough for Today's Web Services." 18 Jun 2002. URL: http://www.forumsys.com/papers/ssl_not_enough.pdf (19 Jul 2002).

Fowler, Dennis. Virtual Private Networks: Making the Right Connection. San Francisco: Morgan Kaufmann Publishers, Inc., 1999.

International Engineering Consortium. "Virtual Private Networks (VPNs)." © 2002. URL: <http://www.iec.org/online/tutorials/vpn/> (08 Sep 2002).

Nortel Networks Ltd. "Virtual Private Networks and IPSec: Ensuring the Security of Enterprise VPNs." 28 Mar 2002. URL: <http://www.nortelnetworks.com/products/library/collateral/12002.25-03-02.pdf> (07 Sep 2002).

Schneier, Bruce. Applied Cryptography: Protocols, Algorithms, and Source Code in C. New York; Toronto: John Wiley and Sons, Inc., 1994.

Schneier, Bruce. Secrets and Lies: Digital Security in a Networked World. New York; Toronto: John Wiley and Sons, Inc., 2000.

Scott, Charlie, Paul Wolfe, and Mike Erwin. Virtual Private Networks (2nd Ed.). Beijing; Sebastopol: O'Reilly and Associates, Inc., 1999.

Smith, Jeremy. "Understanding the Win2K implementation of IPSec." 05 Apr 2002. URL: <http://www.techrepublic.com/article.jhtml?id=r00220020322jsm01.htm> (17 Aug 2002).

Product Datasheets

BorderWare Technologies Inc. "Sentinel IPSec VPN Client." © 1998-2002. URL: <http://www.borderware.com/products/ipsec/ipsec.html> (07 Sep 2002).

Check Point Software Technologies Ltd. "VPN-1 SecureClient". URL: http://www.checkpoint.com/products/protect/vpn-1_sc.html (07 Sep 2002).

Cisco Systems. "Understanding the Cisco VPN Client." © 1992-2002 URL: http://www.cisco.com/univercd/cc/td/doc/product/vpn/client/3_6/user_gd/vc1.htm (07 Sep 2002).

NetScreen Technologies Inc. "NetScreen-Remote Product Page." ©1998-2001. URL: <http://www.netscreen.com/products/nsremote.html> (07 Sep 2002).

WatchGuard Technologies, Inc. "WatchGuard Mobile User VPN Datasheet." © 1996-2002 . URL: http://www.watchguard.com/docs/html/muvpn_ds.asp (07 Sep 2002).

© SANS Institute 2000 - 2005, Author retains full rights.