



# **SANS Institute**

## Information Security Reading Room

# **BUSINESS PARTNER VPN: NEEDED NOW**

---

Karen Duncanson

Copyright SANS Institute 2020. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

## Version 1.4 GSEC

### BUSINESS PARTNER VPN: NEEDED NOW

#### Introduction

This document is intended to take a look at Business Partner VPN and focus on challenges now being dealt with in the face of requirements for a VPN that promises end to end security between two separate business entities and even between the users within those entities. With Web Services technologies evolving quickly the need for end to end security with Business partners is accelerating. While emphasis will be on what is yet needed, appreciation for and discussion of, what is, and why, will be offered.

Business is pushing for methods to cut costs while improving performance. The promise of a cost effective Secure VPN has been with us for awhile, but when it comes to actual implementation between business partners, the promise often breaks down. This document considers what may be needed to deliver on the promise.

#### What is a VPN in the general sense?

As we know, VPN stands for Virtual Private Network. What does that mean? The preferred definition describes the VPN to be a private data network that makes use of the public telecommunication infrastructure, maintaining privacy through the use of tunneling protocol and security features.<sup>6</sup> This is the definition that normally comes to mind when the discussion topic becomes VPN. In examining the concept of VPN more closely it is necessary to understand the broad definition of what a VPN can be and then drill down to the features that are at issue relative to the partner net application of VPN.

The VPN consortium offers some thought on the characteristics of the VPN. Accordingly, a VPN must be a known entity to those who administrate it. Those who deploy and maintain it must know its extent and what data can use it.<sup>6</sup> this helps us get our baseline thought about function and will lead to understanding of the level of security a VPN provides. A VPN is defined by the consortium to be either SECURE, TRUSTED, or HYBRID.

#### Trusted VPN

A trusted VPN is created and maintained by the service provider(s). Only the service provider can control the path the data takes. Only the service provider controls the quality of service. The VPN owners *trust* the service provider. This architecture provides a single source for pathing and the customer *trusts* that these paths will be maintained according to an agreement. This concept is similar to the leased line. However, the customer has no way to test that the agreement is being upheld and

must *trust* the provider. This type of VPN is about bandwidth and quality of service. It is not about security. This type of VPN is based on RFC 3031: MPLS, Multi Protocol Label Switching architecture. It is described here for awareness. We are interested in security and will focus on a discussion of Secure VPN.

## **Secure VPN**

The acronym VPN in general brings to mind the features usually associated with a Secure VPN. Experience talking with users in most corporations indicates that when the topic of building a VPN is brought up, it is with the idea that SECURITY will be gained through encryption. The VPN consortium specifies that the Secure VPN be encrypted and authenticated using the IPsec (secure IP) protocol.

## **Hybrid VPN**

Trusted VPNs exist that also contain within their boundaries one or more Secure VPNs. This is a Hybrid VPN. This VPN provides performance and security. Security exists only between the encryption endpoints. Performance is predictable throughout the entire VPN.

## **Secure VPN and IPsec**

The VPN consortium has specified IPsec protocol as the technology for Secure VPN. IPsec, also known as secure Internet protocol is in the IP layer and is connection oriented due to the security associations needed to maintain an encryption tunnel. All communication layers above the IP layer are protected by the IPsec protocol suite. The suite consists of three protocols: IKE or Key Exchange, AH or Authentication Header (51) and ESP or Encapsulating Security Payload (50).

Before looking at the benefits and drawbacks of layer three or network layer security, it is interesting to understand the rationale for placing security at this level. As mentioned above, IPsec can protect all layers of communication above it in the stack between the security endpoints. There is benefit to be gained from this strategy. However, not all designers of the IPsec protocol supported this rationale because there are drawbacks.

## **Motivation for security at the network layer (IP layer)**

Tanenbaum, whose Computer Networks text is now in its 4<sup>th</sup> edition, has published a brief, but essential section in his book as follows. The essence of this passage is that security was designed with the user community in mind.

*"IETF has known for years that security was lacking in the Internet. Adding it was not easy because a war broke out about where to put it. Most security experts believe that to be really secure, encryption and integrity checks have to end to end (i.e. in the*

*application layer). That is, the source process encrypts and/or integrity protects the data and sends that to the destination process where it is decrypted and verified. "The trouble with this approach is that it requires changing all the applications to make them security aware. In this view, the next best approach is putting encryption in the transport layer or in a new layer between the application and transport layer. Making it still end to end, but not requiring applications to be changed."*<sup>1</sup>

*"The opposite view is that users do not understand security and will not be capable of using it correctly and nobody wants to modify an existing programs in any way, so the network layer should authenticate and/or encrypt without the users being involved."*<sup>1</sup>

This last view prevailed with the argument that network layer encryption does not prevent security savvy users from doing it right and it does somewhat help those who do not understand security. The result of the outcome of that debate is the IETF (Internet Engineering Task Force) series of RFCs that define the IPsec protocol suite.

### **Transport Mode and Tunnel Mode**

Secure VPNs are set up in Tunnel mode or Transport mode. IPsec supports both of these modes. In Transport mode the original IP header is maintained. The IPsec security header is inserted after the original header. The encryption does not include the original header, but everything behind the original header. This is a security concern, The payload is encrypted and the VPN end points can include devices inside the network perimeter (i.e. devices that are not gateways or firewalls) This means the very sensitive data remains encrypted as it traverses the internal network as well as the Internet. This is an IPsec host to IPsec host encrypted connection or peer to peer. The transport implementation offers this end to end level of security which is good for sensitive data. However, this strategy blurs the definition of the security perimeter and exposes internal addresses and traffic patterns.

Transport mode is generally used for remote access from an external workstation to a server within the network. Telecommuters and road warriors are typical users of the transport mode IPsec. As of this writing, the strategy for supporting telecommuters is to use L2TP/IPsec (RFC3193). This authenticates machines and users separately. The IETF drafts that address UDP encapsulation and address NAT issues, are in final working group and once approved should enable the L2TP/IPsec strategy for those who want to use remote access clients including Microsoft native VPN client. Transport mode IPsec is between two hosts. This type of remote access is being addressed.

IPsec Tunnel Mode is between IPsec gateways. The new IP security header is placed in front of the entire packet. The entire original packet including the original IP header is encrypted. The tunnel exists from gateway to gateway. This is used for connecting a remote network to the existing network. The data is encrypted while in the tunnel, but once it leaves the tunnel end point it is in the clear. Once in the clear the original IP header takes the packet to its destination. This means that if the remote site has inbound access into the tunnel and the SA (Security Association) can be established

and upheld, the remote site has access to the far side local network as it leaves the tunnel. It is at this point that security is not adequate for business partner Secure VPNs. The same situation is similar for other Secure VPNs. However, in other situations such as remote networks within the same company, the same level of paranoia may not exist. Perhaps it should.

## **RFCs**

The primary protocols for supporting Secure VPN standards are available and are specified in a series of IPsec RFCs published by the Internet Engineering Task Force (IETF). These are referred to as Secure IP or IPsec. Most of the RFCs that vendors design to today have been in place since 1998, having made obsolete, the first set of RFCs that were published in 1995. There is still networking equipment in production supporting the 1995 vintage RFCs. There will soon be further productive evolution of the IPsec RFCs. Drafts are under consideration proposing solutions to issues that have been chronic problems preventing the deployment of business partner net solutions among other things.

## **Practical Applications for Secure VPN**

What are the things we want to do with Secure VPNs? It is simple to make a list:

- 1.) Secure VPN supporting secure business partner nets and,
- 2.) Secure VPN supporting corporate wan expansion or VPN supporting a private intranet within the corporate net, and
- 3.) The very popular road warrior secure remote support or telecommuter secure remote support.

When we start thinking each of these various applications, the reference to the VPNC baseline requirement regarding control over the “extent” and “what data travels” begins to take on real meaning. The question of where the end points are and how the data is secured beyond the endpoints, as well as how the networking architecture beyond is secured, come into focus.

## **A large hurdle has been, and is interoperability**

An important use of a Secure VPN is to connect two companies or business partners as if they were using a private leased line, but at a much lower cost. A Secure business partner VPN can provide better security than the legacy private leased line making the technology very attractive. Optimum price, performance, and security are possible. When there is so much capability possible, why are there not more Secure VPNs? What causes hesitation in deployment? Why are costly leased lines still maintained? Experience and observation tell me that there are hurdles on the road to what seemingly should be a silver bullet solution.

As mentioned above, most of the RFCs that specify security at the IP layer were created in 1998 and obsolete the earlier RFCs from 1995. RFC 2401 overviews the

security architecture for the IP layer. It addresses the security protocols (AH and ESP), security associations (SA), key management (IKE) and algorithms for authentication and encryption without going into detail. Detail is handled in other RFCs. Since IPsec protocols, as defined by RFC 2401 are specified by the VPN Consortium as being used in all Secure VPNs, it would seem that interoperability would not be a problem. To the credit of the working group of the IETF, IPsec does a decent job handling authentication and encryption between the end points when the hardware is in compliance. The standards were written to allow cryptographic flexibility (the capability to select the best algorithms) amongst tunnel vendors as long as the vendors wrote in hooks for flexibility. Interoperability has been largely achieved in this respect. The most difficult area seems to be the key exchange where perhaps too much flexibility is possible.

Unfortunately, the problem still exists despite RFCs and ongoing testing and compatibility matrix guides. Not all products that are reported to interoperate do so in all modes.<sup>3</sup> As an aside, it should be noted that authentication is handled by the ESP specification as well as AH specification and encryption is handled by ESP, but not AH, making the use of AH optional. The earlier RFCs, vintage 1995, required both AH and ESP to get a tunnel, since ESP previously did not authenticate. As sited earlier, older equipment is sometimes still in use and cannot be upgraded.

The difficulty in building partner nets is related to the difficulty in establishing a Security Association (SA). The difficulty with the SA is the degree of latitude left to developers/vendors. This is true even with the 1998 RFCs. At least one draft standard states that the nonspecific nature of the fields that specify the security parameters for building the IKE key are causing vendors to design products that are difficult to use. The key is necessary to establish the shared secret and thereby bring a tunnel up with a valid SA. This is particularly difficult for partner nets. Looking at the previous list of Practical Applications of Secure VPNs, it becomes clear that the partner net is the one situation where those who administer the end points of the tunnel, have little or no control over the hardware vendor equipment at the other end. There has been more success when deploying Secure VPN for the other two types of applications: remote access and across corporate nets.

Within a single corporation where the security gateways (normally firewalls or routers) can be specified, there should be no interoperability problems. Nor should Secure VPNs from office to office within a corporation pose user authentication issues, since this use of VPN is only expanding an "internal" network. Real world companies have been moving ahead with Secure VPNs supporting gateway to gateway connections when the partners are remote offices of the same company or when the end users are road warriors for well over a year. In these situations there is control over the hardware and software and interoperability is not a concern. Additionally security is a lesser concern since the users of the connection are all on the same team and internal controls are easier to apply. It is probably safe to assume that if you can select matching equipment for both ends of the tunnel, you will be able to create a successful

Security Association. For the other applications listed above, single businesses can control the selection of both end points.

As an example of the difficulties encountered when bringing two separate business entities or partners onto the same secure VPN we can look at the ANX network. The ANX network is a hybrid VPN that supports Secure VPN between business partners. Although interoperability testing to determine a short list of approved VPN products continues today, the ANX partner network has recently taken steps to move to a single device vendor in order to overcome interoperability issues. The fallout from this decision has been substantial considering the size and reach of ANX. Tunnel end points are generally network gateways or firewalls, operating at the network layer of the protocol stack. The vendors who were not selected have most certainly been negatively impacted even if their products were of equal or better quality than those chosen. The other difficulty with this is the challenge to maintenance and development of the IETF standards process.

### **About doing it right**

The partner net VPN is still largely lost in the wilderness of interoperability and security issues. Assuming the Interoperability issues are to be resolved in the near future by rigid adherence to existing and new standards, there will continue to be a large concern for the space between the VPN endpoints and the corporate data. If it is convenient for a *business partner* to connect to another company's network via Secure VPN, then access to specific applications and files should be limited to authorized users rather than all users on a Secure VPN.<sup>4</sup> Tunnel support does not address this space. The alternative of requiring each application on each server to handle security is not administratively effective nor does it suggest efficient performance even if done properly. It is generally agreed that end to end user security for applications utilizing Secure IP is a worthy goal and that security should be handled in a global way rather than a local way.

It should be noted IPsec is based on symmetric key cryptography with the sender and receiver negotiating a shared key before setting up the SA (Security Agreement). The shared key is used in the signature computation and must not be transmitted. Both AH and ESP headers negotiate and provide authentication. AH embeds the authentication information after the IP header. ESP places it at the end of the packet after the payload. Once the SA is in place, communication can begin. The key exchange portion of IPsec, which results in the SA establishment, is complex, causes many interoperability problems and is being reviewed. Some security specialists consider the IKEv1 specification it to be flawed because it allows too many options without enough guidance, hence interoperability problems. Given this, and in light of Industry business leaders taking it upon themselves to impose standardization of equipment to resolve issues, perhaps there will be more momentum to move forward with standards. Improvements to the key exchange specification that will resolve the difficulty of establishing the SA. Improvements to IKE might provide guidance for establishing end user authentication. Such guidelines would be very useful to the Internet Community.

Currently there are works in progress to be taken up at future IETF WG meetings. One in particular deals with IKE. Perhaps we can be optimistic that key exchange among vendors can become simplified enabling most interoperability issues to be cleared up. The intent of the work in progress is to replace RFCs 2407, 2408 and 2409. Among other things, the proposed improvements provide support for a certificate payload and thereby offer opportunity to work on integrating support for PKI certificates into Secure VPN transactions.

### **What is PKI and How might it help?**

Another work in progress at this time, The IP Security PKI Profile of ISAKMP and PKIX, defines explicitly how PKI (Public Key Infrastructure) would be used with IPsec to provide end to end security. The use of PKI to provide end to end security with Secure IPsec has been a topic of discussion for a while. Business partner Secure VPNs would be completely supported by having global security beyond the VPN end points. If PKI certificates were to interoperate across platforms and gateways, the solution would be celebrated. Vendors have created products to support PKI and it has been demonstrated to work in single vendor environments. PKI Interoperability becomes an issue in multi vendor environments.

PKI implementation in an organization is a very large undertaking. PKI provides the complete system of software, hardware, policies, standards and technologies for life cycle management of digital identities and keys. Digital identities or certificates are the mechanism for establishing and verifying the identity of a network *user* or resource. All of this information is stored in a directory structure based on the x.500 (PKI) standard and uses LDAP (lightweight directory access protocol) to access and search. The system is scalable and secure. Cross certification can be used to support business partner relationships providing there are no proprietary issues. Standards are essential to business partner relationships. Standards for PKI are in progress.

To accomplish what is necessary for Secure VPNs to support end to end user security, there will need to be a VPN Policy manager that issues attribute certificates and stores them in an LDAP directory along with the associated identity certificates. The identity certificate is standard a PKI certificate and establishes the identity of a device or user. Attribute certificates provide information about what an entity can access. Both of these are signed certificates. The attribute certificates become the central source for policies. When a user is connected to a network, their identity is authenticated using the usual digital certificate. Next the attribute certificate is checked to grant the appropriate access to any number of resources. The attributes are associated to a VPN policy grouping. The potential is here to give a single user access to specific applications, servers and other resources. Policy can further designate time of day, day of month, or protocol. If for whatever reason, the permissions needed for access as requested can not be verified, the tunnel would not be established.



To digress for a moment, it is important to note that these are the kinds of things that firewalls have attempted to do relative to secure VPNs. Firewalls fail mostly because they cannot do these things globally with enough granularity. Some firewall vendors have made a grand attempt to address the need and partially succeeded, but again the need is global within and across organizations. Even if firewalls succeeded here, the requirement for user specific authentication across the organization would still need addressing.

The added advantage of the PKI based VPN policy is that policies are authenticated by trusted third parties and now they are scalable. All user access is controlled in a central location and does not depend on a system administrator for a specific workstation or server. The VPN policy manager needs to be standards compliant if there is to be interoperability with business partners.

### **Another Idea: Trust Management**

A recent paper by Ioannidis, Blaze, Keromytis, entitled Trust Management for IPsec, is focused on the problem of handling protected traffic at the tunnel end points. The point is made that the space between the tunnel end point and the final destination is unsecured. The paper introduces a policy management scheme that utilizes efficient filters and complex credentials by way of a two tiered policy check added to the IPsec protocol. The general idea is to provide an API to applications and to modify IPsec to contain a hook. Rather than being only theoretical or a work in progress, this scheme has been implemented in open source on OpenBSD using the Keynote Language (RFC2704) and will probably run on other Unixes.

### **How does it work?**

Packet based filtering is performed at the tunnel endpoint. This is true for incoming and outgoing packets and is a normal part of IPsec. While much more could be done with this, it usually is not and the endpoint grants all or nothing access to whatever comes its way based on the SA. The SAs themselves must be set up with a priori knowledge in order to make them work. As mentioned earlier, this does not work well when business partner Secure VPNs are being deployed. At the heart of this scheme is the policy manager concept called "Trust Management".

"Trust Management system provides a standard interface that applications can use to test whether potentially dangerous actions comply with local security policies." <sup>4</sup> The standard interface has an actions list it responds to. It is able to identify the entities that can be associated with the actions. It has a language to describe policies and another to describe credentials. Lastly, the trust management system has a Compliance Checker to pull all of this together. It is possible to write very sophisticated security policies. In the case of IPsec, endpoint control becomes simple. There need only be a checking process that talks to the compliance checker inserted into the IPsec code as it creates the SA. In fact the IKE code that builds the SA is swapped out and the trust management piece is swapped in. This is the code in IKE that supports passphrase

and x.509 certificates. The “trust” code continues to support passphrase and x.509. This strategy supports pki, but does not depend upon it.

The policy checking process does two things. It supports packet filtering based on rules for all packets and trust management, which negotiates and decides which rules to install. The two-phase nature of this process results in performance gains. Because Secure VPN end points typically take large performance hits when asked to process detail beyond the basic access rules this is an important improvement. This offloads the responsibility for managing trust from the application and from the gateway device. An IETF draft that expired in March 2000, titled Compliance Checking and IPsec Management described the deployment of trust management. Although this strategy is not on the IETF standards track at this time, its authors have made a valid point about the need for security beyond the Secure VPN endpoint. It is worth taking a look.

## **Web Services**

The promise of Web Services, accessed across a partner net or across the open Internet, is loosely coupled applications on different servers combined to provide a lightning fast and flexible web based e-business solution. The Web Service security model expects authentication not only of immediate clients, but also of indirect clients whose requests have passed through many intermediate gateways or firewalls. The authentication information is necessary for the client and the server and must be transmitted so that it can be understood by both. The security standard that defines authentication methodology is not established as of this writing. The security model is not developed sufficiently to support the application’s requirements. As promising as Web Services are, most implementers agree it is still not safe to let them outside the firewall. Forty five per cent of enterprise implementers who are planning to implement Web Services in the next two years site security and authentication as the element that will hold back projects.<sup>9</sup>

The good news is that there are two emerging standards: SAML and WS-Security. The W3 consortium will eventually back one of them, putting forth a standard. Presently, it seems that vendors are lining up evenly on both “teams” which will make standardization of Web Services security architecture very difficult. The better news is that both “teams” are supporting the use of PKI as an option for authentication. Perhaps this means that someday we will have end to end security across Secure Business Partner VPNs.

## **Advocacy: Move the Standards Forward Sooner instead of Later**

This discussion has been about the need for a Secure Business Partner VPN, why we can not build one easily, and what we need to remedy the situation. Today the standards process is essential for accomplishing this. As this summary is being written, version 3 of the IKEv2 draft has been posted and progress is being made toward resolving the NAT firewall traversal problems. The IKEv2 draft defines the certificate payload and leaves the door open for deployment of PKI certificates as a

means of authentication and authorization. This need for security policy between business partner nets is underscored by the growing need for supportable security architecture to support Web Services. Ideally the unified security policies can be maintained in a central location adjacent to web services and VPN gateways. Ideally, we will build an architecture that is supportable across platforms and dynamic enough to scale across the corporation and across the secure business partner net. In this way we can achieve end to end security and realize the promise of Secure Business Partner VPN.

## References

1. Tanenbaum, Andrew S. Computer Networks. 4<sup>th</sup> ed. Upper Saddle River, New Jersey: Prentice Hall PTR, 2003. 772-776.
2. Hondo, M., Nagaratnam, N., Nadalin, A., "Securing Web Services." IBM Systems Journal 41 (2002): 228-241.
3. Wilson, Tim. "VPNs Don't Fly Outside Firewalls." Internet Week 28 May 2001. URL: <http://www.internetwk.com/newslead01/lead052501.htm> (October 25, 2002)
4. Blaze, Ioannidis, Keromytis. Trust Management for IPsec. NDSS 2001 paper. NEC Research Index. 2001. URL: <http://citeseer.nj.nec.com/blaze01trust.html>. (October 25, 2002)
5. Alcatel. Enabling Security in an Increasingly Networked World, Technical Report, May 2000. URL: <http://cnscenter.future.co.kr/resource/rsc-center/vendor-wp/alcatel/pki-tp.pdf> (October 25, 2002)
6. VPN Consortium, VPN Technologies: Definitions and Requirements. Jun 2000. URL: <http://www.vpnc.org> (select VPN technologies) (October 25, 2002)
7. Masica, Ken. "Understanding the IP Security Protocol: Encryption and Authentication for IP packets". Internet Security Advisor October 2000: 38-42.
8. Avolio, Fred. "Ipsec and VPNs: The Sad/Glad State of Affairs". Avolio Consulting Inc., February 10, 2001. URL: <http://www.avolio.com/columns/ipsec+vpns.html> (October 25, 2002)
9. Fontanna, John. "Top Web Services Worry: Security". Network World 21 January 2002. URL: <http://www.nwfusion.com/news/2002/0121webservices.html> (October 25, 2002)



# Upcoming SANS Training

[Click here to view a list of all SANS Courses](#)

<b>SANS October Singapore 2020</b>	<b>Singapore, SG</b>	<b>Oct 12, 2020 - Oct 24, 2020</b>	<b>Live Event</b>
<b>SANS Community CTF</b>	<b>,</b>	<b>Oct 15, 2020 - Oct 16, 2020</b>	<b>Self Paced</b>
<b>SANS SEC504 Rennes 2020 (In French)</b>	<b>Rennes, FR</b>	<b>Oct 19, 2020 - Oct 24, 2020</b>	<b>Live Event</b>
<b>SANS SEC560 Lille 2020 (In French)</b>	<b>Lille, FR</b>	<b>Oct 26, 2020 - Oct 31, 2020</b>	<b>Live Event</b>
<b>SANS Tel Aviv November 2020</b>	<b>Tel Aviv, IL</b>	<b>Nov 01, 2020 - Nov 06, 2020</b>	<b>Live Event</b>
<b>SANS Sydney 2020</b>	<b>Sydney, AU</b>	<b>Nov 02, 2020 - Nov 14, 2020</b>	<b>Live Event</b>
<b>SANS Secure Thailand</b>	<b>Bangkok, TH</b>	<b>Nov 09, 2020 - Nov 14, 2020</b>	<b>Live Event</b>
<b>APAC ICS Summit &amp; Training 2020</b>	<b>Singapore, SG</b>	<b>Nov 13, 2020 - Nov 21, 2020</b>	<b>Live Event</b>
<b>SANS FOR508 Rome 2020 (in Italian)</b>	<b>Rome, IT</b>	<b>Nov 16, 2020 - Nov 21, 2020</b>	<b>Live Event</b>
<b>SANS Community CTF</b>	<b>,</b>	<b>Nov 19, 2020 - Nov 20, 2020</b>	<b>Self Paced</b>
<b>SANS Local: Oslo November 2020</b>	<b>Oslo, NO</b>	<b>Nov 23, 2020 - Nov 28, 2020</b>	<b>Live Event</b>
<b>SANS Wellington 2020</b>	<b>Wellington, NZ</b>	<b>Nov 30, 2020 - Dec 12, 2020</b>	<b>Live Event</b>
<b>SANS OnDemand</b>	<b>OnlineUS</b>	<b>Anytime</b>	<b>Self Paced</b>
<b>SANS SelfStudy</b>	<b>Books &amp; MP3s OnlyUS</b>	<b>Anytime</b>	<b>Self Paced</b>