



# Virtual Private Networks Solutions for Remote Access

## Comparison of IPSEC and SSL

*Radia is a registered trademark of Novadigm, Inc.  
SNA is a trademark of International Business Machines Corporation.  
DECnet is a trademark of the Digital Equipment Corporation.  
Windows is a registered trademark of Microsoft Corporation.  
All other marks are the property of their respective owners.  
© 2004 Schlumberger Information Solutions, Houston, Texas. All rights reserved.*

### Summary

Since their emergence a few years ago, Virtual Private Networks (VPNs) have become increasingly useful to enterprises. Initially, all remote access deployments used IP Secure (IPSEC) technology, but VPNs using the Secure Socket Layer (SSL) are now an alternative for this type of use. This white paper compares the technical and business merits of these two types of VPN and discusses the future of VPN technology.

### Introduction

Before the invention of Virtual Private Network (VPN) technology, dedicated connections were required to communicate data between an enterprise's various offices, or to connect the individual user with the company's resources:

- Leased data circuits (either actual leased lines or "private virtual circuits" that used a telecommunication company's infrastructure) connected multiple sites.
- Telephone dial-up permitted individual employees to remotely access the company's e-mail server and web sites at speeds up to 56 kbps, or 112 or 128 kbps via Integrated Services Digital Networks (ISDN).

Once most organizations got access to the Internet, it was logical to also use it to carry both types of intracompany traffic (site-to-site and remote access). By the late 1990s, company employees often had faster Internet access at home through their Internet Service Providers (ISPs) than they had to their company's network via dial-up. A remote worker often had to disconnect from a fast Digital Subscriber Line (DSL) or cable modem connection in order to dial in to his company's modem bank at a considerably slower speed.

Two issues had to be addressed before the Internet could carry site-to-site or remote access traffic: reliability and security. Reliability has improved steadily over the years: connections are rarely dropped, and the performance of DSL and cable services is usually consistent enough for remote users. However, due to the architecture of the Internet (a shared network that passes through many independent sites), security could only be achieved through some form of encryption of the data exchanged. This is where today's two VPN technologies—IPSEC and SSL—came into play.


### How IPSEC and SSL VPNs Work

Using the IPSEC protocol, IPSEC VPNs encrypt all Internet Protocol (IP) traffic over an encrypted tunnel and send it to a gateway, whose role is to decrypt the traffic and pass it to the internal network (see Figure 1). IPSEC tunnels traffic at the packet level, i.e., at the network layer of the Open Systems Interconnect (OSI) seven-layer model, and is indifferent to which higher-level protocol the packets represent (TCP, UDP, ICMP, etc.). Because an IPSEC VPN encapsulates all IP packets regardless of their function, it automatically supports all applications that communicate using IP. IPSEC usually requires the use of an installed program on the client machine to handle the encryption.



SSL VPNs also encapsulate data traffic over an encrypted tunnel to a gateway, but they do so by invoking Secure Sockets Layer (SSL) technology when communicating over an "HTTP secure" (https) web link. The receiving SSL gateway decrypts the traffic and passes it to the internal network. SSL VPNs tunnel traffic at the session layer of the OSI model, not the network layer, so by default they only support some specific IP applications—typically web access and e-mail. On the other hand, SSL support is built into web browsers and most e-mail client programs (Outlook, Eudora, etc.), therefore a separate program is not required on the user's PC to support these applications.

**IPSEC VPN**

All traffic is encrypted on a packet by packet level and sent to the gateway




Gateway decrypts packet and forwards it to original destination




**Pure SSL VPNs**


Web and mail traffic are encapsulated into an https connection





Gateway decrypts and proxies web and email access to 'real' servers




Web and mail traffic are directly encapsulated into an https connection




Select application traffic is passed to the plug-in for https encapsulation (TCP only, outbound traffic)



**SSL VPNs with Application Tunneling**



SSL gateway decrypts as usual and proxies application traffic



**Potential Benefits of SSL over IPSEC**

Because SSL VPNs can reduce or eliminate the need for IPSEC VPNs, there has been a lot of interest in them since about 2000. Some of the benefits of SSL VPNs are:

- The costs of deployment and support are lower than for IPSEC VPNs, for multiple reasons: a software client is not needed, users require less support, and troubleshooting tends to be easier.
- While first-generation SSL VPNs supported only web browsing and electronic mail, today's commercial products support many more applications, and can be extended to support almost any IP-based application.

- Because SSL VPNs impose no specific software or hardware requirements, they can be used from any computer connected to the Internet (the user's home PC, an Internet kiosk at an airport, a colleague's machine, etc.) Therefore, they are a more versatile solution for the frequent traveller or "road warrior." On the other hand, this flexibility heightens some security risks.

The following sections address each of these issues in more detail.

**Cost Comparison**

Several factors contribute to the costs of a VPN solution. When evaluating the cost of any information technology solution, the actual hardware and software investment is only the tip of the

iceberg. To accurately compare solution costs, one should employ a Total Cost of Ownership (TCO) approach.

*Hardware Costs*

Significant hardware costs exist for both IPSEC and SSL VPNs. However, since most IPSEC gateways evolved from firewalls, these gateways can be constructed from existing perimeter firewalls (for example, by enabling the included VPN feature). In contrast, SSL VPNs traditionally involve dedicated hardware and software that cannot be reused for other functions. This is the one component of the cost that is in favor of IPSEC over SSL

*Deployment Costs*

To deploy an IPSEC VPN, every machine that will access the VPN must have specific client software installed to intercept every IP packet handled by the network layer. This software encrypts outgoing packets before passing them to the next lower layer and decrypts incoming packets. Deploying VPN software to hundreds or thousands of machines can be handled by a standard PC configuration management service such as Microsoft’s SMS, Novadigm’s Radia®, etc. However, if no such solution is in place, deployment can be very costly.

In contrast, SSL VPNs have no initial deployment costs. Support of the “https” protocol is already built into standard web browsers and e-mail clients, and other application traffic tunneling is done by applets that are loaded on-demand rather than pre-installed.

*License Costs*

Earlier versions of IPSEC VPNs charged fees for the clients that were used on the remote computer. These fees have disappeared, although suppliers still charge fees for premium clients with policy enforcement and sophisticated personal firewalls.

With SSL VPNs, specific charges are not assessed for each remote machine accessing the VPN, since there is no software client to charge for. However, vendors can impose complicated licensing schemes involving the plug-ins that are used in conjunction with the SSL tunneling of application data. The impact of these schemes must be carefully analyzed.

In both cases, license costs depend on the number of advanced features deployed. An initial, bare-bones installation may appear to be free of license costs, but pressure to add system administration features or new application support may cause these costs to rise over time.

*Support Costs*

VPNs are fairly complicated to support, since they involve clients or applets on remote computers (except in minimal SSL VPN configurations), an optional security policy enforcement engine, encryption between the client computer and the gateway, and optional firewalls in front of the gateway or between the gateway and the organization’s intranet. User-level troubleshooting can be a significant cost (including the indirect cost of lost user time), especially when considering that, by definition, the user and his computer are not immediately accessible to the support technicians.

Because of the complexity of the IPSEC protocol, and because the software client on each machine represents one more possible source of misconfiguration or malfunction, IPSEC VPNs are slightly harder to troubleshoot than SSL VPNs, which use the well-known and well-understood https protocol. However, due to all the newer features that are being added to SSL VPNs, this gap is narrowing. In a Gartner Dataquest study conducted in November 2002, over 50% of the 176 respondents indicated that “difficulty in managing IP VPNs” was an inhibitor to VPN adoption [Harris 2002].

*User Training*

User inertia is always hard to overcome, and the training that users need to become comfortable with new software can be expensive. When an IPSEC VPN is deployed, the users must learn how to operate the client software. If this investment is not made, troubleshooting costs can be expected to rise. When an SSL VPN is deployed, the users only need to master the one screen used to sign on, since most users are already familiar with web browsers. However, the addition of new functionality is making the user interface for SSL VPNs more complex, so the gap in training costs is getting narrower.

The following table summarizes cost comparisons between IPSEC VPNs and SSL VPNs.

Cost	IPSEC VPN	SSL VPN
Hardware	Costly, but can re-use existing firewall hardware	Costly—separate hardware required
Initial deployment cost	Costly	Inexpensive
Licensing	Free basic clients Costly with premium (but necessary) features	Inexpensive for pure SSL tunneling Can be costly with additional application tunneling features
Support	Costly—specialized expertise needed	Inexpensive—web browser troubleshooting is often enough
Training	Costly—new software system to learn	Inexpensive—only web browser knowledge is required

**Application Support**

IPSEC VPNs encapsulate all IP packets without differentiating between the higher-level protocols they relate to (HTTP for web access, FTP for file transfer, SMTP for e-mail, remote login, etc.). By definition, they support all IP applications. This covers nearly all applications in the enterprise, since most organizations have retired applications that used non-Internet protocols like SNA™ or DECnet™.

However, IPSEC VPNs have some issues with NAT (Network Address Translation). NAT is used when an organization or an Internet Service Provider (ISP) cannot or does not want to allocate

a public IP address to each connected device. For an ISP, it may make sense because only a small percentage of its subscribers are actually on-line at any given time. Inside the organization's network, devices communicate using IP addresses from a "private" address pool. At the point where the enterprise network communicates with the outside, the NAT software in the router rewrites all packet headers to replace private return addresses with public ones (for outgoing traffic) or vice versa (for incoming traffic). This header modification interferes with IPSEC. Although most IPSEC VPN vendors support workarounds for NAT, interoperability remains a concern.

By default, SSL VPNs support all web-based applications (those using the HTTP protocol), but only those. Fortunately, most applications introduced recently for mobile workers are web-based, for several reasons:

- Less user training is required.
- Application client software does not need to be installed on the users' PCs.
- The applications are available from any computer (including a shared station) connected to the company network.
- Since these applications are not used intensively for hours on end, it matters little that they are typically slower than traditional client/server software.

However, the restriction of SSL VPNs to web applications can still be a problem, because:

- Some IP-based applications are not web applications; they may generate IP traffic under protocols other than HTTP.
- Even in a "thin client" configuration (access via a browser), some functionality may be provided by Java applets that communicate using a protocol that the SSL VPN cannot handle, causing the application to "hang" when the applet is launched.
- E-mail is critical, and most people access their e-mail via a thick client (Outlook, Eudora, etc.), which communicates with a server using POP or IMAP, not HTTP.

Three solutions can be envisioned to permit the use of legacy applications via SSL VPNs:

- **Web interfaces to legacy applications.** Once an application is accessible via the web, it can be used via an SSL VPN. This approach is effective if the number of applications to be modified is small, and they are mostly TCP-based. Moving to a web interface has other benefits, such as lowering support costs by removing thick clients, or lowering training requirements by embedding more explanatory text in the web interface.
- **Plug-ins to support specific applications.** The administrator specifies a set of applications to be supported, and the VPN supplier (or sometimes the customer) develops plug-ins to support them. When users log in to the SSL VPN gateway, or when they first click on the link to one of the concerned applications, the plug-in is downloaded to the client machine. The plug-in then transparently tunnels all network traffic from

supported applications across the VPN. It can be configured to uninstall on demand, or when the user logs out of the gateway. This solution offers excellent flexibility and high security, if the list of supported applications is small and it is possible to tell users that other applications are not supported.

- **Network shims to support arbitrary applications.** This method sounds similar to the above, but it is more general. A network shim is a very small piece of code that opens a socket and translates between network protocols and internal data formats. In the case of SSL VPNs, shims are inserted into the network stack to provide IPSEC-type functionality. This technique is appealing because of its general applicability, but it is not yet mature: it still lacks stability and enough functionality. It is also unreasonable to expect a network shim to emulate the entire functionality of an IPSEC client.

VPN vendors such as Jupiter currently support browser sessions, e-mail, terminal sessions, and also plug-ins for application tunneling. This seems sufficient for the vast majority of users and does not create excessive complexity.

### SSL VPN Selection Criteria

All SSL VPN products enable remote access to web sites, e-mail, and selected applications for which plug-ins exist or can be developed. On the other hand, these features may be discriminators for some customers:

- **Support for multiple platforms.** Most SSL VPN vendors concentrate on supporting Microsoft Windows® clients first. The enterprise may need to provide access to Macintosh home computers or to Unix or Linux workstations used by engineers working remotely (perhaps at a customer site). These are usually well supported for basic functions, such as accessing web sites via HTTPS, but application plug-ins can be platform-specific.
- **"Backwards" protocol support.** Most SSL VPNs do not support protocols in which client and server communication is reversed, such as X11. This is another concern for remote users of Unix.
- **Support for incoming network connections.** Most SSL VPNs do not support incoming network connections such as active FTP.

### Security Concerns

Most users of IPSEC VPNs employ company standard PCs, on which the "thick" IPSEC client is usually installed and configured by the corporate IT support group. In addition, corporate systems typically have a standard set of policy settings, properly configured and updated anti-virus software, a personal firewall (often bundled with the VPN client), and require strong user authentication (perhaps with a hardware token that supplies a private key or a one-time password). These standard measures provide a certain level of security and lower the risk that the client system will be compromised.

However, there is no implied security level for VPN over SSL: any computer with a browser can access the intranet if the user knows the gateway URL and has a password. There is no assurance that the machine runs a personal firewall or anti-virus software. Since the enterprise cannot control the setup of the client machine, it must address this security limitation in another way.

The following policies can compensate for providing access to uncontrolled client computers:

- **Using built-in policy primitives** that perform security checks on the client machine (such as checking for anti-virus software, or for the presence of a personal firewall). This solution will work best if most users of the VPN service will frequently use standard company-issued computers, otherwise the checks will fail too often and access will be denied.
- **Integrating third-party software**, downloaded on demand to the client machine, that performs an initial security check, then remains on the machine for the duration of the session. This software typically monitors network traffic and blocks anomalous or malicious traffic based on rule-sets and behavioral patterns. If the client environment cannot be trusted, this is an ideal solution for general use. However, it is not the equivalent of a personal firewall.
- **Restricting enterprise access** to users of trusted clients. The enterprise may consider a split scheme in which users of “standard” computers are given full access, while users of non-standard computers receive only limited access (the policy primitives mentioned above may be used to enforce this scheme).
- **Restricting enterprise access** to a small number of hosts/ports for all users of the service. This can be done by setting up packet filtering rules at the gateway. This solution is practical if users access few applications and servers, using fixed port numbers. However, it cannot be used if the intranet services accessed through the gateway are dynamic in any way (e.g. they use remote procedure calls, or dynamically assigned ports, like active FTP).
- **Requiring periodic username/password re-entry** (or the presence of a hardware authentication token) so that a user who walks away without logging out of the gateway session does not leave the door wide open to an intruder for an indefinite period, especially on a public computer.

### Maturity

A certain “convergence” is now taking place between the two types of VPNs: SSL VPNs are adding new features such as plug-ins to improve their functionality and security. This evolution allows them to provide IPSEC-type capability, but it also makes them more complex.

These add-ons may still present defects and interoperability limitations, they require more administrator training, and they cause more support work than one might expect from a technology that, after all, was supposed to simplify the SSL architecture. However, the technology is maturing rapidly and the initial issues are progressively getting ironed out.

### Recommendations

Due to the rapid evolution of the SSL VPN products, enterprises with unmet VPN needs for remote access should consider the following:

- If you are planning to introduce a remote access VPN for the first time, perform an analysis to determine which applications are used, at what frequency, and from where, by your remote users. Find out which applications are critical, and which ones your users could live without. This will determine if IPSEC is required, or if SSL is sufficient, and what application plug-ins may be required for an SSL solution.
- If you run an IPSEC VPN, ask the product supplier about support for SSL VPN using your current infrastructure. At least two vendors (Checkpoint and Cisco) offer SSL VPN support on their flagship IPSEC VPN products. This can be a cost-effective, simple way to implement both IPSEC and SSL. New remote users can then be added without installing more IPSEC clients. One may even consider migrating “mainstream” users (those without esoteric application support requirements) from IPSEC to SSL to reduce support costs, free up IPSEC client licenses, or allow access from a greater variety of computers.
- If the preceding option is not available, implement a pilot SSL VPN with appropriately configured application tunneling to determine whether this technology can replace IPSEC—at least for “mainstream” users.
- If you already run both types of VPN (IPSEC and SSL), place all upgrades of IPSEC VPN “on hold.” You may now be able to support your additional needs using the SSL product, which may lower your costs.
- In all cases, use strong authentication mechanisms, run integrity checks on executable files, and consider on-demand security for SSL VPN clients that are not trusted.

### Conclusion—SSL VPNs Now and in the Future

Most major analysts have voiced cautionary opinions about SSL VPNs in the past. However, if the precautions listed above are observed, SSL VPNs are now ready for prime time. In actual implementation projects, between 90% and 95% of an organization’s remote users only require web access and e-mail. The percentage of users whose needs are met is even higher if some “standard” applications (as defined by the enterprise) are supported through plug-ins.

There will be a continued need for IPSEC-based remote access VPNs in order to tunnel communication to unsupported applications, to serve clients on Macintosh and Linux, and for site-to-site VPNs. However, SSL VPNs offer near-equivalent functionality for most enterprise remote access users. By offering SSL as the default option to all remote users, and providing IPSEC VPN clients only to the few users who need non-standard application support, the enterprise can reduce the complexity of the overall remote access infrastructure, while enabling access from more places, including airport Internet kiosks and web-enabled wireless Personal Digital Assistants (PDAs).

Major IPSEC VPN vendors are starting to offer integrated SSL VPN functionality in their existing products at little or no initial cost to the customer. In the future, all mainstream VPN systems will offer both IPSEC and SSL functionality, making the issue of “choice” moot.

### Glossary and Abbreviations

**DECnet:** A proprietary communication protocol implemented by Digital Equipment Corporation on its family of mini-computers running the RSX-11 and VMS operating systems.

**FTP (File Transfer Protocol):** The Internet protocol for exchanging files, using TCP/IP for the lower layers.

**HTTP (Hyper Text Transfer Protocol):** The protocol used by the World Wide Web, HTTP defines how messages are transmitted, and what actions Web servers and browsers take in response to various commands.

**HTTPS (HTTP Secure):** a version of HTTP that uses SSL to encrypt the data, and also allows the user’s browser to verify that it is communicating with a legitimate web server.

**IP (Internet Protocol):** The protocol that specifies the format of packets and the addressing scheme. Most networks combine IP with a higher-level protocol called Transmission Control Protocol (TCP), which establishes a virtual connection between a destination and a source.

**IPSEC (IP Secure):** A set of protocols that support secure exchange of packets at the IP layer. The sending and receiving devices must share a secret key. IPsec supports two encryption

modes: Transport and Tunnel. Transport mode encrypts only the data portion of each packet; Tunnel mode encrypts both the header and the data.

**SMTP (Simple Mail Transfer Protocol):** A communications protocol that sends e-mail messages from one server to another. The messages can then be retrieved from a server with generally either POP or IMAP.

**SSL (Secure Socket Layer):** A protocol developed by Netscape to transmit data in encrypted form, using a public/private key pair.

**SNA (System Network Architecture):** A proprietary communication protocol implemented by IBM on its family of mainframe and minicomputers.

**TCO (Total Cost of Ownership):** A cost evaluation methodology that considers all components of the cost of an enterprise’s information technology solution, not only the initial purchase of hardware or software, but also the cost of deploying, operating and supporting the solution.

**Tunnel:** An encrypted connection that securely carries traffic across a public network.

**VPN (Virtual Private Network):** The use of a public network, such as the Internet, to carry private traffic from a remote employee to an organization’s servers, or between two sites of the same organization by encrypting it so that it would be useless to anyone who intercepts it. A VPN removes the need for a fixed communication link between the two sites, or between the employee and the organization, and can provide remote access at cable/DSL speed, instead of dial-up modem speed.

---

### References

[Ferraro 03] Ferraro, Crystal L.: “Choosing between SSL and IPSEC VPNs” (an interview of David Passmore, from Burton Group). SearchSecurity.Com, December 2003.  
[http://searchsecurity.techtarget.com/qna/0,289202,sid14\\_gci940324,00.html](http://searchsecurity.techtarget.com/qna/0,289202,sid14_gci940324,00.html)

[Girard 03] Girard, John: “SSL VPN 1H03 Magic Quadrant Evaluation Criteria.” Gartner Group research note M-19-6271. April 8, 2003. 5 pages.

[Harris 02] Harris, Mike and Eric Goodness: “Professional Service Needs for Emerging Enterprise Networks—User Wants and Needs” Gartner Group report ITNI-WW-UW-0112. November 1, 2002. 54 pages.

[ISM 03] Information Security Magazine: Special Issue on VPN. August 2003.  
[http://infosecuritymag.techtarget.com/ss/0,295804,sid6\\_iss21,00.html](http://infosecuritymag.techtarget.com/ss/0,295804,sid6_iss21,00.html)

[Netscreen 04] Netscreen Technologies: “VPN Decision Guide: IPsec or SSL VPN Decision Criteria.” February 1, 2004. 8 pages.  
[http://www.netscreen.com/dm/techpubs/downloads/wp\\_vpn\\_decision\\_criteria.pdf](http://www.netscreen.com/dm/techpubs/downloads/wp_vpn_decision_criteria.pdf)

[Slaby 03] Slaby, Jim: “Choosing between IPSEC and SSL Remote Access VPNs.” Forrester Research “Giga Planning Assumption,” July 31, 2003. 5 pages.

[Smeaton 03] Smeaton, Jo: “Remote Access Trends for a Mobile World.” Intranet Journal, August 2003.  
[http://www.intranetjournal.com/articles/200308/pij\\_08\\_18\\_03a.html](http://www.intranetjournal.com/articles/200308/pij_08_18_03a.html)

[Warden 03] Warden, Waheed: “An Introduction to SSL VPN.” Networknewz.Com, December 2003.  
<http://www.networknewz.com/networknewz-10-20031201AnIntrotoSSLVPN.html>

[Whiteley 04] Whiteley, Robert and Stan Schatt: “Making SSL VPNs a Strategic Part of Your Network.” Forrester Research “TechChoice,” March 19, 2004. 14 pages.

Note: most product suppliers also have TCO comparison documents on their web sites.

---