



NISCC

NATIONAL INFRASTRUCTURE SECURITY CO-ORDINATION CENTRE

NISCC Viewpoint 03/2006 Issued 20 April 2006

Virtual Private Networks (VPNs)

Virtual Private Networks (VPNs) are now widely recognised as a cost-effective method for protecting communications for organisations large and small. In this paper, we define what is meant by a VPN, and describe the main benefits and advantages, including some novel and less well-known applications. We list the technologies that implement VPNs, and mention the main security issues and other obstacles to consider when deploying one.

NISCC Viewpoint papers are intended to provide an overview of emerging technologies and other issues facing the IT sector. A Viewpoint will not necessarily offer mitigation advice; other NISCC products will do this.

**National Infrastructure
Security Co-ordination Centre**
PO Box 832
London, SW1P 1BG

Tel: 0870 487 0748
Fax: 0870 487 0749
Email: enquiries@nisc.gov.uk
Web: www.nisc.gov.uk

Introduction

1. Virtual Private Networks (VPNs) are now widely recognised as a cost-effective method for protecting communications for organisations large and small. In this paper, we define what is meant by a VPN, and describe the main benefits and advantages, including some novel and less well-known applications. We list the technologies that implement VPNs, and mention the main security issues and other obstacles to consider when deploying one.

Background: why VPNs?

2. The 'Internet revolution' has transformed the way people access information and do business. The success of the Internet can be put down to three main things – all-pervasiveness, cheapness, and ease of use. One of the effects of the Internet is to make people think differently about their corporate networks. People began to wonder if it might be possible to achieve the same easy access to their own corporate systems and applications that they enjoy for Internet-based applications.

3. The main technical challenge to such a vision was, and still remains, *security*. Corporate applications are invariably highly sensitive; who can access them and what they can do on them must be tightly controlled. Traditional approaches to security depended to a large extent on the fact that the corporate applications ran on a private network. In more formal terms, information confidentiality, integrity, and the authenticity of those who can access are of paramount importance. Any approach to extending the corporate network beyond its traditional physical boundaries must preserve these characteristics. The class of technology solutions that meet these requirements is referred to as the *Virtual Private Network (VPN)*.

What is a VPN?

4. A Virtual Private Network, as its name suggests, is a network whose security is maintained by virtual or logical measures on the computers connecting to it, rather than by measures implemented by the owner of the communications system - such as protected links or a private site installation. The original term was coined over a decade ago, and since then many products have been branded as VPN solutions of one type or another.

5. Most VPNs use cryptography to protect the data packets and authenticate the endpoints of communication. This is needed because the networks across which the data travels are often open public networks, such as the Internet, a PSTN (telephone system), or a broadband network. Because such networks have large populations and are not well controlled, the risk of compromise from eavesdropping or masquerade attacks is high. Therefore, high-quality cryptography and key management are required to achieve the security requirements discussed above: maintaining the confidentiality and integrity of data in transit, and authenticating the endpoints. The kind of cryptography that is required for these purposes is discussed in more detail later in this paper.

6. Cryptographic VPNs are the most common form of VPN, and are what most people would understand by the term. However certain other

approaches are sometimes marketed as VPN solutions. Techniques such as VLANs (Virtual Local Area Networks) or MPLS (Multiprotocol Label Switching) routing to separate traffic are sometimes also branded as VPN solutions. These techniques separate the traffic into different Communities of Interest (Cols) for security or performance reasons. In order to make them work, the network infrastructure must itself be managed by a trusted party. Consequently, the network's security is more real than virtual, and therefore these do not strictly constitute a VPN solution. Most of the rest of this paper concentrates on "true" VPNs where the network itself is not trusted, and the protection is implemented using cryptography.

Uses and benefits of a VPN

7. If we consider the purposes to which typical current-day organisations put their private network infrastructure, the list usually includes secure file sharing, e-mail, and intranet browsing. For some organisations, the list will also include secure voice, collaborative working, teleconferencing, or Instant Messaging.

8. A VPN allows these private communications to be extended across whatever network resources happen to be available at a particular location and time. This could be a telephone line from a hotel room, a Public Wireless Access Point, or an Internet connection from home. They can also be used to connect securely between different sites of an organisation. PCs, Laptops, PDAs, and some mobile phones can all participate in a VPN.

9. In each case, secure communication is extended beyond the geographical boundaries of the organisation. Therefore, a VPN brings great flexibility and enormous cost savings when compared to the other options for secure communication, such as using a leased line.

Usage of VPNs

10. VPNs can be used in a number of different situations, as illustrated in Figure 1 below. These include clients connecting to a server, and also for client to client (peer-to-peer) communications.

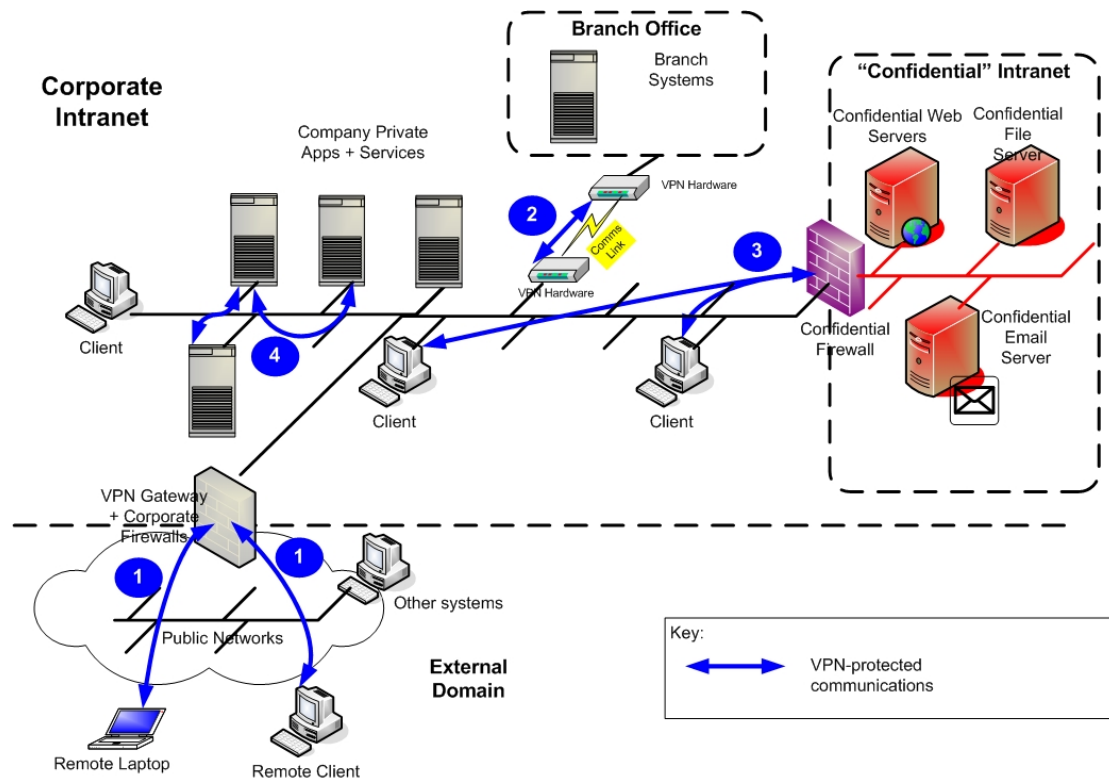


Figure 1 Some Ways to Deploy a VPN

11. Using one or other of these modes, a VPN can protect external communications, labelled as (1) and (2) in Figure 1. They can also create more trusted communication paths across the organisation's Intranet. This allows small groups to communicate at higher sensitivity levels than the intranet normally allows. These are labelled as (3) and (4) in Figure 1. These usages are discussed below.

Protecting external clients

12. Remote users can work from home or on the move, and at the same time use all the services of the corporate network that are available while on site. (See (1) in Figure 1). VPNs can protect voice, tele-conferencing or video-conferencing, interactive whiteboarding, and instant messaging as well as more usual data communications. In each case, protection is provided using the same VPN technology that protects data services. Wired or wireless connections are given the same protection. The encryption services considerably improve their security, as well as reducing cost when compared with other approaches such as arranging specific protection for each kind of communication.

13. The central termination point of a VPN is often referred to as a *VPN Gateway*. The figure shows a VPN Gateway co-located with the organisation's firewall. Another option would be to site it inside the firewall at a termination point inside the corporate boundary, such as for access to a particular set of services and applications. The placement of VPNs with respect to firewalls is discussed later in this paper.

Protecting site to site communications

14. VPNs can also be used to securely connect different parts of an organisation that are dispersed across several different sites, as a cost-effective alternative to a private leased line. This is labelled as (2) in Figure 1.

15. The VPN endpoints authenticate each other at the start of the connection and data authenticity is preserved through the lifetime of the connection by means of a shared key. This ensures that all data exchanged is authentic Intranet data, and prevents third-party traffic from being injected into the connection. As a result, unlike above, a firewall is not normally needed to screen for malicious data, because only legitimate intranet traffic is being exchanged. It is generally assumed that other measures will detect and prevent anomalous activity on the Intranet.

Protected communities within an Intranet

16. Organisations are increasingly recognising the need to partition their Intranets into separate communities of interest, and VPNs can be used to meet this requirement. This usage of VPNs is emerging as an important requirement in organisations who wish to exchange highly sensitive information such as financial data, along with normal company private communications. The need for good protection of such data is increasingly important in satisfying regulatory requirements.

17. Figure 1 shows some examples of how this can be implemented. Label (3) in the Figure shows how it can protect communication to a smaller, more sensitive Intranet segregated from the main Intranet, - a "Confidential" Intranet in Figure 1. Label (4) shows a VPN protecting sensitive peer-to-peer communications across the intranet.

VPN technology

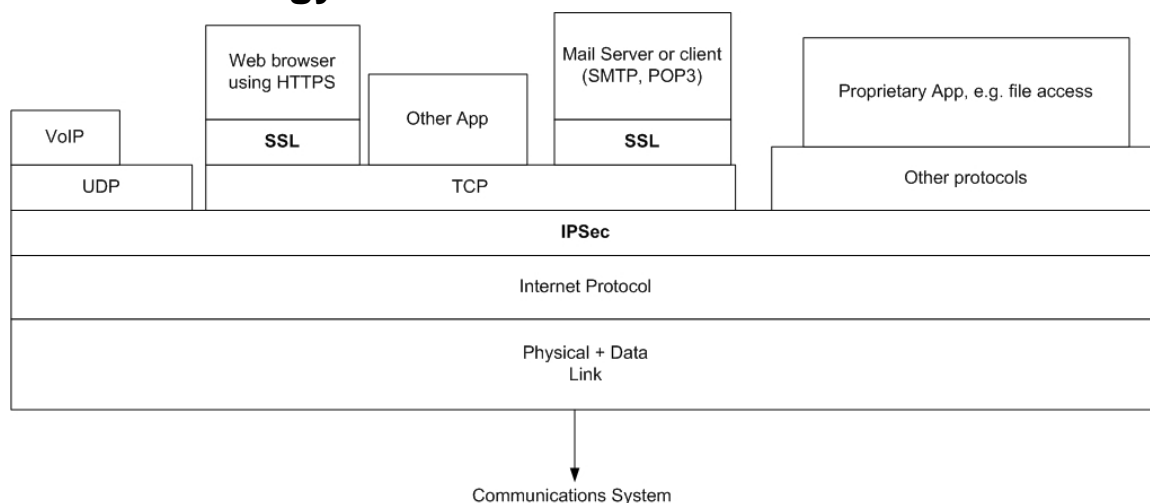


Figure 2 IPsec and SSL VPNs

18. There are two groups of protocols that lead the VPN market. These are solutions based on either IPsec (Internet Protocol Security) or SSL (Secure Sockets Layer) security protocols. The placement of IPsec and SSL VPNs in the protocol stack is illustrated in Figure 2 above¹. Both these protocols have

¹ There are also many proprietary or semi-proprietary approaches in current use. For some of these, steps have been taken to make them open standards: a prime example is SSH (Secure Shell, not to be confused with SSL), which is often used for remote terminal services, particularly to Unix applications. SSH has not achieved the market penetration of IPsec or SSL.

undergone rigorous review and analysis by a large number of security experts across the world. Therefore, we would recommend their use in preference to a proprietary approach for most commercial applications².

19. IPsec operates at the IP layer of the protocol stack and transparently protects all types of communication including TCP (Transport Connection Protocol), UDP (User Datagram Protocol) and other protocols that communicate over IP. UDP support means that applications such as VoIP (Voice over Internet Protocol) are also protected. Thus all corporate applications are immediately protected without the need for any change. As such, this represents a huge benefit of IPsec, from both cost and security perspectives.

20. SSL VPNs secure particular applications, such as communication with a Web browser, or secure e-mail. SSL VPNs can be used in what is referred to as a "clientless" mode, though this is a slightly misleading phrase since a client must still be used. What it really means is that no installation or configuration is required to set them up, whereas IPsec almost always require some steps on a client machine to be enabled. Because SSL VPNs can be clientless, they can be established from any available system, such as a public Internet terminal, or an Internet Café. This turns out to be a somewhat "mixed blessing": on one hand the solution allows great flexibility of working, but it is also susceptible to a number of vulnerabilities that a more tightly managed solution avoids. These issues are discussed further in the next section.

21. A Web browser can be used as the client end of an SSL VPN. Most corporate applications are already Web-enabled, and will therefore operate with an SSL VPN without difficulty. However in some cases, re-engineering will be required for legacy applications that do not natively support web protocols³. In addition, applications such as VoIP or other proprietary protocols cannot be easily operated over an SSL VPN. For these protocols, IPsec is a significantly easier implementation option than SSL.

VPN products

22. The technology to implement VPNs comes in a number of forms, ranging from rack-mounted hardware, through firewall add-ons, to software running on conventional computer platforms. Software solutions can be subdivided into those that come as part of normal Operating Systems and applications, and specialist third-party plug-ins. The choice of approach depends on the balance between cost, performance, and security.

23. Hardware VPN appliances generally have much better performance as well as improved security compared to software implementations. They are therefore generally recommended for the following situations:

- i. the Server end of client-server implementations;
- ii. for clients handling high security data; and
- iii. for site-to-site communications.

² UK Central Government VPN implementations are governed by slightly different rules that are not discussed in this paper.

³ This usually entails building a web front-end that translates between web protocols (http) and the legacy application access protocols.

24. Virtually all modern computer systems have the native capability to support the client side of VPN connections, including both SSL and IPsec. These come either in the form of third-party plug-ins or in many cases from the operating system itself. VPN support is now possible from most PDAs, handheld computers, and some mobile phones.

Security issues to consider

25. In this section, we highlight the main areas of potential risk that can arise from using a VPN. Because VPN endpoints may be outside the physical boundaries of the organisation, and they are connected to networks where there are unknown and potentially hostile agents, they must be considered at high risk of attack. Therefore, it is essential that a VPN deployment be carefully managed and implemented to avoid exposing an organisation to increased risk of compromise. The main issues to consider are in the following areas:

- i. end-point protection;
- ii. authentication and key management;
- iii. protocol tunnelling; and
- iv. product selection and configuration.

Endpoint protection

26. Endpoint protection is essential to the security of any VPN deployment. The impact of compromising an endpoint can be understood by looking at Figure 1. The figure shows that a VPN client can communicate simultaneously with both the other agents on the public network, and the corporate network. Therefore, it can become a route for remotely controlled attack on the organisation's network.

27. To mitigate this risk, we recommend that VPN endpoints should be properly owned and managed, either by the organisation itself or some other party with which it has a relationship. This offers significant advantages over the unmanaged endpoints of the clientless approach (see below).

28. Managing the endpoints allows additional client side security measures to be put in place to reduce the risk of remote attack. At a minimum, antivirus and patch levels should be kept up-to-date. Organisations may also wish to consider other security measures that are now gaining wide acceptance: client security monitors that gather information on patch levels, software versions, running processes, and other system health indicators are now offered by several vendors. These send this information during VPN establishment to allow the server to make a decision on whether access should be granted, or the machine quarantined until further remedial steps are undertaken. State-of-the-art client side Intrusion Detection Systems are also becoming increasingly effective.

29. A final step in securing the endpoints is to prevent secure and insecure communications directly from the client. That is, when connected to the VPN, it should not be possible to simultaneously communicate directly with the Internet from the client machine. This ensures that clients can communicate only via the organisation's intranet. As a result, remote-controlled Trojans become much harder to exploit. For communications that require the Internet,

such as Web browsing, a client can route via the organisation's firewall. This will in most cases be configured to block known viruses and other attacks and offer better protection than a client-side firewall.

Securing clientless solutions

30. A fully clientless approach where there is no organisational control or management of the endpoints represents a high-risk strategy for VPN deployment. As we discussed above, clientless solutions allow flexible working, and allows connection from 'unmanaged' endpoints, such as public Internet terminals. However, the risk from attacks such as keyboard sniffers, screen monitors, and other Trojan Horse activity is high.

31. Even so, for some situations, unmanaged endpoints may be a necessity. In these cases, steps should be taken to mitigate any security risk. The most important of these is to restrict the services that can be used by locking down the server end of the VPN. This means blocking all unnecessary remote access to machines on the intranet, and blocking access to any unneeded applications. An application content screening firewall and Intrusion Detection Services (IDS) are also recommended between the application server and the VPN Gateway.

Authentication and key management

32. Properly authenticating VPN endpoints is essential in avoiding a number of kinds of attack, including masquerade and man-in-the-middle (MITM) attacks. A masquerade is where an attacker impersonates one of the VPN endpoints and tries to connect to other machines in the VPN. If successful, a masquerade attack allows an attacker to access the applications on the machines protected by the VPN: this could include the stealing or modifying of stored data or submitting bogus transactions to an application. A man-in-the-middle attack is where an attacking machine inserts itself into the communication path between two legitimate VPN machines. A successful MITM attack allows the attacker to eavesdrop on the communication between the two machines. It also allows an attacker to modify transactions as they proceed, or insert additional bogus transactions.

33. Defeating these attacks requires a good strategy for the authentication of the endpoints of the VPN. Therefore, authentication and key management are an important part of a VPN deployment, and the various approaches will now be discussed.

34. Password-based authentication is a rather weak approach to identifying clients, and as such is not a strong defence against masquerade or MITM attacks. This is because passwords are easily stolen in a number of ways, including eavesdropping on the network, setting up a bogus site for a 'phishing' attack, a client keyboard sniffer, or social engineering attack.

35. Number generation tokens are a significant improvement over passwords. The user types in the number shown on the token at the start of each session, which changes frequently to prevent subsequent use by an attacker. However, there are some MITM attacks that can defeat number generation tokens. These involve intercepting communications at the start of a session, such as by setting up a bogus site that looks identical to the organisation's VPN portal or by direct interception of communications. Thereafter, the

attacker passes communications between the client and the true VPN server, allowing him to eavesdrop or change the communication as it proceeds.

36. Digital certificate based authentication represents a further improvement in security and is the best available method of authenticating VPN end points known today. A digital certificate allows the endpoints to be authenticated at the start of a session, and also supports the establishment of keys which preserve authenticity and confidentiality throughout the session.

37. Nevertheless, certificate-based communication is not impervious to attack, particularly when SSL is being used. Attacks are made possible by the fact that some browsers prompt the user for a decision if an invalid certificate is detected, such as one containing the wrong name. In most cases the user is likely to click 'Yes' to proceed, allowing a MITM attack to be enacted. More sophisticated attacks are also possible, such as the installation of a bogus 'trust root' certificate in the client browser which then validates the attacker's certificate. A full description of how these work is beyond the scope of this report.

38. Note that SSL VPNs implemented as a separate appliance, and IPsec VPNs in general, do not suffer these problems, since they generally have less interaction with the user and use a simpler approach to certificate validation. These points make them more resilient to attack such as Man-in-the-Middle.

39. Final points to note with respect to authentication are that private keys should be stored in a suitably secure device, and the certificates issued from a suitably trusted Certification Authority (CA).

Protocol tunnelling risks

40. Protocol tunnelling is covered extensively in NISCC Viewpoint 04/06. That document discusses what is meant by the term, and lists the issues and risks that they can raise. One of the chief issues discussed is how protocol tunnels can cause security compromise: this could be either the leakage of sensitive data out through the tunnel, or the introduction of malicious code. VPNs can actually make these problems worse, since they encrypt the data, and add a kind of "steel cladding" to the pipe.

41. The main way to address this issue is to terminate VPNs at network boundaries. Doing this allows screening software to inspect and audit communications as they pass between networks. Terminating the VPN at the network boundary also has an additional benefit: namely, it overcomes the technical issue that Network Address Translation can cause for IPsec VPNs. This point is further discussed at Annex A of this document.

42. Some applications require data to be protected by encryption at all times when it is in transit. For these situations, it is often possible to break the VPN into two parts, terminating one part at the firewall, and re-forming it on the other side of the boundary. This allows the data to be inspected in its plaintext (i.e. decrypted) form at the firewall while maintaining protection over networks. Certain products in the market implement such a solution for SSL

VPNs⁴. The accompanying paper on protocol tunnelling discusses these issues in greater detail.

Product selection and configuration

43. VPNs are only as good as the products that implement them. Choose a product with good security "credentials": this means either a product which has undergone some formal evaluation, or a market leading solution with a good track record.

44. Secondly, ensure the product is properly configured, particularly with respect to choice of algorithm suite and key management. For commercial applications, recommended algorithms are Triple-DES or AES for the data encryption algorithm, SHA-1 for the data integrity algorithm, and DSA/DSS, EC or RSA for signature and key exchange algorithms. Also, ensure the product is configured *not* to fall back to a less secure suite in the protocol negotiation.

45. Thirdly, product testing is strongly advised, ideally with the help of specialist penetration testers. Testing needs to be done prior to initial installation, and repeated periodically throughout the lifetime of the VPN. This gives confidence that it is resistant to the latest attacks, and also to account for any changes to end systems that may have been made since it was last tested.

46. Finally, any VPN will inevitably be a trade-off between security, cost, and functionality. This paper has attempted to describe the main areas to be considered when choosing the right approach for a particular deployment.

⁴ It should be noted that such an approach raises policy problems for selected applications such as electronic banking. For these applications, accessing them through what is effectively a MITM is likely to violate the terms and conditions under which they can be used.

Annex A: Network address translation and IPsec

47. An often-cited issue of IPsec is the complications that arise when it is used with NAT (Network Address Translation) firewalls. Address translation is the replacement of an address in a protocol header with a mapped address on the other side of the boundary. This is useful for organisations who wish to use a private address space on their intranet, which is invisible beyond the firewall. This makes address management easier, and is better for security. In order to make NAT work correctly, various checksums in the data packets must be recalculated. However, IPsec renders this impossible since the packets are encrypted.

48. There are several ways of overcoming the problem of using NAT with IPsec. The simplest is to terminate the VPN at the NAT firewall. Then the firewall can perform the necessary address translation on the decrypted packets. Many firewall vendors offer a VPN plug-in which can be used in precisely this way. (This overcomes many of the risks associated with protocol tunnelling, as discussed elsewhere in this paper.)

49. In some cases, organisations may wish to extend the VPN through the firewall and terminate it further inside the intranet – such as at the “confidential” Intranet boundary (see Figure 1). Therefore the simple approach of terminating the VPN at the firewall will not work. To meet this requirement, a more sophisticated approach is to use some extensions that have been recently added to the IPsec protocol standard. ([RFC 3947](#) and [RFC 3948](#) describe the extensions.) These allow detection of NAT devices during VPN establishment, and make the necessary adjustments for the rest of the session. Note that only certain recent IPsec implementations support this mode of working.

Annex B: VPNs and performance

50. In some circumstances, a VPN may not perform fast enough to support networked applications that use the high bandwidth available on a typical intranet. This may be a particularly serious problem for site-to-site communication where large volumes of data are exchanged. Client applications may also perform unacceptably slowly across a VPN. These may put serious strain on the capacity of the connection and cause unacceptable delays.

51. There are two main reasons why this may occur

- i. The VPN appliance is slowing down the data transfer, particularly if encryption is being performed in software.
- ii. The network itself being used for data transfer has narrow bandwidth.

52. Some simple diagnostic testing can reveal the cause of the bottleneck. If the network itself is the problem, then the solution is clearly to increase the network bandwidth. This could mean migrating from dial-up to broadband networking or using a higher capacity broadband link. In addition, there are several other measures that can help improve the performance of a VPN. The following measures are most applicable for client-to-site connectivity.

- Where possible, use special interfaces for remote users that offer more limited functionality than for locally connected users. A good example is certain e-mail servers that allow messages to be downloaded up to a maximum size or download messages without attachments.
- Block unnecessary "heavyweight" applications such as those offering combined e-mail, scheduling, and calendars, or those that perform file and database synchronisation.
- Train users to use applications efficiently. For example, e-mail and files can be synchronised between client and server while on site, which means that only changes to folders need be synchronised remotely.
- Use special-purpose hardware appliances for the VPN endpoints, such as a security accelerator to improve the performance of the cryptography.

53. The following measures are most applicable for improving the performance of site-to-site connectivity

- Use protocol concentrators at the endpoints to compress data and make more efficient use of the bandwidth.
- Use file caching servers to store temporary copies of information such as frequently accessed web pages.
- Use several VPN connections in parallel.
- Construct the network to ensure the VPN is used as efficiently as possible. For example, replicating e-mail and Web servers so that most clients connect to the local server rather than across the VPN.

Annex C: Glossary of Terms

AES	Advanced Encryption Standard
Col	Community of Interest
DSA/DSS	Digital Signature Algorithm/Digital Signature Standard
EC	Elliptic Curve
HTTPS	HyperText Transfer Protocol over SSL
IDS	Intrusion Detection Service
IPsec	Internet Protocol Security
MITM	Man in the Middle
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
PDA	Personal Digital Assistant
POP3	Post Office Protocol version 3
PSTN	Public Switched Telephone Network
RSA	Rivest-Shamir-Adleman
SHA-1	Secure Hash Algorithm 1
SMTP	Simple Mail Transfer Protocol
SSH	Secure SHell
SSL	Secure Sockets Layer
TCP	Transport Connection Protocol
Triple-DES	Triple – Data Encryption Standard
UDP	User Datagram Protocol
VoIP	Voice over Internet Protocol
VLAN	Virtual Local Area Network
VPN	Virtual Private Network

This paper was produced for NISCC by QinetiQ

About NISCC

The role of NISCC is to minimise the risk to the Critical National Infrastructure from electronic attack. NISCC was set up in 1999 and is an interdepartmental

centre drawing on contributions from across government. Defence, Central Government Policy, Trade, the Intelligence Agencies and Law Enforcement all contribute expertise and effort. Further information can be found at www.niscc.gov.uk



About QinetiQ

QinetiQ is one of the world's leading defence technology and security companies.

For further information please call us on 08700 100 942 or refer to our website: www.QinetiQ.com

