



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 (Security 401)"
at <http://www.giac.org/registration/gsec>

Abstract

This paper identifies some of the key terms and components of a VPN. The topics covered are IPSEC, Internet Key Exchange, Security Association, Data Encryption Standard, Triple DES, Rijndael, Diffie-Hellman, Message Digest 5, Secure Hash Algorithm-1, Hash Message Authentication Codes, Rivest-Shamir-Adelman signatures, Certificate Authorities, Layer Two Tunneling Protocol, and Point-to-Point Tunneling Protocol. The advantages and disadvantages of each component are discussed and compared. The paper concludes with recommended technologies for remote site Internet VPNs, remote user Internet VPNs, intranet VPNs, and extranet VPNs.

Introduction

Virtual Private Networks or VPNs seem to be buzz words lately in the world of IT security. A VPN is designed to provide secure access to resources behind a perimeter firewall from outside of the firewall. Designing and implementing a secure VPN can be a daunting task, because there are several different technologies and encryption options available.

Basic Components of a VPN

Virtual Private Networks (VPNs) can be created by combining several different security components. The combination chosen depends on the specific requirements and goals of the implementation. The following is a list of some of the most common VPN components: IP Security Protocol (IPSec), Internet Key Exchange (IKE), Security Association (SA), Data Encryption Standard (DES), Triple DES (3DES), Rijndael, Diffie-Hellman (D-H), Message Digest 5 (MD5), Secure Hash Algorithm-1 (SHA-1), Hash Message Authentication Codes (HMAC), Rivest-Shamir-Adelman signatures (RSA), Certificate Authorities (CA), Layer Two Tunneling Protocol (L2TP), and Point-to-Point Tunneling Protocol (PPTP).

IPSec

IPSec is the framework for secure communications using IPv4 and IPv6. The services offered by IPSec are access control, connectionless integrity, data origin authentication, protection against replays, confidentiality, and limited traffic flow confidentiality (Security Architecture, 2). These services are integrated into IP and upper layers of the Open Systems Interconnection (OSI) model. IPSec

utilizes Authentication Header (AH) and Encapsulating Security Payload (ESP) to protect traffic flows.

AH provides authentication for the entire packet by creating a digital signature of the entire packet, except for fields that change while traversing the network, such as checksum and time to live. AH was assigned IP Protocol number 51. AH consists of a six fields: 8-bits for the Next Header, 8-bits for the Payload Length, 16-bits for the Reserved field, 32-bits for the Security Parameter Index, 32-bits for the Sequence Number, and a variable length field for the Authentication Data.

- The 8-bit Next Header field is used to identify the protocol number of the payload directly following the Authentication Header. The protocol number can be any 8-bit number defined by the Internet Assigned Numbers Authority (IANA).
- The next field, Payload Length, identifies the length of the AH in 32-bit words minus "2".
- The 16-bit reserved field is saved for future use. This field must be set to zero.
- The Security Parameter Index (SPI) is a 32-bit arbitrary value that is used in combination with the destination IP address and Authentication Header to uniquely identify the Security Association for the datagram. SPI values of 1-255 are reserved by IANA for future use. The SPI value of zero is reserved for local use and should never appear on the wire.
- The 32-bit Sequence Number field contains a monotonically increasing counter. This mandatory counter provides anti-replay protection. While this field must always be present, the receiver has the discretion of whether or not to process the sequence number. When a SA is established, the sender's and receiver's counters are both set to zero. The first packet sent has a sequence number of one.
- The last field in the AH is the Authentication Data field. This field has a variable length that contains the Integrity Check Value (ICV) for the packet. This field is always a multiple of 32-bits in length. The ICV is computed over the IP header fields that do not change in transit, the Authentication Header, and the upper level protocol data (IP Authentication, 1-7).

ESP provides data confidentiality and optional authentication. ESP was assigned IP Protocol number 50. The ESP header directly follows the IP Header. The ESP header is made up of seven separate fields. The fields are the Security Parameters Index, the Sequence Number, the Payload Data, the Padding for Encryption, the Pad Length, the Next Header, and the Authentication Data.

- The Security Parameters Index value used in ESP follows the same rules as the SPI in the Authentication Header.

- The Sequence Number field also follows the same guidelines as the Sequence Number in the Authentication Header.
- The Payload Data field is a variable length containing the data described by the Next Header field. The Payload Data field must be an integral number of bytes in length.
- The Padding for Encryption field is an optional field used to add padding to plaintext for encryption algorithms that require the plaintext to be a multiple of some number of bytes. The Padding may also be used to ensure that the Pad Length and Next Header fields are right aligned within a 4-byte word. The sender may add from 0-255 bytes of padding.
- The Pad Length field is a mandatory field that specifies the number of pad bytes immediately preceding the Pad Length field. The Pad Length value can be between 0 and 255 inclusive.
- The Next Header field is an 8-bit field that identifies the type of data contained within the Payload Data field. The value of this field can be any IANA assigned protocol number.
- The Authentication Data field is a variable-length field that contains the Integrity Check Value. The ICV is computed using the ESP packet minus the Authentication Data. The authentication function selected determines the length of this field. The authentication field is optional, and its presence is determined by the Security Association (IP Encapsulating, 1-7).

AH and ESP can be used by themselves or in combination with one another. If they are used in combination, ESP is encapsulated within AH.

IKE

Internet key exchange is a hybrid of two protocols: ISAKMP and Oakley. IKE is used to authenticate IPSec peers, negotiate IKE and IPSec Security Associations, and to establish keys for encryption (Chapman, 198). IKE exchanges take place in two phases. Phase 1 establishes a secure authenticated channel, and phase 2 negotiates IPSec policies.

Phase 1 exchanges can take place in two different modes “Main Mode” or “Aggressive Mode”. Both modes generate keying material from a Diffie-Hellman exchange. Phase 1 also establishes the following ISAKMP Security Association information: encryption algorithm, hash algorithm, authentication method, and which Diffie-Hellman group to use.

The first two messages in “Main Mode” negotiate policy. The next two messages are used to exchange Diffie-Hellman public values. The last two messages authenticate the Diffie-Hellman exchange.

“Aggressive Mode” performs the same functionality as “Main Mode”, but it does so with fewer messages. The first two messages in “Aggressive Mode” negotiate policy and also exchange Diffie-Hellman public values. The second message also authenticates the responder. The third message authenticates the initiator and provides proof of the exchange.

“Main Mode” and “Aggressive Mode” may be authenticated by digital signatures, two forms of authentication with public key encryption, or pre-shared keys.

“Quick Mode” is used to accomplish a phase 2 exchange. “Quick Mode” is not considered a complete exchange because it is bound to a phase 1 exchange. The ISAKMP SA protects “Quick Mode” exchanges. “Quick Mode” negotiates non-ISAKMP SAs such as IPsec SAs and exchanges nonces to provide replay protection (Internet Key, 4-9).

SA

A Security Association is a database of available services between IPsec peers. Each peer maintains its own SA database. SAs are uniquely defined by peer address, security protocol (AH or ESP), and a 32-bit Security Parameter Index. RFC2401 describes how SAs process IPsec traffic.

The following SAD fields are used in doing IPsec processing:

- Sequence Number Counter: a 32-bit value used to generate the Sequence Number field in AH or ESP headers.
[REQUIRED for all implementations, but used only for outbound traffic.]
- Sequence Counter Overflow: a flag indicating whether overflow of the Sequence Number Counter should generate an auditable event and prevent transmission of additional packets on the SA.
[REQUIRED for all implementations, but used only for outbound traffic.]
- Anti-Replay Window: a 32-bit counter and a bit-map (or equivalent) used to determine whether an inbound AH or ESP packet is a replay.
[REQUIRED for all implementations but used only for inbound traffic.
NOTE: If anti-replay has been disabled by the receiver, e.g., in the case of a manually keyed SA, then the Anti-Replay Window is not used.]
- AH Authentication algorithm keys, etc. [REQUIRED for AH implementations]
- ESP Encryption algorithm, keys, IV mode, IV, etc. [REQUIRED for ESP implementations]
- ESP authentication algorithm, keys, etc. If the authentication service is not selected, this field will be null. [REQUIRED for ESP implementations]

- Lifetime of this Security Association: a time interval after which an SA must be replaced with a new SA (and new SPI) or terminated, plus an indication of which of these actions should occur. This may be expressed as a time or byte count, or a simultaneous use of both, the first lifetime to expire taking precedence. A compliant implementation MUST support both types of lifetimes, and must support a simultaneous use of both. If time is employed, and if IKE employs X.509 certificates for SA establishment, the SA lifetime must be constrained by the validity intervals of the certificates, and the NextIssueDate of the CRLs used in the IKE exchange for the SA. Both initiator and responder are responsible for constraining SA lifetime in this fashion. [REQUIRED for all implementations]
- IPsec protocol mode: tunnel, transport or wildcard. Indicates which mode of AH or ESP is applied to traffic on this SA. Note that if this field is “wildcard” at the sending end of the SA, then the application has to specify the mode to the IPsec implementation. This use of wildcard allows the same SA to be used for either tunnel or transport mode traffic on a per packet basis, e.g., by different sockets. The receiver does not need to know the mode in order to properly process the packet’s IPsec headers. [REQUIRED as follows, unless implicitly defined by context:
 - host implementations must support all modes
 - gateway implementations must support tunnel mode] (22-23)

Security Associations are simplex in nature; meaning to secure traffic in both directions two SAs must be created (“Security Architecture for the Internet Protocol” 21-22). There are two types of SAs: transport mode and tunnel mode.

Tunnel mode is used whenever either endpoint of a SA is a security gateway. In tunnel mode, there is an outer IP header that specifies the address of the security gateway and an inner IP header that specifies the actual destination of the packet. Within tunnel mode, there are two options: split tunneling or tunnel everything. Split tunneling is an option, in which only packets bound for the corporate LAN are protected. This option is less secure because it is possible for an attacker to connect to a split tunnel machine and stay connected to that machine while it accesses the corporate VPN. The disadvantage to tunneling everything is that it requires additional overhead for traffic that does not need to be encrypted.

Transport mode is used when the stations that are talking IPsec are the actual destinations. In this situation, the original IP header is the destination and there is no need for an inner IP header.

DES

The data encryption standard was first published in 1977. DES uses a 64-bit key, of which 56-bits are randomly generated, to encrypt data. The 8-bits of the key that are not randomly generated are used for error detection. These bits are set to create odd parity within each 8-bit byte (Data Encryption). DES is a symmetrical encryption algorithm, meaning that data encrypted with a unique key can only be decrypted with the exact same key. DES can be used in combination with ESP and IKE to provide encryption. With recent advances in computer processing power, DES is no longer considered a strong encryption algorithm, because it can be cracked within a short period of time.

3DES

Triple DES is based on the DES algorithm, but instead of performing a single iteration, 3DES performs three separate iterations with three separate DES keys. This creates a total key length of 168-bits (Data Encryption). 3DES is considerably more secure than DES.

Rijndael

Rijndael is a block cipher based off the 128-bit block cipher Square. Rijndael can be configured to use key lengths of 128, 192, or 256 bits (block cipher). Rijndael uses symmetric keys to cipher and decipher information. The Advanced Encryption Standard (AES), published by the National Institute of Standards and Technology (NIST), identifies Rijndael as the standard cipher (Announcing).

Diffie-Hellman

Diffie-Hellman is a public-key protocol that provides a method for two network peers to establish a shared secret key over a public channel. Each peer creates its own public and private key to secure the exchange. The peers then exchange public keys. A mathematical function is performed using the remote side's public key and the local private key to generate asymmetrical keys. The asymmetrical keys are then used to exchange symmetric keys. Symmetric keys are needed, because bulk encryption with symmetric keys is much faster than with asymmetric keys (Diffie-Hellman).

MD5

Message digest version 5 is a hash algorithm that is used to authenticate packet data. A hash is a one-way operation that takes a message of arbitrary length and creates a 128-bit message digest. The hash is then transmitted with the packet. The receiving peer performs its own hash on the received messages

and compares the two hashes to verify nothing as changed in transit. If even one bit is changed, the two hashes will not be identical and the packet will be discarded. The likely hood of two different messages producing the same message digest is on the order of 2^{64} operations (MD5, 1-6).

SHA-1

Secure hash algorithm is similar to MD5 in that they both perform a one-way hash on the data. SHA-1 is considered more secure than MD5 because it produces a 160-bit message digest instead of the 128-bit message that MD5 produces. The SHA-1 algorithm processes blocks of 512-bytes, this may require some messages to be padded so the entire message length is a multiple of 512 (SECURE HASH).

HMAC

Hash message authentication codes are a method for authenticating the integrity of a message based on a cryptographic hash function. HMAC can be used with any iterated cryptographic hash function, such as MD5 or SHA-1 (HMAC, 1).

RSA Keys

Rivest-Shamir-Adelman keys are a public-key cryptographic system. In a public-key system each peer has its own public and private key. The public-keys can be shared with anyone, but each peer protects its private key. RSA keys can be used for encryption and decryption or signature and verification.

- Encryption is performed by using the peer's public-key to encrypt data. The peer then receives the data encrypted with its own public key and uses its private key to decrypt the packet. This method of cryptography is often known as asymmetrical encryption.
- A signature is created with a party's private key. The receiver then uses the sender's public key to verify the signature (PKCS).

CA

A certificate authority provides a highly scalable method for peers to authenticate each other. Each peer requests a digital certificate from the certificate authority. The certificate authority is responsible for guaranteeing that the party that requested the certificate is really who he or she claims to be. The peers then exchange these certificates to prove their identity to each other (Certificate). For additional security, CAs can advertise a Certificate Revocation List (CRL). CRLs are used to inform machines of issued certificates that are no longer considered secure. One example of when a certificate should be revoked is in the instance

of a laptop theft. It is no longer desirable to allow that laptop to access the corporate LAN, so the CA administrator can simply revoke the laptop's certificate.

L2TP

Layer Two Tunneling Protocol (L2TP) was originally developed for use over Point-to-Point Protocol (PPP) connections. It provides a method for PPP and layer 2 endpoints to reside on different devices that are connected by a packet-switched network (Layer, 3). L2TP can be combined with IPsec to provide a Remote User VPN. L2TP uses user-level authentication and IPsec for computer-level authentication and encryption. Most VPN vendors that support L2TP will support the following user-level authentication protocols:

- Password Authentication Protocol (PAP): PAP sends the username and password requested by the authentication server in clear text. PAP is not considered a secure authentication protocol because it does not protect against password sniffing, replay attacks, remote client impersonation, or remote server impersonation.
- Challenge-Handshake Authentication Protocol (CHAP): CHAP requires three messages to be exchanged between the client and the authentication server. The first message sent by the authentication server is composed of a CHAP Challenge message containing a session ID and an arbitrary challenge string. The client then returns a CHAP Response message containing the username in plain text and a MD5 hash of the challenge string, session ID, and the password. The authentication server then performs the same hash and compares it to the value sent by the client. If the hashes match, the authentication server sends a CHAP Success message. If the hashes don't match, the server sends a CHAP Failure message.
- Microsoft Challenge Handshake Authentication Protocol (MS-CHAP): MS-CHAP is an encrypted authentication protocol largely based on CHAP. Like CHAP, an MS-CHAP authentication server sends a challenge to the remote client. The client responds with a username and a MD4 hash of the challenge string, session ID, and the password. MS-CHAP is more secure than CHAP because the password is never sent in clear text. MS-CHAP is still vulnerable to server impersonation though.
- MS-CHAP version 2 (MS-CHAP v2): MS-CHAP v2 provides the added protection of mutual authentication. The server verifies the identity of the client, and the client verifies the identity of the server. This process requires three messages to be exchanged between the client and server. The first message, from the server, consists of a MS-CHAP v2 Challenge message, a session identifier, and an arbitrary challenge string. The client then sends an MS-CHAP v2 response that contains the username, an arbitrary peer challenge string, and a SHA hash of the received challenge string, the peer challenge string, the session identifier, and the MD4-

hashed version of the user's password. Finally, the server checks the response and sends back a MS-CHAP v2 response composed of an indication of success or failure of the connection and an authentication response based on the sent challenge string, the peer challenge string, the client's encrypted response, and the user's password. The client verifies the response and if everything is correct, the client uses the connection.

- Extensible Authentication Protocol-Message Digest 5 (EAP-MD5): EAP-MD5 uses the CHAP authentication method within the EAP framework. The server sends an EAP-Request message requesting the client ID. The client responds in the form of an EAP-Response message with its user ID. The server then sends an EAP-Request message containing the MD5 challenge string. The client responds to this message with an MD5 hash of its user ID and password. If the response is correct the server sends a Success message to the client.
- Extensible Authentication Protocol-Transport Layer Security (EAP-TLS): EAP-TLS utilizes certificates to validate peers. EAP-TLS also provides data integrity and data confidentiality services (Microsoft Windows, 307-313).

User certificates or smart cards are considered the most secure methods for user authentication, but if those are not available MS-CHAP v2 should be used. MS-CHAP v2 is more secure because both the user and the VPN server are authenticated.

PPTP

Point-to-Point Tunneling Protocol (PPTP) allows the Point-to-Point Protocol to be tunneled through an IP network (Point-to-Point, 1). When incorporated into a VPN PPTP uses user-level authentication and Microsoft Point-to-Point Encryption (MPPE). PPTP user-level authentication may use any of the following protocols:

- MS-CHAP
- MS-CHAP v2
- EAP-TLS

PPTP is limited to these three user authentication protocols because MPPE requires the authentication protocol to generate an encryption key (Virtual).

VPN Types

Several different types of VPNs can be created with the previously mentioned components. Different types of VPNs are needed to meet differing security needs. The four basic types of VPNs that will be discussed are: Remote Site Internet VPNs, Remote User Internet VPNs, Intranet VPNs, and Extranet VPNs.

Remote Site Internet VPNs

Internet VPNs are often used to connect remote offices to a corporate headquarters LAN. This type of VPN often utilizes hardware devices to perform the tunnel creation and encryption. The secure connection between sites is established through two security gateways, one located at each site. Because the two sites appear to be logically connected this configuration is often called a site-to-site VPN. In this configuration, the secure tunnel is transparent to the end users. Meaning that there is no password or special software required to access the VPN.

Remote sites connected by hardware device are easy to configure, because IPSec only needs to be configured on two devices. A remote site VPN is also a lot cheaper to implement than dedicated lines connecting sites. This configuration does introduce some security concerns though. If a hacker compromises a computer at a remote site, he or she then has access to the corporate LAN.

There are several devices that can be used as a security gateway in a site-to-site VPN connection. Any device that can communicate using the IPSec standard and supports tunnel mode should be able to be used as a site-to-site endpoint. Routers, firewalls, or stand-alone VPN devices are the most common endpoints, but there are other options.

A site-to-site configuration should implement IPSec, IKE and ESP. If the number of remote sites is minimal pre-shared keys can be used for IKE authentication, with each connection pair having its own unique pre-shared key. If there are a large number of sites that will be connected, it is more practical to use IKE with RSA signatures for authentication. ESP should be used in combination with IKE to provide encryption. At a minimum DES encryption should be used, with 3DES or Rijndael being the preferred method.

Remote User Internet VPNs

Remote user Internet VPNs provide similar connectivity to a remote site Internet VPNs, but they do so in a slightly different manner. In a remote user VPN the user's computer itself is one of the endpoints of the connection. This allows users to be more mobile. Remote user VPNs often require that additional software be installed on the user's computer. Although, some operating systems have built in support for VPN connectivity.

Remote user VPNs can be used to reduce or completely eliminate the need for dial-in systems. Instead of dialing up to the corporate LAN, users can just

access a local ISP and connect to the corporate LAN through the VPN. This can save companies a lot of money in long distance phone charges. It also can provide faster connection speeds for remote users. Remote user VPNs are often require more time to configure, because every computer that is going to connect to the VPN needs to be configured.

Remote VPNs can be implemented using a variety of different devices. The most common are servers, routers, firewalls, or stand-alone VPN devices. These devices are commonly configured to offer one or a combination of the following services: IPSec, PPTP, or L2TP

Intranet VPNs

Intranet VPNs are becoming more and more common. Intranet VPNs can be used in a number of different situations where a secure channel is required. Some examples include: client-to-server communications, server-to-server communication, and wireless communications.

- A client-to-server VPN can be used to prevent network sniffing of confidential information such as: e-mails, personnel information, payroll information, or research.
- Server-to-server VPNs can be used to connect front-end servers in the DMZ to backend servers that reside on the protected network. They can also be used to protect confidential communications between servers such as DNS zone transfers.
- Wireless VPNs are gaining a lot of popularity because corporations want the mobility of wireless, but require more security than what is built into the 802.11 standards. In a wireless VPN, the wireless clients establish an encrypted tunnel with a security gateway. Even if someone is able to capture the wireless packets, they are unable to decipher the contents because it is encrypted.

Intranet VPNs are usually configured to talk pure IPSec in the case of server-to-server VPNs and L2TP or PPTP in the case of client-to-server VPNs or wireless VPNs.

Extranet VPNs

An extranet allows a company to share its local network resources with suppliers, vendors, customers, partners, or other businesses. Extranets are especially helpful for sharing large amounts of data, collaborating on joint ventures, and for sharing exclusive knowledge that does not need to be made available to the general public.

An extranet has basically the same configuration as a remote access VPN. It can be configured as a site-to-site connection or as a remote user VPN depending on what the needs of the company dictate.

Conclusion

Virtual Private Networks can be implemented with a number of different technologies and several different configurations. Each technology or configuration has its own key benefits and possible disadvantages. What works in one area may not be ideal in another. The key to success is a thorough understanding of the underlying technology that makes VPNs possible.

© SANS Institute 2000 - 2002, Author retains full rights.

References

- “The block cipher Rijndael.” 12 Sept. 2002
<<http://www.esat.kuleuven.ac.be/~rijmen/rijndael/>>.
- Chapman, David W. and Andy Fox, eds. Cisco Secure PIX Firewalls.
Indianapolis: Cisco Press, 2002.
- Cisco Systems, “Certificate Authority Support for IPsec Overview.” 10 Aug. 2002
<http://www.cisco.com/warp/public/cc/so/neso/sqso/eqso/821_pp.pdf>.
- “Diffie-Hellman Key Exchange – A Non-Mathematician’s Explanation.” 9 Aug.
2002
<http://networking.earthweb.com/netsecur/article/0,,12340_624441,00.html>.
- The Internet Engineering Task Force. “HMAC: Keyed-Hashing for Message
Authentication.” 9 Aug. 2002 <<http://www.ietf.org/rfc/rfc2104.txt>>.
- The Internet Engineering Task Force. “The Internet Key Exchange (IKE).” 9 Aug.
2002 <<http://www.ietf.org/rfc/rfc2409.txt>>.
- The Internet Engineering Task Force. “IP Authentication Header.” 7 Aug. 2002
<<http://www.ietf.org/rfc/rfc2402.txt>>.
- The Internet Engineering Task Force. “IP Encapsulating Security Payload
(ESP).” 7 Aug. 2002 <<http://www.ietf.org/rfc/rfc2406.txt>>.
- The Internet Engineering Task Force. “Layer Two Tunneling Protocol ‘L2TP’.” 10
Aug. 2002 <<http://www.ietf.org/rfc/rfc2661.txt>>.
- The Internet Engineering Task Force. “The MD5 Message-Digest Algorithm.” 9
Aug. 2002 <<http://www.ietf.org/rfc/rfc1321.txt>>.
- The Internet Engineering Task Force. “Point-to-Point Tunneling Protocol
(PPTP).” 10 Aug. 2002 <<http://www.ietf.org/rfc/rfc2637.txt?number=2637>>
- The Internet Engineering Task Force. “Security Architecture for the Internet
Protocol.” 4 Aug. 2002 <<http://www.ietf.org/rfc/rfc2401.txt>>.
- Microsoft Corporation. Microsoft Windows 2000 Server Internetworking Guide.
Redmond: Microsoft, 2000.

Microsoft Corporation. "Virtual Private Networking with Windows 2000: Deploying Remote Access VPNs." 12 Aug. 2002
<<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/network/deploy/depovg/vpndeploy.asp>>.

The National Institute of Standards and Technology. "Announcing the ADVANCED ENCRYPTION STANDARD (AES)." 12 Sept. 2002
<<http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>>

The National Institute of Standards and Technology. "Data Encryption Standard (DES)." 9 Aug. 2002 <<http://csrc.nist.gov/publications/fips/fips46-3.pdf>>.

The National Institute of Standards and Technology. "Secure Hash Standard." 9 Aug. 2002 <<http://www.itl.nist.gov/fipspubs/fip180-1.htm>>.

RSA. "PKCS#1 v2.1: RSA Cryptography Standard" 10 Aug. 2002
<<ftp://ftp.rsasecurity.com/pub/pkcs/pkcs-1/pkcs-1v2-1.pdf>>.

© SANS Institute 2000 - 2002, All rights reserved. Author retains full rights.