

Network Working Group  
INTERNET-DRAFT  
[draft-ietf-secsh-transport-05.txt](#)  
Expires in six months

T. Ylonen  
T. Kivinen  
M. Saarinen  
T. Rinne  
S. Lehtinen  
SSH  
22 February 1999

## SSH Transport Layer Protocol

### Status of This memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

### Abstract

SSH is a protocol for secure remote login and other secure network services over an insecure network. This document describes the SSH transport layer protocol which typically runs on top of TCP/IP. The protocol can be used as a basis for a number of secure network services. It provides strong encryption, server authentication, and integrity protection. It may also provide compression. Key exchange method, public key algorithm, symmetric encryption algorithm, message authentication algorithm, and hash algorithm are all negotiated. This document also describes the Diffie-Hellman key exchange method and the minimal set of algorithms that are needed to implement the SSH transport layer protocol.



## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Conventions Used in This Document</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Connection Setup</a>	<a href="#">3</a>
<a href="#">3.1.</a>	<a href="#">Use over TCP/IP</a>	<a href="#">3</a>
<a href="#">3.2.</a>	<a href="#">Protocol Version Exchange</a>	<a href="#">3</a>
<a href="#">3.3.</a>	<a href="#">Compatibility with Old SSH Versions</a>	<a href="#">4</a>
<a href="#">3.3.1.</a>	<a href="#">Old Client, New Server</a>	<a href="#">4</a>
<a href="#">3.3.2.</a>	<a href="#">New Client, Old Server</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Binary Packet Protocol</a>	<a href="#">4</a>
<a href="#">4.1.</a>	<a href="#">Maximum Packet Length</a>	<a href="#">5</a>
<a href="#">4.2.</a>	<a href="#">Compression</a>	<a href="#">5</a>
<a href="#">4.3.</a>	<a href="#">Encryption</a>	<a href="#">6</a>
<a href="#">4.4.</a>	<a href="#">Data Integrity</a>	<a href="#">7</a>
<a href="#">4.5.</a>	<a href="#">Key Exchange Methods</a>	<a href="#">8</a>
<a href="#">4.6.</a>	<a href="#">Public Key Algorithms</a>	<a href="#">8</a>
<a href="#">5.</a>	<a href="#">Key Exchange</a>	<a href="#">10</a>
<a href="#">5.1.</a>	<a href="#">Algorithm Negotiation</a>	<a href="#">10</a>
<a href="#">5.2.</a>	<a href="#">Output from Key Exchange</a>	<a href="#">12</a>
<a href="#">5.3.</a>	<a href="#">Taking Keys into Use</a>	<a href="#">13</a>
<a href="#">6.</a>	<a href="#">Diffie-Hellman Key Exchange</a>	<a href="#">14</a>
<a href="#">6.1.</a>	<a href="#">diffie-hellman-group1-sha1</a>	<a href="#">15</a>
<a href="#">7.</a>	<a href="#">Key Re-Exchange</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">Service Request</a>	<a href="#">16</a>
<a href="#">9.</a>	<a href="#">Additional Messages</a>	<a href="#">17</a>
<a href="#">9.1.</a>	<a href="#">Disconnection Message</a>	<a href="#">17</a>
<a href="#">9.2.</a>	<a href="#">Ignored Data Message</a>	<a href="#">17</a>
<a href="#">9.3.</a>	<a href="#">Debug Message</a>	<a href="#">18</a>
<a href="#">9.4.</a>	<a href="#">Reserved Messages</a>	<a href="#">18</a>
<a href="#">10.</a>	<a href="#">Summary of Message Numbers</a>	<a href="#">18</a>
<a href="#">11.</a>	<a href="#">Security Considerations</a>	<a href="#">19</a>
<a href="#">12.</a>	<a href="#">References</a>	<a href="#">19</a>
<a href="#">13.</a>	<a href="#">Authors' Addresses</a>	<a href="#">20</a>

**[1. Introduction](#)**

The SSH transport layer is a secure low level transport protocol. It provides strong encryption, cryptographic host authentication, and integrity protection.

Authentication in this protocol level is host-based; this protocol does not perform user authentication. A higher level protocol for user authentication can be designed on top of this protocol.

The protocol has been designed to be simple, flexible, to allow parameter negotiation, and to minimize the number of round-trips. Key

exchange method, public key algorithm, symmetric encryption algorithm, message authentication algorithm, and hash algorithm are all negotiated. It is expected that in most environments, only 2 round-trips will be needed for full key exchange, server authentication, service request,

and acceptance notification of service request. The worst case is 3 round-trips.

## **2. Conventions Used in This Document**

The keywords "MUST", "MUST NOT", "REQUIRED", "SHOULD", "SHOULD NOT", and "MAY" that appear in this document are to be interpreted as described in [\[RFC-2119\]](#).

The used data types and terminology are specified in the architecture document [\[SSH-ARCH\]](#).

The architecture document also discusses the algorithm naming conventions that MUST be used with the SSH protocols.

## **3. Connection Setup**

SSH works over any 8-bit clean, binary-transparent transport. The underlying transport SHOULD protect against transmission errors as such errors cause the SSH connection to terminate.

The client initiates the connection.

### **3.1. Use over TCP/IP**

When used over TCP/IP, the server normally listens for connections on port 22. This port number has been registered with the IANA, and has been officially assigned for SSH.

### **3.2. Protocol Version Exchange**

When the connection has been established, both sides MUST send an identification string of the form "SSH-protoversion-softwareversion comments", followed by carriage return and newline characters (ascii 13 and 10, respectively). Both sides MUST be able to process identification strings without carriage return character. No null character is sent. The maximum length of the string is 255 characters, including the carriage return and newline.

The part of the identification string preceding carriage return and newline is used in the Diffie-Hellman key exchange (Section ``Diffie-Hellman Key Exchange'').

The server MAY send other lines of data before sending the version string. Each line SHOULD be terminated by a carriage return and newline. Such lines MUST NOT begin with "SSH-", and SHOULD be encoded in ISO-10646 UTF-8 [\[RFC-2044\]](#) (language is not specified). Clients MUST be able to process such lines; they MAY be silently ignored, or MAY be displayed to the client user; if they are displayed, control character filtering discussed in [\[SSH-ARCH\]](#) SHOULD be used. The primary use of this feature is to allow TCP-wrappers to display an error message before

disconnecting.

Version strings MUST consist of printable US-ASCII characters, not

[I. Ylonen, T. Kivinen, M. Saarinen, T. Rinne and S. Lehtinen](#) [page 3]

including whitespaces or a minus sign (-). The version string is primarily used to trigger compatibility extensions and to indicate the capabilities of an implementation. The comment string should contain additional information that might be useful in solving user problems.

The protocol version described in this document is 2.0.

Key exchange will begin immediately after sending this identifier. All packets following the identification string SHALL use the binary packet protocol, to be described below.

### **3.3. Compatibility with Old SSH Versions**

During a transition period, it is important to be able to work compatibly with installed SSH clients and servers using an older version of the protocol. Information in this section is only relevant for implementations supporting compatibility with SSH versions 1.x.

#### **3.3.1. Old Client, New Server**

Server implementations MAY support a configurable "compatibility" flag that enables compatibility with old versions. When this flag is on, the server SHOULD identify its protocol version as "1.99". Clients using protocol 2.0 MUST be able to identify this as identical to "2.0". In this mode the server SHOULD NOT send carriage return character (ascii 13) after the version identification string.

In the compatibility mode the server SHOULD NOT send any further data after its initialization string until it has received an identification string from the client. The server can then determine whether the client is using an old protocol, and can revert to the old protocol if desired. The server MUST NOT send additional data before the version string, in the compatibility mode.

When compatibility with old clients is not needed, the server MAY send its initial key exchange data immediately after the identification string.

#### **3.3.2. New Client, Old Server**

Since the new client MAY immediately send additional data after its identification string (before receiving server's identification), the old protocol may already have been corrupted when the client learns that the server is old. When this happens, the client SHOULD close the connection to the server, and reconnect using the old protocol.

## **4. Binary Packet Protocol**

Each packet is of the following format.

```
uint32    packet_length
```

```
byte      padding_length
byte[n1]  payload; n1 = packet_length - padding_length - 1
byte[n2]  random padding; n2 = padding_length
```



byte[m]    mac (message authentication code); m = mac\_length

packet\_length

The length of the packet (bytes), not including MAC or the packet\_length field itself.

padding\_length

Length of padding (bytes).

payload

The useful contents of the packet. If compression has been negotiated, this field is compressed. Initially, compression MUST be "none".

Padding

Arbitrary-length padding, such that the total length of (packet\_length || padding\_length || payload || padding) is a multiple of the cipher block size or 8, whichever is larger. There MUST be at least four bytes of padding. The padding SHOULD consist of random bytes. The maximum amount of padding is 255 bytes.

MAC

Message authentication code. If message authentication has been negotiated, this field contains the MAC bytes. Initially, the MAC algorithm MUST be "none".

Note that length of the concatenation of packet length, padding length, payload, and padding MUST be a multiple of the cipher block size or 8, whichever is larger. This constraint MUST be enforced even when using stream ciphers. Note that the packet length field is also encrypted, and processing it requires special care when sending or receiving packets.

The minimum size of a packet is 16 (or the cipher block size, whichever is larger) bytes (plus MAC); implementations SHOULD decrypt the length after receiving the first 8 (or cipher block size, whichever is larger) bytes of a packet.

#### **4.1. Maximum Packet Length**

All implementations MUST be able to process packets with uncompressed payload length of 32768 bytes or less and total packet size of 35000 bytes or less (including length, padding length, payload, padding, and MAC). Implementations SHOULD support longer packets, where they might be needed e.g. if an implementation wants to send a very large number of certificates. Such packets MAY be sent if the version string indicates that the other party is able to process them. However, implementations SHOULD check that the packet length is reasonable for the implementation to avoid denial-of-service and/or buffer overflow attacks.

## 4.2. Compression

I. Ylonen, T. Kivinen, M. Saarinen, T. Rinne and S. Lehtinen [page 5]

If compression has been negotiated, the payload field (and only it) will be compressed using the negotiated algorithm. The length field and MAC will be computed from the compressed payload.

Compression MAY be stateful, depending on the method. Compression MUST be independent for each direction, and implementations MUST allow independently choosing the algorithm for each direction.

The following compression methods are currently defined:

none	REQUIRED	no compression
zlib	OPTIONAL	GNU ZLIB (LZ77) compression

The "zlib" compression is described in [[RFC-1950](#)] and in [[RFC-1951](#)]. The compression context is initialized after each key exchange, and is passed from one packet to the next with only a partial flush being performed at the end of each packet. A partial flush means that all data will be output, but the next packet will continue using compression tables from the end of the previous packet.

Additional methods may be defined as specified in [[SSH-ARCH](#)].

### **[4.3.](#) Encryption**

An encryption algorithm and a key will be negotiated during the key exchange. When encryption is in effect, the packet length, padding length, payload and padding fields of each packet MUST be encrypted with the given algorithm.

The encrypted data in all packets sent in one direction SHOULD be considered a single data stream. For example, initialization vectors SHOULD be passed from the end of one packet to the beginning of the next packet. All ciphers SHOULD use keys with an effective key length of 128 bits or more.

The ciphers in each direction MUST run independently of each other, and implementations MUST allow independently choosing the algorithm for each direction (if multiple algorithms are allowed by local policy).

The following ciphers are currently defined:

3des-cbc	REQUIRED	three-key 3DES in CBC mode
blowfish-cbc	RECOMMENDED	Blowfish in CBC mode
twofish-cbc	RECOMMENDED	Twofish in CBC mode
arcfour	OPTIONAL	the ARCFOUR stream cipher
idea-cbc	OPTIONAL	IDEA in CBC mode
cast128-cbc	OPTIONAL	CAST-128 in CBC mode
none	OPTIONAL	no encryption; NOT RECOMMENDED

The "3des-cbc" cipher is three-key triple-DES (encrypt-decrypt-encrypt),

where the first 8 bytes of the key are used for the first encryption, the next 8 bytes for the decryption, and the following 8 bytes for the final encryption. This requires 24 bytes of key data (of which 168 bits

are actually used). To implement CBC mode, outer chaining **MUST** be used (i.e., there is only one initialization vector). This is a block cipher with 8 byte blocks. This algorithm is defined in [[Schneier](#)].

The "blowfish-cbc" cipher is Blowfish in CBC mode, with 128 bit keys [[Schneier](#)]. This is a block cipher with 8 byte blocks.

The "twofish-cbc" cipher is Twofish in CBC mode, with 256 bit keys as described in Twofish AES submission. This is a block cipher with 16 byte blocks.

The "arcfour" is the Arcfour stream cipher with 128 bit keys. The Arcfour cipher is believed to be compatible with the RC4 cipher [[Schneier](#)]. RC4 is a registered trademark of RSA Data Security Inc.

The "idea-cbc" cipher is the IDEA cipher in CBC mode [[Schneier](#)]. IDEA is patented by Ascom AG.

The "cast128-cbc" cipher is the CAST-128 cipher in CBC mode [[RFC-2144](#)].

The "none" algorithm specifies that no encryption is to be done. Note that this method provides no confidentiality protection, and it is not recommended. Some functionality (e.g. password authentication) may be disabled for security reasons if this cipher is chosen.

Additional methods may be defined as specified in [[SSH-ARCH](#)].

#### **4.4. Data Integrity**

Data integrity is protected by including with each packet a message authentication code (MAC) that is computed from a shared secret, packet sequence number, and the contents of the packet.

The message authentication algorithm and key are negotiated during key exchange. Initially, no MAC will be in effect, and its length **MUST** be zero. After key exchange, the selected MAC will be computed before encryption from the concatenation of packet data:

```
mac = MAC(key, sequence_number || unencrypted_packet)
```

where `unencrypted_packet` is the entire packet without MAC (the length fields, payload and padding), and `sequence_number` is an implicit packet sequence number represented as `uint32`. The sequence number is initialized to zero for the first packet, and is incremented after every packet (regardless of whether encryption or MAC is in use). It is never reset, even if keys/algorithms are renegotiated later. It wraps around to zero after every  $2^{32}$  packets. The packet sequence number itself is not included in the packet sent over the wire.

The MAC algorithms for each direction **MUST** run independently, and implementations **MUST** allow choosing the algorithm independently for both

directions.

The MAC bytes resulting from the MAC algorithm MUST be transmitted

[I. Ylonen](#), T. Kivinen, M. Saarinen, T. Rinne and S. Lehtinen [page 7]

without encryption as the last part of the packet. The number of MAC bytes depends on the algorithm chosen.

The following MAC algorithms are currently defined:

hmac-sha1	REQUIRED	HMAC-SHA1 (length = 20)
hmac-sha-96	RECOMMENDED	first 96 bits of HMAC-SHA1 (length = 12)
hmac-md5	OPTIONAL	HMAC-MD5 (length = 16)
hmac-md5-96	OPTIONAL	first 96 bits of HMAC-MD5 (length = 12)
none	OPTIONAL	no MAC; NOT RECOMMENDED

The "hmac-\*" algorithms are described in [[RFC-2104](#)]. The "\*-n" MACs use only the first n bits of the resulting value.

The hash algorithms are described in [[Schneier](#)].

The "none" method is NOT RECOMMENDED. An active attacker may be able to modify transmitted data if this is used.

Additional methods may be defined as specified in [[SSH-ARCH](#)].

#### [4.5.](#) Key Exchange Methods

The key exchange method specifies how one-time session keys are generated for encryption and for authentication, and how the server authentication is done.

Only one REQUIRED key exchange method has been defined:

diffie-hellman-group1-sha1	REQUIRED
----------------------------	----------

This method is described later in this document.

Additional methods may be defined as specified in [[SSH-ARCH](#)].

#### [4.6.](#) Public Key Algorithms

This protocol has been designed to be able to operate with almost any public key format, encoding, and algorithm (signature and/or encryption).

There are several aspects that define a public key type:

- o Key format: how is the key encoded and how are certificates represented. The key blobs in this protocol MAY contain certificates in addition to keys.
- o Signature and/or encryption algorithms. Some key types may not support both signing and encryption. Key usage may also be restricted by policy statements in e.g. certificates. In this case,

different key types SHOULD be defined for the different policy alternatives.



- o Encoding of signatures and/or encrypted data. This includes but is not limited to padding, byte order, and data formats.

The following public key and/or certificate formats are currently defined:

ssh-dss	REQUIRED	sign	Simple DSS
x509v3	RECOMMENDED	sign	X.509 certificates
spki	OPTIONAL	sign	SPKI certificates
pgp	OPTIONAL	sign	OpenPGP certificates

Additional key types may be defined as specified in [[SSH-ARCH](#)].

The key type **MUST** always be explicitly known (from algorithm negotiation or some other source). It is not normally included in the key blob.

Certificates and public keys are encoded as:

```
uint32  combined length of the format identifier and associated data
string  certificate or public key format identifier
byte[n] key/certificate data
```

The certificate part may have be a zero length string, but a public key is required. This is the public key that will be used for authentication; the certificate sequence contained in the certificate blob can be used to provide authorization.

The "ssh-dss" key format has the following specific encoding:

```
uint32  length
string  "ssh-dss"
mpint   p
mpint   q
mpint   g
mpint   y
```

Here the "p", "q", "g", and "y" parameters form the signature key blob.

Signing and verifying using this key format are done according to the Digital Signature Standard [[FIPS-186](#)] using the SHA-1 hash. A description can also be found in [[Schneier](#)].

The resulting signature is encoded as:

```
uint32  length
string  "ssh-dss"
string  dss_signature_blob
```

The "x509v3" method indicates that the certificates, the public key, and the resulting signature are in X.509v3 compatible DER-encoded format.

The formats used in X.509v3 is described in [[PKIX-Part1](#)].

The "spki" method indicates that the certificate blob contains a

sequence of SPKI certificates. The format of SPKI certificates is described in [SPKI].

The "pgp" method indicates the the certificates, the public key, and the signature are in OpenPGP compatible binary format. [[PGP](#)]

## 5. Key Exchange

Key exchange begins by each side sending lists of supported algorithms. Each side has a preferred algorithm in each category, and it is assumed that most implementations at any given time will use the same preferred algorithm. Each side MAY guess which algorithm the other side is using, and MAY send an initial key exchange packet according to the algorithm if appropriate for the preferred method. If all algorithms were guessed right, the optimistically sent packet MUST be handled as the first key exchange packet. However, if the guess was wrong, and a packet was optimistically sent by one or both parties, such packets MUST be ignored (even if the error in the guess would not affect the contents of the initial packet(s)), and the appropriate side MUST send the correct initial packet.

Server authentication in the key exchange MAY be implicit. After a key exchange with implicit server authentication, the client MUST wait for response to its service request message before sending any further data.

### 5.1. Algorithm Negotiation

Key exchange begins by each side sending the following packet:

```
byte      SSH_MSG_KEXINIT
byte[16]  cookie (random bytes)
string    kex_algorithms
string    server_host_key_algorithms
string    encryption_algorithms_client_to_server
string    encryption_algorithms_server_to_client
string    mac_algorithms_client_to_server
string    mac_algorithms_server_to_client
string    compression_algorithms_client_to_server
string    compression_algorithms_server_to_client
string    languages_client_to_server
string    languages_server_to_client
boolean   first_kex_packet_follows
uint32    0 (reserved for future extension)
```

Each of the algorithm strings MUST be a comma-separated list of algorithm names (see ``Algorithm Naming'' in [[SSH-ARCH](#)]). Each supported (allowed) algorithm MUST be listed, in order of preference.

The first algorithm in each list MUST be the preferred (guessed) algorithm. Each string MUST contain at least one algorithm name.

cookie

The cookie MUST be a random value generated by the sender. Its

purpose is to make it impossible for either side to fully determine the keys and the session identifier.

#### kex\_algorithms

Key exchange algorithms were defined above. The first algorithm MUST be the preferred (and guessed) algorithm. If both sides make the same guess, that algorithm MUST be used. Otherwise, the following algorithm MUST be used to choose a key exchange method: iterate over client's kex algorithms, one at a time. Choose the first algorithm that satisfies the following conditions:

- o the server also supports the algorithm,
- o if the algorithm requires an encryption-capable host key, there is an encryption-capable algorithm on the server's `server_host_key_algorithms` that is also supported by the client, and
- o if the algorithm requires a signature-capable host key, there is a signature-capable algorithm on the server's `server_host_key_algorithms` that is also supported by the client.

If no algorithm satisfying all these conditions can be found, the connection fails, and both sides MUST disconnect.

#### server\_host\_key\_algorithms

List of the algorithms supported for the server host key. The server lists the algorithms for which it has host keys; the client lists the algorithms that it is willing to accept. (There MAY be multiple host keys for a host, possibly with different algorithms.)

Some host keys may not support both signatures and encryption (this can be determined from the algorithm), and thus not all host keys are valid for all key exchange methods.

Algorithm selection depends on whether the chosen key exchange algorithm requires a signature- or encryption capable host key. It MUST be possible to determine this from the public key algorithm name. The first algorithm on the client's list that satisfies the requirements and is also supported by the server MUST be chosen. If there is no such algorithm, both sides MUST disconnect.

#### encryption\_algorithms

Lists the acceptable symmetric encryption algorithms in order of preference. The chosen encryption algorithm to each direction MUST be the first algorithm on the client's list that is also on the server's list. If there is no such algorithm, both sides MUST disconnect.

Note that "none" must be explicitly listed if it is to be acceptable. The defined algorithm names are listed in Section

```Encryption''`.

#### `mac_algorithms`

Lists the acceptable MAC algorithms in order of preference. The chosen MAC algorithm MUST be the first algorithm on the client's list that is also on the server's list. If there is no such algorithm, both sides MUST disconnect.

Note that "none" must be explicitly listed if it is to be acceptable. The MAC algorithm names are listed in Section ```Data Integrity''`.

#### `compression_algorithms`

Lists the acceptable compression algorithms in order of preference. The chosen compression algorithm MUST be the first algorithm on the client's list that is also on the server's list. If there is no such algorithm, both sides MUST disconnect.

Note that "none" must be explicitly listed if it is to be acceptable. The compression algorithm names are listed in Section ```Compression''`.

#### `languages`

This is a comma-separated list of language tags in order of preference [[RFC-1766](#)]. Both parties MAY ignore this list. If there are no language preferences, this list SHOULD be empty.

#### `first_kex_packet_follows`

Indicates whether a guessed key exchange packet follows. If a guessed packet will be sent, this MUST be true. If no guessed packet will be sent, this MUST be false.

After receiving the SSH\_MSG\_KEXINIT packet from the other side, each party will know whether their guess was right. If the other party's guess was wrong, and this field was true, the next packet MUST be silently ignored, and both sides MUST then act as determined by the negotiated key exchange method. If the guess was right, key exchange MUST continue using the guessed packet.

After the KEXINIT packet exchange, the key exchange algorithm is run. It may involve several packet exchanges, as specified by the key exchange method.

## **[5.2](#). Output from Key Exchange**

The key exchange produces two values: a shared secret K, and an exchange hash H. Encryption and authentication keys are derived from these. The exchange hash H from the first key exchange is additionally used as the session identifier, which is a unique identifier for this connection. It is used by authentication methods as a part of the data that is

signed as a proof of possession of a private key. Once computed, the session identifier is not changed, even if keys are later re-exchanged.



Each key exchange method specifies a hash function that is used in the key exchange. The same hash algorithm **MUST** be used in key derivation. Here, we'll call it HASH.

Encryption keys **MUST** be computed as HASH of a known value and K as follows:

- o Initial IV client to server: HASH(K || H || "A" || session\_id) ("A" means the single character A, ascii 65).
- o Initial IV server to client: HASH(K || H || "B" || session\_id)
- o Encryption key client to server: HASH(K || H || "C" || session\_id)
- o Encryption key server to client: HASH(K || H || "D" || session\_id)
- o Integrity key client to server: HASH(K || H || "E" || session\_id)
- o Integrity key server to client: HASH(K || H || "F" || session\_id)

Key data **MUST** be taken from the beginning of the hash output. 128 bits (16 bytes) **SHOULD** be used for algorithms with variable-length keys. For other algorithms, as many bytes as are needed are taken from the beginning of the hash value. If the key length is longer than the output of the HASH, the key is extended by computing HASH of the concatenation of K and H and the entire key so far, and appending the resulting bytes (as many as HASH generates) to the key. This process is repeated until enough key material is available; the key is taken from the beginning of this value. In other words,

```
K1 = HASH(K || H || X || session_id)  (X is e.g. "A")
K2 = HASH(K || H || K1)
K3 = HASH(K || H || K1 || K2)
...
key = K1 || K2 || K3 || ...
```

### **5.3. Taking Keys into Use**

Key exchange ends by each side sending an SSH\_MSG\_NEWKEYS message. This message is sent with the old keys and algorithms. All messages sent after this message **MUST** use the new keys and algorithms.

When this message is received, the new keys and algorithms **MUST** be taken into use for receiving.

This message is the only valid message after key exchange, in addition to SSH\_MSG\_DEBUG, SSH\_MSG\_DISCONNECT and SSH\_MSG\_IGNORE messages. The purpose of this message is to ensure that a party is able to respond with a disconnect message that the other party can understand if something goes wrong with the key exchange. Implementations **MUST NOT**

accept any other messages after key exchange before receiving  
SSH\_MSG\_NEWKEYS.

byte        SSH\_MSG\_NEWKEYS

## 6. Diffie-Hellman Key Exchange

The Diffie-Hellman key exchange provides a shared secret that can not be determined by either party alone. The key exchange is combined with a signature with the host key to provide host authentication.

In the following description (C is the client, S is the server; p is a large safe prime, g is a generator for a subgroup of GF(p), and q is the order of the subgroup; V\_S is S's version string; V\_C is C's version string; K\_S is S's public host key; I\_C is C's KEXINIT message and I\_S S's KEXINIT message which have been exchanged before this part begins):

1. C generates a random number x ( $1 < x < q$ ) and computes  $e = g^x \bmod p$ . C sends "e" to S.
2. S generates a random number y ( $0 < y < q$ ) and computes  $f = g^y \bmod p$ . S receives "e". It computes  $K = e^y \bmod p$ ,  $H = \text{hash}(V_C || V_S || I_C || I_S || K_S || e || f || K)$ , and signature s on H with its private host key. S sends "K\_S || f || s" to C. The signing operation may involve a second hashing operation.
3. C verifies that K\_S really is the host key for S (e.g. using certificates or a local database). C is also allowed to accept the key without verification; however, doing so will render the protocol insecure against active attacks (but may be desirable for practical reasons in the short term in many environments). C then computes  $K = f^x \bmod p$ ,  $H = \text{hash}(V_C || V_S || I_C || I_S || K_S || e || f || K)$ , and verifies the signature s on H.

Either side MUST NOT send or accept e or f values that are not in the range [1, p-1]. If this condition is violated, the key exchange fails.

This is implemented with the following messages. The hash algorithm for computing the exchange hash is defined by the method name, and is called HASH. The public key algorithm for signing is negotiated with the KEXINIT messages.

First, the client sends:

byte        SSH\_MSG\_KEXDH\_INIT  
mpint       e

The server responds with:

byte        SSH\_MSG\_KEXDH\_REPLY  
string       server public host key and certificates (K\_S)  
mpint       f  
string       signature of H

The hash  $H$  is computed as the HASH hash of the concatenation of the following:

[I. Ylonen](#), [T. Kivinen](#), [M. Saarinen](#), [T. Rinne](#) and [S. Lehtinen](#) [page 14]

```

string    V_C, the client's version string (CR and NL excluded)
string    V_S, the server's version string (CR and NL excluded)
string    I_C, the payload of the client's SSH_MSG_KEXINIT
string    I_S, the payload of the server's SSH_MSG_KEXINIT
string    K_S, the host key
mpint     e, exchange value sent by the client
mpint     f, exchange value sent by the server
mpint     K, the shared secret

```

This value is called the exchange hash, and it is used to authenticate the key exchange.

The signature algorithm MUST be applied over H, not the original data. Most signature algorithms include hashing and additional padding. For example, a "ssh-dss" specifies SHA-1 hashing; in that case, the data is first hashed with HASH to compute H, and H is then hashed with SHA-1 as part of the signing operation.

### [6.1. diffie-hellman-group1-sha1](#)

The "diffie-hellman-group1-sha1" method specifies Diffie-Hellman key exchange with SHA-1 as HASH, and the following group:

The prime  $p$  is equal to  $2^{1024} - 2^{960} - 1 + 2^{64} * \text{floor}(2^{894} \text{ Pi} + \text{129093})$ . Its hexadecimal value is

```

FFFFFFFF FFFFFFFF C90FDAA2 2168C234 C4C6628B 80DC1CD1
29024E08 8A67CC74 020BBEA6 3B139B22 514A0879 8E3404DD
EF9519B3 CD3A431B 302B0A6D F25F1437 4FE1356D 6D51C245
E485B576 625E7EC6 F44C42E9 A637ED6B 0BFF5CB6 F406B7ED
EE386BFB 5A899FA5 AE9F2411 7C4B1FE6 49286651 ECE65381
FFFFFFFF FFFFFFFF.

```

In decimal, this value is

```

179769313486231590770839156793787453197860296048756011706444
423684197180216158519368947833795864925541502180565485980503
646440548199239100050792877003355816639229553136239076508735
759914822574862575007425302077447712589550957937778424442426
617334727629299387668709205606050270810842907692932019128194
467627007.

```

The generator used with this prime is  $g = 2$ . The group order  $q$  is  $(p - 1) / 2$ .

This group was taken from the ISAKMP/Oakley specification, and was originally generated by Richard Schroepel at the University of Arizona. Properties of this prime are described in [[Orm96](#)].

## [7. Key Re-Exchange](#)

Key re-exchange is started by sending a SSH\_MSG\_KEXINIT packet when not already doing a key exchange (as described in Section ``Algorithm

Negotiation'''). When this message is received, a party MUST respond with its own SSH\_MSG\_KEXINIT message except when the received SSH\_MSG\_KEXINIT already was a reply. Either party MAY initiate the re-exchange, but roles MUST NOT be changed (i.e., the server remains the server, and the client remains the client).

Key re-exchange is performed using whatever encryption was in effect when the exchange was started. Encryption, compression, and MAC methods are not changed before a new SSH\_MSG\_NEWKEYS is sent after the key exchange (as in the initial key exchange). Re-exchange is processed identically to the initial key exchange, except for the session identifier that will remain unchanged. It is permissible to change some or all of the algorithms during the re-exchange. Host keys can also change. All keys and initialization vectors are recomputed after the exchange. Compression and encryption contexts are reset.

It is recommended that the keys are changed after each gigabyte of transmitted data or after each hour of connection time, whichever comes sooner. However, since the re-exchange is a public key operation, it requires a fair amount of processing power and should not be performed too often.

More application data may be sent after the SSH\_MSG\_NEWKEYS packet has been sent; key exchange does not affect the protocols that lie above the SSH transport layer.

## **8. Service Request**

After the key exchange, the client requests a service. The service is identified by a name. The format of names and procedures for defining new names are defined in [[SSH-ARCH](#)].

Currently, the following names have been reserved:

```
ssh-userauth
ssh-connection
```

Similar local naming policy is applied to the service names that is applied to the algorithm names; a local service should use the servicename@domain syntax.

```
byte      SSH_MSG_SERVICE_REQUEST
string    service name
```

If the server rejects the service request, it SHOULD send an appropriate SSH\_MSG\_DISCONNECT message and MUST disconnect.

When the service starts, it may have access to the session identifier generated during the key exchange.

If the server supports the service (and permits the client to use it),  
it MUST respond with



```
byte      SSH_MSG_SERVICE_ACCEPT
string    service name
```

Message numbers used by services should be in the area reserved for them (see Section ``Summary of Message Numbers''). The transport level will continue to process its own messages.

Note that after a key exchange with implicit server authentication, the client MUST wait for response to its service request message before sending any further data.

## **9. Additional Messages**

Either party may send any of the following messages at any time.

### **9.1. Disconnection Message**

```
byte      SSH_MSG_DISCONNECT
uint32    reason code
string    description [RFC-2044]
string    language tag [RFC-1766]
```

This message causes immediate termination of the connection. All implementations MUST be able to process this message; they SHOULD be able to send this message.

The sender MUST NOT send or receive any data after this message, and the recipient MUST NOT accept any data after receiving this message. The description field gives a more specific explanation in a human-readable form. The error code gives the reason in a more machine-readable format (suitable for localization), and can have the following values:

```
#define SSH_DISCONNECT_HOST_NOT_ALLOWED_TO_CONNECT      1
#define SSH_DISCONNECT_PROTOCOL_ERROR                   2
#define SSH_DISCONNECT_KEY_EXCHANGE_FAILED              3
#define SSH_DISCONNECT_HOST_AUTHENTICATION_FAILED       4
#define SSH_DISCONNECT_MAC_ERROR                        5
#define SSH_DISCONNECT_COMPRESSION_ERROR                6
#define SSH_DISCONNECT_SERVICE_NOT_AVAILABLE           7
#define SSH_DISCONNECT_PROTOCOL_VERSION_NOT_SUPPORTED   8
#define SSH_DISCONNECT_HOST_KEY_NOT_VERIFIABLE          9
#define SSH_DISCONNECT_CONNECTION_LOST                  10
#define SSH_DISCONNECT_BY_APPLICATION                   11
```

If the description string is displayed, control character filtering discussed in [[SSH-ARCH](#)] should be used to avoid attacks by sending terminal control characters.

### **9.2. Ignored Data Message**

byte	SSH_MSG_IGNORE
string	data

All implementations MUST understand (and ignore) this message at any time (after receiving the protocol version). No implementation is required to send them. This message can be used as an additional protection measure against advanced traffic analysis techniques.

### **9.3. Debug Message**

byte	SSH_MSG_DEBUG
boolean	always_display
string	message [ <a href="#">RFC-2044</a> ]
string	language tag [ <a href="#">RFC-1766</a> ]

All implementations MUST understand this message, but they are allowed to ignore it. This message is used to pass information to the other side that may help debugging. If `always_display` is true, the message SHOULD be displayed. Otherwise, it SHOULD NOT be displayed unless debugging information has been explicitly requested by the user.

The message doesn't need to contain a newline. It is, however, allowed to consist of multiple lines separated by CRLF (carriage return - line feed) pairs.

If the message string is displayed, terminal control character filtering discussed in [[SSH-ARCH](#)] should be used to avoid attacks by sending terminal control characters.

### **9.4. Reserved Messages**

An implementation MUST respond to all unrecognized messages with an `SSH_MSG_UNIMPLEMENTED` message in the order in which they were received. Such messages MUST be otherwise ignored. Later protocol versions may define other meanings for these message types.

byte	SSH_MSG_UNIMPLEMENTED
uint32	packet sequence number of rejected message

## **10. Summary of Message Numbers**

The following message numbers have been defined in this protocol.

<code>#define SSH_MSG_DISCONNECT</code>	1
<code>#define SSH_MSG_IGNORE</code>	2
<code>#define SSH_MSG_UNIMPLEMENTED</code>	3
<code>#define SSH_MSG_DEBUG</code>	4
<code>#define SSH_MSG_SERVICE_REQUEST</code>	5
<code>#define SSH_MSG_SERVICE_ACCEPT</code>	6
<code>#define SSH_MSG_KEXINIT</code>	20
<code>#define SSH_MSG_NEWKEYS</code>	21

```
/* Numbers 30-49 used for kex packets.  
   Different kex methods may reuse message numbers in  
   this range. */
```

```
#define SSH_MSG_KEXDH_INIT          30
#define SSH_MSG_KEXDH_REPLY        31
```

## **11. Security Considerations**

This protocol provides a secure encrypted channel over an unsecure network. It performs server host authentication, key exchange, encryption, and integrity protection. It also derives a unique session id that may be used by higher-level protocols.

It is expected that this protocol will sometimes be used without insisting on reliable association between the server host key and the server host name. Such use is inherently insecure, but may be necessary in non-security critical environments, and still provides protection against passive attacks. However, implementors of protocols running on top of this protocol should keep this possibility in mind.

This protocol is designed to be used over a reliable transport. If transmission errors or message manipulation occur, the connection is closed. The connection SHOULD be re-established if this occurs. Denial of service attacks of this type ("wire cutter") are almost impossible to avoid.

The protocol was not designed to eliminate covert channels. For example, the padding, SSH\_MSG\_IGNORE messages, and several other places in the protocol can be used to pass covert information, and the recipient has no reliable way to verify whether such information is being sent.

## **12. References**

[FIPS-186] Federal Information Processing Standards Publication (FIPS PUB) 186, Digital Signature Standard, 18 May 1994.

[Orm96] Orman, H., "The Oakley Key Determination Protocol", version 1, TR97-92, Department of Computer Science Technical Report, University of Arizona.

[RFC-1034] Mockapetris, P., "Domain Names - Concepts and Facilities", November 1987.

[RFC-1766] Alvestrand, H., "Tags for the Identification of Languages", March 1995.

[RFC-1950] Deutch, P. and Gailly, J-L., "ZLIB Compressed Data Format Specification version 3.3", May 1996.

[RFC-1951] Deutch, P., "DEFLATE Compressed Data Format Specification version 1.3", May 1996.

[RFC-2044] Yergeau, F., "UTF-8, a Transformation Format of Unicode and ISO 10646", October 1996.

[RFC-2104] Krawczyk, H., Bellare, M., and Canetti, R., "HMAC: Keyed-Hashing for Message Authentication", February 1997

[RFC-2119] Bradner, S., "Key words for use in RFCs to indicate Requirement Levels", March 1997.

[RFC-2144] Adams, C., "The CAST-128 Encryption Algorithm", May 1997.

[Schneier] Schneier, B., "Applied Cryptography Second Edition: protocols, algorithms, and source code in C", 2nd edition, John Wiley & Sons, New York, NY, 1996.

[SSH-ARCH] Ylonen, T., et al, "SSH Protocol Architecture", Internet Draft, [draft-ietf-secsh-architecture-03.txt](#)

[SSH-USERAUTH] Ylonen, T., et al, "SSH Authentication Protocol", Internet Draft, [draft-ietf-secsh-userauth-05.txt](#)

[SSH-CONNECT] Ylonen, T., et al, "SSH Connection Protocol", Internet Draft, [draft-ietf-secsh-connect-05.txt](#)

[PGP] (XXX n.a.)

[PKIX-Part1] (XXX n.a.)

### **13. Authors' Addresses**

Tatu Ylonen  
SSH Communications Security Ltd.  
Tekniikantie 12  
FIN-02150 ESP00  
Finland  
E-mail: ylo@ssh.fi

Tero Kivinen  
SSH Communications Security Ltd.  
Tekniikantie 12  
FIN-02150 ESP00  
Finland  
E-mail: kivinen@ssh.fi

Markku-Juhani O. Saarinen  
SSH Communications Security Ltd.  
Tekniikantie 12  
FIN-02150 ESP00  
Finland  
E-mail: mjos@ssh.fi





INTERNET-DRAFT

6 August 1998

Timo J. Rinne  
SSH Communications Security Ltd.  
Tekniikantie 12  
FIN-02150 ESP00  
Finland  
E-mail: tri@ssh.fi

Sami Lehtinen  
SSH Communications Security Ltd.  
Tekniikantie 12  
FIN-02150 ESP00  
Finland  
E-mail: sjl@ssh.fi

