SSH Protocol Architecture

T. Ylonen, T. Kivinen, M. Saarinen, T. Rinne and S. Lehtinen 6 August 1998

SSH is a protocol for secure remote login and other secure network services over an insecure network. This document describes the architecture of the SSH protocol, and the notation and terminology used in SSH protocol documents. It also discusses the SSH algorithm naming system that allows local extensions. The SSH protocol consists of three major components: Transport layer protocol provides server authentication, confidentiality, and integrity with perfect forward secrecy. User authentication protocol authenticates the client to the server. Connection protocol multiplexes the encrypted tunnel into several logical channels. Details of these protocols are described in separate documents.

Contents

1 Introduction

SSH is a protocol for secure remote login and other secure network services over an insecure network. It consists of three major components:

- Transport layer protocol [SSH-TRANS] provides server authentication, confidentiality, and integrity. It may optionally also provide compression. The transport layer will typically be run over a TCP/IP connection, but might also be used on top of any other reliable data stream.
- User authentication protocol [SSH-USERAUTH] authenticates the client-side user to the server. It runs over the transport layer protocol.
- Connection protocol [SSH-CONN] multiplexes the encrypted tunnel into several logical channels. It runs over the user authentication protocol.

The client sends a service request once a secure transport layer connection has been established. A second service request is sent after user authentication is complete. This allows new protocols to be defined and coexist with the protocols listed above.

The connection protocol provides channels that can be used for a wide range of purposes. Standard methods are provided for setting up secure interactive shell sessions and for forwarding (”tunneling”) arbitrary TCP/IP ports and X11 connections.

2 Specification of Requirements

All of the documents related to the SSH protocols shall use the keywords ”MUST”, ”MUST NOT”, ”REQUIRED”, ”SHALL”, ”SHALL NOT”, ”SHOULD”, ”SHOULD NOT”, ”RECOMMENDED, ”MAY”, and ”OPTIONAL” to describe requirements. They are to be interpreted as described in [RFC-2119].

3 Architecture

3.1 Host Keys

Each server host MUST have a host key. Hosts MAY have multiple host keys using multiple different algorithms. Multiple hosts MAY share the same host key. Every host MUST have at least one key using each REQUIRED public key algorithm (currently DSS [FIPS-186]).

The server host key is used during key exchange to verify that the client is really talking to the correct server. For this to be possible, the client must have a priori knowledge of the server's public host key.

Two different trust models can be used:

- The client has a local database that associates each host name (as typed by the user) with the corresponding public host key. This method requires no centrally administered infrastructure, and no third-party coordination. The downside is that the database of name-key associations may become burdensome to maintain.
- The host name - key association is certified by some trusted certification authority. The client only knows the CA root key, and can verify the validity of all host keys certified by accepted CAs.

  The second alternative eases the maintenance problem, since ideally only a single CA key needs to be securely stored on the client. On the other hand, each host key must be appropriately certified by a central authority before authorization is possible. Also, a lot of trust is placed on the central infrastructure.

The protocol provides the option that the server name - host key association is not checked when connecting the host for the first time. This allows communication without prior communication of host keys or certification. The connection still provides protection against passive listening; however, it becomes vulnerable to active man-in-the-middle attacks. Implementations SHOULD NOT normally allow such connections by default, as they pose a potential security problem. However, as there is no widely deployed key infrastructure available on the Internet yet, this option makes the protocol much more usable during the transition time until such an infrastructure emerges, while still providing a much higher level of security than that offered by older solutions (e.g. telnet [RFC-854] and rlogin [RFC-1282]).

Implementations SHOULD try to make best effort to check host keys. An example of a possible strategy is to only accept a host key without checking the first time a host is connected, save the key in a local database, and compare against that key on all future connections to that host.

Implementations MAY provide additional methods for verifying the correctness of host keys, e.g. a hexadecimal fingerprint derived from the SHA-1 hash of the public key. Such fingerprints can easily be verified by using telephone or other external communication channels.

All implementations SHOULD provide an option to not accept host keys that cannot be verified.

We believe that ease of use is critical to end-user acceptance of security solutions, and no improvement in security is gained if the new solutions are not used. Thus, providing the option not to check the the server host key is believed to improve overall security of the Internet, even though it reduces the security of the protocol in configurations where it is allowed.

3.2 Extensibility

We believe that the protocol will evolve over time, and some organizations will want to use their own encryption, authentication and/or key exchange methods. Central registration of all extensions is cumbersome, especially for experimental or classified features. On the other hand, having no central registration leads to conflicts in method identifiers, making interoperability difficult.

We have chosen to identify algorithms, methods, formats, and extension protocols with textual names that are of a specific format. DNS names are used to create local namespaces where experimental or classified extensions can be defined without fear of conflicts with other implementations.

One design goal has been to keep the base protocol as simple as possible, and to require as few algorithms as possible. However, all implementations MUST support a minimal set of algorithms to ensure interoperability (this does not imply that the local policy on all hosts would necessary allow these algorithms). The mandatory algorithms are specified in the relevant protocol documents.

Additional algorithms, methods, formats, and extension protocols can be defined in separate drafts. See Section 5 (Algorithm Naming) for more information.

3.3 Policy Issues

The protocol allows full negotiation of encryption, integrity, key exchange, compression, and public key algorithms and formats. Encryption, integrity, public key, and compression algorithms can be different for each direction.

The following policy issues SHOULD be addressed in the configuration mechanisms of each implementation:

- Encryption, integrity, and compression algorithms, separately for each direction. The policy MUST specify which is the preferred algorithm (e.g. the first algorithm listed in each category).
- Public key algorithms and key exchange method to be used for host authentication. The existence of trusted host keys for different public key algorithms also affects this choice.
- The authentication methods that are to be required by the server for each user. The server's policy MAY require multiple authentication for some or all users. The required algorithms MAY depend on the location from where the user is trying to log in from.
- The operations that the user is allowed to perform using the connection protocol. Some issues are related to security; for example, the policy SHOULD NOT allow the server to start sessions or run commands on the client machine, and MUST NOT allow connections to the authentication agent unless forwarding it has been requested. Other issues, such as which TCP/IP ports can be forwarded and by whom, are clear local policy issues. Many of these issues may involve traversing or bypassing firewalls, and are interrelated with the local security policy.

3.4 Security Properties

The primary goal of the SSH protocols is improved security on the Internet. It attempts to do this in a way that is easy to deploy, even at the cost of absolute security.

- All encryption, integrity, and public key algorithms used are well-known, well-established algorithms.
- All algorithms are used with cryptographically sound key sizes that are believed to provide protection against even the strongest cryptanalytic attacks for decades.
- All algorithms are negotiated, and in case some algorithm is broken, it is easy to switch to some other algorithm without modifying the base protocol.

Specific concessions were made to make wide-spread fast deployment easier. The particular case where this comes up is verifying that the server host key really belongs to the desired host; the protocol allows the verification to be left out (but this is NOT RECOMMENDED). This is believed to significantly improve usability in the short term, until widespread Internet public key infrastructures emerge.

3.5 Packet Size and Overhead

Some readers will worry about the increase in packet size due to new headers, padding, and MAC. The minimum packet size in the order of 28 bytes (depending on negotiated algorithms). The increase is negligible for large packets, but very significant for one-byte packets (telnet-type sessions). There are, however, several factors that make this a non-issue in almost all cases:

- The minimum size of a TCP/IP header is 32 bytes. Thus, the increase is actually from 33 to 51 bytes (roughly).
- The minimum size of the data field of an ethernet packet is 46 bytes [RFC-894]. Thus, the increase is by no more than 5 bytes. When ethernet headers are considered, the increase is by less than 10 percent.
- The total fraction of telnet-type data in the Internet is negligible, even with increased packet sizes.

The only environment where the packet size increase is likely to have significant effect is PPP [RFC-1134] over slow modem lines (PPP compresses the TCP/IP headers, emphasizing the increase in packet size). However, with modern modems, the time needed to transfer is on the order of 2ms, which is a lot faster than people can type.

There are also issues related to the maximum packet size. To minimize delays in screen updates, one does not want excessively large packets for interactive sessions. The maximum packet size is negotiated separately for each channel.

3.6 Localization and Character Set Support

For the most part, the SSH protocols do not directly pass text that would be displayed to the user. However, there are some places where such data might be passed. When applicable, the character set for the data MUST be explicitly specified. In most places, ISO 10646 with UTF-8 encoding is used [RFC-2044]. When applicable, a field is also be provided for a language tag [RFC-1766].

One big issue is the character set of the interactive session. There is no clear solution, as different applications may display data in different formats. Different types of terminal emulation may also be employed in the client, and the character set to be used is effectively determined by the terminal emulation. Thus, no place is provided for specifying the character set or encoding for terminal session data directly. However, the terminal emulation type (e.g. "vt100") is transmitted to the remote site, and it implicitly specifies the character set and encoding. Applications typically use the terminal type to determine what character set they use, or the character set is determined using some external means. The terminal emulation may also allow configuring the default character set. In any case, character set for the terminal session is considered primarily a client local issue.

Internal names used to identify algorithms or protocols are normally never displayed to users, and must be in US-ASCII.

The client and server user names are inherently constrained by what the server is prepared to accept. They might, however, occasionally be displayed in logs, reports, etc. They MUST be encoded using ISO 10646 UTF-8, but other encodings may be required in some cases. It is up to the server to decide how to map user names to accepted user names. Straight bit-wise binary comparison is RECOMMENDED.

For localization purposes, the protocol attempts to minimize the number of textual messages transmitted. When present, such messages typically relate to errors, debugging information, or some externally configured data. For data that is normally displayed, it SHOULD be possible to fetch a localized message instead of the transmitted by using a numeric code. The remaining messages SHOULD be configurable.

4 Data Type Representations Used in the SSH Protocols

`byte`

A `byte` represents an arbitrary 8-bit value (octet) [RFC1700]. Fixed length data is sometimes represented as an array of bytes, written `byte[n]`, where `n` is the number of bytes in the array.

`boolean`

A boolean value is stored as a single byte. The value 0 represents `false`, and the value 1 represents `true`. All non-zero values MUST be interpreted as true; however, applications MUST not store values other than 0 and 1.

`uint32`

Represents a 32-bit unsigned integer. Stored as four bytes in the order of decreasing significance (network byte order).

For example, the value 699921578 (0x29b7f4aa) is stored as `29 b7 f4 aa`.

`string`

Arbitrary length binary string. Strings are allowed to contain arbitrary binary data, including null characters and 8-bit characters. They are stored as a `uint32` containing its length (number of bytes that follow) and zero (= empty string) or more bytes that are the value of the string. Terminating null characters are not used.

`Strings` are also used to store text. In that case, US-ASCII is used for internal names, and ISO-10646 UTF-8 for text that might be displayed to the user. Terminating null character SHOULD normally not be stored in the string.

For example, the US-ASCII string "testing" is represented as `00 00 00 07 t e s t i n g`. The UTF8 mapping does not alter the encoding of US-ASCII characters.

`mpint`

Represents multiple precision integers in two's complement format, stored as a `string`, 8 bits per byte, MSB first. Negative numbers have one in the most significant bit of the first byte of the data partition of. If the most significant bit would be set for a positive number, the number MUST be preceded by a zero byte. Unnecessary leading zero or 255 bytes MUST NOT be included. The value zero MUST be stored as a `string` with zero bytes of data.

By convention, a number that is used in modular computations in Z_n SHOULD be represented in the range 0 <= x < n.

Examples:

```
value (hex)        representation (hex)
---------------------------------------------------------------
0                  00 00 00 00
9a378f9b2e332a7    00 00 00 08 09 a3 78 f9 b2 e3 32 a7
80                 00 00 00 02 00 80
-1234              00 00 00 02 ed cc
-deadbeef          00 00 00 05 ff 21 52 41 11
```

4.1 Encoding of Network Addresses

Network addresses are encoded as `strings`. DNS names MUST NOT be used, as DNS is an insecure protocol.

If an address contains a colon (':', ascii 58), it is interpreted as an IPv6 address. The encoding of IPv6 addresses is described in [RFC-1884]. IPv4 addresses are expressed in the standard dot-separated decimal format (e.g. 127.0.0.1).

5 Algorithm Naming

The SSH protocols refer to particular hash, encryption, integrity, compression, and key exchange algorithms or protocols by names. There are some standard algorithms that all implementations MUST support. There are also algorithms that are defined in the protocol specification but are OPTIONAL. Furthermore, it is expected that some organizations will want to use their own algorithms.

In this protocol, all algorithm identifiers MUST be printable US-ASCII strings no longer than 64 characters. Names MUST be case-sensitive.

There are two formats for algorithm names:

- Names that do not contain an at-sign (@) are reserved to be assigned by IANA (Internet Assigned Numbers Authority). Examples include '3des-cbc', 'sha-1', 'hmac-sha1', and 'zlib' (the quotes are not part of the name). Additional names of this format may be registered with IANA; see Section 7 (IANA Considerations). Names of this format MUST NOT be used without first registering with IANA. Registered names MUST NOT contain an at-sign (@) or a comma (,).
- Anyone can define additional algorithms by using names in the format name@domainname, e.g. "ourcipher-cbc@ssh.fi". The format of the part preceding the at sign is not specified; it MUST consist of US-ASCII characters except at-sign and comma. The part following the at-sign MUST be a valid fully qualified internet domain name [RFC-1034] controlled by the person or organization defining the name. It is up to each domain how it manages its local namespace.

6 Message Numbers

SSH packets have message numbers in the range 1-255. These numbers have been allocated as follows:

```
Transport layer protocol:

  1-19     Transport layer generic (e.g. disconnect, ignore, debug,
           etc)
  20-29    Algorithm negotiation
  30-49    Key exchange method specific (numbers can be reused for
           different authentication methods)

User authentication protocol:

  50-59    User authentication generic
  60-79    User authentication method specific (numbers can be reused
           for different authentication methods)

Connection protocol:

  80-89    Connection protocol generic
  90-127   Channel related messages

Reserved for client protocols:

  128-191  Reserved

Local extensions:

  192-255  Local extensions
```

7 IANA Considerations

Allocation of the following types of names in the SSH protocols is assigned to IANA:

- encryption algorithm names,
- MAC algorithm names,
- public key algorithm names (public key algorithm also implies encoding and signature/encryption capability),
- key exchange method names, and
- protocol (service) names.

The IANA-allocated names MUST be printable US-ASCII strings, and MUST NOT contain the characters at-sign ('@'), comma (','), or whitespace or control characters (ascii codes 32 or less). Names are case-sensitive, and MUST not be longer than 64 characters.

Each category of names listed above has a separate namespace. However, using the same name in multiple categories SHOULD be avoided to minimize confusion.

8 Security Considerations

Special care should be taken to ensure that all of the random numbers are of good quality. The random numbers SHOULD be produced with safe mechanisms discussed in [RFC1750].

When displaying text, such as error or debug messages to the user, the client software SHOULD replace any control characters (except tab, carriage return and newline) with safe sequences to avoid attacks by sending terminal control characters.

Not using MAC or encryption SHOULD be avoided. The user authentication protocol is subject to man-in-the-middle attacks if the encryption is disabled. The SSH protocol does not protect against message alteration if no MAC is used.

9 References

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

[RFC-854] Postel, J. and Reynolds, J., "Telnet Protocol Specification", May 1983.

[RFC-894] Hornig, C., "A Standard for the Transmission of IP Datagrams over Ethernet Networks", April 1984.

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

[RFC-1134] Perkins, D., "The Point-to-Point Protocol: A Proposal for Multi-Protocol Transmission o Datagrams Over Point-to-Point Links", November 1989.

[RFC-1282] Kantor, B., "BSD Rlogin", December 1991.

[RFC-1700] Reynolds, J. and Postel, J., "Assigned Numbers", October 1994 (also STD 2).

[RFC-1750] Eastlake, D., Crocker, S., and Schiller, J., "Randomness Recommendations for Security", December 1994.

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

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

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

[Schneier] Schneier, B., "Applied Cryptography Second Edition", John Wiley & Sons, New York, NY, 1995.

[SSH-TRANS] Ylonen, T., Kivinen, T, and Saarinen, M., "SSH Transport Layer Protocol", Internet Draft, draft-ietf-secsh-transport-02.txt

[SSH-USERAUTH] Ylonen, T., Kivinen, T, and Saarinen, M., "SSH Authentication Protocol", Internet Draft, draft-ietf-secsh-userauth-02.txt

[SSH-CONNECT] Ylonen, T., Kivinen, T, and Saarinen, M., "SSH Connection Protocol", Internet Draft, draft-ietf-secsh-connect-02.txt

10 Authors' Addresses

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

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

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

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

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