

The WebSocket protocol
draft-ietf-hybi-thewebsocketprotocol-03

Abstract

The WebSocket protocol enables two-way communication between a user agent running untrusted code running in a controlled environment to a remote host that has opted-in to communications from that code. The security model used for this is the Origin-based security model commonly used by Web browsers. The protocol consists of an initial handshake followed by basic message framing, layered over TCP. The goal of this technology is to provide a mechanism for browser-based applications that need two-way communication with servers that does not rely on opening multiple HTTP connections (e.g. using XMLHttpRequest or <iframe>s and long polling).

Please send feedback to the hybi@ietf.org mailing list.

Note

This draft is meant to reflect changes in direction in the HyBi working group. There is not yet consensus on everything in this draft. Specifically, details about the framing are still under discussion, however this draft is much closer to what the group is discussing than the previous draft. There have also been proposals to change the handshake, so the handshake is also not in a final form.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 20, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

<u>1.</u>	Introduction	<u>5</u>
<u>1.1.</u>	Background	<u>5</u>
<u>1.2.</u>	Protocol overview	<u>5</u>
<u>1.3.</u>	Opening handshake	<u>8</u>
<u>1.4.</u>	Closing handshake	<u>11</u>
<u>1.5.</u>	Design philosophy	<u>12</u>
<u>1.6.</u>	Security model	<u>12</u>
<u>1.7.</u>	Relationship to TCP and HTTP	<u>13</u>
<u>1.8.</u>	Establishing a connection	<u>13</u>
<u>1.9.</u>	Subprotocols using the WebSocket protocol	<u>14</u>
<u>2.</u>	Conformance requirements	<u>15</u>
<u>2.1.</u>	Terminology	<u>15</u>
<u>3.</u>	WebSocket URLs	<u>17</u>
<u>3.1.</u>	Parsing WebSocket URLs	<u>17</u>
<u>3.2.</u>	Constructing WebSocket URLs	<u>18</u>
<u>3.3.</u>	Valid WebSocket URLs	<u>18</u>
<u>4.</u>	Data Framing	<u>20</u>
<u>4.1.</u>	Overview	<u>20</u>
<u>4.2.</u>	Base Framing Protocol	<u>20</u>
<u>4.3.</u>	Fragmentation	<u>22</u>
<u>4.4.</u>	Control Frames	<u>23</u>
<u>4.4.1.</u>	Close	<u>23</u>
<u>4.4.2.</u>	Ping	<u>24</u>
<u>4.4.3.</u>	Pong	<u>24</u>
<u>4.5.</u>	Data Frames	<u>24</u>
<u>4.6.</u>	Examples	<u>24</u>
<u>4.7.</u>	Extensibility	<u>25</u>
<u>5.</u>	Opening Handshake	<u>26</u>
<u>5.1.</u>	Client Requirements	<u>26</u>
<u>5.2.</u>	Server-side requirements	<u>36</u>
<u>5.2.1.</u>	Reading the client's opening handshake	<u>36</u>
<u>5.2.2.</u>	Sending the server's opening handshake	<u>39</u>
<u>6.</u>	Error Handling	<u>44</u>
<u>6.1.</u>	Handling errors in UTF-8 from the server	<u>44</u>
<u>6.2.</u>	Handling errors in UTF-8 from the client	<u>44</u>
<u>7.</u>	Closing the connection	<u>45</u>
<u>7.1.</u>	Client-initiated closure	<u>45</u>
<u>7.2.</u>	Server-initiated closure	<u>45</u>
<u>7.3.</u>	Closure	<u>45</u>
<u>8.</u>	Known extensions	<u>47</u>
<u>8.1.</u>	Compression	<u>47</u>
<u>9.</u>	Security considerations	<u>48</u>
<u>10.</u>	IANA considerations	<u>49</u>
<u>10.1.</u>	Registration of ws: scheme	<u>49</u>
<u>10.2.</u>	Registration of wss: scheme	<u>50</u>
<u>10.3.</u>	Registration of the "WebSocket" HTTP Upgrade keyword	<u>51</u>

10.4.	Sec-WebSocket-Key1 and Sec-WebSocket-Key2	51
10.5.	Sec-WebSocket-Location	52
10.6.	Sec-WebSocket-Origin	53
10.7.	Sec-WebSocket-Protocol	54
10.8.	Sec-WebSocket-Draft	54
11.	Using the WebSocket protocol from other specifications	56
12.	Acknowledgements	57
13.	Normative References	58
	Author's Address	60

[1.](#) Introduction

[1.1.](#) Background

`_This section is non-normative._`

Historically, creating an instant messenger chat client as a Web application has required an abuse of HTTP to poll the server for updates while sending upstream notifications as distinct HTTP calls.

This results in a variety of problems:

- `o The server is forced to use a number of different underlying TCP connections for each client: one for sending information to the client, and a new one for each incoming message.`
- `o The wire protocol has a high overhead, with each client-to-server message having an HTTP header.`
- `o The client-side script is forced to maintain a mapping from the outgoing connections to the incoming connection to track replies.`

A simpler solution would be to use a single TCP connection for traffic in both directions. This is what the WebSocket protocol provides. Combined with the WebSocket API, it provides an alternative to HTTP polling for two-way communication from a Web page to a remote server. [[WSAPI](#)]

The same technique can be used for a variety of Web applications: games, stock tickers, multiuser applications with simultaneous editing, user interfaces exposing server-side services in real time, etc.

[1.2.](#) Protocol overview

`_This section is non-normative._`

The protocol has two parts: a handshake, and then the data transfer.

The handshake from the client looks as follows:

```
GET /demo HTTP/1.1
Host: example.com
Connection: Upgrade
Sec-WebSocket-Key2: 12998 5 Y3 1 .P00
Sec-WebSocket-Protocol: sample
Upgrade: WebSocket
Sec-WebSocket-Key1: 4 @1 46546xW%0l 1 5
Origin: http://example.com
```

```
^n:ds[4U
```

The handshake from the server looks as follows:

```
HTTP/1.1 101 WebSocket Protocol Handshake
Upgrade: WebSocket
Connection: Upgrade
Sec-WebSocket-Origin: http://example.com
Sec-WebSocket-Location: ws://example.com/demo
Sec-WebSocket-Protocol: sample
```

```
8jKS'y:G*Co,Wxa-
```

The leading line from the client follows the Request-Line format. The leading line from the server follows the Status-Line format. The Request-Line and Status-Line productions are defined in the HTTP specification.

After the leading line in both cases come an unordered ASCII case-insensitive set of fields, one per line, that each match the following non-normative ABNF: [[RFC5234](#)] [[ANSI.X3-4.1986](#)]

```
field      = 1*name-char colon [ space ] *any-char cr lf
colon      = %x003A ; U+003A COLON (:)
space      = %x0020 ; U+0020 SPACE
cr         = %x000D ; U+000D CARRIAGE RETURN (CR)
lf         = %x000A ; U+000A LINE FEED (LF)
name-char  = %x0000-0009
            / %x000B-000C
            / %x000E-0039
            / %x003B-10FFFF
            ; a Unicode character other than
            ; U+000A LINE FEED (LF),
            ; U+000D CARRIAGE RETURN (CR),
            ; or U+003A COLON (:)
any-char   = %x0000-0009 / %x000B-000C / %x000E-10FFFF
            ; a Unicode character other than
            ; U+000A LINE FEED (LF)
            ; or U+000D CARRIAGE RETURN (CR)
```

NOTE: The character set for the above ABNF is Unicode. The fields themselves are encoded as UTF-8.

Lines that don't match the above production cause the connection to be aborted.

Finally, after the last field, the client sends 10 bytes starting with 0x0D 0x0A and followed by 8 random bytes, part of a challenge, and the server sends 18 bytes starting with 0x0D 0x0A and followed by 16 bytes consisting of a challenge response. The details of this challenge and other parts of the handshake are described in the next section.

Once the client and server have both sent their handshakes, and if the handshake was successful, then the data transfer part starts. This is a two-way communication channel where each side can, independently from the other, send data at will.

Clients and servers, after a successful handshake, transfer data back and forth in conceptual units referred to in this specification as "messages". A message is a complete unit of data at an application level, with the expectation that many or most applications implementing this protocol (such as web user agents) provide APIs in terms of sending and receiving messages. The websocket message does not necessarily correspond to a particular network layer framing, as a fragmented message may be coalesced, or vice versa, e.g. by an intermediary.

Data is sent on the wire in the form of frames that have an

associated type. Broadly speaking, there are types for textual data, which is interpreted as UTF-8 text, binary data (whose interpretation is left up to the application), and control frames, which are not intended to carry data for the application, but instead for protocol-level signalling, such as to signal that the connection should be closed. This version of the protocol defines six frame types and leaves ten reserved for future use.

The WebSocket protocol uses this framing so that specifications that use the WebSocket protocol can expose such connections using an event-based mechanism instead of requiring users of those specifications to implement buffering and piecing together of messages manually.

[1.3.](#) Opening handshake

`_This section is non-normative._`

The opening handshake is intended to be compatible with HTTP-based server-side software, so that a single port can be used by both HTTP clients talking to that server and WebSocket clients talking to that server. To this end, the WebSocket client's handshake appears to HTTP servers to be a regular GET request with an Upgrade offer:

```
GET / HTTP/1.1
Upgrade: WebSocket
Connection: Upgrade
```

Fields in the handshake are sent by the client in a random order; the order is not meaningful.

Additional fields are used to select options in the WebSocket protocol. The only options available in this version are the subprotocol selector, `|Sec-WebSocket-Protocol|`, and `|Cookie|`, which can be used for sending cookies to the server (e.g. as an authentication mechanism). The `|Sec-WebSocket-Protocol|` request-header field can be used to indicate what subprotocols (application-level protocols layered over the WebSocket protocol) are acceptable to the client. The server selects one of the acceptable protocols and echoes that value in its handshake to indicate that it has selected that protocol.

```
Sec-WebSocket-Protocol: chat superchat
```

The other fields in the handshake are all security-related. The `|Host|` field is used to protect against DNS rebinding attacks and to allow multiple domains to be served from one IP address.

Host: example.com

The server includes the hostname in the |Sec-WebSocket-Location| field of its handshake, so that both the client and the server can verify that they agree on which host is in use.

The |Origin| field is used to protect against unauthorized cross-origin use of a WebSocket server by scripts using the |WebSocket| API in a Web browser. The server specifies which origin it is willing to receive requests from by including a |Sec-WebSocket-Origin| field with that origin. If multiple origins are authorized, the server echoes the value in the |Origin| field of the client's handshake.

Origin: http://example.com

Finally, the server has to prove to the client that it received the client's WebSocket handshake, so that the server doesn't accept connections that are not WebSocket connections. This prevents an attacker from tricking a WebSocket server by sending it carefully-crafted packets using |XMLHttpRequest| or a |form| submission.

To prove that the handshake was received, the server has to take three pieces of information and combine them to form a response. The first two pieces of information come from the |Sec-WebSocket-Key1| and |Sec-WebSocket-Key2| fields in the client handshake:

```
Sec-WebSocket-Key1: 18x 6]8vM;54 *(5: { U1]8 z [ 8
Sec-WebSocket-Key2: 1_ tx7X d < nw 334J702) 7]o}` 0
```

For each of these fields, the server has to take the digits from the value to obtain a number (in this case 1868545188 and 1733470270 respectively), then divide that number by the number of spaces characters in the value (in this case 12 and 10) to obtain a 32-bit number (155712099 and 173347027). These two resulting numbers are then used in the server handshake, as described below.

The counting of spaces is intended to make it impossible to smuggle this field into the resource name; making this even harder is the presence of two such fields, and the use of a newline as the only reliable indicator that the end of the key has been reached. The use of random characters interspersed with the spaces and the numbers ensures that the implementor actually looks for spaces and newlines, instead of being treating any character like a space, which would make it again easy to smuggle the fields into the path and trick the server. Finally, dividing by this number of spaces is intended to make sure that even the most naive of implementations will check for spaces, since if the server does not verify that there are some

spaces, the server will try to divide by zero, which is usually fatal (a correct handshake will always have at least one space).

The third piece of information is given after the fields, in the last eight bytes of the handshake, expressed here as they would be seen if interpreted as UTF-8:

```
Tm[K T2u
```

The concatenation of the number obtained from processing the |Sec-WebSocket-Key1| field, expressed as a big-endian 32 bit number, the number obtained from processing the |Sec-WebSocket-Key2| field, again expressed as a big-endian 32 bit number, and finally the eight bytes at the end of the handshake, form a 128 bit string whose MD5 sum is then used by the server to prove that it read the handshake.

The handshake from the server is much simpler than the client handshake. The first line is an HTTP Status-Line, with the status code 101 (the HTTP version and reason phrase aren't important):

```
HTTP/1.1 101 WebSocket Protocol Handshake
```

The fields follow. Two of the fields are just for compatibility with HTTP:

```
Upgrade: WebSocket  
Connection: Upgrade
```

Two of the fields are part of the security model described above, echoing the origin and stating the exact host, port, resource name, and whether the connection is expected to be encrypted:

```
Sec-WebSocket-Origin: http://example.com  
Sec-WebSocket-Location: ws://example.com/
```

These fields are checked by the Web browser when it is acting as a |WebSocket| client for scripted pages. A server that only handles one origin and only serves one resource can therefore just return hard-coded values and does not need to parse the client's handshake to verify the correctness of the values.

Option fields can also be included. In this version of the protocol, the main option field is |Sec-WebSocket-Protocol|, which indicates the subprotocol that the server has selected. Web browsers verify that the server included one of the values as was specified in the

|WebSocket| constructor. A server that speaks multiple subprotocols has to make sure it selects one based on the client's handshake and specifies it in its handshake.

Sec-WebSocket-Protocol: chat

The server can also set cookie-related option fields to `_set_cookies`, as in HTTP.

After the fields, the server sends the aforementioned MD5 sum, a 16 byte (128 bit) value, shown here as if interpreted as UTF-8:

fQJ,fN/4F4!~K~MH

This value depends on what the client sends, as described above. If it doesn't match what the client is expecting, the client would disconnect.

Having part of the handshake appear after the fields ensures that both the server and the client verify that the connection is not being interrupted by an HTTP intermediary such as a man-in-the-middle cache or proxy.

[1.4.](#) Closing handshake

`_This section is non-normative._`

The closing handshake is far simpler than the opening handshake.

Either peer can send a control frame with data containing a specified control sequence to begin the closing handshake. Upon receiving such a frame, the other peer sends an identical frame in acknowledgement, if it hasn't already sent one. Upon receiving `_that_` control frame, the first peer then closes the connection, safe in the knowledge that no further data is forthcoming.

After sending a control frame indicating the connection should be closed, a peer does not send any further data; after receiving a control frame frame indicating the connection should be closed, a peer discards any further data received.

It is safe for both peers to initiate this handshake simultaneously.

The closing handshake is intended to replace the TCP closing handshake (FIN/ACK), on the basis that the TCP closing handshake is not always reliable end-to-end, especially in the presence of man-in-the-middle proxies and other intermediaries.

[1.5.](#) Design philosophy

This section is non-normative.

The WebSocket protocol is designed on the principle that there should be minimal framing (the only framing that exists is to make the protocol frame-based instead of stream-based, and to support a distinction between Unicode text and binary frames). It is expected that metadata would be layered on top of WebSocket by the application layer, in the same way that metadata is layered on top of TCP by the application layer (HTTP).

Conceptually, WebSocket is really just a layer on top of TCP that adds a Web "origin"-based security model for browsers; adds an addressing and protocol naming mechanism to support multiple services on one port and multiple host names on one IP address; layers a framing mechanism on top of TCP to get back to the IP packet mechanism that TCP is built on, but without length limits; and reimplements the closing handshake in-band. Other than that, it adds nothing. Basically it is intended to be as close to just exposing raw TCP to script as possible given the constraints of the Web. It's also designed in such a way that its servers can share a port with HTTP servers, by having its handshake be a valid HTTP Upgrade handshake also.

The protocol is intended to be extensible; future versions will likely introduce a mechanism to compress data and might support sending binary data.

[1.6.](#) Security model

This section is non-normative.

The WebSocket protocol uses the origin model used by Web browsers to restrict which Web pages can contact a WebSocket server when the WebSocket protocol is used from a Web page. Naturally, when the WebSocket protocol is used by a dedicated client directly (i.e. not from a Web page through a Web browser), the origin model is not useful, as the client can provide any arbitrary origin string.

This protocol is intended to fail to establish a connection with servers of pre-existing protocols like SMTP or HTTP, while allowing HTTP servers to opt-in to supporting this protocol if desired. This is achieved by having a strict and elaborate handshake, and by limiting the data that can be inserted into the connection before the handshake is finished (thus limiting how much the server can be influenced).

It is similarly intended to fail to establish a connection when data from other protocols, especially HTTP, is sent to a WebSocket server, for example as might happen if an HTML `|form|` were submitted to a WebSocket server. This is primarily achieved by requiring that the server prove that it read the handshake, which it can only do if the handshake contains the appropriate parts which themselves can only be sent by a WebSocket handshake; in particular, fields starting with `|Sec-|` cannot be set by an attacker from a Web browser, even when using `|XMLHttpRequest|`.

[1.7.](#) Relationship to TCP and HTTP

`_This section is non-normative._`

The WebSocket protocol is an independent TCP-based protocol. Its only relationship to HTTP is that its handshake is interpreted by HTTP servers as an Upgrade request.

Based on the expert recommendation of the IANA, the WebSocket protocol by default uses port 80 for regular WebSocket connections and port 443 for WebSocket connections tunneled over TLS.

[1.8.](#) Establishing a connection

`_This section is non-normative._`

There are several options for establishing a WebSocket connection.

On the face of it, the simplest method would seem to be to use port 80 to get a direct connection to a WebSocket server. Port 80 traffic, however, will often be intercepted by man-in-the-middle HTTP proxies, which can lead to the connection failing to be established.

The most reliable method, therefore, is to use TLS encryption and port 443 to connect directly to a WebSocket server. This has the advantage of being more secure; however, TLS encryption can be computationally expensive.

When a connection is to be made to a port that is shared by an HTTP server (a situation that is quite likely to occur with traffic to ports 80 and 443), the connection will appear to the HTTP server to be a regular GET request with an Upgrade offer. In relatively simple setups with just one IP address and a single server for all traffic to a single hostname, this might allow a practical way for systems based on the WebSocket protocol to be deployed. In more elaborate setups (e.g. with load balancers and multiple servers), a dedicated set of hosts for WebSocket connections separate from the HTTP servers is probably easier to manage.

[1.9.](#) Subprotocols using the WebSocket protocol

`_This section is non-normative._`

The client can request that the server use a specific subprotocol by including the `|Sec-WebSocket-Protocol|` field in its handshake. If it is specified, the server needs to include the same field and one of the selected subprotocol values in its response for the connection to be established.

These subprotocol names do not need to be registered, but if a subprotocol is intended to be implemented by multiple independent WebSocket servers, potential clashes with the names of subprotocols defined independently can be avoided by using names that contain the domain name of the subprotocol's originator. For example, if Example Corporation were to create a Chat subprotocol to be implemented by many servers around the Web, they could name it `"chat.example.com"`. If the Example Organisation called their competing subprotocol `"example.org's chat protocol"`, then the two subprotocols could be implemented by servers simultaneously, with the server dynamically selecting which subprotocol to use based on the value sent by the client.

Subprotocols can be versioned in backwards-incompatible ways by changing the subprotocol name, eg. going from `"bookings.example.net"` to `"v2.bookings.example.net"`. These subprotocols would be considered completely separate by WebSocket clients. Backwards-compatible versioning can be implemented by reusing the same subprotocol string but carefully designing the actual subprotocol to support this kind of extensibility.

2. Conformance requirements

All diagrams, examples, and notes in this specification are non-normative, as are all sections explicitly marked non-normative. Everything else in this specification is normative.

The key words "MUST", "MUST NOT", "REQUIRED", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in the normative parts of this document are to be interpreted as described in [RFC2119](#). For readability, these words do not appear in all uppercase letters in this specification. [[RFC2119](#)]

Requirements phrased in the imperative as part of algorithms (such as "strip any leading space characters" or "return false and abort these steps") are to be interpreted with the meaning of the key word ("must", "should", "may", etc) used in introducing the algorithm.

Conformance requirements phrased as algorithms or specific steps may be implemented in any manner, so long as the end result is equivalent. (In particular, the algorithms defined in this specification are intended to be easy to follow, and not intended to be performant.)

Implementations may impose implementation-specific limits on otherwise unconstrained inputs, e.g. to prevent denial of service attacks, to guard against running out of memory, or to work around platform-specific limitations.

The conformance classes defined by this specification are user agents and servers.

2.1. Terminology

ASCII shall mean the character-encoding scheme defined in [[ANSI.X3-4.1986](#)].

Converting a string to ASCII lowercase means replacing all characters in the range U+0041 to U+005A (i.e. LATIN CAPITAL LETTER A to LATIN CAPITAL LETTER Z) with the corresponding characters in the range U+0061 to U+007A (i.e. LATIN SMALL LETTER A to LATIN SMALL LETTER Z).

Comparing two strings in an *ASCII case-insensitive* manner means comparing them exactly, code point for code point, except that the characters in the range U+0041 to U+005A (i.e. LATIN CAPITAL LETTER A to LATIN CAPITAL LETTER Z) and the corresponding characters in the range U+0061 to U+007A (i.e. LATIN SMALL LETTER A to LATIN SMALL LETTER Z) are considered to also match.

The term "URL" is used in this section in a manner consistent with the terminology used in HTML, namely, to denote a string that might or might not be a valid URI or IRI and to which certain error handling behaviors will be applied when the string is parsed. [[HTML](#)]

When an implementation is required to `_send_` data as part of the WebSocket protocol, the implementation may delay the actual transmission arbitrarily, e.g. buffering data so as to send fewer IP packets.

[3.](#) WebSocket URLs

[3.1.](#) Parsing WebSocket URLs

The steps to **parse a WebSocket URL's components** from a string `/url/` are as follows. These steps return either a `/host/`, a `/port/`, a `/resource name/`, and a `/secure/` flag, or they fail.

1. If the `/url/` string is not an absolute URL, then fail this algorithm. [[RFC3986](#)] [[RFC3987](#)]
2. Resolve the `/url/` string using the resolve a Web address algorithm defined by the Web addresses specification, with the URL character encoding set to UTF-8. [[RFC3629](#)] [[RFC3986](#)] [[RFC3987](#)]

NOTE: It doesn't matter what it is resolved relative to, since we already know it is an absolute URL at this point.

3. If `/url/` does not have a `<scheme>` component whose value, when converted to ASCII lowercase, is either "ws" or "wss", then fail this algorithm.
4. If `/url/` has a `<fragment>` component, then fail this algorithm.
5. If the `<scheme>` component of `/url/` is "ws", set `/secure/` to false; otherwise, the `<scheme>` component is "wss", set `/secure/` to true.
6. Let `/host/` be the value of the `<host>` component of `/url/`, converted to ASCII lowercase.
7. If `/url/` has a `<port>` component, then let `/port/` be that component's value; otherwise, there is no explicit `/port/`.
8. If there is no explicit `/port/`, then: if `/secure/` is false, let `/port/` be 80, otherwise let `/port/` be 443.
9. Let `/resource name/` be the value of the `<path>` component (which might be empty) of `/url/`.
10. If `/resource name/` is the empty string, set it to a single character U+002F SOLIDUS (/).
11. If `/url/` has a `<query>` component, then append a single U+003F QUESTION MARK character (?) to `/resource name/`, followed by the value of the `<query>` component.

12. Return `/host/`, `/port/`, `/resource name/`, and `/secure/`.

[3.2.](#) Constructing WebSocket URLs

The steps to *construct a WebSocket URL* from a `/host/`, a `/port/`, a `/resource name/`, and a `/secure/` flag, are as follows:

1. Let `/url/` be the empty string.
2. If the `/secure/` flag is false, then append the string `"ws://"` to `/url/`. Otherwise, append the string `"wss://"` to `/url/`.
3. Append `/host/` to `/url/`.
4. If the `/secure/` flag is false and port is not 80, or if the `/secure/` flag is true and port is not 443, then append the string `":"` followed by `/port/` to `/url/`.
5. Append `/resource name/` to `/url/`.
6. Return `/url/`.

[3.3.](#) Valid WebSocket URLs

For a WebSocket URL to be considered valid, the following conditions MUST hold.

- o The `/host/` must be ASCII-only (i.e. it must have been punycode-encoded already if necessary, and MUST NOT contain any characters above U+007E).
- o The `/origin/` must not contain characters in the range U+0041 to U+005A (i.e. LATIN CAPITAL LETTER A to LATIN CAPITAL LETTER Z).
- o The `/resource name/` string must be a non-empty string of characters in the range U+0021 to U+007E that starts with a U+002F SOLIDUS character (/).
- o The various strings in `/protocols/` MUST all be non-empty strings with characters in the range U+0021 to U+007E and MUST all be unique.

Any WebSocket URLs not meeting the above criteria are considered invalid, and a client MUST NOT attempt to make a connection to an invalid WebSocket URL. A client SHOULD attempt to parse a URL obtained from any external source (such as a web site or a user) using the steps specified in [Section 3.1](#) to obtain a valid WebSocket URL, but MUST NOT attempt to connect with such an unparsed URL, and

instead only use the parsed version and only if that version is considered valid by the criteria above.

Opcode: 4 bits

Defines the interpretation of the payload data

Payload length: 7 bits

The length of the payload: if 0-125, that is the payload length. If 126, the following 2 bytes interpreted as a 16 bit unsigned integer are the payload length. If 127, the following 8 bytes interpreted as a 64-bit unsigned integer (the high bit must be 0) are the payload length. Multibyte length quantities are expressed in network byte order. The payload length is the length of the Extension data + the length of the Application Data. The length of the Extension data may be zero, in which case the Payload length is the length of the Application data.

Extension data: n bytes

The extension data is 0 bytes unless there is a reserved op-code or reserved bit present in the frame which indicates an extension has been negotiated. Any extension MUST specify the length of the extension data, or how that length may be calculated, and its use MUST be negotiated during the handshake. If present, the extension data is included in the total payload length.

Application data: n bytes

Arbitrary application data, taking up the remainder of the frame after any extension data. The length of the Application data is equal to the payload length minus the length of the Extension data.

The base framing protocol is formally defined by the following ABNF [[RFC5234](#)]:

ws-frame	= frame-more frame-rsv1 frame-rsv2 frame-rsv3 frame-opcode frame-rsv4 frame-length frame-extension application-data;
frame-more	= %x0 ; final frame of message / %x1 ; more frames of this message follow
frame-rsv1	= %x0 ; 1 bit, must be 0
frame-rsv2	= %x0 ; 1 bit, must be 0
frame-rsv3	= %x0 ; 1 bit, must be 0
frame-opcode	= %x0 ; continuation frame / %x1 ; connection close / %x2 ; ping / %x3 ; pong / %x4 ; text frame / %x5 ; binary frame / %x6-F ; reserved
frame-rsv4	= %x0 ; 1 bit, must be 0
frame-length	= %x00-7D / %x7E frame-length-16 / %x7F frame-length-63
frame-length-16	= %x0000-FFFF
frame-length-63	= %x0000000000000000-7FFFFFFFFFFFFFFF
frame-extension	= *(%x00-FF) ; to be defined later
application-data	= *(%x00-FF)

[4.3.](#) Fragmentation

The following rules apply to fragmentation:

- o An unfragmented message consists of a single frame with the MORE bit clear and an opcode other than 0.

- o A fragmented message consists of a single frame with the MORE bit set and an opcode other than 0, followed by zero or more frames with the MORE bit set and the opcode set to 0, and terminated by a single frame with the MORE bit clear and an opcode of 0. Its content is the concatenation of the application data from each of those frames in order.
- o Note: There is an open question as to whether control frames be interjected in the middle of a fragmented message. If so, it must be decided whether they be fragmented (which would require keeping a stack of "in-progress" messages).__
- o A sender MAY create fragments of any size for non control messages.
- o Clients and servers MUST support receiving both fragmented and unfragmented messages.
- o An intermediary MAY change the fragmentation of a message if the message uses only opcode and reserved bit values known to the intermediary.

[4.4.](#) Control Frames

Control frames have opcodes of 0x01 (Close), 0x02 (Ping), or 0x03 (Pong). Control frames are used to communicate state about the websocket.

All control frames MUST be 125 bytes or less in length and MUST NOT be fragmented.

[4.4.1.](#) Close

The Close message contains an opcode of 0x01.

The application MUST NOT send any more data messages after sending a close message.

A received close message is deemed to be an acknowledgement if the message body matches the body of a close message previously sent by the receiver. Otherwise the close message is a close initiated by the sender.

Upon receipt of an initiated close the endpoint MUST send a close acknowledgment. It should do so as soon as is practical.

The websocket is considered fully closed when an endpoint has either received a close acknowledgment or sent a close acknowledgment.

[4.4.2.](#) Ping

The Ping message contains an opcode of 0x02.

Upon receipt of a Ping message, an endpoint MUST send a Pong message in response. It SHOULD do so as soon as is practical. The message bodies of the Ping and Pong MUST be the same.

[4.4.3.](#) Pong

The Pong message contains an opcode of 0x03.

Upon receipt of a Ping message, an endpoint MUST send a Pong message in response. It SHOULD do so as soon as is practical. The message bodies of the Ping and Pong MUST be the same.

[4.5.](#) Data Frames

All frame types not listed in [Section 4.4](#) are data frames, which transport application-layer data. The opcode determines the interpretation of the application data:

Text

The payload data is text data encoded as UTF-8.

Binary

The payload data is arbitrary binary data whose interpretation is solely up to the application layer.

[4.6.](#) Examples

This section is non-normative.

- o A single-frame text message

- * 0x04 0x05 "Hello"

- o A fragmented text message

- * 0x84 0x03 "Hel"

- * 0x00 0x02 "lo"

- o Ping request and response

- * 0x02 0x05 "Hello"
- * 0x03 0x05 "Hello"
- o 256 bytes binary message in a single frame
 - * 0x05 0x7E 0x0100 [256 bytes of binary data]
- o 64KiB binary message in a single frame
 - * 0x05 0x7F 0x000000000000010000 [65536 bytes of binary data]

4.7. Extensibility

The protocol is designed to allow for extensions, which will add capabilities to the base protocols. The endpoints of a connection MUST negotiate the use of any extensions during the handshake. This specification provides opcodes 0x6 through 0xF, the extension data field, and the frame-rsv1, frame-rsv2, frame-rsv3, and frame-rsv4 bits of the frame header for use by extensions. Below are some anticipated uses of extensions. This list is neither complete nor proscriptive.

- o Extension data may be placed in the payload before the application data.
- o Reserved bits can be allocated for per-frame needs.
- o Reserved opcode values can be defined.
- o Reserved bits can be allocated to the opcode field if more opcode values are needed.
- o A reserved bit or an "extension" opcode can be defined which allocates additional bits out of the payload area to define larger opcodes or more per-frame bits.

[5.](#) Opening Handshake

[5.1.](#) Client Requirements

User agents running in controlled environments, e.g. browsers on mobile handsets tied to specific carriers, may offload the management of the connection to another agent on the network. In such a situation, the user agent for the purposes of conformance is considered to include both the handset software and any such agents.

When the user agent is to *establish a WebSocket connection* to a host */host/*, on a port */port/*, from an origin whose ASCII serialization is */origin/*, with a flag */secure/*, with a string giving a */resource name/*, with a (possibly empty) list of strings giving the */protocols/*, and optionally with a */defer cookies/* flag, it must run the following steps. [[ORIGIN](#)]

1. Verify that the WebSocket URL and its components are valid according to [Section 3.3](#). If any of the requirements are not met, the client **MUST** fail the WebSocket connection and abort these steps.
2. If the user agent already has a WebSocket connection to the remote host (IP address) identified by */host/*, even if known by another name, wait until that connection has been established or for that connection to have failed. If multiple connections to the same IP address are attempted simultaneously, the user agent must serialize them so that there is no more than one connection at a time running through the following steps.

If the user agent cannot determine the IP address of the remote host (for example because all communication is being done through a proxy server that performs DNS queries itself), then the user agent must assume for the purposes of this step that each host name refers to a distinct remote host, but should instead limit the total number of simultaneous connections that are not established to a reasonably low number (e.g., in a Web browser, to the number of tabs the user has open).

NOTE: This makes it harder for a script to perform a denial of service attack by just opening a large number of WebSocket connections to a remote host. A server can further reduce the load on itself when attacked by making use of this by pausing before closing the connection, as that will reduce the rate at which the client reconnects.

NOTE: There is no limit to the number of established WebSocket connections a user agent can have with a single remote host.

Servers can refuse to connect users with an excessive number of connections, or disconnect resource-hogging users when suffering high load.

3. `_Connect_`: If the user agent is configured to use a proxy when using the WebSocket protocol to connect to host `/host/` and/or port `/port/`, then connect to that proxy and ask it to open a TCP connection to the host given by `/host/` and the port given by `/port/`.

EXAMPLE: For example, if the user agent uses an HTTP proxy for all traffic, then if it was to try to connect to port 80 on server `example.com`, it might send the following lines to the proxy server:

```
CONNECT example.com:80 HTTP/1.1
Host: example.com
```

If there was a password, the connection might look like:

```
CONNECT example.com:80 HTTP/1.1
Host: example.com
Proxy-authorization: Basic ZWRuYW1vZGU6bm9jYXBlcY=
```

Otherwise, if the user agent is not configured to use a proxy, then open a TCP connection to the host given by `/host/` and the port given by `/port/`.

NOTE: Implementations that do not expose explicit UI for selecting a proxy for WebSocket connections separate from other proxies are encouraged to use a SOCKS proxy for WebSocket connections, if available, or failing that, to prefer the proxy configured for HTTPS connections over the proxy configured for HTTP connections.

For the purpose of proxy autoconfiguration scripts, the URL to pass the function must be constructed from `/host/`, `/port/`, `/resource name/`, and the `/secure/` flag using the steps to construct a WebSocket URL.

NOTE: The WebSocket protocol can be identified in proxy autoconfiguration scripts from the scheme (`"ws:"` for unencrypted connections and `"wss:"` for encrypted connections).

4. If the connection could not be opened, then fail the WebSocket connection and abort these steps.

5. If `/secure/` is true, perform a TLS handshake over the connection. If this fails (e.g. the server's certificate could not be verified), then fail the WebSocket connection and abort these steps. Otherwise, all further communication on this channel must run through the encrypted tunnel. [[RFC2246](#)]

User agents must use the Server Name Indication extension in the TLS handshake. [[RFC4366](#)]

6. Send the UTF-8 string "GET" followed by a UTF-8-encoded U+0020 SPACE character to the remote side (the server).

Send the `/resource name/` value, encoded as UTF-8.

Send another UTF-8-encoded U+0020 SPACE character, followed by the UTF-8 string "HTTP/1.1", followed by a UTF-8-encoded U+000D CARRIAGE RETURN U+000A LINE FEED character pair (CRLF).

7. Let `/fields/` be an empty list of strings.
8. Add the string "Upgrade: WebSocket" to `/fields/`.
9. Add the string "Connection: Upgrade" to `/fields/`.
10. Let `/hostport/` be an empty string.
11. Append the `/host/` value, converted to ASCII lowercase, to `/hostport/`.
12. If `/secure/` is false, and `/port/` is not 80, or if `/secure/` is true, and `/port/` is not 443, then append a U+003A COLON character (:) followed by the value of `/port/`, expressed as a base-ten integer, to `/hostport/`.
13. Add the string consisting of the concatenation of the string "Host:", a U+0020 SPACE character, and `/hostport/`, to `/fields/`.
14. Add the string consisting of the concatenation of the string "Origin:", a U+0020 SPACE character, and the `/origin/` value, to `/fields/`.
15. Add the string "Sec-WebSocket-Draft: 2" to `/fields/`.
16. If there is no `/protocols/`, then skip this step.

Otherwise, generate the acceptable protocol string by joining each protocol in `/protocols/` using a U+0020 SPACE character. Add the string consisting of the concatenation of the string

"Sec-WebSocket-Protocol:", a U+0020 SPACE character, and the acceptable protocol string generated above, to /fields/.

17. If the client has any cookies that would be relevant to a resource accessed over HTTP, if /secure/ is false, or HTTPS, if it is true, on host /host/, port /port/, with /resource name/ as the path (and possibly query parameters), then add to /fields/ any HTTP headers that would be appropriate for that information. [\[RFC2616\]](#) [\[RFC2109\]](#) [\[RFC2965\]](#)

This includes "HttpOnly" cookies (cookies with the http-only-flag set to true); the WebSocket protocol is not considered a non-HTTP API for the purpose of cookie processing.

18. When one or more HTTP headers are to be added to /fields/ for this step, each header must be added separately, and each header must be added as one entry consisting of the header's name in its canonical case, followed by a U+003A COLON character (:) and a U+0020 SPACE character, followed by the value with no use of continuation lines (e.g. containing no U+000A LINE FEED characters.)
19. Let /spaces_1/ be a random integer from 1 to 12 inclusive.
Let /spaces_2/ be a random integer from 1 to 12 inclusive.
EXAMPLE: For example, 5 and 9.
20. Let /max_1/ be the largest integer not greater than 4,294,967,295 divided by /spaces_1/.
Let /max_2/ be the largest integer not greater than 4,294,967,295 divided by /spaces_2/.
EXAMPLE: Continuing the example, 858,993,459 and 477,218,588.
21. Let /number_1/ be a random integer from 0 to /max_1/ inclusive.
Let /number_2/ be a random integer from 0 to /max_2/ inclusive.
EXAMPLE: For example, 777,007,543 and 114,997,259.
22. Let /product_1/ be the result of multiplying /number_1/ and /spaces_1/ together.
Let /product_2/ be the result of multiplying /number_2/ and /spaces_2/ together.

EXAMPLE: Continuing the example, 3,885,037,715 and 1,034,975,331.

23. Let /key_1/ be a string consisting of /product_1/, expressed in base ten using the numerals in the range U+0030 DIGIT ZERO (0) to U+0039 DIGIT NINE (9).

Let /key_2/ be a string consisting of /product_2/, expressed in base ten using the numerals in the range U+0030 DIGIT ZERO (0) to U+0039 DIGIT NINE (9).

EXAMPLE: Continuing the example, "3885037715" and "1034975331".

24. Insert between one and twelve random characters from the ranges U+0021 to U+002F and U+003A to U+007E into /key_1/ at random positions.

Insert between one and twelve random characters from the ranges U+0021 to U+002F and U+003A to U+007E into /key_2/ at random positions.

NOTE: This corresponds to random printable ASCII characters other than the digits and the U+0020 SPACE character.

EXAMPLE: Continuing the example, this could lead to "P3880503D&ul7{K%gX(%715" and "1N?|kUT0or3o4I97N5-S3031".

25. Insert /spaces_1/ U+0020 SPACE characters into /key_1/ at random positions other than the start or end of the string.

Insert /spaces_2/ U+0020 SPACE characters into /key_2/ at random positions other than the start or end of the string.

EXAMPLE: Continuing the example, this could lead to "P388 0503D&ul7 {K%gX(%7 15" and "1 N ?|k UT0or 3o 4 I97N 5-S30 31".

26. Add the string consisting of the concatenation of the string "Sec-WebSocket-Key1:", a U+0020 SPACE character, and the /key_1/ value, to /fields/.

Add the string consisting of the concatenation of the string "Sec-WebSocket-Key2:", a U+0020 SPACE character, and the /key_2/ value, to /fields/.

27. For each string in /fields/, in a random order: send the string, encoded as UTF-8, followed by a UTF-8-encoded U+000D CARRIAGE RETURN U+000A LINE FEED character pair (CRLF). It is important that the fields be output in a random order so that servers not

depend on the particular order used by any particular client.

28. Send a UTF-8-encoded U+000D CARRIAGE RETURN U+000A LINE FEED character pair (CRLF).
29. Let `/key_3/` be a string consisting of eight random bytes (or equivalently, a random 64 bit unsigned integer encoded in big-endian order).

EXAMPLE: For example, 0x47 0x30 0x22 0x2D 0x5A 0x3F 0x47 0x58.

30. Send `/key_3/` to the server.
31. Read bytes from the server until either the connection closes, or a 0x0A byte is read. Let `/field/` be these bytes, including the 0x0A byte.

If `/field/` is not at least seven bytes long, or if the last two bytes aren't 0x0D and 0x0A respectively, or if `/field/` contains any 0x0D bytes other than the penultimate byte, or if `/field/` does not contain at least two 0x20 bytes, then `_fail` the `WebSocket connection_` and abort these steps.

User agents may apply a timeout to this step, failing the `WebSocket connection` if the server does not send back data in a suitable time period.

32. Let `/code/` be the substring of `/field/` that starts from the byte after the first 0x20 byte, and ends with the byte before the second 0x20 byte.
33. If `/code/`, interpreted as UTF-8, is "101", then move to the next step.

If `/code/`, interpreted as UTF-8, is "407", then either close the connection and jump back to step 2, providing appropriate authentication information, or fail the `WebSocket connection`. 407 is the code used by HTTP meaning "Proxy Authentication Required". User agents that support proxy authentication must interpret the response as defined by HTTP (e.g. to find and interpret the `|Proxy-Authenticate|` header).

Otherwise, fail the `WebSocket connection` and abort these steps.

34. Let `/fields/` be a list of name-value pairs, initially empty.
35. `_Field_`: Let `/name/` and `/value/` be empty byte arrays.

36. Read a byte from the server.

If the connection closes before this byte is received, then fail the WebSocket connection and abort these steps.

Otherwise, handle the byte as described in the appropriate entry below:

- > If the byte is 0x0D (UTF-8 CR)
If the /name/ byte array is empty, then jump to the fields processing step. Otherwise, fail the WebSocket connection and abort these steps.
- > If the byte is 0x0A (UTF-8 LF)
Fail the WebSocket connection and abort these steps.
- > If the byte is 0x3A (UTF-8 :)
Move on to the next step.
- > If the byte is in the range 0x41 to 0x5A (UTF-8 A-Z)
Append a byte whose value is the byte's value plus 0x20 to the /name/ byte array and redo this step for the next byte.
- > Otherwise
Append the byte to the /name/ byte array and redo this step for the next byte.

NOTE: This reads a field name, terminated by a colon, converting upper-case letters in the range A-Z to lowercase, and aborting if a stray CR or LF is found.

37. Let /count/ equal 0.

NOTE: This is used in the next step to skip past a space character after the colon, if necessary.

38. Read a byte from the server and increment /count/ by 1.

If the connection closes before this byte is received, then fail the WebSocket connection and abort these steps.

Otherwise, handle the byte as described in the appropriate entry below:

- > If the byte is 0x20 (UTF-8 space) and /count/ equals 1
Ignore the byte and redo this step for the next byte.

- > If the byte is 0x0D (UTF-8 CR)
Move on to the next step.
- > If the byte is 0x0A (UTF-8 LF)
Fail the WebSocket connection and abort these steps.
- > Otherwise
Append the byte to the /value/ byte array and redo this step
for the next byte.

NOTE: This reads a field value, terminated by a CRLF, skipping past a single space after the colon if there is one.

39. Read a byte from the server.

If the connection closes before this byte is received, or if the byte is not a 0x0A byte (UTF-8 LF), then fail the WebSocket connection and abort these steps.

NOTE: This skips past the LF byte of the CRLF after the field.

40. Append an entry to the /fields/ list that has the name given by the string obtained by interpreting the /name/ byte array as a UTF-8 byte stream and the value given by the string obtained by interpreting the /value/ byte array as a UTF-8 byte stream.

41. Return to the "Field" step above.

42. _Fields processing_: Read a byte from the server.

If the connection closes before this byte is received, or if the byte is not a 0x0A byte (UTF-8 LF), then fail the WebSocket connection and abort these steps.

NOTE: This skips past the LF byte of the CRLF after the blank line after the fields.

43. Let the /list of cookies/ be empty.

44. If there is not exactly one entry in the /fields/ list whose name is "upgrade", or if there is not exactly one entry in the /fields/ list whose name is "connection", or if there is not exactly one entry in the /fields/ list whose name is "sec-websocket-origin", or if there is not exactly one entry in the /fields/ list whose name is "sec-websocket-location", or if the /protocol/ was specified but there is not exactly one entry in the /fields/ list whose name is "sec-websocket-protocol", or if there are any entries in the /fields/ list whose names are the

empty string, then fail the WebSocket connection and abort these steps. Otherwise, handle each entry in the /fields/ list as follows:

- > If the entry's name is "upgrade"
If the value, converted to ASCII lowercase, is not exactly equal to the string "websocket", then fail the WebSocket connection and abort these steps.
- > If the entry's name is "connection"
If the value, converted to ASCII lowercase, is not exactly equal to the string "upgrade", then fail the WebSocket connection and abort these steps.
- > If the entry's name is "sec-websocket-origin"
If the value is not exactly equal to /origin/, then fail the WebSocket connection and abort these steps. [[ORIGIN](#)]
- > If the entry's name is "sec-websocket-location"
If the value is not exactly equal to a string obtained from the steps to construct a WebSocket URL from /host/, /port/, /resource name/, and the /secure/ flag, then fail the WebSocket connection and abort these steps.
- > If the entry's name is "sec-websocket-protocol"
If there was a /protocols/ string specified, and the value is not exactly equal to one of the items in /protocols/, then fail the WebSocket connection and abort these steps. (If no /protocols/ was specified, the field is ignored.)
- > If the entry's name is "set-cookie" or "set-cookie2" or another cookie-related field name
If the relevant specification is supported by the user agent, add the cookie, interpreted as defined by the appropriate specification, to the /list of cookies/, with the resource being the one with the host /host/, the port /port/, the path (and possibly query parameters) /resource name/, and the scheme |http| if /secure/ is false and |https| if /secure/ is true. [[RFC2109](#)] [[RFC2965](#)]

If the relevant specification is not supported by the user agent, then the field must be ignored.

The cookies added to the /list of cookies/ are discarded if the connection fails to be established. Only if and when the connection is established do the cookies actually get applied.

-> Any other name
Ignore it.

45. Let /challenge/ be the concatenation of /number_1/, expressed as a big-endian 32 bit integer, /number_2/, expressed as a big-endian 32 bit integer, and the eight bytes of /key_3/ in the order they were sent on the wire.

EXAMPLE: Using the examples given earlier, this leads to the 16 bytes 0x2E 0x50 0x31 0xB7 0x06 0xDA 0xB8 0x0B 0x47 0x30 0x22 0x2D 0x5A 0x3F 0x47 0x58.

46. Let /expected/ be the MD5 fingerprint of /challenge/ as a big-endian 128 bit string. [[RFC1321](#)]

EXAMPLE: Using the examples given earlier, this leads to the 16 bytes 0x30 0x73 0x74 0x33 0x52 0x6C 0x26 0x71 0x2D 0x32 0x5A 0x55 0x5E 0x77 0x65 0x75. In UTF-8, these bytes correspond to the string "0st3Rl&q-2ZU^weu".

47. Read sixteen bytes from the server. Let /reply/ be those bytes.

If the connection closes before these bytes are received, then fail the WebSocket connection and abort these steps.

48. If /reply/ does not exactly equal /expected/, then fail the WebSocket connection and abort these steps.

49. If the /defer cookies/ flag is not set, apply the cookies in the /list of cookies/.

50. The *WebSocket connection is established*. Now the user agent must send and receive to and from the connection as described in the next section.

51. If the /defer cookies/ flag is set, store the /list of cookies/ for use by the component that invoked this algorithm.

Where the algorithm above requires that a user agent fail the WebSocket connection, the user agent may first read an arbitrary number of further bytes from the connection (and then discard them) before actually *failing the WebSocket connection*. Similarly, if a user agent can show that the bytes read from the connection so far are such that there is no subsequent sequence of bytes that the server can send that would not result in the user agent being required to *fail the WebSocket connection*, the user agent may immediately *fail the WebSocket connection* without waiting for those

bytes.

NOTE: The previous paragraph is intended to make it conforming for user agents to implement the algorithm in subtly different ways that are equivalent in all ways except that they terminate the connection at earlier or later points. For example, it enables an implementation to buffer the entire handshake response before checking it, or to verify each field as it is received rather than collecting all the fields and then checking them as a block.

When the user agent is to "apply the cookies" in a /list of cookies/, it must handle each cookie in the /list of cookies/ as defined by the appropriate specification. [[RFC2109](#)] [[RFC2965](#)]

[5.2.](#) Server-side requirements

This section only applies to servers.

Servers may offload the management of the connection to other agents on the network, for example load balancers and reverse proxies. In such a situation, the server for the purposes of conformance is considered to include all parts of the server-side infrastructure from the first device to terminate the TCP connection all the way to the server that processes requests and sends responses.

EXAMPLE: For example, a data center might have a server that responds to Web Socket requests with an appropriate handshake, and then passes the connection to another server to actually process the data frames. For the purposes of this specification, the "server" is the combination of both computers.

[5.2.1.](#) Reading the client's opening handshake

When a client starts a WebSocket connection, it sends its part of the opening handshake. The server must parse at least part of this handshake in order to obtain the necessary information to generate the server part of the handshake.

The client handshake consists of the following parts. If the server, while reading the handshake, finds that the client did not send a handshake that matches the description below, the server should abort the WebSocket connection.

1. The three-character UTF-8 string "GET".
2. A UTF-8-encoded U+0020 SPACE character (0x20 byte).

3. A string consisting of all the bytes up to the next UTF-8-encoded U+0020 SPACE character (0x20 byte). The result of decoding this string as a UTF-8 string is the name of the resource requested by the server. If the server only supports one resource, then this can safely be ignored; the client verifies that the right resource is supported based on the information included in the server's own handshake. The resource name will begin with U+002F SOLIDUS character (/) and will only include characters in the range U+0021 to U+007E.
4. A string of bytes terminated by a UTF-8-encoded U+000D CARRIAGE RETURN U+000A LINE FEED character pair (CRLF). All the characters from the second 0x20 byte up to the first 0x0D 0x0A byte pair in the data from the client can be safely ignored. (It will probably be the string "HTTP/1.1".)
5. A series of fields.

Each field is terminated by a UTF-8-encoded U+000D CARRIAGE RETURN U+000A LINE FEED character pair (CRLF). The end of the fields is denoted by the terminating CRLF pair being followed immediately by another CRLF pair.

NOTE: In other words, the fields start with the first 0x0D 0x0A byte pair, end with the first 0x0D 0x0A 0x0D 0x0A byte sequence, and are separate from each other by 0x0D 0x0A byte pairs.

The fields are encoded as UTF-8.

Each field consists of a name, consisting of one or more characters in the ranges U+0021 to U+0039 and U+003B to U+007E, followed by a U+003A COLON character (:) and a U+0020 SPACE character, followed by zero or more characters forming the value.

The expected field names, the meaning of their corresponding values, and the processing servers are required to apply to those fields, are described below, after the description of the client handshake.

6. After the first 0x0D 0x0A 0x0D 0x0A byte sequence, indicating the end of the fields, the client sends eight random bytes. These are used in constructing the server handshake.

The expected field names, and the meaning of their corresponding values, are as follows. Field names must be compared in an ASCII case-insensitive manner.

|Upgrade|

Invariant part of the handshake. Will always have a value that is an ASCII case-insensitive match for the string "WebSocket".

Can be safely ignored, though the server should abort the WebSocket connection if this field is absent or has a different value, to avoid vulnerability to cross-protocol attacks.

|Connection|

Invariant part of the handshake. Will always have a value that is an ASCII case-insensitive match for the string "Upgrade".

Can be safely ignored, though the server should abort the WebSocket connection if this field is absent or has a different value, to avoid vulnerability to cross-protocol attacks.

|Host|

The value gives the hostname that the client intended to use when opening the WebSocket. It would be of interest in particular to virtual hosting environments, where one server might serve multiple hosts, and might therefore want to return different data.

Can be safely ignored, though the server should abort the WebSocket connection if this field is absent or has a value that does not match the server's host name, to avoid vulnerability to cross-protocol attacks and DNS rebinding attacks.

|Origin|

The value gives the scheme, hostname, and port (if it's not the default port for the given scheme) of the page that asked the client to open the WebSocket. It would be interesting if the server's operator had deals with operators of other sites, since the server could then decide how to respond (or indeed, whether to respond) based on which site was requesting a connection.

[\[ORIGIN\]](#)

Can be safely ignored, though the server should abort the WebSocket connection if this field is absent or has a value that does not match one of the origins the server is expecting to communicate with, to avoid vulnerability to cross-protocol attacks and cross-site scripting attacks.

|Sec-WebSocket-Protocol|

The value gives the name of a subprotocol that the client is intending to select. It would be interesting if the server supports multiple protocols or protocol versions.

Can be safely ignored, though the server may abort the WebSocket

connection if the field is absent but the conventions for communicating with the server are such that the field is expected; and the server should abort the WebSocket connection if the field has a value that does not match one of the subprotocols that the server supports, to avoid integrity errors once the connection is established.

|Sec-WebSocket-Key1|

|Sec-WebSocket-Key2|

The values provide the information required for computing the server's handshake, as described in the next section.

|Sec-WebSocket-Draft|

The value provides the version of this draft protocol that the client is attempting to establish a connection using. If this value is not equal to a version of the draft protocol that the server understands, the server **MUST** abort the WebSocket connection.

Other fields

Other fields can be used, such as "Cookie", for authentication purposes. Their semantics are equivalent to the semantics of the HTTP headers with the same names.

Unrecognized fields can be safely ignored, and are probably either the result of intermediaries injecting fields unrelated to the operation of the WebSocket protocol, or clients that support future versions of the protocol offering options that the server doesn't support.

[5.2.2](#). Sending the server's opening handshake

When a client establishes a WebSocket connection to a server, the server must run the following steps.

1. If the server supports encryption, perform a TLS handshake over the connection. If this fails (e.g. the client indicated a host name in the extended client hello "server_name" extension that the server does not host), then close the connection; otherwise, all further communication for the connection (including the server handshake) must run through the encrypted tunnel.
[\[RFC2246\]](#)
2. Establish the following information:

/host/

The host name or IP address of the WebSocket server, as it is to be addressed by clients. The host name must be punycode-encoded if necessary. If the server can respond to requests to multiple hosts (e.g. in a virtual hosting environment), then the value should be derived from the client's handshake, specifically from the "Host" field. The /host/ value must be lowercase (not containing characters in the range U+0041 LATIN CAPITAL LETTER A to U+005A LATIN CAPITAL LETTER Z).

/port/

The port number on which the server expected and/or received the connection.

/resource name/

An identifier for the service provided by the server. If the server provides multiple services, then the value should be derived from the resource name given in the client's handshake.

/secure flag/

True if the connection is encrypted or if the server expected it to be encrypted; false otherwise.

/origin/

The ASCII serialization of the origin that the server is willing to communicate with, converted to ASCII lowercase. If the server can respond to requests from multiple origins (or indeed, all origins), then the value should be derived from the client's handshake, specifically from the "Origin" field. [[ORIGIN](#)]

/subprotocol/

Either null, or a string representing the subprotocol the server is ready to use. If the server supports multiple subprotocols, then the value should be derived from the client's handshake, specifically by selecting one of the values from the "Sec-WebSocket-Protocol" field. The absence of such a field is equivalent to the null value. The empty string is not the same as the null value for these purposes.

/key_1/

The value of the "Sec-WebSocket-Key1" field in the client's handshake.

/key_2/

The value of the "Sec-WebSocket-Key2" field in the client's handshake.

/key_3/

The eight random bytes sent after the first 0x0D 0x0A 0x0D 0x0A sequence in the client's handshake.

3. Let /location/ be the string that results from constructing a WebSocket URL from /host/, /port/, /resource name/, and /secure flag/.
4. Let /key-number_1/ be the digits (characters in the range U+0030 DIGIT ZERO (0) to U+0039 DIGIT NINE (9)) in /key_1/, interpreted as a base ten integer, ignoring all other characters in /key_1/.

Let /key-number_2/ be the digits (characters in the range U+0030 DIGIT ZERO (0) to U+0039 DIGIT NINE (9)) in /key_2/, interpreted as a base ten integer, ignoring all other characters in /key_2/.

EXAMPLE: For example, assume that the client handshake was:

```
GET / HTTP/1.1
Connection: Upgrade
Host: example.com
Upgrade: WebSocket
Sec-WebSocket-Key1: 3e6b263 4 17 80
Origin: http://example.com
Sec-WebSocket-Key2: 17 9 G`ZD9 2 2b 7X 3 /r90

WjN}|M(6
```

The /key-number_1/ would be the number 3,626,341,780, and the /key-number_2/ would be the number 1,799,227,390.

In this example, incidentally, /key_3/ is "WjN}|M(6", or 0x57 0x6A 0x4E 0x7D 0x7C 0x4D 0x28 0x36.

5. Let /spaces_1/ be the number of U+0020 SPACE characters in /key_1/.

Let /spaces_2/ be the number of U+0020 SPACE characters in /key_2/.

If either /spaces_1/ or /spaces_2/ is zero, then abort the WebSocket connection. This is a symptom of a cross-protocol attack.

EXAMPLE: In the example above, /spaces_1/ would be 4 and /spaces_2/ would be 10.

6. If /key-number_1/ is not an integral multiple of /spaces_1/, then abort the WebSocket connection.

If /key-number_2/ is not an integral multiple of /spaces_2/, then abort the WebSocket connection.

NOTE: This can only happen if the client is not a conforming WebSocket client.

7. Let /part_1/ be /key-number_1/ divided by /spaces_1/.

Let /part_2/ be /key-number_2/ divided by /spaces_2/.

EXAMPLE: In the example above, /part_1/ would be 906,585,445 and /part_2/ would be 179,922,739.

8. Let /challenge/ be the concatenation of /part_1/, expressed as a big-endian 32 bit integer, /part_2/, expressed as a big-endian 32 bit integer, and the eight bytes of /key_3/ in the order they were sent on the wire.

EXAMPLE: In the example above, this would be the 16 bytes 0x36 0x09 0x65 0x65 0x0A 0xB9 0x67 0x33 0x57 0x6A 0x4E 0x7D 0x7C 0x4D 0x28 0x36.

9. Let /response/ be the MD5 fingerprint of /challenge/ as a big-endian 128 bit string. [[RFC1321](#)]

EXAMPLE: In the example above, this would be the 16 bytes 0x6E 0x60 0x39 0x65 0x42 0x6B 0x39 0x7A 0x24 0x52 0x38 0x70 0x4F 0x74 0x56 0x62, or "n`9eBk9z\$R8p0tVb" in UTF-8.

10. Send the following line, terminated by the two characters U+000D CARRIAGE RETURN and U+000A LINE FEED (CRLF) and encoded as UTF-8, to the client:

HTTP/1.1 101 WebSocket Protocol Handshake

This line may be sent differently if necessary, but must match the Status-Line production defined in the HTTP specification, with the Status-Code having the value 101.

11. Send the following fields to the client. Each field must be sent as a line consisting of the field name, which must be an ASCII case-insensitive match for the field name in the list

below, followed by a U+003A COLON character (:) and a U+0020 SPACE character, followed by the field value as specified in the list below, followed by the two characters U+000D CARRIAGE RETURN and U+000A LINE FEED (CRLF). The lines must be encoded as UTF-8. The lines may be sent in any order.

|Upgrade|

The value must be the string "WebSocket".

|Connection|

The value must be the string "Upgrade".

|Sec-WebSocket-Location|

The value must be /location/

|Sec-WebSocket-Origin|

The value must be /origin/

|Sec-WebSocket-Protocol|

This field must be included if /subprotocol/ is not null, and must not be included if /subprotocol/ is null.

If included, the value must be /subprotocol/

Optionally, include "Set-Cookie", "Set-Cookie2", or other cookie-related fields, with values equal to the values that would be used for the identically named HTTP headers. [[RFC2109](#)] [[RFC2965](#)]

12. Send two bytes 0x0D 0x0A (UTF-8 CRLF).

13. Send /response/.

This completes the server's handshake. If the server finishes these steps without aborting the WebSocket connection, and if the client does not then fail the WebSocket connection, then the connection is established and the server may begin sending and receiving data, as described in the next section.

[6.](#) Error Handling

[6.1.](#) Handling errors in UTF-8 from the server

When a client is to interpret a byte stream as UTF-8 but finds that the byte stream is not in fact a valid UTF-8 stream, then any bytes or sequences of bytes that are not valid UTF-8 sequences must be interpreted as a U+FFFD REPLACEMENT CHARACTER.

[6.2.](#) Handling errors in UTF-8 from the client

When a server is to interpret a byte stream as UTF-8 but finds that the byte stream is not in fact a valid UTF-8 stream, behavior is undefined. A server could close the connection, convert invalid byte sequences to U+FFFD REPLACEMENT CHARACTERS, store the data verbatim, or perform application-specific processing. Subprotocols layered on the WebSocket protocol might define specific behavior for servers.

[7.](#) Closing the connection

[7.1.](#) Client-initiated closure

Certain algorithms require the user agent to **fail the WebSocket connection**. To do so, the user agent must close the WebSocket connection, and may report the problem to the user (which would be especially useful for developers).

Except as indicated above or as specified by the application layer (e.g. a script using the WebSocket API), user agents should not close the connection.

User agents must not convey any failure information to scripts in a way that would allow a script to distinguish the following situations:

- o A server whose host name could not be resolved.
- o A server to which packets could not successfully be routed.
- o A server that refused the connection on the specified port.
- o A server that did not complete the opening handshake (e.g. because it was not a WebSocket server).
- o A WebSocket server that sent a correct opening handshake, but that specified options that caused the client to drop the connection (e.g. the server specified an origin that differed from the script's).
- o A WebSocket server that abruptly closed the connection after successfully completing the opening handshake.

[7.2.](#) Server-initiated closure

Certain algorithms require or recommend that the server **abort the WebSocket connection** during the opening handshake. To do so, the server must simply close the WebSocket connection.

[7.3.](#) Closure

To **close the WebSocket connection**, the user agent or server must close the TCP connection, using whatever mechanism possible (e.g. either the TCP RST or FIN mechanisms). When a user agent notices that the server has closed its connection, it must immediately close its side of the connection also. Whether the user agent or the server closes the connection first, it is said that the **WebSocket*

connection is closed*. If the connection was closed after the client finished the WebSocket closing handshake, then the WebSocket connection is said to have been closed cleanly.

Servers may close the WebSocket connection whenever desired. User agents should not close the WebSocket connection arbitrarily.

[8.](#) Known extensions

Extensions provide a mechanism for implementations to opt-in to additional protocol features. This section defines the meaning of well-known extensions but implementations may use extensions defined separately as well.

[8.1.](#) Compression

The registered extension token for this compression extension is "deflate-stream".

The extension does not have any per message extension data and it does not define the use of any WebSocket reserved bits or op codes.

Senders using this extension MUST apply [RFC 1951](#) encodings to all bytes of the data stream following the handshake including both data and control messages. The data stream MAY include multiple blocks of both compressed and uncompressed types as defined by [RFC 1951](#).
[RFC1951]

Senders MUST NOT delay the transmission of any portion of a WebSocket message because the deflate encoding of the message does not end on a byte boundary. The encodings for adjacent messages MAY appear in the same byte if no delay in transmission is occurred by doing so.

9. Security considerations

While this protocol is intended to be used by scripts in Web pages, it can also be used directly by hosts. Such hosts are acting on their own behalf, and can therefore send fake "Origin" fields, misleading the server. Servers should therefore be careful about assuming that they are talking directly to scripts from known origins, and must consider that they might be accessed in unexpected ways. In particular, a server should not trust that any input is valid.

EXAMPLE: For example, if the server uses input as part of SQL queries, all input text should be escaped before being passed to the SQL server, lest the server be susceptible to SQL injection.

Servers that are not intended to process input from any Web page but only for certain sites should verify the "Origin" field is an origin they expect, and should only respond with the corresponding "Sec-WebSocket-Origin" if it is an accepted origin. Servers that only accept input from one origin can just send back that value in the "Sec-WebSocket-Origin" field, without bothering to check the client's value.

If at any time a server is faced with data that it does not understand, or that violates some criteria by which the server determines safety of input, or when the server sees a handshake that does not correspond to the values the server is expecting (e.g. incorrect path or origin), the server should just disconnect. It is always safe to disconnect.

The biggest security risk when sending text data using this protocol is sending data using the wrong encoding. If an attacker can trick the server into sending data encoded as ISO-8859-1 verbatim (for instance), rather than encoded as UTF-8, then the attacker could inject arbitrary frames into the data stream.

[10.](#) IANA considerations

[10.1.](#) Registration of ws: scheme

A |ws:| URL identifies a WebSocket server and resource name.

URI scheme name.

ws

Status.

Permanent.

URI scheme syntax.

In ABNF terms using the terminals from the URI specifications:
[[RFC5234](#)] [[RFC3986](#)]

"ws" ":" hier-part ["?" query]

The path and query components form the resource name sent to the server to identify the kind of service desired. Other components have the meanings described in [RFC3986](#).

URI scheme semantics.

The only operation for this scheme is to open a connection using the WebSocket protocol.

Encoding considerations.

Characters in the host component that are excluded by the syntax defined above must be converted from Unicode to ASCII by applying the IDNA ToASCII algorithm to the Unicode host name, with both the AllowUnassigned and UseSTD3ASCIIRules flags set, and using the result of this algorithm as the host in the URI. [[RFC3490](#)]

Characters in other components that are excluded by the syntax defined above must be converted from Unicode to ASCII by first encoding the characters as UTF-8 and then replacing the corresponding bytes using their percent-encoded form as defined in the URI and IRI specification. [[RFC3986](#)] [[RFC3987](#)]

Applications/protocols that use this URI scheme name.

WebSocket protocol.

Interoperability considerations.

None.

Security considerations.

See "Security considerations" section above.

Contact.

Ian Hickson <ian@hixie.ch>

Author/Change controller.

Ian Hickson <ian@hixie.ch>

References.

This document.

[10.2.](#) Registration of wss: scheme

A |wss:| URL identifies a WebSocket server and resource name, and indicates that traffic over that connection is to be encrypted.

URI scheme name.

wss

Status.

Permanent.

URI scheme syntax.

In ABNF terms using the terminals from the URI specifications:
[[RFC5234](#)] [[RFC3986](#)]

"wss" ":" hier-part ["?" query]

The path and query components form the resource name sent to the server to identify the kind of service desired. Other components have the meanings described in [RFC3986](#).

URI scheme semantics.

The only operation for this scheme is to open a connection using the WebSocket protocol, encrypted using TLS.

Encoding considerations.

Characters in the host component that are excluded by the syntax defined above must be converted from Unicode to ASCII by applying the IDNA ToASCII algorithm to the Unicode host name, with both the AllowUnassigned and UseSTD3ASCIIRules flags set, and using the result of this algorithm as the host in the URI. [[RFC3490](#)]

Characters in other components that are excluded by the syntax defined above must be converted from Unicode to ASCII by first encoding the characters as UTF-8 and then replacing the corresponding bytes using their percent-encoded form as defined in

the URI and IRI specification. [[RFC3986](#)] [[RFC3987](#)]

Applications/protocols that use this URI scheme name.
WebSocket protocol over TLS.

Interoperability considerations.
None.

Security considerations.
See "Security considerations" section above.

Contact.
Ian Hickson <ian@hixie.ch>

Author/Change controller.
Ian Hickson <ian@hixie.ch>

References.
This document.

[10.3.](#) Registration of the "WebSocket" HTTP Upgrade keyword

Name of token.
WebSocket

Author/Change controller.
Ian Hickson <ian@hixie.ch>

Contact.
Ian Hickson <ian@hixie.ch>

References.
This document.

[10.4.](#) Sec-WebSocket-Key1 and Sec-WebSocket-Key2

This section describes two header fields for registration in the Permanent Message Header Field Registry. [[RFC3864](#)]

Header field name
Sec-WebSocket-Key1

Applicable protocol
http

Status

reserved; do not use outside WebSocket handshake

Author/Change controller

IETF

Specification document(s)

This document is the relevant specification.

Related information

None.

Header field name

Sec-WebSocket-Key2

Applicable protocol

http

Status

reserved; do not use outside WebSocket handshake

Author/Change controller

IETF

Specification document(s)

This document is the relevant specification.

Related information

None.

The |Sec-WebSocket-Key1| and |Sec-WebSocket-Key2| headers are used in the WebSocket handshake. They are sent from the client to the server to provide part of the information used by the server to prove that it received a valid WebSocket handshake. This helps ensure that the server does not accept connections from non-Web-Socket clients (e.g. HTTP clients) that are being abused to send data to unsuspecting WebSocket servers.

[10.5.](#) Sec-WebSocket-Location

This section describes a header field for registration in the Permanent Message Header Field Registry. [[RFC3864](#)]

Header field name

Sec-WebSocket-Location

Applicable protocol
http

Status
reserved; do not use outside WebSocket handshake

Author/Change controller
IETF

Specification document(s)
This document is the relevant specification.

Related information
None.

The |Sec-WebSocket-Location| header is used in the WebSocket handshake. It is sent from the server to the client to confirm the URL of the connection. This enables the client to verify that the connection was established to the right server, port, and path, instead of relying on the server to verify that the requested host, port, and path are correct.

[10.6.](#) Sec-WebSocket-Origin

This section describes a header field for registration in the Permanent Message Header Field Registry. [[RFC3864](#)]

Header field name
Sec-WebSocket-Origin

Applicable protocol
http

Status
reserved; do not use outside WebSocket handshake

Author/Change controller
IETF

Specification document(s)
This document is the relevant specification.

Related information
None.

The |Sec-WebSocket-Origin| header is used in the WebSocket handshake. It is sent from the server to the client to confirm the origin of the script that opened the connection. This enables user agents to

verify that the server is willing to serve the script that opened the connection.

[10.7.](#) Sec-WebSocket-Protocol

This section describes a header field for registration in the Permanent Message Header Field Registry. [[RFC3864](#)]

Header field name

Sec-WebSocket-Protocol

Applicable protocol

http

Status

reserved; do not use outside WebSocket handshake

Author/Change controller

IETF

Specification document(s)

This document is the relevant specification.

Related information

None.

The |Sec-WebSocket-Protocol| header is used in the WebSocket handshake. It is sent from the client to the server and back from the server to the client to confirm the subprotocol of the connection. This enables scripts to both select a subprotocol and be sure that the server agreed to serve that subprotocol.

[10.8.](#) Sec-WebSocket-Draft

This section describes a header field for registration in the Permanent Message Header Field Registry. [[RFC3864](#)]

Header field name

Sec-WebSocket-Draft

Applicable protocol

http

Status

reserved; do not use outside WebSocket handshake

Author/Change controller
IETF

Specification document(s)
This document is the relevant specification.

Related information
None.

The |Sec-WebSocket-Draft| header is used in the WebSocket handshake. It is sent from the client to the server to indicate the draft protocol version of the connection. This enables servers to correctly interpret the handshake and subsequent data being sent from the data, and close the connection if the server cannot interpret that data in a safe manner.

[11.](#) Using the WebSocket protocol from other specifications

The WebSocket protocol is intended to be used by another specification to provide a generic mechanism for dynamic author-defined content, e.g. in a specification defining a scripted API.

Such a specification first needs to "establish a WebSocket connection", providing that algorithm with:

- o The destination, consisting of a /host/ and a /port/.
- o A /resource name/, which allows for multiple services to be identified at one host and port.
- o A /secure/ flag, which is true if the connection is to be encrypted, and false otherwise.
- o An ASCII serialization of an origin that is being made responsible for the connection. [[ORIGIN](#)]
- o Optionally a string identifying a protocol that is to be layered over the WebSocket connection.

The /host/, /port/, /resource name/, and /secure/ flag are usually obtained from a URL using the steps to parse a WebSocket URL's components. These steps fail if the URL does not specify a WebSocket.

If a connection can be established, then it is said that the "WebSocket connection is established".

If at any time the connection is to be closed, then the specification needs to use the "close the WebSocket connection" algorithm.

When the connection is closed, for any reason including failure to establish the connection in the first place, it is said that the "WebSocket connection is closed".

While a connection is open, the specification will need to handle the cases when "a WebSocket message has been received" with text /data/.

To send some text /data/ to an open connection, the specification needs to "send /data/ using the WebSocket".

12. Acknowledgements

Special thanks are due to Ian Hickson, who was the original author and editor of this protocol. The initial design of this specification benefitted from the participation of many people in the WHATWG and WHATWG mailing list. Contributions to that specification are not tracked by section, but a list of all who contributed to that specification is given in the WHATWG HTML specification. [[HTML](#)]

Special thanks also to John Tamplin for providing a significant amount of text for the Data Framing section of this specification.

13. Normative References

- [HTML] Hickson, I., "HTML", August 2010, <<http://whatwg.org/html5>>.
- [ORIGIN] Barth, A., Jackson, C., and I. Hickson, "The HTTP Origin Header", [draft-abarth-origin](http://tools.ietf.org/html/draft-abarth-origin) (work in progress), September 2009, <<http://tools.ietf.org/html/draft-abarth-origin>>.
- [ANSI.X3-4.1986] American National Standards Institute, "Coded Character Set - 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.
- [RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", [RFC 1321](http://tools.ietf.org/html/rfc1321), April 1992.
- [RFC1951] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", [RFC 1951](http://tools.ietf.org/html/rfc1951), May 1996.
- [RFC2109] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", [RFC 2109](http://tools.ietf.org/html/rfc2109), February 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](http://tools.ietf.org/html/bcp14), [RFC 2119](http://tools.ietf.org/html/rfc2119), March 1997.
- [RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", [RFC 2246](http://tools.ietf.org/html/rfc2246), January 1999.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](http://tools.ietf.org/html/rfc2616), June 1999.
- [RFC2965] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", [RFC 2965](http://tools.ietf.org/html/rfc2965), October 2000.
- [RFC3490] Faltstrom, P., Hoffman, P., and A. Costello, "Internationalizing Domain Names in Applications (IDNA)", [RFC 3490](http://tools.ietf.org/html/rfc3490), March 2003.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](http://tools.ietf.org/html/rfc3629), November 2003.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](http://tools.ietf.org/html/bcp90), [RFC 3864](http://tools.ietf.org/html/rfc3864), September 2004.

- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", [RFC 3987](#), January 2005.
- [RFC4366] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", [RFC 4366](#), April 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [WSAPI] Hickson, I., "The Web Sockets API", August 2010, <<http://dev.w3.org/html5/websockets/>>.

Author's Address

Ian Fette
Google, Inc.

Email: ifette+ietf@google.com
URI: <http://www.ianfette.com/>