

HyBi Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 24, 2011

I. Fette
Google, Inc.
April 22, 2011

The WebSocket protocol
draft-ietf-hybi-thewebsocketprotocol-07

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.

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 October 24, 2011.

Copyright Notice

Copyright (c) 2011 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

1.	Introduction	4
1.1.	Background	4
1.2.	Protocol overview	4
1.3.	Opening handshake	6
1.4.	Closing handshake	8
1.5.	Design philosophy	9
1.6.	Security model	10
1.7.	Relationship to TCP and HTTP	10
1.8.	Establishing a connection	11
1.9.	Subprotocols using the WebSocket protocol	11
2.	Conformance requirements	13
2.1.	Terminology	13
3.	WebSocket URIs	15
3.1.	Parsing WebSocket URIs	15
3.2.	Constructing WebSocket URIs	16
3.3.	Valid WebSocket URIs	16
4.	Data Framing	17
4.1.	Overview	17
4.2.	Base Framing Protocol	17
4.3.	Client-to-Server Masking	20
4.4.	Fragmentation	21
4.5.	Control Frames	23
4.5.1.	Close	23
4.5.2.	Ping	24
4.5.3.	Pong	24
4.6.	Data Frames	24
4.7.	Examples	25
4.8.	Extensibility	25
5.	Opening Handshake	27
5.1.	Client Requirements	27
5.2.	Server-side requirements	31
5.2.1.	Reading the client's opening handshake	32

5.2.2.	Sending the server's opening handshake	32
6.	Error Handling	36
6.1.	Handling errors in UTF-8 from the server	36
6.2.	Handling errors in UTF-8 from the client	36
7.	Closing the connection	37

7.1.	Definitions	37
7.1.1.	Close the WebSocket Connection	37
7.1.2.	Start the WebSocket Closing Handshake	37
7.1.3.	The WebSocket Connection Is Closed	37
7.1.4.	Fail the WebSocket Connection	37
7.2.	Abnormal closures	37
7.2.1.	Client-initiated closure	38
7.2.2.	Server-initiated closure	38
7.3.	Normal closure of connections	38
7.4.	Status codes	38
7.4.1.	Defined Status Codes	38
7.4.2.	Reserved status code ranges	39
8.	Extensions	41
8.1.	Negotiating extensions	41
8.2.	Known extensions	42
8.2.1.	Compression	42
9.	Security considerations	44
10.	IANA considerations	46
10.1.	Registration of ws: scheme	46
10.2.	Registration of wss: scheme	47
10.3.	Registration of the "WebSocket" HTTP Upgrade keyword	48
10.4.	Sec-WebSocket-Key	48
10.5.	Sec-WebSocket-Extensions	49
10.6.	Sec-WebSocket-Accept	50
10.7.	Sec-WebSocket-Origin	50
10.8.	Sec-WebSocket-Protocol	51
10.9.	Sec-WebSocket-Version	51
11.	Using the WebSocket protocol from other specifications	53
12.	Acknowledgements	54
13.	References	55
13.1.	Normative References	55
13.2.	Informative References	56
	Author's Address	57

[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.[\[RFC6202\]](#)

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:

Fette	Expires October 24, 2011	[Page 4]
-------	--------------------------	----------

Internet-Draft	The WebSocket protocol	April 2011
----------------	------------------------	------------

```
GET /chat HTTP/1.1
Host: server.example.com
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Key: dGhlIHNhbXBsZSBub25jZQ==
Sec-WebSocket-Origin: http://example.com
Sec-WebSocket-Protocol: chat, superchat
Sec-WebSocket-Version: 7
```

The handshake from the server looks as follows:

```
HTTP/1.1 101 Switching Protocols
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Accept: s3pPLMBiTxaQ9kYGzzhZRbK+x0o=
Sec-WebSocket-Protocol: chat
```

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 [[RFC2616](#)].

After the leading line in both cases come an unordered set of header

fields. The meaning of these header fields is specified in [Section 5](#) of this document. Additional header fields may also be present, such as cookies [[I-D.ietf-httpstate-cookie](#)] required to identify the user. The format and parsing of headers is as defined in [[RFC2616](#)].

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. A message is composed of one or more frames, all of which contain the same type of data. Broadly speaking, there are types for textual data, which is interpreted as UTF-8 [[RFC3629](#)] 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 signaling, 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 and intermediaries, 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 is an HTTP Upgrade request:

```
GET /chat HTTP/1.1
Host: server.example.com
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Key: dGhlIHNhbXBsZSBub25jZQ==
Sec-WebSocket-Origin: http://example.com
Sec-WebSocket-Protocol: chat, superchat
Sec-WebSocket-Version: 7
```

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

The "Request-URI" of the GET method [[RFC2616](#)] is used to identify the endpoint of the WebSocket connection, both to allow multiple domains to be served from one IP address and to allow multiple WebSocket endpoints to be served by a single server.

The client includes the hostname in the Host header of its handshake as per [[RFC2616](#)], so that both the client and the server can verify that they agree on which host is in use.

Additional headers are used to select options in the WebSocket protocol. 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
```

The |Sec-WebSocket-Origin| header 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 is informed of the script origin generating the WebSocket connection request. If the server does not wish to accept connections from this origin, it can choose to abort the connection. This header is sent by browser clients, for non-browser clients this header may be sent if it makes sense in the context of those clients.

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 two pieces of information and combine them to form a response. The first piece of information comes from the |Sec-WebSocket-Key| header in the client handshake:

Sec-WebSocket-Key: dGhliHNhbXBsZSBub25jZQ==

For this header, the server has to take the value (as present in the header, e.g. the base64-encoded [[RFC4648](#)] version minus leading and trailing whitespace), and concatenate this with the GUID "258EAFA5-E914-47DA-95CA-C5AB0DC85B11" in string form, which is unlikely to be used by network endpoints that do not understand the WebSocket protocol. A SHA-1 hash, base64-encoded, of this concatenation is then returned in the server's handshake [[FIPS.180-2.2002](#)].

Concretely, if as in the example above, header |Sec-WebSocket-Key| had the value "dGhliHNhbXBsZSBub25jZQ==", the server would concatenate the string "258EAFA5-E914-47DA-95CA-C5AB0DC85B11" to form the string "dGhliHNhbXBsZSBub25jZQ==258EAFA5-E914-47DA-95CA-C5AB0DC85B11". The server would then take the SHA-1 hash of this, giving the value 0xb3 0x7a 0x4f 0x2c 0xc0 0x62 0x4f 0x16 0x90 0xf6 0x46 0x06 0xcf 0x38 0x59 0x45 0xb2 0xbe 0xc4 0xea. This value is then base64-encoded, to give the value "s3pPLMBiTxaQ9kYGzzhZRbK+x0o=". This value would then be echoed in the header |Sec-WebSocket-Accept|.

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:

```
HTTP/1.1 101 Switching Protocols
```

Any status code other than 101 indicates that the WebSocket handshake has not completed, and that the semantics of HTTP still apply. The headers follow the status code.

The `|Connection|` and `|Upgrade|` headers complete the HTTP Upgrade. The `|Sec-WebSocket-Accept|` header indicates whether the server is willing to accept the connection. If present, this header must include a hash of the client's nonce sent in `|Sec-WebSocket-Key|` along with a predefined GUID. Any other value must not be interpreted as an acceptance of the connection by the server.

```
HTTP/1.1 101 Switching Protocols
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Accept: s3pPLMBiTxaQ9kYGzzhZRbK+x0o=
```

These fields are checked by the Web browser when it is acting as a `|WebSocket|` client for scripted pages. If the `|Sec-WebSocket-Accept|` value does not match the expected value, or if the header is missing, or if the HTTP status code is not 101, the connection will not be established and WebSockets frames will not be sent.

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 client's handshake. 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.

[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 (detailed in [Section 4.5.1](#)). Upon receiving such a frame, the other peer sends a close frame in response, 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 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.

By sending a close frame and waiting for a close frame in response, certain cases are avoided where data may be unnecessarily lost. For instance, on some platforms, if a socket is closed with data in the receive queue, a RST packet is sent, which will then cause `recv()` to fail for the party that received the RST, even if there was data waiting to be read.

[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 re-implements 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 request mechanism also.

The protocol is intended to be extensible; future versions will likely introduce additional concepts such as multiplexing.

[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 [[RFC5321](#)] and 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, at the time of writing of this specification, fields starting with `|Sec-|` cannot be set by an attacker from a Web browser using only HTML and JavaScript APIs such as `|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.

By default the WebSocket protocol uses port 80 for regular WebSocket connections and port 443 for WebSocket connections tunneled over TLS [[RFC2818](#)].

Fette

Expires October 24, 2011

[Page 10]

Internet-Draft

The WebSocket protocol

April 2011

[1.8.](#) Establishing a connection

This section is non-normative.

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. At the time of writing of this specification, it should be noted that connections on port 80 and 443 have significantly different success rates, with connections on port 443 being significantly more likely to succeed, though this may change with time.

[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 Organization 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, e.g. 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](#). [[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 "URI" is used in this document as defined in [\[RFC3986\]](#).

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 URIs

[3.1.](#) Parsing WebSocket URIs

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

1. If the `/uri/` string is not an absolute URI, then fail this algorithm. [[RFC3986](#)]
2. Resolve the `/uri/` string using the resolve a Web address algorithm defined by the Web addresses specification, with the URI 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 URI at this point.

3. If `/uri/` does not have a `<scheme>` component whose value, when converted to ASCII lowercase, is either `"ws"` or `"wss"`, then fail this algorithm.
4. If `/uri/` has a `<fragment>` component, then fail this algorithm.
5. If the `<scheme>` component of `/uri/` is `"ws"`, set `/secure/` to false; otherwise, if the `<scheme>` component is `"wss"`, set `/secure/` to true; if neither of the above apply, fail this algorithm.
6. Let `/host/` be the value of the `<host>` component of `/uri/`, converted to ASCII lowercase.
7. If `/uri/` 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 `/uri/`.
10. If `/resource name/` is the empty string, set it to a single character U+002F SOLIDUS (`/`).
11. If `/uri/` 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 URIs

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

1. Let /uri/ be the empty string.
2. If the /secure/ flag is false, then append the string "ws://" to /uri/. Otherwise, append the string "wss://" to /uri/.
3. Append /host/ to /uri/.
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 /uri/.
5. Append /resource name/ to /uri/.
6. Return /uri/.

[3.3.](#) Valid WebSocket URIs

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

- o The /host/ MUST be ASCII-only (i.e. it MUST have been punycode-encoded [[RFC3492](#)] already if necessary, and MUST NOT contain any characters above U+007E).
- o The /resource name/ string MUST be a non-empty string of characters in the range U+0021 to U+007E and MUST start with a U+002F SOLIDUS character (/).

Any WebSocket URIs not meeting the above criteria are considered invalid. A client MUST NOT attempt to make a connection to an invalid WebSocket URI. A client SHOULD attempt to parse a URI 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 URI, but MUST NOT attempt to connect with such an unparsed URI, and instead only use the parsed version and only if that version is considered valid by the criteria above.

0									1									2									3								
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1				
+ - + - + - + - - - - - + - + - - - - +									-----+-----+-----+-----+									-----+-----+-----+-----+																	
F R R R				opcode M					Payload len					Extended payload length																					
I S S S				(4)					A	(7)					(16/63)																				
N V V V									S						(if payload len==126/127)																				
1 2 3									K																										
+ - + - + - + - - - - - + - + - - - - +									-----+-----+-----+-----+									- +																	
Extended payload length continued, if payload len == 127																																			
+ - - - - - + - - - - - + - - - - - + - - - - - +									-----+-----+-----+-----+									-----+-----+-----+-----+																	
																Masking-key, if MASK set to 1																			
-----+-----+-----+-----+																																			

Masking-key (continued)	Payload Data
-----	-----
: Payload Data continued ... :	
-----	-----
Payload Data continued ...	
-----	-----

FIN: 1 bit

Indicates that this is the final fragment in a message. The first fragment MAY also be the final fragment.

RSV1, RSV2, RSV3: 1 bit each

Must be 0 unless an extension is negotiated which defines meanings for non-zero values

Opcode: 4 bits

Defines the interpretation of the payload data

Mask: 1 bit

Defines whether the payload data is masked. If set to 1, a masking key is present in Masking-key, and this is used to unmask the payload data as per [Section 4.3](#). All frames sent from client to server have this bit set to 1.

Payload length: 7 bits, 7+16 bits, or 7+64 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 most significant 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. The length of this field is always at least 7 bits. If the value of the

first 7 bits is 0-125, the length of this field is 7 bits. If the value is 126, there exist 16 additional bits with a 16-bit length. If the value is 127, there exist 64 additional bits with a 63-bit length (the most significant bit MUST be 0).

Masking-key: 0 or 4 bytes

All frames sent from the client to the server are masked by a 32-bit value that is contained within the frame. This field is present if the Mask bit is set to 1, and is absent if the Mask bit is set to 0. See [Section 4.3](#) for further information on client-to-server masking.

Fette

Expires October 24, 2011

[Page 18]

Internet-Draft

The WebSocket protocol

April 2011

Payload data: n bytes

The payload data is defined as Extension Data concatenated with Application Data.

Extension data: n bytes

The extension data is 0 bytes unless an extension has been negotiated. Any extension MUST specify the length of the extension data, or how that length may be calculated, and how the extension 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-fin
 frame-rsv1

```

frame-rsv2
frame-rsv3
frame-opcode
frame-masked
frame-payload-length
[ frame-masking-key ]
frame-payload-data

frame-fin          = %x0 ; more frames of this message follow
                   / %x1 ; final frame of message

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 ; text frame
                   / %x2 ; binary frame
                   / %3-7 ; reserved for further non-control frames
                   / %x8 ; connection close

```

```

                   / %x9 ; ping
                   / %xA ; pong
                   / %xB-F ; reserved for further control frames

frame-masked       = %x0 ; frame is not masked, no frame-masking-key
                   = %x1 ; frame is masked, frame-masking-key present

frame-payload-length = %x00-7D
                   / %x7E frame-payload-length-16
                   / %x7F frame-payload-length-63

frame-payload-length-16 = %x0000-FFFF

frame-payload-length-63 = %x0000000000000000-7FFFFFFFFFFFFFFF

frame-masking-key   = <4>( %x00-FF ) ; present only if frame-masked is

frame-payload-data  = (frame-masked-extension-data
                      frame-masked-application-data) ; frame-masked 1

```

```

/ (frame-unmasked-extension-data
  frame-unmasked-application-data) ; frame-masked 0

frame-masked-extension-data      = *( %x00-FF ) ; to be defined later

frame-masked-application-data    = *( %x00-FF )

frame-unmasked-extension-data    = *( %x00-FF ) ; to be defined later

frame-unmasked-application-data = *( %x00-FF )

```

4.3. Client-to-Server Masking

The client MUST mask all frames sent to the server. A server MUST close the connection upon receiving a frame with the MASK bit set to 0. In this case, a server MAY send a close frame with a status code of 1002 (protocol error) as defined in [Section 7.4.1](#).

A masked frame MUST have the field frame-masked set to 1, as defined in [Section 4.2](#).

The masking key is contained completely within the frame, as defined in [Section 4.2](#) as frame-masking-key. It is used to mask the payload data defined in the same section as frame-payload-data, which includes extension and application data.

The masking key is a 32-bit value chosen at random by the client. The masking key MUST be derived from a strong source of entropy, and the masking key for a given frame MUST NOT make it simple for a

server to predict the masking key for a subsequent frame.

The masking does not affect the length of the payload data. To convert masked data into unmasked data, or vice versa, the following algorithm is applied. The same algorithm applies regardless of the direction of the translation - e.g. the same steps are applied to mask the data as to unmask the data.

Octet *i* of the transformed data ("transformed-octet-*i*") is the XOR of octet *i* of the original data ("original-octet-*i*") with octet *i* modulo 4 of the masking key ("masking-key-octet-*j*"):

$j = i \text{ MOD } 4$
 $\text{transformed-octet-}i = \text{original-octet-}i \text{ XOR masking-key-octet-}j$

When preparing a masked frame, the client **MUST** pick a fresh masking key uniformly at random from the set of allowed 32-bit values. The unpredictability of the masking-nonce is essential to prevent the author of malicious application data from selecting the bytes that appear on the wire.

The payload length, indicated in the framing as `frame-payload-length`, does NOT include the length of the masking key. It is the length of the payload data, e.g. the number of bytes following the masking key.

[4.4.](#) Fragmentation

The primary purpose of fragmentation is to allow sending a message that is of unknown size when the message is started without having to buffer that message. If messages couldn't be fragmented, then an endpoint would have to buffer the entire message so its length could be counted before first byte is sent. With fragmentation, a server or intermediary may choose a reasonable size buffer, and when the buffer is full write a fragment to the network.

A secondary use-case for fragmentation is for multiplexing, where it is not desirable for a large message on one logical channel to monopolize the output channel, so the MUX needs to be free to split the message into smaller fragments to better share the output channel.

The following rules apply to fragmentation:

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

- o A fragmented message consists of a single frame with the FIN bit clear and an opcode other than 0, followed by zero or more frames with the FIN bit clear and the opcode set to 0, and terminated by a single frame with the FIN bit set and an opcode of 0. Its content is the concatenation of the application data (and any

extension data that may be present) from each of those frames in order. As an example, for a text message sent as three fragments, the first fragment would have an opcode of 0x4 and a FIN bit clear, the second fragment would have an opcode of 0x0 and a FIN bit clear, and the third fragment would have an opcode of 0x0 and a FIN bit that is set.

- o Control frames MAY be injected in the middle of a fragmented message. Control frames themselves MUST NOT be fragmented.
- o An endpoint MUST be capable of handling control frames in the middle of a fragmented message.
- o _Note: if control frames could not be interjected, the latency of a ping, for example, would be very long if behind a large message. Hence, the requirement of handling control frames in the middle of a fragmented message._
- 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 As control frames cannot be fragmented, an intermediary MUST NOT attempt to change the fragmentation of a control frame.
- o An intermediary MUST NOT change the fragmentation of a message if any reserved bit values are used and the meaning of these values is not known to the intermediary.
- o An intermediary MUST NOT change the fragmentation of any message in the context of a connection where extensions have been negotiated and the intermediary is not aware of the semantics of the negotiated extensions.
- o As a consequence of these rules, all fragments of a message are of the same type, as set by the first fragment's opcode. Since Control frames cannot be fragmented, the type for all fragments in a message MUST be either text or binary, or one of the reserved opcodes.

[4.5.](#) Control Frames

Control frames are identified by opcodes where the most significant bit of the opcode is 1. Currently defined opcodes for control frames include 0x8 (Close), 0x9 (Ping), and 0xA (Pong). Opcodes 0xB–0xF are reserved for further control frames yet to be defined.

Control frames are used to communicate state about the websocket. Control frames can be interjected in the middle of a fragmented message.

All control frames MUST have a payload length of 125 bytes or less and MUST NOT be fragmented.

[4.5.1.](#) Close

The Close message contains an opcode of 0x8.

The Close message MAY contain a body (the "application data" portion of the frame) that indicates a reason for closing, such as an endpoint shutting down, an endpoint having received a message too large, or an endpoint having received a message that does not conform to the format expected by the other endpoint. If there is a body, the first two bytes of the body MUST be a 2-byte integer (in network byte order) representing a status code defined in [Section 7.4](#). Following the 2-byte integer the body MAY contain UTF-8 encoded data, the interpretation of which is not defined by this specification. This data is not necessarily human readable, but may be useful for debugging or passing information relevant to the script that opened the connection.

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

If an endpoint receives a Close message and that endpoint did not previously send a Close message, the endpoint MUST send a Close message in response. It SHOULD do so as soon as is practical.

After both sending and receiving a close message, an endpoint considers the websocket connection closed, and SHOULD close the underlying TCP connection.

If a client and server both send a Close message at the same time, both endpoints will have sent and received a Close message and should consider the websocket connection closed and close the underlying TCP connection.

[4.5.2.](#) Ping

The Ping message contains an opcode of 0x9.

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 (i.e. both the Extension data (if any) and the Application data) of the Ping and Pong **MUST** be the same.

An endpoint **MAY** send a Ping message any time after the connection is established and before the connection is closed. **NOTE:** A ping message may serve either as a keepalive, or to verify that the remote endpoint is still responsive.

[4.5.3.](#) Pong

The Pong message contains an opcode of 0xA.

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 (i.e. both the Extension data (if any) and the Application data) of the Ping and Pong **MUST** be the same. In the case multiple Pings have been received, a Pong **MUST** be issued only in response to the most recent Ping.

A Pong message **MAY** be sent unsolicited. This serves as a unidirectional heartbeat. A response to an unsolicited pong is not expected.

[4.6.](#) Data Frames

Data frames (e.g. non control frames) are identified by opcodes where the most significant bit of the opcode is 0. Currently defined opcodes for data frames include 0x1 (Text), 0x2 (Binary). Opcodes 0x3-0x7 are reserved for further non-control frames yet to be defined.

Data frames carry application-layer or extension-layer data. The opcode determines the interpretation of the 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.

Fette

Expires October 24, 2011

[Page 24]

Internet-Draft

The WebSocket protocol

April 2011

[4.7.](#) Examples

This section is non-normative.

- o A single-frame unmasked text message
 - * 0x81 0x05 0x48 0x65 0x6c 0x6c 0x6f (contains "Hello")
- o A single-frame masked text message
 - * 0x81 0x85 0x37 0xfa 0x21 0x3d 0x7f 0x9f 0x4d 0x51 0x58 (contains "Hello")
- o A fragmented unmasked text message
 - * 0x01 0x03 0x48 0x65 0x6c (contains "Hel")
 - * 0x80 0x02 0x6c 0x6f (contains "lo")
- o Ping request and response
 - * 0x89 0x05 0x48 0x65 0x6c 0x6c 0x6f (contains a body of "Hello", but the contents of the body are arbitrary)
 - * 0x8a 0x05 0x48 0x65 0x6c 0x6c 0x6f (contains a body of "Hello", matching the body of the ping)
- o 256 bytes binary message in a single unmasked frame
 - * 0x82 0x7E 0x0100 [256 bytes of binary data]
- o 64KiB binary message in a single unmasked frame
 - * 0x82 0x7F 0x00000000000010000 [65536 bytes of binary data]

[4.8.](#) 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 0x3 through 0x7 and 0xB through 0xF, the extension data field, and the frame-rsv1, frame-rsv2, and frame-rsv3 bits of the frame header for use by extensions. The negotiation of extensions is discussed in further detail in [Section 8.1](#). 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* given either a WebSocket URI `/uri/` or the constituent components of a URI as specified in [Section 11](#), it MUST meet the following requirements. In the following text, we will use terms from [Section 3](#) such as `/host/` and `/secure/` flag as defined in that section.

1. The WebSocket URI and its components derived by applying the steps defined in [Section 3.3](#), or if the following algorithm was supplied with the constituent components as defined in [Section 11](#) then those components provided, MUST be 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/ and port /port/ pair, even if the remote host is known by another name, the user agent MUST wait until that connection has been established or for that connection to have failed. There MUST be no more than one connection in a CONNECTING state. 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 accept connections from hosts with an excessive number of existing connections, or disconnect resource-hogging connections when suffering high load.

3. Proxy Usage: 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 the user agent SHOULD 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 ZWRuYW1vZGU6bm9jYXB1cyE=
```

If the user agent is not configured to use a proxy, then a direct TCP connection SHOULD be opened 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 URI to pass the function MUST be constructed from /host/, /port/, /resource name/, and the /secure/ flag using the steps to construct a WebSocket URI as given in [Section 3.2](#).

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, either because a direct connection failed or because any proxy used returned an error, then the user agent MUST fail the WebSocket connection and abort the connection attempt.
5. If /secure/ is true, the user agent MUST perform a TLS handshake over the connection [[RFC2818](#)]. If this fails (e.g. the server's

certificate could not be verified), then the user agent MUST fail the WebSocket connection and abort the connection. Otherwise, all further communication on this channel MUST run through the encrypted tunnel. [[RFC5246](#)]

User agents MUST use the Server Name Indication extension in the TLS handshake. [[RFC6066](#)]

Once a connection to the server has been established (including a connection via a proxy or over a TLS-encrypted tunnel), the client MUST send a handshake to the server. The handshake consists of an HTTP upgrade request, along with a list of required and optional headers. The requirements for this handshake are as follows.

1. The handshake MUST be a valid HTTP request as specified by [[RFC2616](#)].
2. The Method of the request MUST be GET and the HTTP version MUST be at least 1.1.

For example, if the WebSocket URI is "ws://example.com/chat", The first line sent should be "GET /chat HTTP/1.1"

3. The request MUST contain a "Request-URI" as part of the GET method. This MUST match the /resource name/ [Section 3](#).
4. The request MUST contain a "Host" header whose value is equal to /host/
5. The request MUST contain an "Upgrade" header whose value is equal to "websocket".
6. The request MUST contain a "Connection" header whose value MUST include the "Upgrade" token.
7. The request MUST include a header with the name "Sec-WebSocket-Key". The value of this header MUST be a nonce consisting of a randomly selected 16-byte value that has been base64-encoded [[RFC3548](#)]. The nonce MUST be selected randomly for each connection.

NOTE: As an example, if the randomly selected value was the sequence of bytes 0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f 0x10, the value of the header would be "AQIDBAUGBwgJCgsMDQ4PEC=="

8. The request MUST include a header with the name "Sec-WebSocket-Origin" if the request is coming from a browser client. If the connection is from a non-browser client, the request MAY include this header if the semantics of that client match the use-case described here for browser clients. The value of this header MUST be the ASCII serialization of origin of the context in which the code establishing the connection is running, and MUST be lower-case. The value MUST NOT contain letters in the range U+0041 to U+005A (i.e. LATIN CAPITAL LETTER A to LATIN CAPITAL LETTER Z) [[I-D.ietf-websec-origin](#)].

As an example, if code is running on `www.example.com` attempting to establish a connection to `ww2.example.com`, the value of the header would be `"http://www.example.com"`.

9. The request MUST include a header with the name "Sec-WebSocket-Version". The value of this header MUST be 7.
10. The request MAY include a header with the name "Sec-WebSocket-Protocol". If present, this value indicates the subprotocol(s) the client wishes to speak. The elements that comprise this value MUST be non-empty strings with characters in the range U+0021 to U+007E and MUST all be unique strings. The ABNF for the value of this header is `1#(token | quoted-string)`, where the definitions of constructs and rules are as given in [[RFC2616](#)].
11. The request MAY include a header with the name "Sec-WebSocket-Extensions". If present, this value indicates the protocol-level extension(s) the client wishes to speak. The interpretation and format of this header is described in [Section 8.1](#).
12. The request MAY include headers associated with sending cookies, as defined by the appropriate specifications [[I-D.ietf-httpstate-cookie](#)].

Once the client's opening handshake has been sent, the client MUST wait for a response from the server before sending any further data. The client MUST validate the server's response as follows:

- o If the status code received from the server is not 101, the client handles the response per HTTP procedures. Otherwise, proceed as follows.

- o If the response lacks an Upgrade header or the Upgrade header contains a value that is not an ASCII case-insensitive match for the value "websocket", the client MUST fail the WebSocket connection.
- o If the response lacks a Connection header or the Connection header contains a value that is not an ASCII case-insensitive match for the value "Upgrade", the client MUST fail the WebSocket connection.
- o If the response lacks a Sec-WebSocket-Accept header or the Sec-WebSocket-Accept contains a value other than the base64-encoded SHA-1 of the concatenation of the Sec-WebSocket-Key (as a string, not base64-decoded) with the string "258EAF5E-E914-47DA-95CA-C5AB0DC85B11", the client MUST fail the WebSocket connection.

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.

[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 WebSocket 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 **MUST** abort the WebSocket connection.

1. An HTTP/1.1 or higher GET request, including a "Request-URI" [[RFC2616](#)] that should be interpreted as a /resource name/ [Section 3](#).
2. A "Host" header containing the server's authority.
3. A "Sec-WebSocket-Key" header with a base64-encoded value that, when decoded, is 16 bytes in length.
4. A "Sec-WebSocket-Version" header, with a value of 7.
5. Optionally, a "Sec-WebSocket-Origin" header. This header is sent by all browser clients. A connection attempt lacking this header **SHOULD NOT** be interpreted as coming from a browser client.
6. Optionally, a "Sec-WebSocket-Protocol" header, with a list of values indicating which protocols the client would like to speak, ordered by preference.
7. Optionally, a "Sec-WebSocket-Extensions" header, with a list of

values indicating which extensions the client would like to speak. The interpretation of this header is discussed in [Section 8.1](#).

8. Optionally, other headers, such as those used to send cookies to a server. Unknown headers MUST be ignored.

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

When a client establishes a WebSocket connection to a server, the server MUST complete the following steps to accept the connection and send the server's opening handshake.

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.
[\[RFC5246\]](#)
2. Establish the following information:

/origin/

The |Sec-WebSocket-Origin| header in the client's handshake indicates the origin of the script establishing the connection. The origin is serialized to ASCII and converted to lowercase. The server MAY use this information as part of a determination of whether to accept the incoming connection. If the server does not validate the origin, it will accept connections from anywhere. For more detail, refer to [Section 9](#).

/key/

The |Sec-WebSocket-Key| header in the client's handshake includes a base64-encoded value that, if decoded, is 16 bytes in length. This (encoded) value is used in the creation of the server's handshake to indicate an acceptance of the connection. It is not necessary for the server to base64-decode the Sec-WebSocket-Key value.

/version/

The |Sec-WebSocket-Version| header in the client's handshake includes the version of the WebSocket protocol the client is attempting to communicate with. If this version does not match a version understood by the server, the server **MUST** abort the WebSocket connection. The server **MAY** send a non-200 response code with a |Sec-WebSocket-Version| header indicating the version(s) the server is capable of understanding.

/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 from the Request-URI [[RFC2616](#)] of the GET method.

/subprotocol/

Either a single value or null, representing the subprotocol the server is ready to use. If the server supports multiple subprotocols, then the value **MUST** 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, and is not a legal value for this field. The ABNF for the value of this header is (token | quoted-string), where the definitions of constructs and rules are as given in [[RFC2616](#)].

/extensions/

A (possibly empty) list representing the protocol-level extensions the server is ready to use. If the server supports multiple extensions, then the value **MUST** be derived from the client's handshake, specifically by selecting one or more of the values from the "Sec-WebSocket-Extensions" 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. Extensions not listed by the client **MUST NOT** be listed. The method by which these values should be selected and interpreted is discussed in [Section 8.1](#).

3. If the server chooses to accept the incoming connection, it **MUST** reply with a valid HTTP response indicating the following.

1. A 101 response code. Such a response could look like "HTTP/1.1 101 Switching Protocols"
2. A "Sec-WebSocket-Accept" header. The value of this header is constructed by concatenating /key/, defined above in Paragraph 2 of [Section 5.2.2](#), with the string "258EAF5E-E914-47DA-95CA-C5AB0DC85B11", taking the SHA-1 hash of this concatenated value to obtain a 20-byte value, and base64-encoding this 20-byte hash.

NOTE: As an example, if the value of the "Sec-WebSocket-Key" header in the client's handshake were "dGh1IHhXbXBsZSBub25jZQ==", the server would append the string "258EAF5E-E914-47DA-95CA-C5AB0DC85B11" to form the string "dGh1IHhXbXBsZSBub25jZQ==258EAF5E-E914-47DA-95CA-C5AB0DC85B11". The server would then take the SHA-1 hash of this string, giving the value 0xb3 0x7a 0x4f 0x2c 0xc0 0x62 0x4f 0x16 0x90 0xf6 0x46 0x06 0xcf 0x38 0x59 0x45 0xb2 0xbe 0xc4 0xea. This value is then base64-encoded, to give the value "s3pPLMBiTxaQ9kYGzzhZRbK+x0o=", which would be returned in the "Sec-WebSocket-Accept" header.

3. Optionally, a "Sec-WebSocket-Protocol" header, with a value /subprotocol/ as defined in Paragraph 2 of [Section 5.2.2](#).

4. Optionally, a "Sec-WebSocket-Extensions" header, with a value /extensions/ as defined in Paragraph 2 of [Section 5.2.2](#).

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.

[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.](#) Definitions

[7.1.1.](#) Close the WebSocket Connection

To `_Close the WebSocket Connection_`, an endpoint closes the underlying TCP connection. An endpoint **SHOULD** use a method that cleanly closes the TCP connection, as well as the TLS session, if applicable, discarding any trailing bytes that may be received. An endpoint **MAY** close the connection via any means available when necessary, such as when under attack.

As an example of how to obtain a clean closure in C using Berkeley sockets, one would call `shutdown()` with `SHUT_WR` on the socket, call `recv()` until obtaining a return value of 0 indicating that the peer has also performed an orderly shutdown, and finally calling `close()` on the socket.

[7.1.2.](#) Start the WebSocket Closing Handshake

To `_start the WebSocket closing handshake_`, an endpoint **MUST** send a Close control frame, as described in [Section 4.5.1](#). Once an endpoint has both sent and received a Close control frame, that endpoint **SHOULD** `_Close the WebSocket Connection_` as defined in [Section 7.1.1](#).

[7.1.3.](#) The WebSocket Connection Is Closed

When the underlying TCP connection is closed, it is said that `_the WebSocket connection is closed_`. If the tcp connection was closed after the WebSocket closing handshake was completed, the WebSocket connection is said to have been closed `_cleanly_`.

[7.1.4.](#) Fail the WebSocket Connection

Certain algorithms and specifications require a 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) in an appropriate manner.

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.

[7.2.](#) Abnormal closures

[7.2.1.](#) Client-initiated closure

Certain algorithms, namely during the initial handshake, require the user agent to **fail the WebSocket connection**. To do so, the user agent **MUST** *_Close the WebSocket connection_* as previously defined, and **MAY** report the problem to the user via an appropriate mechanism (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.

[7.2.2.](#) 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_* ([Section 7.1.1](#)).

[7.3.](#) Normal closure of connections

Servers **MAY** close the WebSocket connection whenever desired. User agents **SHOULD NOT** close the WebSocket connection arbitrarily. In either case, an endpoint initiates a closure by following the procedures to *_start the WebSocket closing handshake_* ([Section 7.1.2](#)).

[7.4.](#) Status codes

When closing an established connection (e.g. when sending a Close frame, after the handshake has completed), an endpoint **MAY** indicate a reason for closure. The interpretation of this reason by an endpoint, and the action an endpoint should take given this reason, are left undefined by this specification. This specification defines a set of pre-defined status codes, and specifies which ranges may be used by extensions, frameworks, and end applications. The status code and any associated textual message are optional components of a Close frame.

[7.4.1.](#) Defined Status Codes

Endpoints **MAY** use the following pre-defined status codes when sending a Close frame.

1000

1000 indicates a normal closure, meaning whatever purpose the

connection was established for has been fulfilled.

1001

1001 indicates that an endpoint is "going away", such as a server going down, or a browser having navigated away from a page.

1002

1002 indicates that an endpoint is terminating the connection due to a protocol error.

1003

1003 indicates that an endpoint is terminating the connection because it has received a type of data it cannot accept (e.g. an endpoint that understands only text data MAY send this if it receives a binary message).

1004

1004 indicates that an endpoint is terminating the connection because it has received a message that is too large.

[7.4.2.](#) Reserved status code ranges

0-999

Status codes in the range 0-999 are not used.

1000-1999

Status codes in the range 1000-1999 are reserved for definition by this protocol.

2000-2999

Status codes in the range 2000-2999 are reserved for use by extensions.

3000-3999

Status codes in the range 3000-3999 MAY be used by libraries and frameworks. The interpretation of these codes is undefined by this protocol. End applications MUST NOT use status codes in this range.

Fette

Expires October 24, 2011

[Page 39]

Internet-Draft

The WebSocket protocol

April 2011

4000-4999

Status codes in the range 4000-4999 MAY be used by application code. The interpretation of these codes is undefined by this protocol.

[8.](#) Extensions

WebSocket clients MAY request extensions to this specification, and WebSocket servers MAY accept some or all extensions requested by the client. A server MUST NOT respond with any extension not requested by the client. If extension parameters are included in negotiations between the client and the server, those parameters MUST be chosen in accordance with the specification of the extension to which the parameters apply.

[8.1.](#) Negotiating extensions

A client requests extensions by including a "Sec-WebSocket-Extensions" header, which follows the normal rules for HTTP headers (see [\[RFC2616\] section 4.2](#)) and the value of the header is defined by the following ABNF. Note that unlike other section of the document this section is using ABNF syntax/rules from [\[RFC2616\]](#).

```
extension-list = 1#extension
extension      = extension-token *( ";" extension-param )
extension-token = registered-token | private-use-token
registered-token = token
private-use-token = "x-" token
extension-param = token [ "=" ( token | quoted-string ) ]
```

Note that like other HTTP headers, this header MAY be split or combined across multiple lines. Ergo, the following are equivalent:

```
Sec-WebSocket-Extensions: foo
Sec-WebSocket-Extensions: bar; baz=2
```

is exactly equivalent to

```
Sec-WebSocket-Extensions: foo, bar; baz=2
```

Any extension-token used MUST either be a registered token (registration TBD), or have a prefix of "x-" to indicate a private-use token. The parameters supplied with any given extension MUST be defined for that extension. Note that the client is only offering to use any advertised extensions, and MUST NOT use them unless the server indicates that it wishes to use the extension.

Note that the order of extensions is significant. Any interactions between multiple extensions MAY be defined in the documents defining the extensions. In the absence of such definition, the interpretation is that the headers listed by the client in its request represent a preference of the headers it wishes to use, with the first options listed being most preferable. The extensions

listed by the server in response represent the extensions actually in use for the connection. Should the extensions modify the data and/or framing, the order of operations on the data should be assumed to be the same as the order in which the extensions are listed in the server's response in the opening handshake.

For example, if there are two extensions "foo" and "bar", if the header |Sec-WebSocket-Extensions| sent by the server has the value "foo, bar" then operations on the data will be made as bar(foo(data)), be those changes to the data itself (such as compression) or changes to the framing they may "stack".

Non-normative examples of acceptable extension headers:

```
Sec-WebSocket-Extensions: deflate-stream
Sec-WebSocket-Extensions: mux; max-channels=4; flow-control, deflate-stream
Sec-WebSocket-Extensions: x-private-extension
```

A server accepts one or more extensions by including a |Sec-WebSocket-Extensions| header containing one or more extensions which were requested by the client. The interpretation of any extension parameters, and what constitutes a valid response by a server to a requested set of parameters by a client, will be defined by each such extension.

[8.2.](#) 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.

[8.2.1.](#) 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 [\[RFC1951\]](#).

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.

Historically there have been some confusion and interoperability problems around the specification of compression algorithms. In this specification "deflate-stream" requires a [\[RFC1951\]](#) deflate encoding. It MUST NOT be wrapped in any of the header formats often associated with [RFC 1951](#) such as "zlib" [\[RFC1950\]](#). This requirement is given special attention with this note because of confusion in this area, the presence of some popular open source libraries that create both formats under a single API call with confusing naming conventions, and the fact that the popular HTTP [\[RFC2616\]](#) specification defines

"deflate" compression differently than this specification.

[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.

In addition to endpoints being the target of attacks via WebSockets, other parts of web infrastructure, such as proxies, may be the subject of an attack. In particular, an intermediary may interpret a WebSocket message from a client as a request, and a message from the server as a response to that request. For instance, an attacker could get a browser to establish a connection to its server, get the browser to send a message that looks to an intermediary like a GET request for a common piece of JavaScript on another domain, and send

back a message that is interpreted as a cacheable response to that request, thus poisoning the cache for other users. To prevent this attack, messages sent from clients are masked on the wire with a 32-bit value, to prevent an attacker from controlling the bits on the wire and thus lessen the probability of an attacker being able to construct a message that can be misinterpreted by a proxy as a non-WebSocket request.

As mentioned in [Section 6.2](#), servers must be extremely cautious interpreting invalid UTF-8 data from the client. A naive UTF-8 parsing implementation can result in buffer overflows in the case of invalid input data.

[10.](#) IANA considerations

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

A |ws:| URI 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> [[RFC3986](#)] 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 specifications. [[RFC3986](#)] [[RFC3987](#)]

Applications/protocols that use this URI scheme name.

WebSocket protocol.

Interoperability considerations.

None.

Fette

Expires October 24, 2011

[Page 46]

Internet-Draft

The WebSocket protocol

April 2011

Security considerations.

See "Security considerations" section above.

Contact.

HYBI WG <hybi@ietf.org>

Author/Change controller.

IETF <iesg@ietf.org>

References.

RFC XXXX

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

A |wss:| URI 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

Fette

Expires October 24, 2011

[Page 47]

Internet-Draft

The WebSocket protocol

April 2011

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.
HYBI WG <hybi@ietf.org>

Author/Change controller.
IETF <iesg@ietf.org>

References.
RFC XXXX

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

Name of token.
WebSocket

Author/Change controller.

IETF <iesg@ietf.org>

Contact.

HYBI <hybi@ietf.org>

References.

RFC XXXX

[10.4.](#) Sec-WebSocket-Key

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

Header field name

Sec-WebSocket-Key

Applicable protocol

http

Fette

Expires October 24, 2011

[Page 48]

Internet-Draft

The WebSocket protocol

April 2011

Status

standard

Author/Change controller

IETF

Specification document(s)

RFC XXXX

Related information

This header field is only used for WebSocket handshake.

The |Sec-WebSocket-Key| header is used in the WebSocket handshake. It is 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-WebSocket clients (e.g. HTTP clients) that are being abused to send data to unsuspecting WebSocket servers.

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

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

Header field name

Sec-WebSocket-Extensions

Applicable protocol

http

Status

standard

Author/Change controller

IETF

Specification document(s)

RFC XXXX

Related information

This header field is only used for WebSocket handshake.

The |Sec-WebSocket-Extensions| header is used in the WebSocket handshake. It is initially sent from the client to the server, and then subsequently sent from the server to the client, to agree on a set of protocol-level extensions to use for the duration of the connection.

Fette

Expires October 24, 2011

[Page 49]

Internet-Draft

The WebSocket protocol

April 2011

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

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

Header field name

Sec-WebSocket-Accept

Applicable protocol

http

Status

standard

Author/Change controller
IETF

Specification document(s)
RFC XXXX

Related information
This header field is only used for WebSocket handshake.

The |Sec-WebSocket-Accept| header is used in the WebSocket handshake. It is sent from the server to the client to confirm that the server is willing to initiate the connection.

[10.7.](#) 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
standard

Author/Change controller
IETF

Specification document(s)
RFC XXXX

Related information
This header field is only used for WebSocket handshake.

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.

[10.8.](#) 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
standard

Author/Change controller
IETF

Specification document(s)
RFC XXXX

Related information
This header field is only used for WebSocket handshake.

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.

[10.9.](#) Sec-WebSocket-Version

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

Header field name
Sec-WebSocket-Version

Applicable protocol
http

Status
standard

Author/Change controller
IETF

Specification document(s)
RFC XXXX

Related information
This header field is only used for WebSocket handshake.

The |Sec-WebSocket-Version| header is used in the WebSocket handshake. It is sent from the client to the server to indicate the 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. [[I-D.ietf-websec-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 URI using the steps to parse a WebSocket URI's components. These steps fail if the URI 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 at <http://whatwg.org/html5>.

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

Special thanks also to Adam Barth for providing a significant amount of text and background research for the Data Masking section of this specification.

[13.](#) References

[13.1.](#) Normative References

[ANSI.X3-4.1986]

American National Standards Institute, "Coded Character Set - 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.

[FIPS.180-2.2002]

National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-2, August 2002, <<http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>>.

[RFC1951] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", [RFC 1951](#), May 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[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](#), June 1999.

[RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), May 2000.

[RFC3490] Faltstrom, P., Hoffman, P., and A. Costello, "Internationalizing Domain Names in Applications (IDNA)",

[RFC 3490](#), March 2003.

- [RFC3492] Costello, A., "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)", [RFC 3492](#), March 2003.
- [RFC3548] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 3548](#), July 2003.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), September 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.

Fette

Expires October 24, 2011

[Page 55]

Internet-Draft

The WebSocket protocol

April 2011

- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", [RFC 3987](#), January 2005.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.
- [RFC6066] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", [RFC 6066](#), January 2011.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), October 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.

[13.2](#). Informative References

- [WSAPI] Hickson, I., "The Web Sockets API", August 2010, <<http://dev.w3.org/html5/websockets/>>.
- [I-D.ietf-httpstate-cookie]

Barth, A., "HTTP State Management Mechanism",
[draft-ietf-httpstate-cookie-20](#) (work in progress),
December 2010.

[I-D.ietf-websec-origin]

Barth, A., "The Web Origin Concept",
[draft-ietf-websec-origin-00](#) (work in progress),
December 2010.

[RFC1950] Deutsch, L. and J-L. Gailly, "ZLIB Compressed Data Format
Specification version 3.3", [RFC 1950](#), May 1996.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", [RFC 5321](#),
October 2008.

[RFC6202] Loreto, S., Saint-Andre, P., Salsano, S., and G. Wilkins,
"Known Issues and Best Practices for the Use of Long
Polling and Streaming in Bidirectional HTTP", [RFC 6202](#),
April 2011.

Fette

Expires October 24, 2011

[Page 56]

Internet-Draft

The WebSocket protocol

April 2011

Author's Address

Ian Fette
Google, Inc.

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

