

Internet Engineering Task Force (IETF)
Request for Comments: 6455
Category: Standards Track
ISSN: 2070-1721

I. Fette
Google, Inc.
A. Melnikov
Isode Ltd.
December 2011

The WebSocket Protocol

Abstract

The WebSocket Protocol enables two-way communication between a client running untrusted code 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 opening 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).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6455>.

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	5
1.3.	Opening Handshake	6
1.4.	Closing Handshake	9
1.5.	Design Philosophy	9
1.6.	Security Model	10
1.7.	Relationship to TCP and HTTP	11
1.8.	Establishing a Connection	11
1.9.	Subprotocols Using the WebSocket Protocol	12
2.	Conformance Requirements	12
2.1.	Terminology and Other Conventions	13
3.	WebSocket URIs	14
4.	Opening Handshake	14
4.1.	Client Requirements	14
4.2.	Server-Side Requirements	20
4.2.1.	Reading the Client's Opening Handshake	21
4.2.2.	Sending the Server's Opening Handshake	22
4.3.	Collected ABNF for New Header Fields Used in Handshake	25
4.4.	Supporting Multiple Versions of WebSocket Protocol	26
5.	Data Framing	27
5.1.	Overview	27
5.2.	Base Framing Protocol	28
5.3.	Client-to-Server Masking	32
5.4.	Fragmentation	33
5.5.	Control Frames	36
5.5.1.	Close	36
5.5.2.	Ping	37
5.5.3.	Pong	37
5.6.	Data Frames	38
5.7.	Examples	38
5.8.	Extensibility	39
6.	Sending and Receiving Data	39
6.1.	Sending Data	39
6.2.	Receiving Data	40
7.	Closing the Connection	41
7.1.	Definitions	41
7.1.1.	Close the WebSocket Connection	41
7.1.2.	Start the WebSocket Closing Handshake	42
7.1.3.	The WebSocket Closing Handshake is Started	42
7.1.4.	The WebSocket Connection is Closed	42
7.1.5.	The WebSocket Connection Close Code	42

7.1.6.	The WebSocket Connection Close Reason	43
7.1.7.	Fail the WebSocket Connection	43
7.2.	Abnormal Closures	44
7.2.1.	Client-Initiated Closure	44
7.2.2.	Server-Initiated Closure	44
7.2.3.	Recovering from Abnormal Closure	44
7.3.	Normal Closure of Connections	45
7.4.	Status Codes	45
7.4.1.	Defined Status Codes	45
7.4.2.	Reserved Status Code Ranges	47
8.	Error Handling	48
8.1.	Handling Errors in UTF-8-Encoded Data	48
9.	Extensions	48
9.1.	Negotiating Extensions	48
9.2.	Known Extensions	50
10.	Security Considerations	50
10.1.	Non-Browser Clients	50
10.2.	Origin Considerations	50
10.3.	Attacks On Infrastructure (Masking)	51
10.4.	Implementation-Specific Limits	52
10.5.	WebSocket Client Authentication	53
10.6.	Connection Confidentiality and Integrity	53
10.7.	Handling of Invalid Data	53
10.8.	Use of SHA-1 by the WebSocket Handshake	54
11.	IANA Considerations	54
11.1.	Registration of New URI Schemes	54
11.1.1.	Registration of "ws" Scheme	54
11.1.2.	Registration of "wss" Scheme	55
11.2.	Registration of the "WebSocket" HTTP Upgrade Keyword	56
11.3.	Registration of New HTTP Header Fields	57
11.3.1.	Sec-WebSocket-Key	57
11.3.2.	Sec-WebSocket-Extensions	58
11.3.3.	Sec-WebSocket-Accept	58
11.3.4.	Sec-WebSocket-Protocol	59
11.3.5.	Sec-WebSocket-Version	60
11.4.	WebSocket Extension Name Registry	61
11.5.	WebSocket Subprotocol Name Registry	61
11.6.	WebSocket Version Number Registry	62
11.7.	WebSocket Close Code Number Registry	64
11.8.	WebSocket Opcode Registry	65
11.9.	WebSocket Framing Header Bits Registry	66
12.	Using the WebSocket Protocol from Other Specifications	66
13.	Acknowledgements	67
14.	References	68
14.1.	Normative References	68
14.2.	Informative References	69

1. Introduction

1.1. Background

`_This section is non-normative._`

Historically, creating web applications that need bidirectional communication between a client and a server (e.g., instant messaging and gaming applications) 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 [WSAPI], it provides an alternative to HTTP polling for two-way communication from a web page to a remote server.

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.

The WebSocket Protocol is designed to supersede existing bidirectional communication technologies that use HTTP as a transport layer to benefit from existing infrastructure (proxies, filtering, authentication). Such technologies were implemented as trade-offs between efficiency and reliability because HTTP was not initially meant to be used for bidirectional communication (see [RFC6202] for further discussion). The WebSocket Protocol attempts to address the goals of existing bidirectional HTTP technologies in the context of the existing HTTP infrastructure; as such, it is designed to work over HTTP ports 80 and 443 as well as to support HTTP proxies and intermediaries, even if this implies some complexity specific to the current environment. However, the design does not limit WebSocket to HTTP, and future implementations could use a simpler handshake over a

dedicated port without reinventing the entire protocol. This last point is important because the traffic patterns of interactive messaging do not closely match standard HTTP traffic and can induce unusual loads on some components.

1.2. Protocol Overview

This section is non-normative.

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

The handshake from the client looks as follows:

```
GET /chat HTTP/1.1
Host: server.example.com
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Key: dGhlIHNhbXBsZSBub25jZQ==
Origin: http://example.com
Sec-WebSocket-Protocol: chat, superchat
Sec-WebSocket-Version: 13
```

The handshake from the server looks as follows:

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

An unordered set of header fields comes after the leading line in both cases. The meaning of these header fields is specified in Section 4 of this document. Additional header fields may also be present, such as cookies [RFC6265]. 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.

After a successful handshake, clients and servers transfer data back and forth in conceptual units referred to in this specification as "messages". On the wire, a message is composed of one or more

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.