

HTTP Working Group  
INTERNET-DRAFT  
<draft-ietf-http-v11-spec-00.txt>  
Expires May 22, 1996

R. Fielding, UC Irvine  
H. Frystyk, MIT/LCS  
T. Berners-Lee, MIT/LCS  
November 22, 1995

## Hypertext Transfer Protocol -- HTTP/1.1

### Status of this Memo

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

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

To learn the current status of any Internet-Draft, please check the "lid-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this document is unlimited. Please send comments to the HTTP working group at <http-wg@cuckoo.hpl.hp.com>. Discussions of the working group are archived at <URL:http://www.ics.uci.edu/pub/ietf/http/>. General discussions about HTTP and the applications which use HTTP should take place on the <www-talk@w3.org> mailing list.

### Abstract

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. It is a generic, stateless, object-oriented protocol which can be used for many tasks, such as name servers and distributed object management systems, through extension of its request methods (commands). A feature of HTTP is the typing and negotiation of data representation, allowing systems to be built independently of the data being transferred.

HTTP has been in use by the World-Wide Web global information initiative since 1990. This specification defines the protocol referred to as "HTTP/1.1".

### Table of Contents

1. Introduction
  - 1.1 Purpose
  - 1.2 Requirements
  - 1.3 Terminology
  - 1.4 Overall Operation
2. Notational Conventions and Generic Grammar
  - 2.1 Augmented BNF
  - 2.2 Basic Rules

- 3. Protocol Parameters
  - 3.1 HTTP Version
  - 3.2 Uniform Resource Identifiers
    - 3.2.1 General Syntax
    - 3.2.2 http URL
  - 3.3 Date/Time Formats
    - 3.3.1 Full Date
    - 3.3.2 Delta Seconds
  - 3.4 Character Sets
  - 3.5 Content Codings
  - 3.6 Transfer Codings
  - 3.7 Media Types
    - 3.7.1 Canonicalization and Text Defaults
    - 3.7.2 Multipart Types
  - 3.8 Product Tokens
  - 3.9 Quality Values
  - 3.10 Language Tags
  - 3.11 Logic Bags
- 4. HTTP Message
  - 4.1 Message Types
  - 4.2 Message Headers
  - 4.3 General Header Fields
- 5. Request
  - 5.1 Request-Line
    - 5.1.1 Method
    - 5.1.2 Request-URI
  - 5.2 Request Header Fields
- 6. Response
  - 6.1 Status-Line
    - 6.1.1 Status Code and Reason Phrase
  - 6.2 Response Header Fields
- 7. Entity
  - 7.1 Entity Header Fields
  - 7.2 Entity Body
    - 7.2.1 Type
    - 7.2.2 Length
- 8. Method Definitions
  - 8.1 OPTIONS
  - 8.2 GET
  - 8.3 HEAD
  - 8.4 POST
  - 8.5 PUT
  - 8.6 PATCH
  - 8.7 COPY
  - 8.8 MOVE
  - 8.9 DELETE
  - 8.10 LINK
  - 8.11 UNLINK
  - 8.12 TRACE
  - 8.13 WRAPPED
- 9. Status Code Definitions
  - 9.1 Informational 1xx
  - 9.2 Successful 2xx
  - 9.3 Redirection 3xx

- 9.4 Client Error 4xx
- 9.5 Server Error 5xx
- 10. Header Field Definitions
  - 10.1 Accept
  - 10.2 Accept-Charset
  - 10.3 Accept-Encoding
  - 10.4 Accept-Language
  - 10.5 Allow
  - 10.6 Authorization
  - 10.7 Base
  - 10.8 Cache-Control
  - 10.9 Connection
    - 10.9.1 Persistent Connections
  - 10.10 Content-Encoding
  - 10.11 Content-Language
  - 10.12 Content-Length
  - 10.13 Content-MD5
  - 10.14 Content-Range
  - 10.15 Content-Type
  - 10.16 Content-Version
  - 10.17 Date
  - 10.18 Derived-From
  - 10.19 Expires
  - 10.20 Forwarded
  - 10.21 From
  - 10.22 Host
  - 10.23 If-Modified-Since
  - 10.24 Keep-Alive
  - 10.25 Last-Modified
  - 10.26 Link
  - 10.27 Location
  - 10.28 MIME-Version
  - 10.29 Pragma
  - 10.30 Proxy-Authenticate
  - 10.31 Proxy-Authorization
  - 10.32 Public
  - 10.33 Range
  - 10.34 Referer
  - 10.35 Refresh
  - 10.36 Retry-After
  - 10.37 Server
  - 10.38 Title
  - 10.39 Transfer Encoding
  - 10.40 Unless
  - 10.41 Upgrade
  - 10.42 URI
  - 10.43 User-Agent
  - 10.44 WWW-Authenticate
- 11. Access Authentication
  - 11.1 Basic Authentication Scheme
  - 11.2 Digest Authentication Scheme
- 12. Content Negotiation
  - 12.1 Preemptive Negotiation
- 13. Caching
- 14. Security Considerations
  - 14.1 Authentication of Clients

- 14.2 Safe Methods
- 14.3 Abuse of Server Log Information
- 14.4 Transfer of Sensitive Information

## 15. Acknowledgments

## 16. References

## 17. Authors' Addresses

[Appendix A.](#) Internet Media Type message/http

[Appendix B.](#) Tolerant Applications

[Appendix C.](#) Relationship to MIME

- C.1 Conversion to Canonical Form
  - C.1.1 Representation of Line Breaks
  - C.1.2 Default Character Set
- C.2 Conversion of Date Formats
- C.3 Introduction of Content-Encoding
- C.4 No Content-Transfer-Encoding
- C.5 Introduction of Transfer-Encoding

[Appendix D.](#) Changes from HTTP/1.0

## 1. Introduction

### 1.1 Purpose

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. HTTP has been in use by the World-Wide Web global information initiative since 1990. The first version of HTTP, referred to as HTTP/0.9, was a simple protocol for raw data transfer across the Internet. HTTP/1.0, as defined by RFC xxxx [6], improved the protocol by allowing messages to be in the format of MIME-like entities, containing metainformation about the data transferred and modifiers on the request/response semantics. However, HTTP/1.0 does not sufficiently take into consideration the effect of hierarchical proxies and caching, the desire for persistent connections and virtual hosts, and a number of other details that slipped through the cracks of existing implementations. In addition, the proliferation of incompletely-implemented applications calling themselves "HTTP/1.0" has necessitated a protocol version change in order for two communicating applications to determine each other's true capabilities.

This specification defines the protocol referred to as "HTTP/1.1". This protocol is backwards-compatible with HTTP/1.0, but includes more stringent requirements in order to ensure reliable implementation of its features.

Practical information systems require more functionality than simple retrieval, including search, front-end update, and annotation. HTTP allows an open-ended set of methods to be used to indicate the purpose of a request. It builds on the discipline of reference provided by the Uniform Resource Identifier (URI) [3], as a location (URL) [4] or name (URN) [20], for indicating the resource on which a method is to be applied. Messages are passed in a format similar to that used by Internet Mail [9] and the Multipurpose Internet Mail Extensions (MIME) [7].

HTTP is also used as a generic protocol for communication between user agents and proxies/gateways to other Internet protocols, such as SMTP [16], NNTP [13], FTP [18], Gopher [2], and WAIS [10], allowing basic hypermedia access to resources available from diverse applications and simplifying the implementation of user agents.

## 1.2 Requirements

This specification uses the same words as RFC 1123 [8] for defining the significance of each particular requirement. These words are:

must

This word or the adjective "required" means that the item is an absolute requirement of the specification.

should

This word or the adjective "recommended" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.

may

This word or the adjective "optional" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

An implementation is not compliant if it fails to satisfy one or more of the must requirements for the protocols it implements. An implementation that satisfies all the must and all the should requirements for its protocols is said to be "unconditionally compliant"; one that satisfies all the must requirements but not all the should requirements for its protocols is said to be "conditionally compliant".

## 1.3 Terminology

This specification uses a number of terms to refer to the roles played by participants in, and objects of, the HTTP communication.

connection

A transport layer virtual circuit established between two application programs for the purpose of communication.

message

The basic unit of HTTP communication, consisting of a structured sequence of octets matching the syntax defined in Section 4 and transmitted via the connection.

request

An HTTP request message (as defined in Section 5).

# 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.