**HTTP/1.1, part 1: URIs, Connections, and Message Parsing**
**draft-ietf-httpbis-p1-messaging-10**

**Abstract**

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypertext information systems. HTTP has been in use by the World Wide Web global information initiative since 1990. This document is Part 1 of the seven-part specification that defines the protocol referred to as "HTTP/1.1" and, taken together, obsoletes RFC 2616. Part 1 provides an overview of HTTP and its associated terminology, defines the "http" and "https" Uniform Resource Identifier (URI) schemes, defines the generic message syntax and parsing requirements for HTTP message frames, and describes general security concerns for implementations.

**Editorial Note (To be removed by RFC Editor)**

Discussion of this draft should take place on the HTTPBIS working group mailing list (ietf-http-wg@w3.org). The current issues list is at http://tools.ietf.org/wg/httpbis/trac/report/3 and related documents (including fancy diffs) can be found at http://tools.ietf.org/wg/httpbis/.
The changes in this draft are summarized in Appendix D.11 (Since draft-ietf-httpbis-p1-messaging-09).

**Table of Contents**

---

## 1.  Introduction                                            [TOC](#)

The Hypertext Transfer Protocol (HTTP) is an application-level request/
response protocol that uses extensible semantics and MIME-like message
payloads for flexible interaction with network-based hypertext
information systems. HTTP relies upon the Uniform Resource Identifier
(URI) standard [RFC3986] (Berners-Lee, T., Fielding, R., and L.
Masinter, "Uniform Resource Identifier (URI): Generic Syntax,"
January 2005.) to indicate request targets and relationships between
resources. Messages are passed in a format similar to that used by
Internet mail [RFC5322] (Resnick, P., "Internet Message Format,"
October 2008.) and the Multipurpose Internet Mail Extensions (MIME)
[RFC2045] (Freed, N. and N. Borenstein, "Multipurpose Internet Mail
Extensions (MIME) Part One: Format of Internet Message Bodies,"
November 1996.) (see Appendix A of [Part3] (Fielding, R., Ed., Gettys,
J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.,
Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload
and Content Negotiation," July 2010.) for the differences between HTTP
and MIME messages).
HTTP is a generic interface protocol for information systems. It is
designed to hide the details of how a service is implemented by
presenting a uniform interface to clients that is independent of the
types of resources provided. Likewise, servers do not need to be aware
of each client's purpose: an HTTP request can be considered in
isolation rather than being associated with a specific type of client
or a predetermined sequence of application steps. The result is a
protocol that can be used effectively in many different contexts and
for which implementations can evolve independently over time.
HTTP is also designed for use as a generic protocol for translating
communication to and from other Internet information systems. HTTP

proxies and gateways provide access to alternative information services by translating their diverse protocols into a hypertext format that can be viewed and manipulated by clients in the same way as HTTP services. One consequence of HTTP flexibility is that the protocol cannot be defined in terms of what occurs behind the interface. Instead, we are limited to defining the syntax of communication, the intent of received communication, and the expected behavior of recipients. If the communication is considered in isolation, then successful actions should be reflected in corresponding changes to the observable interface provided by servers. However, since multiple clients may act in parallel and perhaps at cross-purposes, we cannot require that such changes be observable beyond the scope of a single response. This document is Part 1 of the seven-part specification of HTTP, defining the protocol referred to as "HTTP/1.1" and obsoleting [RFC2616] (Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.). Part 1 describes the architectural elements that are used or referred to in HTTP, defines the "http" and "https" URI schemes, describes overall network operation and connection management, and defines HTTP message framing and forwarding requirements. Our goal is to define all of the mechanisms necessary for HTTP message handling that are independent of message semantics, thereby defining the complete set of requirements for message parsers and message-forwarding intermediaries.

---

## 1.1.  Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.). An implementation is not compliant if it fails to satisfy one or more of the "MUST" or "REQUIRED" level requirements for the protocols it implements. An implementation that satisfies all the "MUST" or "REQUIRED" level and all the "SHOULD" level requirements for its protocols is said to be "unconditionally compliant"; one that satisfies all the "MUST" level requirements but not all the "SHOULD" level requirements for its protocols is said to be "conditionally compliant".

---

## 1.2.  Syntax Notation

This specification uses the Augmented Backus-Naur Form (ABNF) notation of [RFC5234] (Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," January 2008.).
The following core rules are included by reference, as defined in [RFC5234] (Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," January 2008.), Appendix B.1: ALPHA (letters), CR (carriage return), CRLF (CR LF), CTL (controls), DIGIT (decimal 0-9), DQUOTE (double quote), HEXDIG (hexadecimal 0-9/A-F/a-f), LF (line feed), OCTET (any 8-bit sequence of data), SP (space), VCHAR (any visible [USASCII] (American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange," 1986.) character), and WSP (whitespace).
As a syntactical convention, ABNF rule names prefixed with "obs-" denote "obsolete" grammar rules that appear for historical reasons.

---

### 1.2.1.  ABNF Extension: #rule

The #rule extension to the ABNF rules of [RFC5234] (Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," January 2008.) is used to improve readability.
A construct "#" is defined, similar to "*", for defining comma-delimited lists of elements. The full form is "<n>#<m>element" indicating at least <n> and at most <m> elements, each separated by a single comma (",") and optional whitespace (OWS, Section 1.2.2 (Basic Rules)).
Thus,

        1#element => element *( OWS "," OWS element )

and:

        #element => [ 1#element ]

and for n >= 1 and m > 1:

        <n>#<m>element => element <n-1>*<m-1>( OWS "," OWS element )

For compatibility with legacy list rules, recipients SHOULD accept empty list elements. In other words, consumers would follow the list productions:

        #element => [ ( "," / element ) *( OWS "," [ OWS element ] ) ]

        1#element => *( "," OWS ) element *( OWS "," [ OWS element ] )

Note that empty elements do not contribute to the count of elements present, though.

For example, given these ABNF productions:

```
example-list      = 1#example-list-elmt
example-list-elmt = token ; see Section 1.2.2
```

Then these are valid values for example-list (not including the double quotes, which are present for delimitation only):

```
"foo,bar"
" foo ,bar,"
"  foo , ,bar,charlie   "
"foo ,bar,   charlie "
```

But these values would be invalid, as at least one non-empty element is required:

```
""
","
",     ,"
```

[Appendix C (Collected ABNF)](#) shows the collected ABNF, with the list rules expanded as explained above.

---

### 1.2.2.  Basic Rules

HTTP/1.1 defines the sequence CR LF as the end-of-line marker for all protocol elements except the entity-body (see [Appendix A (Tolerant Applications)](#) for tolerant applications). The end-of-line marker within an entity-body is defined by its associated media type, as described in Section 2.3 of [[Part3] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload and Content Negotiation," July 2010.)](#).
This specification uses three rules to denote the use of linear whitespace: OWS (optional whitespace), RWS (required whitespace), and BWS ("bad" whitespace).
The OWS rule is used where zero or more linear whitespace characters may appear. OWS SHOULD either not be produced or be produced as a single SP character. Multiple OWS characters that occur within field-content SHOULD be replaced with a single SP before interpreting the field value or forwarding the message downstream.
RWS is used when at least one linear whitespace character is required to separate field tokens. RWS SHOULD be produced as a single SP character. Multiple RWS characters that occur within field-content SHOULD be replaced with a single SP before interpreting the field value or forwarding the message downstream.

BWS is used where the grammar allows optional whitespace for historical reasons but senders SHOULD NOT produce it in messages. HTTP/1.1 recipients MUST accept such bad optional whitespace and remove it before interpreting the field value or forwarding the message downstream.

```
OWS            = *( [ obs-fold ] WSP )
               ; "optional" whitespace
RWS            = 1*( [ obs-fold ] WSP )
               ; "required" whitespace
BWS            = OWS
               ; "bad" whitespace
obs-fold       = CRLF
               ; see Section 3.2
```

Many HTTP/1.1 header field values consist of words (token or quoted-string) separated by whitespace or special characters. These special characters MUST be in a quoted string to be used within a parameter value (as defined in [Section 6.2 (Transfer Codings)](#)).

```
word           = token / quoted-string

token          = 1*tchar

tchar          = "!" / "#" / "$" / "%" / "&" / "'" / "*"
               / "+" / "-" / "." / "^" / "_" / "`" / "|" / "~"
               / DIGIT / ALPHA
               ; any VCHAR, except special

special        = "(" / ")" / "<" / ">" / "@" / ","
               / ";" / ":" / "\" / DQUOTE / "/" / "["
               / "]" / "?" / "=" / "{" / "}"
```

A string of text is parsed as a single word if it is quoted using double-quote marks.

```
quoted-string  = DQUOTE *( qdtext / quoted-pair ) DQUOTE
qdtext         = OWS / %x21 / %x23-5B / %x5D-7E / obs-text
               ; OWS / <VCHAR except DQUOTE and "\"> / obs-text
obs-text       = %x80-FF
```

The backslash character ("\") can be used as a single-character quoting mechanism within quoted-string constructs:

```
quoted-pair    = "\" ( WSP / VCHAR / obs-text )
```

Producers SHOULD NOT escape characters that do not require escaping (i.e., other than DQUOTE and the backslash character).

### 1.2.3.  ABNF Rules defined in other Parts of the Specification

The ABNF rules below are defined in other parts:

```
request-header  = <request-header, defined in [Part2], Section 3>
response-header = <response-header, defined in [Part2], Section 5>


entity-body     = <entity-body, defined in [Part3], Section 3.2>
entity-header   = <entity-header, defined in [Part3], Section 3.1>


Cache-Control   = <Cache-Control, defined in [Part6], Section 3.4>
Pragma          = <Pragma, defined in [Part6], Section 3.4>
Warning         = <Warning, defined in [Part6], Section 3.6>
```

---

## 2.  HTTP architecture

HTTP was created for the World Wide Web architecture and has evolved over time to support the scalability needs of a worldwide hypertext system. Much of that architecture is reflected in the terminology and syntax productions used to define HTTP.

---

### 2.1.  Client/Server Operation

HTTP is a request/response protocol that operates by exchanging messages across a reliable transport or session-layer connection. An HTTP client is a program that establishes a connection to a server for the purpose of sending one or more HTTP requests. An HTTP server is a program that accepts connections in order to service HTTP requests by sending HTTP responses.

Note that the terms "client" and "server" refer only to the roles that these programs perform for a particular connection. The same program may act as a client on some connections and a server on others. We use the term "user agent" to refer to the program that initiates a request, such as a WWW browser, editor, or spider (web-traversing robot), and the term "origin server" to refer to the program that can originate authoritative responses to a request.
Most HTTP communication consists of a retrieval request (GET) for a representation of some resource identified by a URI. In the simplest case, this may be accomplished via a single connection (v) between the user agent (UA) and the origin server (O).

```
        request chain ------------------------>
    UA -------------------v------------------ O
        <---------------------- response chain
```

A client sends an HTTP request to the server in the form of a request
message ([Section 4 (Request)](#)), beginning with a method, URI, and
protocol version, followed by MIME-like header fields containing
request modifiers, client information, and payload metadata, an empty
line to indicate the end of the header section, and finally the payload
body (if any).
A server responds to the client's request by sending an HTTP response
message ([Section 5 (Response)](#)), beginning with a status line that
includes the protocol version, a success or error code, and textual
reason phrase, followed by MIME-like header fields containing server
information, resource metadata, and payload metadata, an empty line to
indicate the end of the header section, and finally the payload body
(if any).
The following example illustrates a typical message exchange for a GET
request on the URI "http://www.example.com/hello.txt":
client request:

```
    GET /hello.txt HTTP/1.1
    User-Agent: curl/7.16.3 libcurl/7.16.3 OpenSSL/0.9.7l zlib/1.2.3
    Host: www.example.com
    Accept: */*
```

server response:

```
    HTTP/1.1 200 OK
    Date: Mon, 27 Jul 2009 12:28:53 GMT
    Server: Apache
    Last-Modified: Wed, 22 Jul 2009 19:15:56 GMT
    ETag: "34aa387-d-1568eb00"
    Accept-Ranges: bytes
    Content-Length: 14
    Vary: Accept-Encoding
    Content-Type: text/plain

    Hello World!
```

## 2.2.  Intermediaries

A more complicated situation occurs when one or more intermediaries are
present in the request/response chain. There are three common forms of
intermediary: proxy, gateway, and tunnel. In some cases, a single
intermediary may act as an origin server, proxy, gateway, or tunnel,
switching behavior based on the nature of each request.

```
            request chain --------------------------------------->
        UA -----v----- A -----v----- B -----v----- C -----v----- O
            <--------------------------------- response chain
```
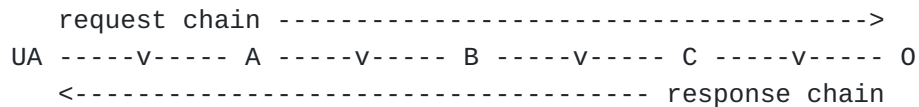
The figure above shows three intermediaries (A, B, and C) between the
user agent and origin server. A request or response message that
travels the whole chain will pass through four separate connections.
Some HTTP communication options may apply only to the connection with
the nearest, non-tunnel neighbor, only to the end-points of the chain,
or to all connections along the chain. Although the diagram is linear,
each participant may be engaged in multiple, simultaneous
communications. For example, B may be receiving requests from many
clients other than A, and/or forwarding requests to servers other than
C, at the same time that it is handling A's request.
We use the terms "upstream" and "downstream" to describe various
requirements in relation to the directional flow of a message: all
messages flow from upstream to downstream. Likewise, we use the terms
"inbound" and "outbound" to refer to directions in relation to the
request path: "inbound" means toward the origin server and "outbound"
means toward the user agent.
 A proxy is a message forwarding agent that is selected by the client,
usually via local configuration rules, to receive requests for some
type(s) of absolute URI and attempt to satisfy those requests via
translation through the HTTP interface. Some translations are minimal,
such as for proxy requests for "http" URIs, whereas other requests may
require translation to and from entirely different application-layer
protocols. Proxies are often used to group an organization's HTTP
requests through a common intermediary for the sake of security,
annotation services, or shared caching.
 A gateway (a.k.a., reverse proxy) is a receiving agent that acts as a
layer above some other server(s) and translates the received requests
to the underlying server's protocol. Gateways are often used for load
balancing or partitioning HTTP services across multiple machines.
Unlike a proxy, a gateway receives requests as if it were the origin
server for the requested resource; the requesting client will not be
aware that it is communicating with a gateway. A gateway communicates
with the client as if the gateway is the origin server and thus is
subject to all of the requirements on origin servers for that
connection. A gateway communicates with inbound servers using any
protocol it desires, including private extensions to HTTP that are
outside the scope of this specification.

A tunnel acts as a blind relay between two connections without changing the messages. Once active, a tunnel is not considered a party to the HTTP communication, though the tunnel may have been initiated by an HTTP request. A tunnel ceases to exist when both ends of the relayed connection are closed. Tunnels are used to extend a virtual connection through an intermediary, such as when transport-layer security is used to establish private communication through a shared firewall proxy.

---

## 2.3.  Caches

Any party to HTTP communication that is not acting as a tunnel may employ an internal cache for handling requests. A cache is a local store of previous response messages and the subsystem that controls its message storage, retrieval, and deletion. A cache stores cacheable responses in order to reduce the response time and network bandwidth consumption on future, equivalent requests. Any client or server may include a cache, though a cache cannot be used by a server while it is acting as a tunnel.

The effect of a cache is that the request/response chain is shortened if one of the participants along the chain has a cached response applicable to that request. The following illustrates the resulting chain if B has a cached copy of an earlier response from O (via C) for a request which has not been cached by UA or A.

```
            request chain ---------->
       UA -----v----- A -----v----- B - - - - - - C - - - - - - O
            <--------- response chain
```

A response is cacheable if a cache is allowed to store a copy of the response message for use in answering subsequent requests. Even when a response is cacheable, there may be additional constraints placed by the client or by the origin server on when that cached response can be used for a particular request. HTTP requirements for cache behavior and cacheable responses are defined in Section 2 of [Part6] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1, part 6: Caching," July 2010.).

There are a wide variety of architectures and configurations of caches and proxies deployed across the World Wide Web and inside large organizations. These systems include national hierarchies of proxy caches to save transoceanic bandwidth, systems that broadcast or multicast cache entries, organizations that distribute subsets of cached data via optical media, and so on.

---

## 2.4.  Transport Independence

HTTP systems are used in a wide variety of environments, from corporate intranets with high-bandwidth links to long-distance communication over low-power radio links and intermittent connectivity.
HTTP communication usually takes place over TCP/IP connections. The default port is TCP 80 (http://www.iana.org/assignments/port-numbers), but other ports can be used. This does not preclude HTTP from being implemented on top of any other protocol on the Internet, or on other networks. HTTP only presumes a reliable transport; any protocol that provides such guarantees can be used; the mapping of the HTTP/1.1 request and response structures onto the transport data units of the protocol in question is outside the scope of this specification.
In HTTP/1.0, most implementations used a new connection for each request/response exchange. In HTTP/1.1, a connection may be used for one or more request/response exchanges, although connections may be closed for a variety of reasons (see Section 7.1 (Persistent Connections)).

---

## 2.5.  HTTP Version                                    TOC

HTTP uses a "<major>.<minor>" numbering scheme to indicate versions of the protocol. The protocol versioning policy is intended to allow the sender to indicate the format of a message and its capacity for understanding further HTTP communication, rather than the features obtained via that communication. No change is made to the version number for the addition of message components which do not affect communication behavior or which only add to extensible field values.
The <minor> number is incremented when the changes made to the protocol add features which do not change the general message parsing algorithm, but which may add to the message semantics and imply additional capabilities of the sender. The <major> number is incremented when the format of a message within the protocol is changed. See [RFC2145] (Mogul, J., Fielding, R., Gettys, J., and H. Nielsen, "Use and Interpretation of HTTP Version Numbers," May 1997.) for a fuller explanation.
The version of an HTTP message is indicated by an HTTP-Version field in the first line of the message. HTTP-Version is case-sensitive.

```
     HTTP-Version   = HTTP-Prot-Name "/" 1*DIGIT "." 1*DIGIT
     HTTP-Prot-Name = %x48.54.54.50 ; "HTTP", case-sensitive
```

Note that the major and minor numbers MUST be treated as separate integers and that each MAY be incremented higher than a single digit. Thus, HTTP/2.4 is a lower version than HTTP/2.13, which in turn is lower than HTTP/12.3. Leading zeros MUST be ignored by recipients and MUST NOT be sent.

An application that sends a request or response message that includes HTTP-Version of "HTTP/1.1" MUST be at least conditionally compliant with this specification. Applications that are at least conditionally compliant with this specification SHOULD use an HTTP-Version of "HTTP/1.1" in their messages, and MUST do so for any message that is not compatible with HTTP/1.0. For more details on when to send specific HTTP-Version values, see [RFC2145] (Mogul, J., Fielding, R., Gettys, J., and H. Nielsen, "Use and Interpretation of HTTP Version Numbers," May 1997.).

The HTTP version of an application is the highest HTTP version for which the application is at least conditionally compliant.

Proxy and gateway applications need to be careful when forwarding messages in protocol versions different from that of the application. Since the protocol version indicates the protocol capability of the sender, a proxy/gateway MUST NOT send a message with a version indicator which is greater than its actual version. If a higher version request is received, the proxy/gateway MUST either downgrade the request version, or respond with an error, or switch to tunnel behavior.

Due to interoperability problems with HTTP/1.0 proxies discovered since the publication of [RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.), caching proxies MUST, gateways MAY, and tunnels MUST NOT upgrade the request to the highest version they support. The proxy/gateway's response to that request MUST be in the same major version as the request.

> Note: Converting between versions of HTTP may involve modification of header fields required or forbidden by the versions involved.

## 2.6.  Uniform Resource Identifiers                           TOC

Uniform Resource Identifiers (URIs) [RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.) are used throughout HTTP as the means for identifying resources. URI references are used to target requests, indicate redirects, and define relationships. HTTP does not limit what a resource may be; it merely defines an interface that can be used to interact with a resource via HTTP. More information on the scope of URIs and resources can be found in [RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.).

This specification adopts the definitions of "URI-reference", "absolute-URI", "relative-part", "port", "host", "path-abempty", "path-absolute", "query", and "authority" from [RFC3986] (Berners-Lee, T.,

[Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.)](). In addition, we define a partial-URI rule for protocol elements that allow a relative URI without a fragment.

```
URI-reference = <URI-reference, defined in [RFC3986], Section 4.1>
absolute-URI  = <absolute-URI, defined in [RFC3986], Section 4.3>
relative-part = <relative-part, defined in [RFC3986], Section 4.2>
authority     = <authority, defined in [RFC3986], Section 3.2>
path-abempty  = <path-abempty, defined in [RFC3986], Section 3.3>
path-absolute = <path-absolute, defined in [RFC3986], Section 3.3>
port          = <port, defined in [RFC3986], Section 3.2.3>
query         = <query, defined in [RFC3986], Section 3.4>
uri-host      = <host, defined in [RFC3986], Section 3.2.2>

partial-URI   = relative-part [ "?" query ]
```

Each protocol element in HTTP that allows a URI reference will indicate in its ABNF production whether the element allows only a URI in absolute form (absolute-URI), any relative reference (relative-ref), or some other subset of the URI-reference grammar. Unless otherwise indicated, URI references are parsed relative to the request target (the default base URI for both the request and its corresponding response).

---

### 2.6.1.  http URI scheme

The "http" URI scheme is hereby defined for the purpose of minting identifiers according to their association with the hierarchical namespace governed by a potential HTTP origin server listening for TCP connections on a given port. The HTTP server is identified via the generic syntax's [authority (Uniform Resource Identifiers)]() component, which includes a host identifier and optional TCP port, and the remainder of the URI is considered to be identifying data corresponding to a resource for which that server might provide an HTTP interface.

```
http-URI = "http:" "//" authority path-abempty [ "?" query ]
```

The host identifier within an [authority (Uniform Resource Identifiers)]() component is defined in [[RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.)](), Section 3.2.2. If host is provided as an IP literal or IPv4 address, then the HTTP server is any listener on the indicated TCP port at that IP address. If host is a registered name, then that name is considered an indirect identifier and the recipient might use a name resolution service, such as DNS, to find the address of a listener for that host. The host MUST NOT be empty; if an "http" URI is received

with an empty host, then it MUST be rejected as invalid. If the port
subcomponent is empty or not given, then TCP port 80 is assumed (the
default reserved port for WWW services).

Regardless of the form of host identifier, access to that host is not
implied by the mere presence of its name or address. The host may or
may not exist and, even when it does exist, may or may not be running
an HTTP server or listening to the indicated port. The "http" URI
scheme makes use of the delegated nature of Internet names and
addresses to establish a naming authority (whatever entity has the
ability to place an HTTP server at that Internet name or address) and
allows that authority to determine which names are valid and how they
might be used.

When an "http" URI is used within a context that calls for access to
the indicated resource, a client MAY attempt access by resolving the
host to an IP address, establishing a TCP connection to that address on
the indicated port, and sending an HTTP request message to the server
containing the URI's identifying data as described in [Section 4
(Request)](). If the server responds to that request with a non-interim
HTTP response message, as described in [Section 5 (Response)](), then that
response is considered an authoritative answer to the client's request.

Although HTTP is independent of the transport protocol, the "http"
scheme is specific to TCP-based services because the name delegation
process depends on TCP for establishing authority. An HTTP service
based on some other underlying connection protocol would presumably be
identified using a different URI scheme, just as the "https" scheme
(below) is used for servers that require an SSL/TLS transport layer on
a connection. Other protocols may also be used to provide access to
"http" identified resources --- it is only the authoritative interface
used for mapping the namespace that is specific to TCP.

---

### 2.6.2.  https URI scheme

The "https" URI scheme is hereby defined for the purpose of minting
identifiers according to their association with the hierarchical
namespace governed by a potential HTTP origin server listening for SSL/
TLS-secured connections on a given TCP port. The host and port are
determined in the same way as for the "http" scheme, except that a
default TCP port of 443 is assumed if the port subcomponent is empty or
not given.

```
     https-URI = "https:" "//" authority path-abempty [ "?" query ]
```

The primary difference between the "http" and "https" schemes is that
interaction with the latter is required to be secured for privacy
through the use of strong encryption. The URI cannot be sent in a
request until the connection is secure. Likewise, the default for

caching is that each response that would be considered "public" under the "http" scheme is instead treated as "private" and thus not eligible for shared caching.

The process for authoritative access to an "https" identified resource is defined in [RFC2818] (Rescorla, E., "HTTP Over TLS," May 2000.).

---

### 2.6.3.  http and https URI Normalization and Comparison

Since the "http" and "https" schemes conform to the URI generic syntax, such URIs are normalized and compared according to the algorithm defined in [RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.), Section 6, using the defaults described above for each scheme.

If the port is equal to the default port for a scheme, the normal form is to elide the port subcomponent. Likewise, an empty path component is equivalent to an absolute path of "/", so the normal form is to provide a path of "/" instead. The scheme and host are case-insensitive and normally provided in lowercase; all other components are compared in a case-sensitive manner. Characters other than those in the "reserved" set are equivalent to their percent-encoded octets (see [RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," January 2005.), Section 2.1): the normal form is to not encode them.

For example, the following three URIs are equivalent:

        http://example.com:80/~smith/home.html
        http://EXAMPLE.com/%7Esmith/home.html
        http://EXAMPLE.com:/%7esmith/home.html

[TODO-not-here] (roy: This paragraph does not belong here.) If path-abempty is the empty string (i.e., there is no slash "/" path separator following the authority), then the "http" URI MUST be given as "/" when used as a request-target (Section 4.1.2 (request-target)). If a proxy receives a host name which is not a fully qualified domain name, it MAY add its domain to the host name it received. If a proxy receives a fully qualified domain name, the proxy MUST NOT change the host name.

---

### 3.  HTTP Message

All HTTP/1.1 messages consist of a start-line followed by a sequence of characters in a format similar to the Internet Message Format [RFC5322] (Resnick, P., "Internet Message Format," October 2008.): zero or more header fields (collectively referred to as the "headers" or the "header

section"), an empty line indicating the end of the header section, and an optional message-body.

An HTTP message can either be a request from client to server or a response from server to client. Syntactically, the two types of message differ only in the start-line, which is either a Request-Line (for requests) or a Status-Line (for responses), and in the algorithm for determining the length of the message-body (Section 3.4 (Message Length)). In theory, a client could receive requests and a server could receive responses, distinguishing them by their different start-line formats, but in practice servers are implemented to only expect a request (a response is interpreted as an unknown or invalid request method) and clients are implemented to only expect a response.

```
    HTTP-message    = start-line
                       *( header-field CRLF )
                       CRLF
                       [ message-body ]
    start-line      = Request-Line / Status-Line
```

Whitespace (WSP) MUST NOT be sent between the start-line and the first header field. The presence of whitespace might be an attempt to trick a noncompliant implementation of HTTP into ignoring that field or processing the next line as a new request, either of which may result in security issues when implementations within the request chain interpret the same message differently. HTTP/1.1 servers MUST reject such a message with a 400 (Bad Request) response.

---

### 3.1.  Message Parsing Robustness

In the interest of robustness, servers SHOULD ignore at least one empty line received where a Request-Line is expected. In other words, if the server is reading the protocol stream at the beginning of a message and receives a CRLF first, it should ignore the CRLF.

Some old HTTP/1.0 client implementations generate an extra CRLF after a POST request as a lame workaround for some early server applications that failed to read message-body content that was not terminated by a line-ending. An HTTP/1.1 client MUST NOT preface or follow a request with an extra CRLF. If terminating the request message-body with a line-ending is desired, then the client MUST include the terminating CRLF octets as part of the message-body length.

The normal procedure for parsing an HTTP message is to read the start-line into a structure, read each header field into a hash table by field name until the empty line, and then use the parsed data to determine if a message-body is expected. If a message-body has been indicated, then it is read as a stream until an amount of OCTETs equal to the message-length is read or the connection is closed. Care must be

taken to parse an HTTP message as a sequence of OCTETs in an encoding that is a superset of US-ASCII. Attempting to parse HTTP as a stream of Unicode characters in a character encoding like UTF-16 may introduce security flaws due to the differing ways that such parsers interpret invalid characters.

---

### 3.2.  Header Fields

Each HTTP header field consists of a case-insensitive field name followed by a colon (":"), optional whitespace, and the field value.

```
header-field   = field-name ":" OWS [ field-value ] OWS
field-name     = token
field-value    = *( field-content / OWS )
field-content  = *( WSP / VCHAR / obs-text )
```

No whitespace is allowed between the header field name and colon. For security reasons, any request message received containing such whitespace MUST be rejected with a response code of 400 (Bad Request). A proxy MUST remove any such whitespace from a response message before forwarding the message downstream.
A field value MAY be preceded by optional whitespace (OWS); a single SP is preferred. The field value does not include any leading or trailing white space: OWS occurring before the first non-whitespace character of the field value or after the last non-whitespace character of the field value is ignored and SHOULD be removed before further processing (as this does not change the meaning of the header field).
The order in which header fields with differing field names are received is not significant. However, it is "good practice" to send header fields that contain control data first, such as Host on requests and Date on responses, so that implementations can decide when not to handle a message as early as possible. A server MUST wait until the entire header section is received before interpreting a request message, since later header fields might include conditionals, authentication credentials, or deliberately misleading duplicate header fields that would impact request processing.
Multiple header fields with the same field name MUST NOT be sent in a message unless the entire field value for that header field is defined as a comma-separated list [i.e., #(values)]. Multiple header fields with the same field name can be combined into one "field-name: field-value" pair, without changing the semantics of the message, by appending each subsequent field value to the combined field value in order, separated by a comma. The order in which header fields with the same field name are received is therefore significant to the interpretation of the combined field value; a proxy MUST NOT change the order of these field values when forwarding a message.

Note: The "Set-Cookie" header as implemented in practice (as opposed to how it is specified in [RFC2109] (Kristol, D. and L. Montulli, "HTTP State Management Mechanism," February 1997.)) can occur multiple times, but does not use the list syntax, and thus cannot be combined into a single line. (See Appendix A.2.3 of [Kri2001] (Kristol, D., "HTTP Cookies: Standards, Privacy, and Politics," November 2001.) for details.) Also note that the Set-Cookie2 header specified in [RFC2965] (Kristol, D. and L. Montulli, "HTTP State Management Mechanism," October 2000.) does not share this problem.

Historically, HTTP header field values could be extended over multiple lines by preceding each extra line with at least one space or horizontal tab character (line folding). This specification deprecates such line folding except within the message/http media type (Section 10.3.1 (Internet Media Type message/http)). HTTP/1.1 senders MUST NOT produce messages that include line folding (i.e., that contain any field-content that matches the obs-fold rule) unless the message is intended for packaging within the message/http media type. HTTP/1.1 recipients SHOULD accept line folding and replace any embedded obs-fold whitespace with a single SP prior to interpreting the field value or forwarding the message downstream.

Historically, HTTP has allowed field content with text in the ISO-8859-1 [ISO-8859-1] (International Organization for Standardization, "Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1," 1998.) character encoding and supported other character sets only through use of [RFC2047] (Moore, K., "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text," November 1996.) encoding. In practice, most HTTP header field values use only a subset of the US-ASCII character encoding [USASCII] (American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange," 1986.). Newly defined header fields SHOULD limit their field values to US-ASCII characters. Recipients SHOULD treat other (obs-text) octets in field content as opaque data.

Comments can be included in some HTTP header fields by surrounding the comment text with parentheses. Comments are only allowed in fields containing "comment" as part of their field value definition.

```
comment         = "(" *( ctext / quoted-cpair / comment ) ")"
ctext           = OWS / %x21-27 / %x2A-5B / %x5D-7E / obs-text
                ; OWS / <VCHAR except "(", ")", and "\"> / obs-text
```

The backslash character ("\") can be used as a single-character quoting mechanism within comment constructs:

```
quoted-cpair    = "\" ( WSP / VCHAR / obs-text )
```

Producers SHOULD NOT escape characters that do not require escaping
(i.e., other than the backslash character "\" and the parentheses "("
and ")").

---

### 3.3.  Message Body

The message-body (if any) of an HTTP message is used to carry the
entity-body associated with the request or response. The message-body
differs from the entity-body only when a transfer-coding has been
applied, as indicated by the Transfer-Encoding header field
(Section 9.7 (Transfer-Encoding)).

```
      message-body = entity-body
                   / <entity-body encoded as per Transfer-Encoding>
```

Transfer-Encoding MUST be used to indicate any transfer-codings applied
by an application to ensure safe and proper transfer of the message.
Transfer-Encoding is a property of the message, not of the entity, and
thus MAY be added or removed by any application along the request/
response chain. (However, Section 6.2 (Transfer Codings) places
restrictions on when certain transfer-codings may be used.)
The rules for when a message-body is allowed in a message differ for
requests and responses.
The presence of a message-body in a request is signaled by the
inclusion of a Content-Length or Transfer-Encoding header field in the
request's header fields. When a request message contains both a
message-body of non-zero length and a method that does not define any
semantics for that request message-body, then an origin server SHOULD
either ignore the message-body or respond with an appropriate error
message (e.g., 413). A proxy or gateway, when presented the same
request, SHOULD either forward the request inbound with the message-
body or ignore the message-body when determining a response.
For response messages, whether or not a message-body is included with a
message is dependent on both the request method and the response status
code (Section 5.1.1 (Status Code and Reason Phrase)). All responses to
the HEAD request method MUST NOT include a message-body, even though
the presence of entity-header fields might lead one to believe they do.
All 1xx (Informational), 204 (No Content), and 304 (Not Modified)
responses MUST NOT include a message-body. All other responses do
include a message-body, although it MAY be of zero length.

---

### 3.4.  Message Length

The transfer-length of a message is the length of the message-body as it appears in the message; that is, after any transfer-codings have been applied. When a message-body is included with a message, the transfer-length of that body is determined by one of the following (in order of precedence):

1. Any response message which "MUST NOT" include a message-body (such as the 1xx, 204, and 304 responses and any response to a HEAD request) is always terminated by the first empty line after the header fields, regardless of the entity-header fields present in the message.

2. If a Transfer-Encoding header field ([Section 9.7 (Transfer-Encoding)](#)) is present and the "chunked" transfer-coding ([Section 6.2 (Transfer Codings)](#)) is used, the transfer-length is defined by the use of this transfer-coding. If a Transfer-Encoding header field is present and the "chunked" transfer-coding is not present, the transfer-length is defined by the sender closing the connection.

3. If a Content-Length header field ([Section 9.2 (Content-Length)](#)) is present, its value in OCTETs represents both the entity-length and the transfer-length. The Content-Length header field MUST NOT be sent if these two lengths are different (i.e., if a Transfer-Encoding header field is present). If a message is received with both a Transfer-Encoding header field and a Content-Length header field, the latter MUST be ignored.

4. If the message uses the media type "multipart/byteranges", and the transfer-length is not otherwise specified, then this self-delimiting media type defines the transfer-length. This media type MUST NOT be used unless the sender knows that the recipient can parse it; the presence in a request of a Range header with multiple byte-range specifiers from a HTTP/1.1 client implies that the client can parse multipart/byteranges responses.

   > A range header might be forwarded by a HTTP/1.0 proxy that does not understand multipart/byteranges; in this case the server MUST delimit the message using methods defined in items 1, 3 or 5 of this section.

5. By the server closing the connection. (Closing the connection cannot be used to indicate the end of a request body, since that would leave no possibility for the server to send back a response.)

For compatibility with HTTP/1.0 applications, HTTP/1.1 requests
containing a message-body MUST include a valid Content-Length header
field unless the server is known to be HTTP/1.1 compliant. If a request
contains a message-body and a Content-Length is not given, the server
SHOULD respond with 400 (Bad Request) if it cannot determine the length
of the message, or with 411 (Length Required) if it wishes to insist on
receiving a valid Content-Length.

All HTTP/1.1 applications that receive entities MUST accept the
"chunked" transfer-coding ([Section 6.2 (Transfer Codings)](#)), thus
allowing this mechanism to be used for messages when the message length
cannot be determined in advance.

Messages MUST NOT include both a Content-Length header field and a
transfer-coding. If the message does include a transfer-coding, the
Content-Length MUST be ignored.

When a Content-Length is given in a message where a message-body is
allowed, its field value MUST exactly match the number of OCTETs in the
message-body. HTTP/1.1 user agents MUST notify the user when an invalid
length is received and detected.

---

### 3.5.  General Header Fields

There are a few header fields which have general applicability for both
request and response messages, but which do not apply to the entity
being transferred. These header fields apply only to the message being
transmitted.

```
general-header = Cache-Control          ; [Part6], Section 3.2
               / Connection             ; Section 9.1
               / Date                   ; Section 9.3
               / Pragma                 ; [Part6], Section 3.4
               / Trailer                ; Section 9.6
               / Transfer-Encoding      ; Section 9.7
               / Upgrade                ; Section 9.8
               / Via                    ; Section 9.9
               / Warning                ; [Part6], Section 3.6
```

General-header field names can be extended reliably only in combination
with a change in the protocol version. However, new or experimental
header fields may be given the semantics of general header fields if
all parties in the communication recognize them to be general-header
fields. Unrecognized header fields are treated as entity-header fields.

---

## 4.  Request

A request message from a client to a server includes, within the first
line of that message, the method to be applied to the resource, the
identifier of the resource, and the protocol version in use.

```
Request         = Request-Line              ; Section 4.1
                  *(( general-header        ; Section 3.5
                   / request-header         ; [Part2], Section 3
                   / entity-header ) CRLF ) ; [Part3], Section 3.1
                  CRLF
                  [ message-body ]          ; Section 3.3
```

---

### 4.1.  Request-Line

The Request-Line begins with a method token, followed by the request-
target and the protocol version, and ending with CRLF. The elements are
separated by SP characters. No CR or LF is allowed except in the final
CRLF sequence.

```
Request-Line  = Method SP request-target SP HTTP-Version CRLF
```

---

### 4.1.1.  Method

The Method token indicates the method to be performed on the resource
identified by the request-target. The method is case-sensitive.

```
Method        = token
```

---

### 4.1.2.  request-target

The request-target identifies the resource upon which to apply the
request.

```
request-target = "*"
               / absolute-URI
               / ( path-absolute [ "?" query ] )
               / authority
```

The four options for request-target are dependent on the nature of the request.
The asterisk "*" means that the request does not apply to a particular resource, but to the server itself, and is only allowed when the method used does not necessarily apply to a resource. One example would be

        OPTIONS * HTTP/1.1

The absolute-URI form is REQUIRED when the request is being made to a proxy. The proxy is requested to forward the request or service it from a valid cache, and return the response. Note that the proxy MAY forward the request on to another proxy or directly to the server specified by the absolute-URI. In order to avoid request loops, a proxy MUST be able to recognize all of its server names, including any aliases, local variations, and the numeric IP address. An example Request-Line would be:

        GET http://www.example.org/pub/WWW/TheProject.html HTTP/1.1

To allow for transition to absolute-URIs in all requests in future versions of HTTP, all HTTP/1.1 servers MUST accept the absolute-URI form in requests, even though HTTP/1.1 clients will only generate them in requests to proxies.
The authority form is only used by the CONNECT method (Section 7.9 of [Part2] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics," July 2010.)).
The most common form of request-target is that used to identify a resource on an origin server or gateway. In this case the absolute path of the URI MUST be transmitted (see Section 2.6.1 (http URI scheme), path-absolute) as the request-target, and the network location of the URI (authority) MUST be transmitted in a Host header field. For example, a client wishing to retrieve the resource above directly from the origin server would create a TCP connection to port 80 of the host "www.example.org" and send the lines:

        GET /pub/WWW/TheProject.html HTTP/1.1
        Host: www.example.org

followed by the remainder of the Request. Note that the absolute path cannot be empty; if none is present in the original URI, it MUST be given as "/" (the server root).
If a proxy receives a request without any path in the request-target and the method specified is capable of supporting the asterisk form of request-target, then the last proxy on the request chain MUST forward the request with "*" as the final request-target.
For example, the request

        OPTIONS http://www.example.org:8001 HTTP/1.1

would be forwarded by the proxy as

        OPTIONS * HTTP/1.1
        Host: www.example.org:8001

after connecting to port 8001 of host "www.example.org".
The request-target is transmitted in the format specified in
Section 2.6.1 (http URI scheme). If the request-target is percent-
encoded ([RFC3986] (Berners-Lee, T., Fielding, R., and L. Masinter,
"Uniform Resource Identifier (URI): Generic Syntax," January 2005.),
Section 2.1), the origin server MUST decode the request-target in order
to properly interpret the request. Servers SHOULD respond to invalid
request-targets with an appropriate status code.
A transparent proxy MUST NOT rewrite the "path-absolute" part of the
received request-target when forwarding it to the next inbound server,
except as noted above to replace a null path-absolute with "/" or "*".

> Note: The "no rewrite" rule prevents the proxy from changing the
> meaning of the request when the origin server is improperly using a
> non-reserved URI character for a reserved purpose. Implementors
> should be aware that some pre-HTTP/1.1 proxies have been known to
> rewrite the request-target.

HTTP does not place a pre-defined limit on the length of a request-
target. A server MUST be prepared to receive URIs of unbounded length
and respond with the 414 (URI Too Long) status if the received request-
target would be longer than the server wishes to handle (see Section
8.4.15 of [Part2] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk,
H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J.
Reschke, Ed., "HTTP/1.1, part 2: Message Semantics," July 2010.)).
Various ad-hoc limitations on request-target length are found in
practice. It is RECOMMENDED that all HTTP senders and recipients
support request-target lengths of 8000 or more OCTETs.

> Note: Fragments ([RFC3986] (Berners-Lee, T., Fielding, R., and L.
> Masinter, "Uniform Resource Identifier (URI): Generic Syntax,"
> January 2005.), Section 3.5) are not part of the request-target and
> thus will not be transmitted in an HTTP request.

---

### 4.2.  The Resource Identified by a Request                         TOC

The exact resource identified by an Internet request is determined by
examining both the request-target and the Host header field.
An origin server that does not allow resources to differ by the
requested host MAY ignore the Host header field value when determining
the resource identified by an HTTP/1.1 request. (But see Appendix B.1.1

for other requirements on Host support in HTTP/1.1.)
An origin server that does differentiate resources based on the host
requested (sometimes referred to as virtual hosts or vanity host names)
MUST use the following rules for determining the requested resource on
an HTTP/1.1 request:

1. If request-target is an absolute-URI, the host is part of the
   request-target. Any Host header field value in the request MUST
   be ignored.

2. If the request-target is not an absolute-URI, and the request
   includes a Host header field, the host is determined by the
   Host header field value.

3. If the host as determined by rule 1 or 2 is not a valid host on
   the server, the response MUST be a 400 (Bad Request) error
   message.

Recipients of an HTTP/1.0 request that lacks a Host header field MAY
attempt to use heuristics (e.g., examination of the URI path for
something unique to a particular host) in order to determine what exact
resource is being requested.

---

### 4.3.  Effective Request URI

HTTP requests often do not carry the absolute URI ([RFC3986] (Berners-
Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier
(URI): Generic Syntax," January 2005.), Section 4.3) for the resource
they are intended for; instead, the value needs to be inferred from the
request-target, Host header and other context. The result of this
process is the "Effective Request URI".
If the request-target is an absolute-URI, then the Effective Request
URI is the request-target.
If the request-target uses the path-absolute (plus optional query)
syntax or if it is just the asterisk "*", then the Effective Request
URI is constructed by concatenating

   *the scheme name: "http" if the request was received over an
    insecure TCP connection, or "https" when received over SSL/TLS-
    secured TCP connection,

   *the character sequence "://",

   *the authority component, as specified in the Host header
    ([Section 9.4 (Host)](#)) and determined by the rules in [Section 4.2
    (The Resource Identified by a Request)](#), [effrequri-nohost] (jre:

> [Do we need to include the handling of missing hosts in HTTP/1.0 messages, as described in Section 4.2?)](#) and

> *the request-target obtained from the Request-Line, unless the request-target is just the asterisk "*".

Otherwise, when request-target uses the authority form, the Effective Request URI is undefined.
Example 1: the Effective Request URI for the message

```
GET /pub/WWW/TheProject.html HTTP/1.1
Host: www.example.org:8080
```

(received over an insecure TCP connection) is "http", plus "://", plus the authority component "www.example.org:8080", plus the request-target "/pub/WWW/TheProject.html", thus "http://www.example.org:8080/pub/WWW/TheProject.html".
Example 2: the Effective Request URI for the message

```
GET * HTTP/1.1
Host: www.example.org
```

(received over an SSL/TLS secured TCP connection) is "https", plus "://", plus the authority component "www.example.org", thus "https://www.example.org".
Effective Request URIs are compared using the rules described in [Section 2.6.3 (http and https URI Normalization and Comparison)](#), except that empty path components MUST NOT be treated as equivalent to an absolute path of "/".

---

## 5.  Response

After receiving and interpreting a request message, a server responds with an HTTP response message.

```
Response        = Status-Line                ; Section 5.1
                  *(( general-header         ; Section 3.5
                   / response-header         ; [Part2], Section 5
                   / entity-header ) CRLF )  ; [Part3], Section 3.1
                  CRLF
                  [ message-body ]           ; Section 3.3
```

---

## 5.1.  Status-Line

The first line of a Response message is the Status-Line, consisting of
the protocol version followed by a numeric status code and its
associated textual phrase, with each element separated by SP
characters. No CR or LF is allowed except in the final CRLF sequence.

    Status-Line = HTTP-Version SP Status-Code SP Reason-Phrase CRLF

---

## 5.1.1.  Status Code and Reason Phrase

The Status-Code element is a 3-digit integer result code of the attempt
to understand and satisfy the request. These codes are fully defined in
Section 8 of [Part2] (Fielding, R., Ed., Gettys, J., Mogul, J.,
Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed.,
and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics,"
July 2010.). The Reason Phrase exists for the sole purpose of providing
a textual description associated with the numeric status code, out of
deference to earlier Internet application protocols that were more
frequently used with interactive text clients. A client SHOULD ignore
the content of the Reason Phrase.
The first digit of the Status-Code defines the class of response. The
last two digits do not have any categorization role. There are 5 values
for the first digit:

    *1xx: Informational - Request received, continuing process

    *2xx: Success - The action was successfully received, understood,
     and accepted

    *3xx: Redirection - Further action must be taken in order to
     complete the request

    *4xx: Client Error - The request contains bad syntax or cannot be
     fulfilled

    *5xx: Server Error - The server failed to fulfill an apparently
     valid request


      Status-Code    = 3DIGIT
      Reason-Phrase  = *( WSP / VCHAR / obs-text )

## 6.  Protocol Parameters

---

### 6.1.  Date/Time Formats: Full Date

HTTP applications have historically allowed three different formats for the representation of date/time stamps.
The first format is preferred as an Internet standard and represents a fixed-length subset of that defined by [RFC1123] (Braden, R., "Requirements for Internet Hosts - Application and Support," October 1989.):

        Sun, 06 Nov 1994 08:49:37 GMT  ; RFC 1123

The other formats are described here only for compatibility with obsolete implementations.

        Sunday, 06-Nov-94 08:49:37 GMT ; obsolete RFC 850 format
        Sun Nov  6 08:49:37 1994       ; ANSI C's asctime() format

HTTP/1.1 clients and servers that parse the date value MUST accept all three formats (for compatibility with HTTP/1.0), though they MUST only generate the RFC 1123 format for representing HTTP-date values in header fields. See Appendix A (Tolerant Applications) for further information.
All HTTP date/time stamps MUST be represented in Greenwich Mean Time (GMT), without exception. For the purposes of HTTP, GMT is exactly equal to UTC (Coordinated Universal Time). This is indicated in the first two formats by the inclusion of "GMT" as the three-letter abbreviation for time zone, and MUST be assumed when reading the asctime format. HTTP-date is case sensitive and MUST NOT include additional whitespace beyond that specifically included as SP in the grammar.

        HTTP-date    = rfc1123-date / obs-date

Preferred format:

```
rfc1123-date = day-name "," SP date1 SP time-of-day SP GMT

day-name     = %x4D.6F.6E ; "Mon", case-sensitive
             / %x54.75.65 ; "Tue", case-sensitive
             / %x57.65.64 ; "Wed", case-sensitive
             / %x54.68.75 ; "Thu", case-sensitive
             / %x46.72.69 ; "Fri", case-sensitive
             / %x53.61.74 ; "Sat", case-sensitive
             / %x53.75.6E ; "Sun", case-sensitive

date1        = day SP month SP year
             ; e.g., 02 Jun 1982

day          = 2DIGIT
month        = %x4A.61.6E ; "Jan", case-sensitive
             / %x46.65.62 ; "Feb", case-sensitive
             / %x4D.61.72 ; "Mar", case-sensitive
             / %x41.70.72 ; "Apr", case-sensitive
             / %x4D.61.79 ; "May", case-sensitive
             / %x4A.75.6E ; "Jun", case-sensitive
             / %x4A.75.6C ; "Jul", case-sensitive
             / %x41.75.67 ; "Aug", case-sensitive
             / %x53.65.70 ; "Sep", case-sensitive
             / %x4F.63.74 ; "Oct", case-sensitive
             / %x4E.6F.76 ; "Nov", case-sensitive
             / %x44.65.63 ; "Dec", case-sensitive
year         = 4DIGIT

GMT   = %x47.4D.54 ; "GMT", case-sensitive

time-of-day  = hour ":" minute ":" second
                ; 00:00:00 - 23:59:59

hour         = 2DIGIT
minute       = 2DIGIT
second       = 2DIGIT
```

The semantics of day-name, day, month, year, and time-of-day are the same as those defined for the RFC 5322 constructs with the corresponding name ([RFC5322] (Resnick, P., "Internet Message Format," October 2008.), Section 3.3).
Obsolete formats:

```
obs-date     = rfc850-date / asctime-date
```

```
rfc850-date  = day-name-l "," SP date2 SP time-of-day SP GMT
date2        = day "-" month "-" 2DIGIT
               ; day-month-year (e.g., 02-Jun-82)

day-name-l   = %x4D.6F.6E.64.61.79 ; "Monday", case-sensitive
       / %x54.75.65.73.64.61.79 ; "Tuesday", case-sensitive
       / %x57.65.64.6E.65.73.64.61.79 ; "Wednesday", case-sensitive
       / %x54.68.75.72.73.64.61.79 ; "Thursday", case-sensitive
       / %x46.72.69.64.61.79 ; "Friday", case-sensitive
       / %x53.61.74.75.72.64.61.79 ; "Saturday", case-sensitive
       / %x53.75.6E.64.61.79 ; "Sunday", case-sensitive


asctime-date = day-name SP date3 SP time-of-day SP year
date3        = month SP ( 2DIGIT / ( SP 1DIGIT ))
               ; month day (e.g., Jun  2)
```

> Note: Recipients of date values are encouraged to be robust in
> accepting date values that may have been sent by non-HTTP
> applications, as is sometimes the case when retrieving or posting
> messages via proxies/gateways to SMTP or NNTP.

> Note: HTTP requirements for the date/time stamp format apply only to
> their usage within the protocol stream. Clients and servers are not
> required to use these formats for user presentation, request
> logging, etc.

## 6.2.  Transfer Codings

Transfer-coding values are used to indicate an encoding transformation
that has been, can be, or may need to be applied to an entity-body in
order to ensure "safe transport" through the network. This differs from
a content coding in that the transfer-coding is a property of the
message, not of the original entity.

```
transfer-coding          = "chunked" ; Section 6.2.1
                         / "compress" ; Section 6.2.2.1
                         / "deflate" ; Section 6.2.2.2
                         / "gzip" ; Section 6.2.2.3
                         / transfer-extension
transfer-extension       = token *( OWS ";" OWS transfer-parameter )
```

Parameters are in the form of attribute/value pairs.

```
transfer-parameter      = attribute BWS "=" BWS value
attribute               = token
value                   = word
```

All transfer-coding values are case-insensitive. HTTP/1.1 uses
transfer-coding values in the TE header field ([Section 9.5 (TE)](#)) and in
the Transfer-Encoding header field ([Section 9.7 (Transfer-Encoding)](#)).
Whenever a transfer-coding is applied to a message-body, the set of
transfer-codings MUST include "chunked", unless the message indicates
it is terminated by closing the connection. When the "chunked"
transfer-coding is used, it MUST be the last transfer-coding applied to
the message-body. The "chunked" transfer-coding MUST NOT be applied
more than once to a message-body. These rules allow the recipient to
determine the transfer-length of the message ([Section 3.4 (Message
Length)](#)).
Transfer-codings are analogous to the Content-Transfer-Encoding values
of MIME, which were designed to enable safe transport of binary data
over a 7-bit transport service ([[RFC2045] (Freed, N. and N. Borenstein,
"Multipurpose Internet Mail Extensions (MIME) Part One: Format of
Internet Message Bodies," November 1996.)](#), Section 6). However, safe
transport has a different focus for an 8bit-clean transfer protocol. In
HTTP, the only unsafe characteristic of message-bodies is the
difficulty in determining the exact body length ([Section 3.4 (Message
Length)](#)), or the desire to encrypt data over a shared transport.
A server which receives an entity-body with a transfer-coding it does
not understand SHOULD return 501 (Not Implemented), and close the
connection. A server MUST NOT send transfer-codings to an HTTP/1.0
client.

---

### 6.2.1.  Chunked Transfer Coding

The chunked encoding modifies the body of a message in order to
transfer it as a series of chunks, each with its own size indicator,
followed by an OPTIONAL trailer containing entity-header fields. This
allows dynamically produced content to be transferred along with the
information necessary for the recipient to verify that it has received
the full message.

```
Chunked-Body   = *chunk
                 last-chunk
                 trailer-part
                 CRLF

chunk          = chunk-size *WSP [ chunk-ext ] CRLF
                 chunk-data CRLF
chunk-size     = 1*HEXDIG
last-chunk     = 1*("0") *WSP [ chunk-ext ] CRLF

chunk-ext      = *( ";" *WSP chunk-ext-name
                    [ "=" chunk-ext-val ] *WSP )
chunk-ext-name = token
chunk-ext-val  = token / quoted-str-nf
chunk-data     = 1*OCTET ; a sequence of chunk-size octets
trailer-part   = *( entity-header CRLF )

quoted-str-nf  = DQUOTE *( qdtext-nf / quoted-pair ) DQUOTE
                 ; like quoted-string, but disallowing line folding
qdtext-nf      = WSP / %x21 / %x23-5B / %x5D-7E / obs-text
                 ; WSP / <VCHAR except DQUOTE and "\"> / obs-text
```

The chunk-size field is a string of hex digits indicating the size of
the chunk-data in octets. The chunked encoding is ended by any chunk
whose size is zero, followed by the trailer, which is terminated by an
empty line.
The trailer allows the sender to include additional HTTP header fields
at the end of the message. The Trailer header field can be used to
indicate which header fields are included in a trailer (see Section 9.6
(Trailer)).
A server using chunked transfer-coding in a response MUST NOT use the
trailer for any header fields unless at least one of the following is
true:

    1. the request included a TE header field that indicates
       "trailers" is acceptable in the transfer-coding of the
       response, as described in Section 9.5 (TE); or,

    2. the server is the origin server for the response, the trailer
       fields consist entirely of optional metadata, and the recipient
       could use the message (in a manner acceptable to the origin
       server) without receiving this metadata. In other words, the
       origin server is willing to accept the possibility that the
       trailer fields might be silently discarded along the path to
       the client.

This requirement prevents an interoperability failure when the message
is being received by an HTTP/1.1 (or later) proxy and forwarded to an
HTTP/1.0 recipient. It avoids a situation where compliance with the

protocol would have necessitated a possibly infinite buffer on the
proxy.
A process for decoding the "chunked" transfer-coding can be represented
in pseudo-code as:

```
length := 0
read chunk-size, chunk-ext (if any) and CRLF
while (chunk-size > 0) {
    read chunk-data and CRLF
    append chunk-data to entity-body
    length := length + chunk-size
    read chunk-size and CRLF
}
read entity-header
while (entity-header not empty) {
    append entity-header to existing header fields
    read entity-header
}
Content-Length := length
Remove "chunked" from Transfer-Encoding
```

All HTTP/1.1 applications MUST be able to receive and decode the
"chunked" transfer-coding, and MUST ignore chunk-ext extensions they do
not understand.

---

### 6.2.2. Compression Codings

The codings defined below can be used to compress the payload of a
message.

> Note: Use of program names for the identification of encoding
> formats is not desirable and is discouraged for future encodings.
> Their use here is representative of historical practice, not good
> design.
>
> Note: For compatibility with previous implementations of HTTP,
> applications SHOULD consider "x-gzip" and "x-compress" to be
> equivalent to "gzip" and "compress" respectively.

---

### 6.2.2.1. Compress Coding

The "compress" format is produced by the common UNIX file compression program "compress". This format is an adaptive Lempel-Ziv-Welch coding (LZW).

---

### 6.2.2.2.  Deflate Coding

The "deflate" format is defined as the "deflate" compression mechanism (described in [RFC1951] (Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3," May 1996.)) used inside the "zlib" data format ([RFC1950] (Deutsch, L. and J-L. Gailly, "ZLIB Compressed Data Format Specification version 3.3," May 1996.)).

> Note: Some incorrect implementations send the "deflate" compressed data without the zlib wrapper.

---

### 6.2.2.3.  Gzip Coding

The "gzip" format is produced by the file compression program "gzip" (GNU zip), as described in [RFC1952] (Deutsch, P., Gailly, J-L., Adler, M., Deutsch, L., and G. Randers-Pehrson, "GZIP file format specification version 4.3," May 1996.). This format is a Lempel-Ziv coding (LZ77) with a 32 bit CRC.

---

### 6.2.3.  Transfer Coding Registry

The HTTP Transfer Coding Registry defines the name space for the transfer coding names.
Registrations MUST include the following fields:

> *Name

> *Description

> *Pointer to specification text

Names of transfer codings MUST NOT overlap with names of content codings (Section 2.2 of [Part3] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload and Content Negotiation," July 2010.)), unless the encoding transformation

is identical (as it is the case for the compression codings defined in
[Section 6.2.2 (Compression Codings)](#)).
Values to be added to this name space require expert review and a
specification (see "Expert Review" and "Specification Required" in
Section 4.1 of [[RFC5226] (Narten, T. and H. Alvestrand, "Guidelines for
Writing an IANA Considerations Section in RFCs," May 2008.)](#)), and MUST
conform to the purpose of transfer coding defined in this section.
The registry itself is maintained at [http://www.iana.org/assignments/
http-parameters](#).

---

## 6.3.  Product Tokens TOC

Product tokens are used to allow communicating applications to identify
themselves by software name and version. Most fields using product
tokens also allow sub-products which form a significant part of the
application to be listed, separated by whitespace. By convention, the
products are listed in order of their significance for identifying the
application.

        product         = token ["/" product-version]
        product-version = token

Examples:

        User-Agent: CERN-LineMode/2.15 libwww/2.17b3
        Server: Apache/0.8.4

Product tokens SHOULD be short and to the point. They MUST NOT be used
for advertising or other non-essential information. Although any token
character MAY appear in a product-version, this token SHOULD only be
used for a version identifier (i.e., successive versions of the same
product SHOULD only differ in the product-version portion of the
product value).

---

## 6.4.  Quality Values TOC

Both transfer codings (TE request header, [Section 9.5 (TE)](#)) and content
negotiation (Section 4 of [[Part3] (Fielding, R., Ed., Gettys, J.,
Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.,
Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload
and Content Negotiation," July 2010.)](#)) use short "floating point"
numbers to indicate the relative importance ("weight") of various
negotiable parameters. A weight is normalized to a real number in the
range 0 through 1, where 0 is the minimum and 1 the maximum value. If a

parameter has a quality value of 0, then content with this parameter is "not acceptable" for the client. HTTP/1.1 applications MUST NOT generate more than three digits after the decimal point. User configuration of these values SHOULD also be limited in this fashion.

```
qvalue          = ( "0" [ "." 0*3DIGIT ] )
                / ( "1" [ "." 0*3("0") ] )
```

> Note: "Quality values" is a misnomer, since these values merely represent relative degradation in desired quality.

---

## 7.  Connections

---

## 7.1.  Persistent Connections

---

### 7.1.1.  Purpose

Prior to persistent connections, a separate TCP connection was established to fetch each URL, increasing the load on HTTP servers and causing congestion on the Internet. The use of inline images and other associated data often requires a client to make multiple requests of the same server in a short amount of time. Analysis of these performance problems and results from a prototype implementation are available [Pad1995] (Padmanabhan, V. and J. Mogul, "Improving HTTP Latency," December 1995.) [Spe] (Spero, S., "Analysis of HTTP Performance Problems," .). Implementation experience and measurements of actual HTTP/1.1 implementations show good results [Nie1997] (Frystyk, H., Gettys, J., Prud'hommeaux, E., Lie, H., and C. Lilley, "Network Performance Effects of HTTP/1.1, CSS1, and PNG," September 1997.). Alternatives have also been explored, for example, T/TCP [Tou1998] (Touch, J., Heidemann, J., and K. Obraczka, "Analysis of HTTP Performance," Aug 1998.).
Persistent HTTP connections have a number of advantages:

   *By opening and closing fewer TCP connections, CPU time is saved in routers and hosts (clients, servers, proxies, gateways, tunnels, or caches), and memory used for TCP protocol control blocks can be saved in hosts.

    \*HTTP requests and responses can be pipelined on a connection.
     Pipelining allows a client to make multiple requests without
     waiting for each response, allowing a single TCP connection to be
     used much more efficiently, with much lower elapsed time.

    \*Network congestion is reduced by reducing the number of packets
     caused by TCP opens, and by allowing TCP sufficient time to
     determine the congestion state of the network.

    \*Latency on subsequent requests is reduced since there is no time
     spent in TCP's connection opening handshake.

    \*HTTP can evolve more gracefully, since errors can be reported
     without the penalty of closing the TCP connection. Clients using
     future versions of HTTP might optimistically try a new feature,
     but if communicating with an older server, retry with old
     semantics after an error is reported.

HTTP implementations SHOULD implement persistent connections.

---

### 7.1.2.  Overall Operation

A significant difference between HTTP/1.1 and earlier versions of HTTP
is that persistent connections are the default behavior of any HTTP
connection. That is, unless otherwise indicated, the client SHOULD
assume that the server will maintain a persistent connection, even
after error responses from the server.
Persistent connections provide a mechanism by which a client and a
server can signal the close of a TCP connection. This signaling takes
place using the Connection header field (Section 9.1 (Connection)).
Once a close has been signaled, the client MUST NOT send any more
requests on that connection.

---

### 7.1.2.1.  Negotiation

An HTTP/1.1 server MAY assume that a HTTP/1.1 client intends to
maintain a persistent connection unless a Connection header including
the connection-token "close" was sent in the request. If the server
chooses to close the connection immediately after sending the response,
it SHOULD send a Connection header including the connection-token
"close".
An HTTP/1.1 client MAY expect a connection to remain open, but would
decide to keep it open based on whether the response from a server
contains a Connection header with the connection-token close. In case

the client does not want to maintain a connection for more than that request, it SHOULD send a Connection header including the connection-token close.

If either the client or the server sends the close token in the Connection header, that request becomes the last one for the connection.

Clients and servers SHOULD NOT assume that a persistent connection is maintained for HTTP versions less than 1.1 unless it is explicitly signaled. See [Appendix B.2 (Compatibility with HTTP/1.0 Persistent Connections)](#) for more information on backward compatibility with HTTP/ 1.0 clients.

In order to remain persistent, all messages on the connection MUST have a self-defined message length (i.e., one not defined by closure of the connection), as described in [Section 3.4 (Message Length)](#).

---

### 7.1.2.2.  Pipelining

A client that supports persistent connections MAY "pipeline" its requests (i.e., send multiple requests without waiting for each response). A server MUST send its responses to those requests in the same order that the requests were received.

Clients which assume persistent connections and pipeline immediately after connection establishment SHOULD be prepared to retry their connection if the first pipelined attempt fails. If a client does such a retry, it MUST NOT pipeline before it knows the connection is persistent. Clients MUST also be prepared to resend their requests if the server closes the connection before sending all of the corresponding responses.

Clients SHOULD NOT pipeline requests using non-idempotent methods or non-idempotent sequences of methods (see Section 7.1.2 of [[Part2] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/ 1.1, part 2: Message Semantics," July 2010.)](#)). Otherwise, a premature termination of the transport connection could lead to indeterminate results. A client wishing to send a non-idempotent request SHOULD wait to send that request until it has received the response status for the previous request.

---

### 7.1.3.  Proxy Servers

It is especially important that proxies correctly implement the properties of the Connection header field as specified in [Section 9.1 (Connection)](#).

The proxy server MUST signal persistent connections separately with its clients and the origin servers (or other proxy servers) that it connects to. Each persistent connection applies to only one transport link.
A proxy server MUST NOT establish a HTTP/1.1 persistent connection with an HTTP/1.0 client (but see Section 19.7.1 of [RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.) for information and discussion of the problems with the Keep-Alive header implemented by many HTTP/1.0 clients).

---

### 7.1.3.1.  End-to-end and Hop-by-hop Headers

For the purpose of defining the behavior of caches and non-caching proxies, we divide HTTP headers into two categories:

   *End-to-end headers, which are transmitted to the ultimate
    recipient of a request or response. End-to-end headers in
    responses MUST be stored as part of a cache entry and MUST be
    transmitted in any response formed from a cache entry.

   *Hop-by-hop headers, which are meaningful only for a single
    transport-level connection, and are not stored by caches or
    forwarded by proxies.

The following HTTP/1.1 headers are hop-by-hop headers:

   *Connection

   *Keep-Alive

   *Proxy-Authenticate

   *Proxy-Authorization

   *TE

   *Trailer

   *Transfer-Encoding

   *Upgrade

All other headers defined by HTTP/1.1 are end-to-end headers.
Other hop-by-hop headers MUST be listed in a Connection header (Section 9.1 (Connection)).

---

**7.1.3.2.  Non-modifiable Headers**

Some features of HTTP/1.1, such as Digest Authentication, depend on the
value of certain end-to-end headers. A transparent proxy SHOULD NOT
modify an end-to-end header unless the definition of that header
requires or specifically allows that.
A transparent proxy MUST NOT modify any of the following fields in a
request or response, and it MUST NOT add any of these fields if not
already present:

    *Content-Location

    *Content-MD5

    *ETag

    *Last-Modified

A transparent proxy MUST NOT modify any of the following fields in a
response:

    *Expires

but it MAY add any of these fields if not already present. If an
Expires header is added, it MUST be given a field-value identical to
that of the Date header in that response.
A proxy MUST NOT modify or add any of the following fields in a message
that contains the no-transform cache-control directive, or in any
request:

    *Content-Encoding

    *Content-Range

    *Content-Type

A non-transparent proxy MAY modify or add these fields to a message
that does not include no-transform, but if it does so, it MUST add a
Warning 214 (Transformation applied) if one does not already appear in
the message (see Section 3.6 of [Part6] (Fielding, R., Ed., Gettys, J.,
Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.,
Lafon, Y., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1,
part 6: Caching," July 2010.)).

> Warning: Unnecessary modification of end-to-end headers might cause
> authentication failures if stronger authentication mechanisms are
> introduced in later versions of HTTP. Such authentication mechanisms
> MAY rely on the values of header fields not listed here.

The Content-Length field of a request or response is added or deleted according to the rules in [Section 3.4 (Message Length)](). A transparent proxy MUST preserve the entity-length (Section 3.2.2 of [[Part3] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/ 1.1, part 3: Message Payload and Content Negotiation," July 2010.)]()) of the entity-body, although it MAY change the transfer-length ([Section 3.4 (Message Length)]()).

---

### 7.1.4.  Practical Considerations

Servers will usually have some time-out value beyond which they will no longer maintain an inactive connection. Proxy servers might make this a higher value since it is likely that the client will be making more connections through the same server. The use of persistent connections places no requirements on the length (or existence) of this time-out for either the client or the server.
When a client or server wishes to time-out it SHOULD issue a graceful close on the transport connection. Clients and servers SHOULD both constantly watch for the other side of the transport close, and respond to it as appropriate. If a client or server does not detect the other side's close promptly it could cause unnecessary resource drain on the network.
A client, server, or proxy MAY close the transport connection at any time. For example, a client might have started to send a new request at the same time that the server has decided to close the "idle" connection. From the server's point of view, the connection is being closed while it was idle, but from the client's point of view, a request is in progress.
This means that clients, servers, and proxies MUST be able to recover from asynchronous close events. Client software SHOULD reopen the transport connection and retransmit the aborted sequence of requests without user interaction so long as the request sequence is idempotent (see Section 7.1.2 of [[Part2] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics," July 2010.)]()). Non-idempotent methods or sequences MUST NOT be automatically retried, although user agents MAY offer a human operator the choice of retrying the request(s). Confirmation by user-agent software with semantic understanding of the application MAY substitute for user confirmation. The automatic retry SHOULD NOT be repeated if the second sequence of requests fails.
Servers SHOULD always respond to at least one request per connection, if at all possible. Servers SHOULD NOT close a connection in the middle of transmitting a response, unless a network or client failure is suspected.

Clients (including proxies) SHOULD limit the number of simultaneous connections that they maintain to a given server (including proxies). Previous revisions of HTTP gave a specific number of connections as a ceiling, but this was found to be impractical for many applications. As a result, this specification does not mandate a particular maximum number of connections, but instead encourages clients to be conservative when opening multiple connections.
In particular, while using multiple connections avoids the "head-of-line blocking" problem (whereby a request that takes significant server-side processing and/or has a large payload can block subsequent requests on the same connection), each connection used consumes server resources (sometimes significantly), and furthermore using multiple connections can cause undesirable side effects in congested networks. Note that servers might reject traffic that they deem abusive, including an excessive number of connections from a client.

---

### 7.2.  Message Transmission Requirements

---

### 7.2.1.  Persistent Connections and Flow Control

HTTP/1.1 servers SHOULD maintain persistent connections and use TCP's flow control mechanisms to resolve temporary overloads, rather than terminating connections with the expectation that clients will retry. The latter technique can exacerbate network congestion.

---

### 7.2.2.  Monitoring Connections for Error Status Messages

An HTTP/1.1 (or later) client sending a message-body SHOULD monitor the network connection for an error status while it is transmitting the request. If the client sees an error status, it SHOULD immediately cease transmitting the body. If the body is being sent using a "chunked" encoding (Section 6.2 (Transfer Codings)), a zero length chunk and empty trailer MAY be used to prematurely mark the end of the message. If the body was preceded by a Content-Length header, the client MUST close the connection.

---

### 7.2.3.  Use of the 100 (Continue) Status

The purpose of the 100 (Continue) status (see Section 8.1.1 of [Part2] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics," July 2010.)) is to allow a client that is sending a request message with a request body to determine if the origin server is willing to accept the request (based on the request headers) before the client sends the request body. In some cases, it might either be inappropriate or highly inefficient for the client to send the body if the server will reject the message without looking at the body.
Requirements for HTTP/1.1 clients:

  *If a client will wait for a 100 (Continue) response before
   sending the request body, it MUST send an Expect request-header
   field (Section 9.2 of [Part2] (Fielding, R., Ed., Gettys, J.,
   Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.,
   Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message
   Semantics," July 2010.)) with the "100-continue" expectation.

  *A client MUST NOT send an Expect request-header field (Section
   9.2 of [Part2] (Fielding, R., Ed., Gettys, J., Mogul, J.,
   Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y.,
   Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics,"
   July 2010.)) with the "100-continue" expectation if it does not
   intend to send a request body.

Because of the presence of older implementations, the protocol allows ambiguous situations in which a client may send "Expect: 100-continue" without receiving either a 417 (Expectation Failed) status or a 100 (Continue) status. Therefore, when a client sends this header field to an origin server (possibly via a proxy) from which it has never seen a 100 (Continue) status, the client SHOULD NOT wait for an indefinite period before sending the request body.
Requirements for HTTP/1.1 origin servers:

  *Upon receiving a request which includes an Expect request-header
   field with the "100-continue" expectation, an origin server MUST
   either respond with 100 (Continue) status and continue to read
   from the input stream, or respond with a final status code. The
   origin server MUST NOT wait for the request body before sending
   the 100 (Continue) response. If it responds with a final status
   code, it MAY close the transport connection or it MAY continue to
   read and discard the rest of the request. It MUST NOT perform the
   requested method if it returns a final status code.

  *An origin server SHOULD NOT send a 100 (Continue) response if the
   request message does not include an Expect request-header field

with the "100-continue" expectation, and MUST NOT send a 100 (Continue) response if such a request comes from an HTTP/1.0 (or earlier) client. There is an exception to this rule: for compatibility with [RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.), a server MAY send a 100 (Continue) status in response to an HTTP/1.1 PUT or POST request that does not include an Expect request-header field with the "100-continue" expectation. This exception, the purpose of which is to minimize any client processing delays associated with an undeclared wait for 100 (Continue) status, applies only to HTTP/1.1 requests, and not to requests with any other HTTP-version value.

*An origin server MAY omit a 100 (Continue) response if it has already received some or all of the request body for the corresponding request.

*An origin server that sends a 100 (Continue) response MUST ultimately send a final status code, once the request body is received and processed, unless it terminates the transport connection prematurely.

*If an origin server receives a request that does not include an Expect request-header field with the "100-continue" expectation, the request includes a request body, and the server responds with a final status code before reading the entire request body from the transport connection, then the server SHOULD NOT close the transport connection until it has read the entire request, or until the client closes the connection. Otherwise, the client might not reliably receive the response message. However, this requirement is not be construed as preventing a server from defending itself against denial-of-service attacks, or from badly broken client implementations.

Requirements for HTTP/1.1 proxies:

*If a proxy receives a request that includes an Expect request-header field with the "100-continue" expectation, and the proxy either knows that the next-hop server complies with HTTP/1.1 or higher, or does not know the HTTP version of the next-hop server, it MUST forward the request, including the Expect header field.

*If the proxy knows that the version of the next-hop server is HTTP/1.0 or lower, it MUST NOT forward the request, and it MUST respond with a 417 (Expectation Failed) status.

*Proxies SHOULD maintain a cache recording the HTTP version numbers received from recently-referenced next-hop servers.

*A proxy MUST NOT forward a 100 (Continue) response if the request
 message was received from an HTTP/1.0 (or earlier) client and did
 not include an Expect request-header field with the "100-
 continue" expectation. This requirement overrides the general
 rule for forwarding of 1xx responses (see Section 8.1 of [Part2]
 (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter,
 L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke,
 Ed., "HTTP/1.1, part 2: Message Semantics," July 2010.)).

---

### 7.2.4.  Client Behavior if Server Prematurely Closes Connection      TOC

If an HTTP/1.1 client sends a request which includes a request body,
but which does not include an Expect request-header field with the
"100-continue" expectation, and if the client is not directly connected
to an HTTP/1.1 origin server, and if the client sees the connection
close before receiving any status from the server, the client SHOULD
retry the request. If the client does retry this request, it MAY use
the following "binary exponential backoff" algorithm to be assured of
obtaining a reliable response:

1. Initiate a new connection to the server

2. Transmit the request-headers

3. Initialize a variable R to the estimated round-trip time to the
   server (e.g., based on the time it took to establish the
   connection), or to a constant value of 5 seconds if the round-
   trip time is not available.

4. Compute T = R * (2**N), where N is the number of previous
   retries of this request.

5. Wait either for an error response from the server, or for T
   seconds (whichever comes first)

6. If no error response is received, after T seconds transmit the
   body of the request.

7. If client sees that the connection is closed prematurely,
   repeat from step 1 until the request is accepted, an error
   response is received, or the user becomes impatient and
   terminates the retry process.

If at any point an error status is received, the client

*SHOULD NOT continue and

*SHOULD close the connection if it has not completed sending the
   request message.

---

## 8.  Miscellaneous notes that may disappear

---

### 8.1.  Scheme aliases considered harmful

[TBD-aliases-harmful] (describe why aliases like webcal are harmful.)

---

### 8.2.  Use of HTTP for proxy communication

[TBD-proxy-other] (Configured to use HTTP to proxy HTTP or other
protocols.)

---

### 8.3.  Interception of HTTP for access control

[TBD-intercept] (Interception of HTTP traffic for initiating access
control.)

---

### 8.4.  Use of HTTP by other protocols

[TBD-profiles] (Profiles of HTTP defined by other protocol. Extensions
of HTTP like WebDAV.)

---

### 8.5.  Use of HTTP by media type specification

[TBD-hypertext] (Instructions on composing HTTP requests via hypertext
formats.)

---

## 9. Header Field Definitions

This section defines the syntax and semantics of HTTP/1.1 header fields related to message framing and transport protocols.
For entity-header fields, both sender and recipient refer to either the client or the server, depending on who sends and who receives the entity.

---

### 9.1. Connection

The "Connection" general-header field allows the sender to specify options that are desired for that particular connection and MUST NOT be communicated by proxies over further connections.
The Connection header's value has the following grammar:

```
Connection        = "Connection" ":" OWS Connection-v
Connection-v      = 1#connection-token
connection-token = token
```

HTTP/1.1 proxies MUST parse the Connection header field before a message is forwarded and, for each connection-token in this field, remove any header field(s) from the message with the same name as the connection-token. Connection options are signaled by the presence of a connection-token in the Connection header field, not by any corresponding additional header field(s), since the additional header field may not be sent if there are no parameters associated with that connection option.
Message headers listed in the Connection header MUST NOT include end-to-end headers, such as Cache-Control.
HTTP/1.1 defines the "close" connection option for the sender to signal that the connection will be closed after completion of the response.
For example,

        Connection: close

in either the request or the response header fields indicates that the connection SHOULD NOT be considered "persistent" (Section 7.1 (Persistent Connections)) after the current request/response is complete.
An HTTP/1.1 client that does not support persistent connections MUST include the "close" connection option in every request message.
An HTTP/1.1 server that does not support persistent connections MUST include the "close" connection option in every response message that does not have a 1xx (Informational) status code.
A system receiving an HTTP/1.0 (or lower-version) message that includes a Connection header MUST, for each connection-token in this field, remove and ignore any header field(s) from the message with the same

name as the connection-token. This protects against mistaken forwarding
of such header fields by pre-HTTP/1.1 proxies. See Appendix B.2
(Compatibility with HTTP/1.0 Persistent Connections).

---

## 9.2.  Content-Length

The "Content-Length" entity-header field indicates the size of the
entity-body, in number of OCTETs. In the case of responses to the HEAD
method, it indicates the size of the entity-body that would have been
sent had the request been a GET.

        Content-Length  = "Content-Length" ":" OWS 1*Content-Length-v
        Content-Length-v = 1*DIGIT

An example is

        Content-Length: 3495

Applications SHOULD use this field to indicate the transfer-length of
the message-body, unless this is prohibited by the rules in Section 3.4
(Message Length).
Any Content-Length greater than or equal to zero is a valid value.
Section 3.4 (Message Length) describes how to determine the length of a
message-body if a Content-Length is not given.
Note that the meaning of this field is significantly different from the
corresponding definition in MIME, where it is an optional field used
within the "message/external-body" content-type. In HTTP, it SHOULD be
sent whenever the message's length can be determined prior to being
transferred, unless this is prohibited by the rules in Section 3.4
(Message Length).

---

## 9.3.  Date

The "Date" general-header field represents the date and time at which
the message was originated, having the same semantics as the
Origination Date Field (orig-date) defined in Section 3.6.1 of
[RFC5322] (Resnick, P., "Internet Message Format," October 2008.). The
field value is an HTTP-date, as described in Section 6.1 (Date/Time
Formats: Full Date); it MUST be sent in rfc1123-date format.

        Date  = "Date" ":" OWS Date-v
        Date-v = HTTP-date

An example is

```
        Date: Tue, 15 Nov 1994 08:12:31 GMT
```

Origin servers MUST include a Date header field in all responses,
except in these cases:

1. If the response status code is 100 (Continue) or 101 (Switching
   Protocols), the response MAY include a Date header field, at
   the server's option.

2. If the response status code conveys a server error, e.g., 500
   (Internal Server Error) or 503 (Service Unavailable), and it is
   inconvenient or impossible to generate a valid Date.

3. If the server does not have a clock that can provide a
   reasonable approximation of the current time, its responses
   MUST NOT include a Date header field. In this case, the rules
   in Section 9.3.1 (Clockless Origin Server Operation) MUST be
   followed.

A received message that does not have a Date header field MUST be
assigned one by the recipient if the message will be cached by that
recipient or gatewayed via a protocol which requires a Date. An HTTP
implementation without a clock MUST NOT cache responses without
revalidating them on every use. An HTTP cache, especially a shared
cache, SHOULD use a mechanism, such as NTP [RFC1305] (Mills, D.,
"Network Time Protocol (Version 3) Specification, Implementation,"
March 1992.), to synchronize its clock with a reliable external
standard.
Clients SHOULD only send a Date header field in messages that include
an entity-body, as in the case of the PUT and POST requests, and even
then it is optional. A client without a clock MUST NOT send a Date
header field in a request.
The HTTP-date sent in a Date header SHOULD NOT represent a date and
time subsequent to the generation of the message. It SHOULD represent
the best available approximation of the date and time of message
generation, unless the implementation has no means of generating a
reasonably accurate date and time. In theory, the date ought to
represent the moment just before the entity is generated. In practice,
the date can be generated at any time during the message origination
without affecting its semantic value.

---

### 9.3.1. Clockless Origin Server Operation

Some origin server implementations might not have a clock available. An
origin server without a clock MUST NOT assign Expires or Last-Modified
values to a response, unless these values were associated with the

resource by a system or user with a reliable clock. It MAY assign an
Expires value that is known, at or before server configuration time, to
be in the past (this allows "pre-expiration" of responses without
storing separate Expires values for each resource).

## 9.4.  Host

The "Host" request-header field specifies the Internet host and port
number of the resource being requested, allowing the origin server or
gateway to differentiate between internally-ambiguous URLs, such as the
root "/" URL of a server for multiple host names on a single IP
address.
The Host field value MUST represent the naming authority of the origin
server or gateway given by the original URL obtained from the user or
referring resource (generally an http URI, as described in
Section 2.6.1 (http URI scheme)).

```
        Host   = "Host" ":" OWS Host-v
        Host-v = uri-host [ ":" port ] ; Section 2.6.1
```

A "host" without any trailing port information implies the default port
for the service requested (e.g., "80" for an HTTP URL). For example, a
request on the origin server for <http://www.example.org/pub/WWW/>
would properly include:

```
        GET /pub/WWW/ HTTP/1.1
        Host: www.example.org
```

A client MUST include a Host header field in all HTTP/1.1 request
messages. If the requested URI does not include an Internet host name
for the service being requested, then the Host header field MUST be
given with an empty value. An HTTP/1.1 proxy MUST ensure that any
request message it forwards does contain an appropriate Host header
field that identifies the service being requested by the proxy. All
Internet-based HTTP/1.1 servers MUST respond with a 400 (Bad Request)
status code to any HTTP/1.1 request message which lacks a Host header
field.
See Sections 4.2 (The Resource Identified by a Request) and B.1.1
(Changes to Simplify Multi-homed Web Servers and Conserve IP Addresses)
for other requirements relating to Host.

## 9.5.  TE

The "TE" request-header field indicates what extension transfer-codings it is willing to accept in the response, and whether or not it is willing to accept trailer fields in a chunked transfer-coding. Its value may consist of the keyword "trailers" and/or a comma-separated list of extension transfer-coding names with optional accept parameters (as described in Section 6.2 (Transfer Codings)).

```
TE        = "TE" ":" OWS TE-v
TE-v      = #t-codings
t-codings = "trailers" / ( transfer-extension [ te-params ] )
te-params = OWS ";" OWS "q=" qvalue *( te-ext )
te-ext    = OWS ";" OWS token [ "=" word ]
```

The presence of the keyword "trailers" indicates that the client is willing to accept trailer fields in a chunked transfer-coding, as defined in Section 6.2.1 (Chunked Transfer Coding). This keyword is reserved for use with transfer-coding values even though it does not itself represent a transfer-coding. Examples of its use are:

```
TE: deflate
TE:
TE: trailers, deflate;q=0.5
```

The TE header field only applies to the immediate connection. Therefore, the keyword MUST be supplied within a Connection header field (Section 9.1 (Connection)) whenever TE is present in an HTTP/1.1 message. A server tests whether a transfer-coding is acceptable, according to a TE field, using these rules:

1. The "chunked" transfer-coding is always acceptable. If the keyword "trailers" is listed, the client indicates that it is willing to accept trailer fields in the chunked response on behalf of itself and any downstream clients. The implication is that, if given, the client is stating that either all downstream clients are willing to accept trailer fields in the forwarded response, or that it will attempt to buffer the response on behalf of downstream recipients.

   Note: HTTP/1.1 does not define any means to limit the size of a chunked response such that a client can be assured of buffering the entire response.

2. If the transfer-coding being tested is one of the transfer-codings listed in the TE field, then it is acceptable unless it is accompanied by a qvalue of 0. (As defined in Section 6.4 (Quality Values), a qvalue of 0 means "not acceptable.")

3. If multiple transfer-codings are acceptable, then the
   acceptable transfer-coding with the highest non-zero qvalue is
   preferred. The "chunked" transfer-coding always has a qvalue of
   1.

If the TE field-value is empty or if no TE field is present, the only
transfer-coding is "chunked". A message with no transfer-coding is
always acceptable.

---

## 9.6.  Trailer

The "Trailer" general-header field indicates that the given set of
header fields is present in the trailer of a message encoded with
chunked transfer-coding.

```
Trailer  = "Trailer" ":" OWS Trailer-v
Trailer-v = 1#field-name
```

An HTTP/1.1 message SHOULD include a Trailer header field in a message
using chunked transfer-coding with a non-empty trailer. Doing so allows
the recipient to know which header fields to expect in the trailer.
If no Trailer header field is present, the trailer SHOULD NOT include
any header fields. See Section 6.2.1 (Chunked Transfer Coding) for
restrictions on the use of trailer fields in a "chunked" transfer-
coding.
Message header fields listed in the Trailer header field MUST NOT
include the following header fields:

   *Transfer-Encoding

   *Content-Length

   *Trailer

---

## 9.7.  Transfer-Encoding

The "Transfer-Encoding" general-header field indicates what transfer-
codings (if any) have been applied to the message body. It differs from
Content-Encoding (Section 2.2 of [Part3] (Fielding, R., Ed., Gettys,
J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.,
Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload
and Content Negotiation," July 2010.)) in that transfer-codings are a

property of the message (and therefore are removed by intermediaries), whereas content-codings are not.

```
Transfer-Encoding   = "Transfer-Encoding" ":" OWS
                       Transfer-Encoding-v
Transfer-Encoding-v = 1#transfer-coding
```

Transfer-codings are defined in <u>Section 6.2 (Transfer Codings)</u>. An example is:

```
Transfer-Encoding: chunked
```

If multiple encodings have been applied to an entity, the transfer-codings MUST be listed in the order in which they were applied. Additional information about the encoding parameters MAY be provided by other entity-header fields not defined by this specification. Many older HTTP/1.0 applications do not understand the Transfer-Encoding header.

---

### 9.8.  Upgrade TOC

The "Upgrade" general-header field allows the client to specify what additional communication protocols it would like to use, if the server chooses to switch protocols. Additionally, the server MUST use the Upgrade header field within a 101 (Switching Protocols) response to indicate which protocol(s) are being switched to.

```
Upgrade   = "Upgrade" ":" OWS Upgrade-v
Upgrade-v = 1#product
```

For example,

```
Upgrade: HTTP/2.0, SHTTP/1.3, IRC/6.9, RTA/x11
```

The Upgrade header field is intended to provide a simple mechanism for transition from HTTP/1.1 to some other, incompatible protocol. It does so by allowing the client to advertise its desire to use another protocol, such as a later version of HTTP with a higher major version number, even though the current request has been made using HTTP/1.1. This eases the difficult transition between incompatible protocols by allowing the client to initiate a request in the more commonly supported protocol while indicating to the server that it would like to use a "better" protocol if available (where "better" is determined by the server, possibly according to the nature of the method and/or resource being requested).
The Upgrade header field only applies to switching application-layer protocols upon the existing transport-layer connection. Upgrade cannot

be used to insist on a protocol change; its acceptance and use by the server is optional. The capabilities and nature of the application-layer communication after the protocol change is entirely dependent upon the new protocol chosen, although the first action after changing the protocol MUST be a response to the initial HTTP request containing the Upgrade header field.

The Upgrade header field only applies to the immediate connection. Therefore, the upgrade keyword MUST be supplied within a Connection header field ([Section 9.1 (Connection)](#)) whenever Upgrade is present in an HTTP/1.1 message.

The Upgrade header field cannot be used to indicate a switch to a protocol on a different connection. For that purpose, it is more appropriate to use a 301, 302, 303, or 305 redirection response.

This specification only defines the protocol name "HTTP" for use by the family of Hypertext Transfer Protocols, as defined by the HTTP version rules of [Section 2.5 (HTTP Version)](#) and future updates to this specification. Additional tokens can be registered with IANA using the registration procedure defined below.

---

### 9.8.1.  Upgrade Token Registry

The HTTP Upgrade Token Registry defines the name space for product tokens used to identify protocols in the Upgrade header field. Each registered token should be associated with one or a set of specifications, and with contact information.

Registrations should be allowed on a First Come First Served basis as described in Section 4.1 of [[RFC5226] (Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," May 2008.)](#). These specifications need not be IETF documents or be subject to IESG review, but should obey the following rules:

1. A token, once registered, stays registered forever.

2. The registration MUST name a responsible party for the registration.

3. The registration MUST name a point of contact.

4. The registration MAY name the documentation required for the token.

5. The responsible party MAY change the registration at any time. The IANA will keep a record of all such changes, and make them available upon request.

6. The responsible party for the first registration of a "product" token MUST approve later registrations of a "version" token

together with that "product" token before they can be
registered.

7. If absolutely required, the IESG MAY reassign the
responsibility for a token. This will normally only be used in
the case when a responsible party cannot be contacted.

It is not required that specifications for upgrade tokens be made
publicly available, but the contact information for the registration
should be.

---

**9.9. Via**

The "Via" general-header field MUST be used by gateways and proxies to
indicate the intermediate protocols and recipients between the user
agent and the server on requests, and between the origin server and the
client on responses. It is analogous to the "Received" field defined in
Section 3.6.7 of [RFC5322] (Resnick, P., "Internet Message Format,"
October 2008.) and is intended to be used for tracking message
forwards, avoiding request loops, and identifying the protocol
capabilities of all senders along the request/response chain.

```
Via               = "Via" ":" OWS Via-v
Via-v             = 1#( received-protocol RWS received-by
                        [ RWS comment ] )
received-protocol = [ protocol-name "/" ] protocol-version
protocol-name     = token
protocol-version  = token
received-by       = ( uri-host [ ":" port ] ) / pseudonym
pseudonym         = token
```

The received-protocol indicates the protocol version of the message
received by the server or client along each segment of the request/
response chain. The received-protocol version is appended to the Via
field value when the message is forwarded so that information about the
protocol capabilities of upstream applications remains visible to all
recipients.
The protocol-name is optional if and only if it would be "HTTP". The
received-by field is normally the host and optional port number of a
recipient server or client that subsequently forwarded the message.
However, if the real host is considered to be sensitive information, it
MAY be replaced by a pseudonym. If the port is not given, it MAY be
assumed to be the default port of the received-protocol.
Multiple Via field values represent each proxy or gateway that has
forwarded the message. Each recipient MUST append its information such
that the end result is ordered according to the sequence of forwarding
applications.

Comments MAY be used in the Via header field to identify the software of the recipient proxy or gateway, analogous to the User-Agent and Server header fields. However, all comments in the Via field are optional and MAY be removed by any recipient prior to forwarding the message.

For example, a request message could be sent from an HTTP/1.0 user agent to an internal proxy code-named "fred", which uses HTTP/1.1 to forward the request to a public proxy at p.example.net, which completes the request by forwarding it to the origin server at www.example.com. The request received by www.example.com would then have the following Via header field:

    Via: 1.0 fred, 1.1 p.example.net (Apache/1.1)

Proxies and gateways used as a portal through a network firewall SHOULD NOT, by default, forward the names and ports of hosts within the firewall region. This information SHOULD only be propagated if explicitly enabled. If not enabled, the received-by host of any host behind the firewall SHOULD be replaced by an appropriate pseudonym for that host.

For organizations that have strong privacy requirements for hiding internal structures, a proxy MAY combine an ordered subsequence of Via header field entries with identical received-protocol values into a single such entry. For example,

    Via: 1.0 ricky, 1.1 ethel, 1.1 fred, 1.0 lucy

could be collapsed to

    Via: 1.0 ricky, 1.1 mertz, 1.0 lucy

Applications SHOULD NOT combine multiple entries unless they are all under the same organizational control and the hosts have already been replaced by pseudonyms. Applications MUST NOT combine entries which have different received-protocol values.

---

## 10.  IANA Considerations

---

## 10.1.  Message Header Registration

The Message Header Registry located at http://www.iana.org/assignments/message-headers/message-header-index.html should be updated with the permanent registrations below (see [RFC3864] (Klyne, G., Nottingham,

M., and J. Mogul, "Registration Procedures for Message Header Fields," September 2004.)):

| Header Field Name | Protocol | Status | Reference |
|---|---|---|---|
| Connection | http | standard | Section 9.1 (Connection) |
| Content-Length | http | standard | Section 9.2 (Content-Length) |
| Date | http | standard | Section 9.3 (Date) |
| Host | http | standard | Section 9.4 (Host) |
| TE | http | standard | Section 9.5 (TE) |
| Trailer | http | standard | Section 9.6 (Trailer) |
| Transfer-Encoding | http | standard | Section 9.7 (Transfer-Encoding) |
| Upgrade | http | standard | Section 9.8 (Upgrade) |
| Via | http | standard | Section 9.9 (Via) |

The change controller is: "IETF (iesg@ietf.org) - Internet Engineering Task Force".

## 10.2.  URI Scheme Registration                    TOC

The entries for the "http" and "https" URI Schemes in the registry located at http://www.iana.org/assignments/uri-schemes.html should be updated to point to Sections 2.6.1 (http URI scheme) and 2.6.2 (https URI scheme) of this document (see [RFC4395] (Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes," February 2006.)).

## 10.3.  Internet Media Type Registrations          TOC

This document serves as the specification for the Internet media types "message/http" and "application/http". The following is to be registered with IANA (see [RFC4288] (Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures," December 2005.)).

## 10.3.1.  Internet Media Type message/http          TOC

The message/http type can be used to enclose a single HTTP request or response message, provided that it obeys the MIME restrictions for all "message" types regarding line length and encodings.

**Type name:**  message

**Subtype name:**  http

**Required parameters:**  none

**Optional parameters:**  version, msgtype

> **version:**  The HTTP-Version number of the enclosed message (e.g., "1.1"). If not present, the version can be determined from the first line of the body.

> **msgtype:**  The message type -- "request" or "response". If not present, the type can be determined from the first line of the body.

**Encoding considerations:**  only "7bit", "8bit", or "binary" are permitted

**Security considerations:**  none

**Interoperability considerations:**  none

**Published specification:**  This specification (see [Section 10.3.1 (Internet Media Type message/http)](#)).

**Applications that use this media type:**
**Additional information:**

> **Magic number(s):**  none

> **File extension(s):**  none

> **Macintosh file type code(s):**  none

**Person and email address to contact for further information:**  See Authors Section.

**Intended usage:**  COMMON

**Restrictions on usage:**  none

**Author/Change controller:**  IESG

**10.3.2. Internet Media Type application/http**

The application/http type can be used to enclose a pipeline of one or more HTTP request or response messages (not intermixed).

**Type name:** application

**Subtype name:** http

**Required parameters:** none

**Optional parameters:** version, msgtype

>   **version:** The HTTP-Version number of the enclosed messages (e.g., "1.1"). If not present, the version can be determined from the first line of the body.

>   **msgtype:** The message type -- "request" or "response". If not present, the type can be determined from the first line of the body.

**Encoding considerations:** HTTP messages enclosed by this type are in "binary" format; use of an appropriate Content-Transfer-Encoding is required when transmitted via E-mail.

**Security considerations:** none

**Interoperability considerations:** none

**Published specification:** This specification (see Section 10.3.2 (Internet Media Type application/http)).

**Applications that use this media type:**
**Additional information:**

>   **Magic number(s):** none

>   **File extension(s):** none

>   **Macintosh file type code(s):** none

**Person and email address to contact for further information:** See Authors Section.

**Intended usage:** COMMON

**Restrictions on usage:** none

**Author/Change controller:** IESG

## 10.4.  Transfer Coding Registry TOC

The registration procedure for HTTP Transfer Codings is now defined by
Section 6.2.3 (Transfer Coding Registry) of this document.
The HTTP Transfer Codings Registry located at http://www.iana.org/
assignments/http-parameters should be updated with the registrations
below:

| Name | Description | Reference |
|---|---|---|
| chunked | Transfer in a series of chunks | Section 6.2.1 (Chunked Transfer Coding) |
| compress | UNIX "compress" program method | Section 6.2.2.1 (Compress Coding) |
| deflate | "deflate" compression mechanism ([RFC1951] (Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3," May 1996.)) used inside the "zlib" data format ([RFC1950] (Deutsch, L. and J-L. Gailly, "ZLIB Compressed Data Format Specification version 3.3," May 1996.)) | Section 6.2.2.2 (Deflate Coding) |
| gzip | Same as GNU zip [RFC1952] (Deutsch, P., Gailly, J-L., Adler, M., Deutsch, L., and G. Randers-Pehrson, "GZIP file format specification version 4.3," May 1996.) | Section 6.2.2.3 (Gzip Coding) |

## 10.5.  Upgrade Token Registration TOC

The registration procedure for HTTP Upgrade Tokens --
previously defined in Section 7.2 of [RFC2817] (Khare, R. and S.
Lawrence, "Upgrading to TLS Within HTTP/1.1," May 2000.) -- is now
defined by Section 9.8.1 (Upgrade Token Registry) of this document.
The HTTP Status Code Registry located at http://www.iana.org/
assignments/http-upgrade-tokens/ should be updated with the
registration below:

| Value | Description | Reference |
|---|---|---|
| HTTP | Hypertext Transfer Protocol | Section 2.5 (HTTP Version) of this specification |

TOC

## 11. Security Considerations

This section is meant to inform application developers, information providers, and users of the security limitations in HTTP/1.1 as described by this document. The discussion does not include definitive solutions to the problems revealed, though it does make some suggestions for reducing security risks.

---

### 11.1. Personal Information

HTTP clients are often privy to large amounts of personal information (e.g., the user's name, location, mail address, passwords, encryption keys, etc.), and SHOULD be very careful to prevent unintentional leakage of this information. We very strongly recommend that a convenient interface be provided for the user to control dissemination of such information, and that designers and implementors be particularly careful in this area. History shows that errors in this area often create serious security and/or privacy problems and generate highly adverse publicity for the implementor's company.

---

### 11.2. Abuse of Server Log Information

A server is in the position to save personal data about a user's requests which might identify their reading patterns or subjects of interest. This information is clearly confidential in nature and its handling can be constrained by law in certain countries. People using HTTP to provide data are responsible for ensuring that such material is not distributed without the permission of any individuals that are identifiable by the published results.

---

### 11.3. Attacks Based On File and Path Names

Implementations of HTTP origin servers SHOULD be careful to restrict the documents returned by HTTP requests to be only those that were intended by the server administrators. If an HTTP server translates HTTP URIs directly into file system calls, the server MUST take special care not to serve files that were not intended to be delivered to HTTP clients. For example, UNIX, Microsoft Windows, and other operating systems use ".." as a path component to indicate a directory level above the current one. On such a system, an HTTP server MUST disallow any such construct in the request-target if it would otherwise allow

access to a resource outside those intended to be accessible via the HTTP server. Similarly, files intended for reference only internally to the server (such as access control files, configuration files, and script code) MUST be protected from inappropriate retrieval, since they might contain sensitive information. Experience has shown that minor bugs in such HTTP server implementations have turned into security risks.

---

### 11.4. DNS Spoofing

Clients using HTTP rely heavily on the Domain Name Service, and are thus generally prone to security attacks based on the deliberate mis-association of IP addresses and DNS names. Clients need to be cautious in assuming the continuing validity of an IP number/DNS name association.

In particular, HTTP clients SHOULD rely on their name resolver for confirmation of an IP number/DNS name association, rather than caching the result of previous host name lookups. Many platforms already can cache host name lookups locally when appropriate, and they SHOULD be configured to do so. It is proper for these lookups to be cached, however, only when the TTL (Time To Live) information reported by the name server makes it likely that the cached information will remain useful.

If HTTP clients cache the results of host name lookups in order to achieve a performance improvement, they MUST observe the TTL information reported by DNS.

If HTTP clients do not observe this rule, they could be spoofed when a previously-accessed server's IP address changes. As network renumbering is expected to become increasingly common [RFC1900] (Carpenter, B. and Y. Rekhter, "Renumbering Needs Work," February 1996.), the possibility of this form of attack will grow. Observing this requirement thus reduces this potential security vulnerability.

This requirement also improves the load-balancing behavior of clients for replicated servers using the same DNS name and reduces the likelihood of a user's experiencing failure in accessing sites which use that strategy.

---

### 11.5. Proxies and Caching

By their very nature, HTTP proxies are men-in-the-middle, and represent an opportunity for man-in-the-middle attacks. Compromise of the systems on which the proxies run can result in serious security and privacy problems. Proxies have access to security-related information, personal information about individual users and organizations, and proprietary

information belonging to users and content providers. A compromised proxy, or a proxy implemented or configured without regard to security and privacy considerations, might be used in the commission of a wide range of potential attacks.

Proxy operators should protect the systems on which proxies run as they would protect any system that contains or transports sensitive information. In particular, log information gathered at proxies often contains highly sensitive personal information, and/or information about organizations. Log information should be carefully guarded, and appropriate guidelines for use should be developed and followed. ([Section 11.2 (Abuse of Server Log Information)](#)).

Proxy implementors should consider the privacy and security implications of their design and coding decisions, and of the configuration options they provide to proxy operators (especially the default configuration).

Users of a proxy need to be aware that proxies are no trustworthier than the people who run them; HTTP itself cannot solve this problem. The judicious use of cryptography, when appropriate, may suffice to protect against a broad range of security and privacy attacks. Such cryptography is beyond the scope of the HTTP/1.1 specification.

---

## 11.6.  Denial of Service Attacks on Proxies [TOC]

They exist. They are hard to defend against. Research continues. Beware.

---

## 12.  Acknowledgments [TOC]

HTTP has evolved considerably over the years. It has benefited from a large and active developer community--the many people who have participated on the www-talk mailing list--and it is that community which has been most responsible for the success of HTTP and of the World-Wide Web in general. Marc Andreessen, Robert Cailliau, Daniel W. Connolly, Bob Denny, John Franks, Jean-Francois Groff, Phillip M. Hallam-Baker, Hakon W. Lie, Ari Luotonen, Rob McCool, Lou Montulli, Dave Raggett, Tony Sanders, and Marc VanHeyningen deserve special recognition for their efforts in defining early aspects of the protocol.

This document has benefited greatly from the comments of all those participating in the HTTP-WG. In addition to those already mentioned, the following individuals have contributed to this specification: Gary Adams, Harald Tveit Alvestrand, Keith Ball, Brian Behlendorf, Paul Burchard, Maurizio Codogno, Josh Cohen, Mike Cowlishaw, Roman Czyborra, Michael A. Dolan, Daniel DuBois, David J. Fiander, Alan Freier, Marc

---

## 13.  References

---

## 13.1. Normative References

| | |
|---|---|
| [ISO-8859-1] | International Organization for Standardization, "Information technology -- 8-bit single-byte coded graphic character sets -- Part 1: Latin alphabet No. 1," ISO/IEC 8859-1:1998, 1998. |
| [Part2] | Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 2: Message Semantics," draft-ietf-httpbis-p2-semantics-10 (work in progress), July 2010. |
| [Part3] | Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload and Content Negotiation," draft-ietf-httpbis-p3-payload-10 (work in progress), July 2010. |
| [Part5] | Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 5: Range Requests and Partial Responses," draft-ietf-httpbis-p5-range-10 (work in progress), July 2010. |
| [Part6] | Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1, part 6: Caching," draft-ietf-httpbis-p6-cache-10 (work in progress), July 2010. |
| [RFC1950] | Deutsch, L. and J-L. Gailly, "ZLIB Compressed Data Format Specification version 3.3," RFC 1950, May 1996. RFC 1950 is an Informational RFC, thus it may be less stable than this specification. On the other hand, this downward reference was present since the publication of RFC 2068 in 1997 ([RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.)), therefore it is unlikely to cause problems in practice. See also [BCP97] (Klensin, J. and S. Hartman, "Handling Normative References to Standards-Track Documents," June 2007.). |
| [RFC1951] | Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3," RFC 1951, May 1996. RFC 1951 is an Informational RFC, thus it may be less stable than this specification. On the other hand, this downward reference was present since the publication of RFC 2068 in 1997 ([RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.)), therefore it is unlikely to cause problems in practice. See also [BCP97] (Klensin, |

| | J. and S. Hartman, "Handling Normative References to Standards-Track Documents," June 2007.). |
|---|---|
| [RFC1952] | Deutsch, P., Gailly, J-L., Adler, M., Deutsch, L., and G. Randers-Pehrson, "GZIP file format specification version 4.3," RFC 1952, May 1996. RFC 1952 is an Informational RFC, thus it may be less stable than this specification. On the other hand, this downward reference was present since the publication of RFC 2068 in 1997 ([RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.)), therefore it is unlikely to cause problems in practice. See also [BCP97] (Klensin, J. and S. Hartman, "Handling Normative References to Standards-Track Documents," June 2007.). |
| [RFC2119] | Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997. |
| [RFC3986] | Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax," RFC 3986, STD 66, January 2005. |
| [RFC5234] | Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," STD 68, RFC 5234, January 2008. |
| [USASCII] | American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange," ANSI X3.4, 1986. |

## 13.2. Informative References

| [BCP97] | Klensin, J. and S. Hartman, "Handling Normative References to Standards-Track Documents," BCP 97, RFC 4897, June 2007. |
|---|---|
| [Kri2001] | Kristol, D., "HTTP Cookies: Standards, Privacy, and Politics," ACM Transactions on Internet Technology Vol. 1, #2, November 2001. |
| [Nie1997] | Frystyk, H., Gettys, J., Prud'hommeaux, E., Lie, H., and C. Lilley, "Network Performance Effects of HTTP/1.1, CSS1, and PNG," ACM Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '97, September 1997. |
| [Pad1995] | Padmanabhan, V. and J. Mogul, "Improving HTTP Latency," Computer Networks and ISDN Systems v. 28, pp. 25-35, December 1995. |
| [RFC1123] | Braden, R., "Requirements for Internet Hosts - Application and Support," STD 3, RFC 1123, October 1989. |

| | |
|---|---|
| [RFC1305] | Mills, D., "Network Time Protocol (Version 3) Specification, Implementation," RFC 1305, March 1992. |
| [RFC1900] | Carpenter, B. and Y. Rekhter, "Renumbering Needs Work," RFC 1900, February 1996. |
| [RFC1945] | Berners-Lee, T., Fielding, R., and H. Nielsen, "Hypertext Transfer Protocol -- HTTP/1.0," RFC 1945, May 1996. |
| [RFC2045] | Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies," RFC 2045, November 1996. |
| [RFC2047] | Moore, K., "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text," RFC 2047, November 1996. |
| [RFC2068] | Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," RFC 2068, January 1997. |
| [RFC2109] | Kristol, D. and L. Montulli, "HTTP State Management Mechanism," RFC 2109, February 1997. |
| [RFC2145] | Mogul, J., Fielding, R., Gettys, J., and H. Nielsen, "Use and Interpretation of HTTP Version Numbers," RFC 2145, May 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. |
| [RFC2817] | Khare, R. and S. Lawrence, "Upgrading to TLS Within HTTP/1.1," RFC 2817, May 2000. |
| [RFC2818] | Rescorla, E., "HTTP Over TLS," RFC 2818, May 2000. |
| [RFC2965] | Kristol, D. and L. Montulli, "HTTP State Management Mechanism," RFC 2965, October 2000. |
| [RFC3864] | Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields," BCP 90, RFC 3864, September 2004. |
| [RFC4288] | Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures," BCP 13, RFC 4288, December 2005. |
| [RFC4395] | Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes," BCP 115, RFC 4395, February 2006. |
| [RFC5226] | Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," BCP 26, RFC 5226, May 2008. |
| [RFC5322] | Resnick, P., "Internet Message Format," RFC 5322, October 2008. |
| [Spe] | Spero, S., "Analysis of HTTP Performance Problems." |
| [Tou1998] | Touch, J., Heidemann, J., and K. Obraczka, "Analysis of HTTP Performance," ISI Research Report ISI/RR-98-463, Aug 1998. (original report dated Aug. 1996) |

Although this document specifies the requirements for the generation of
HTTP/1.1 messages, not all applications will be correct in their
implementation. We therefore recommend that operational applications be
tolerant of deviations whenever those deviations can be interpreted
unambiguously.
Clients SHOULD be tolerant in parsing the Status-Line and servers
SHOULD be tolerant when parsing the Request-Line. In particular, they
SHOULD accept any amount of WSP characters between fields, even though
only a single SP is required.
The line terminator for header fields is the sequence CRLF. However, we
recommend that applications, when parsing such headers, recognize a
single LF as a line terminator and ignore the leading CR.
The character set of an entity-body SHOULD be labeled as the lowest
common denominator of the character codes used within that body, with
the exception that not labeling the entity is preferred over labeling
the entity with the labels US-ASCII or ISO-8859-1. See [Part3]
(Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L.,
Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/
1.1, part 3: Message Payload and Content Negotiation," July 2010.).
Additional rules for requirements on parsing and encoding of dates and
other potential problems with date encodings include:

    *HTTP/1.1 clients and caches SHOULD assume that an RFC-850 date
     which appears to be more than 50 years in the future is in fact
     in the past (this helps solve the "year 2000" problem).

    *Although all date formats are specified to be case-sensitive,
     recipients SHOULD match day, week and timezone names case-
     insensitively.

    *An HTTP/1.1 implementation MAY internally represent a parsed
     Expires date as earlier than the proper value, but MUST NOT
     internally represent a parsed Expires date as later than the
     proper value.

    *All expiration-related calculations MUST be done in GMT. The
     local time zone MUST NOT influence the calculation or comparison
     of an age or expiration time.

    *If an HTTP header incorrectly carries a date value with a time
     zone other than GMT, it MUST be converted into GMT using the most
     conservative possible conversion.

## Appendix B.  Compatibility with Previous Versions

HTTP has been in use by the World-Wide Web global information initiative since 1990. The first version of HTTP, later referred to as HTTP/0.9, was a simple protocol for hypertext data transfer across the Internet with only a single method and no metadata. HTTP/1.0, as defined by [RFC1945] (Berners-Lee, T., Fielding, R., and H. Nielsen, "Hypertext Transfer Protocol -- HTTP/1.0," May 1996.), added a range of request methods and MIME-like messaging that could include metadata about the data transferred and modifiers on the request/response semantics. However, HTTP/1.0 did not sufficiently take into consideration the effects of hierarchical proxies, caching, the need for persistent connections, or name-based virtual hosts. The proliferation of incompletely-implemented applications calling themselves "HTTP/1.0" further necessitated a protocol version change in order for two communicating applications to determine each other's true capabilities.

HTTP/1.1 remains compatible with HTTP/1.0 by including more stringent requirements that enable reliable implementations, adding only those new features that will either be safely ignored by an HTTP/1.0 recipient or only sent when communicating with a party advertising compliance with HTTP/1.1.

It is beyond the scope of a protocol specification to mandate compliance with previous versions. HTTP/1.1 was deliberately designed, however, to make supporting previous versions easy. It is worth noting that, at the time of composing this specification, we would expect general-purpose HTTP/1.1 servers to:

   *understand any valid request in the format of HTTP/1.0 and 1.1;

   *respond appropriately with a message in the same major version
    used by the client.

And we would expect HTTP/1.1 clients to:

   *understand any valid response in the format of HTTP/1.0 or 1.1.

For most implementations of HTTP/1.0, each connection is established by the client prior to the request and closed by the server after sending the response. Some implementations implement the Keep-Alive version of persistent connections described in Section 19.7.1 of [RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.).

**B.1.  Changes from HTTP/1.0**

This section summarizes major differences between versions HTTP/1.0 and
HTTP/1.1.

---

**B.1.1.  Changes to Simplify Multi-homed Web Servers and
Conserve IP Addresses**

The requirements that clients and servers support the Host request-
header, report an error if the Host request-header ([Section 9.4 (Host)](#))
is missing from an HTTP/1.1 request, and accept absolute URIs
([Section 4.1.2 (request-target)](#)) are among the most important changes
defined by this specification.
Older HTTP/1.0 clients assumed a one-to-one relationship of IP
addresses and servers; there was no other established mechanism for
distinguishing the intended server of a request than the IP address to
which that request was directed. The changes outlined above will allow
the Internet, once older HTTP clients are no longer common, to support
multiple Web sites from a single IP address, greatly simplifying large
operational Web servers, where allocation of many IP addresses to a
single host has created serious problems. The Internet will also be
able to recover the IP addresses that have been allocated for the sole
purpose of allowing special-purpose domain names to be used in root-
level HTTP URLs. Given the rate of growth of the Web, and the number of
servers already deployed, it is extremely important that all
implementations of HTTP (including updates to existing HTTP/1.0
applications) correctly implement these requirements:

    *Both clients and servers MUST support the Host request-header.

    *A client that sends an HTTP/1.1 request MUST send a Host header.

    *Servers MUST report a 400 (Bad Request) error if an HTTP/1.1
     request does not include a Host request-header.

    *Servers MUST accept absolute URIs.

---

**B.2.  Compatibility with HTTP/1.0 Persistent Connections**

Some clients and servers might wish to be compatible with some previous
implementations of persistent connections in HTTP/1.0 clients and
servers. Persistent connections in HTTP/1.0 are explicitly negotiated
as they are not the default behavior. HTTP/1.0 experimental

implementations of persistent connections are faulty, and the new facilities in HTTP/1.1 are designed to rectify these problems. The problem was that some existing HTTP/1.0 clients may be sending Keep-Alive to a proxy server that doesn't understand Connection, which would then erroneously forward it to the next inbound server, which would establish the Keep-Alive connection and result in a hung HTTP/1.0 proxy waiting for the close on the response. The result is that HTTP/1.0 clients must be prevented from using Keep-Alive when talking to proxies.

However, talking to proxies is the most important use of persistent connections, so that prohibition is clearly unacceptable. Therefore, we need some other mechanism for indicating a persistent connection is desired, which is safe to use even when talking to an old proxy that ignores Connection. Persistent connections are the default for HTTP/1.1 messages; we introduce a new keyword (Connection: close) for declaring non-persistence. See Section 9.1 (Connection).

The original HTTP/1.0 form of persistent connections (the Connection: Keep-Alive and Keep-Alive header) is documented in Section 19.7.1 of [RFC2068] (Fielding, R., Gettys, J., Mogul, J., Nielsen, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," January 1997.).

---

### B.3.  Changes from RFC 2068                                    TOC

This specification has been carefully audited to correct and disambiguate key word usage; RFC 2068 had many problems in respect to the conventions laid out in [RFC2119] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.).

Transfer-coding and message lengths all interact in ways that required fixing exactly when chunked encoding is used (to allow for transfer encoding that may not be self delimiting); it was important to straighten out exactly how message lengths are computed. (Sections 6.2 (Transfer Codings), 3.4 (Message Length), 9.2 (Content-Length), see also [Part3] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 3: Message Payload and Content Negotiation," July 2010.), [Part5] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 5: Range Requests and Partial Responses," July 2010.) and [Part6] (Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1, part 6: Caching," July 2010.))

The use and interpretation of HTTP version numbers has been clarified by [RFC2145] (Mogul, J., Fielding, R., Gettys, J., and H. Nielsen, "Use and Interpretation of HTTP Version Numbers," May 1997.). Require proxies to upgrade requests to highest protocol version they support to

deal with problems discovered in HTTP/1.0 implementations ([Section 2.5 (HTTP Version)](#))

Quality Values of zero should indicate that "I don't want something" to allow clients to refuse a representation. ([Section 6.4 (Quality Values)](#))

Transfer-coding had significant problems, particularly with interactions with chunked encoding. The solution is that transfer-codings become as full fledged as content-codings. This involves adding an IANA registry for transfer-codings (separate from content codings), a new header field (TE) and enabling trailer headers in the future. Transfer encoding is a major performance benefit, so it was worth fixing [Nie1997] (Frystyk, H., Gettys, J., Prud'hommeaux, E., Lie, H., and C. Lilley, "Network Performance Effects of HTTP/1.1, CSS1, and PNG," September 1997.). TE also solves another, obscure, downward interoperability problem that could have occurred due to interactions between authentication trailers, chunked encoding and HTTP/1.0 clients. (Section [6.2 (Transfer Codings)](#), [6.2.1 (Chunked Transfer Coding)](#), [7.1.3.2 (Non-modifiable Headers)](#), and [9.5 (TE)](#))

Proxies should be able to add Content-Length when appropriate. ([Section 7.1.3.2 (Non-modifiable Headers)](#))

---

### B.4.  Changes from RFC 2616                          [TOC](#)

Empty list elements in list productions have been deprecated. ([Section 1.2.1 (ABNF Extension: #rule)](#))

Rules about implicit linear whitespace between certain grammar productions have been removed; now it's only allowed when specifically pointed out in the ABNF. The NUL character is no longer allowed in comment and quoted-string text. The quoted-pair rule no longer allows escaping control characters other than HTAB. Non-ASCII content in header fields and reason phrase has been obsoleted and made opaque (the TEXT rule was removed) ([Section 1.2.2 (Basic Rules)](#))

Clarify that HTTP-Version is case sensitive. ([Section 2.5 (HTTP Version)](#))

Remove reference to non-existent identity transfer-coding value tokens. (Sections [6.2 (Transfer Codings)](#) and [3.4 (Message Length)](#))

Require that invalid whitespace around field-names be rejected. ([Section 3.2 (Header Fields)](#))

Update use of abs_path production from RFC1808 to the path-absolute + query components of RFC3986. ([Section 4.1.2 (request-target)](#))

Clarification that the chunk length does not include the count of the octets in the chunk header and trailer. Furthermore disallowed line folding in chunk extensions. ([Section 6.2.1 (Chunked Transfer Coding)](#))

Remove hard limit of two connections per server. ([Section 7.1.4 (Practical Considerations)](#))

Clarify exactly when close connection options must be sent.
([Section 9.1 (Connection)](#))

**Appendix C.  Collected ABNF**

```
BWS = OWS

Cache-Control = <Cache-Control, defined in [Part6], Section 3.4>
Chunked-Body = *chunk last-chunk trailer-part CRLF
Connection = "Connection:" OWS Connection-v
Connection-v = *( "," OWS ) connection-token *( OWS "," [ OWS
 connection-token ] )
Content-Length = "Content-Length:" OWS 1*Content-Length-v
Content-Length-v = 1*DIGIT

Date = "Date:" OWS Date-v
Date-v = HTTP-date

GMT = %x47.4D.54 ; GMT

HTTP-Prot-Name = %x48.54.54.50 ; HTTP
HTTP-Version = HTTP-Prot-Name "/" 1*DIGIT "." 1*DIGIT
HTTP-date = rfc1123-date / obs-date
HTTP-message = start-line *( header-field CRLF ) CRLF [ message-body
 ]
Host = "Host:" OWS Host-v
Host-v = uri-host [ ":" port ]

Method = token

OWS = *( [ obs-fold ] WSP )

Pragma = <Pragma, defined in [Part6], Section 3.4>

RWS = 1*( [ obs-fold ] WSP )
Reason-Phrase = *( WSP / VCHAR / obs-text )
Request = Request-Line *( ( general-header / request-header /
 entity-header ) CRLF ) CRLF [ message-body ]
Request-Line = Method SP request-target SP HTTP-Version CRLF
Response = Status-Line *( ( general-header / response-header /
 entity-header ) CRLF ) CRLF [ message-body ]

Status-Code = 3DIGIT
Status-Line = HTTP-Version SP Status-Code SP Reason-Phrase CRLF

TE = "TE:" OWS TE-v
TE-v = [ ( "," / t-codings ) *( OWS "," [ OWS t-codings ] ) ]
Trailer = "Trailer:" OWS Trailer-v
Trailer-v = *( "," OWS ) field-name *( OWS "," [ OWS field-name ] )
Transfer-Encoding = "Transfer-Encoding:" OWS Transfer-Encoding-v
Transfer-Encoding-v = *( "," OWS ) transfer-coding *( OWS "," [ OWS
 transfer-coding ] )

URI-reference = <URI-reference, defined in [RFC3986], Section 4.1>
```

```
Upgrade = "Upgrade:" OWS Upgrade-v
Upgrade-v = *( "," OWS ) product *( OWS "," [ OWS product ] )

Via = "Via:" OWS Via-v
Via-v = *( "," OWS ) received-protocol RWS received-by [ RWS comment
 ] *( OWS "," [ OWS received-protocol RWS received-by [ RWS comment ]
 ] )

Warning = <Warning, defined in [Part6], Section 3.6>

absolute-URI = <absolute-URI, defined in [RFC3986], Section 4.3>
asctime-date = day-name SP date3 SP time-of-day SP year
attribute = token
authority = <authority, defined in [RFC3986], Section 3.2>

chunk = chunk-size *WSP [ chunk-ext ] CRLF chunk-data CRLF
chunk-data = 1*OCTET
chunk-ext = *( ";" *WSP chunk-ext-name [ "=" chunk-ext-val ] *WSP )
chunk-ext-name = token
chunk-ext-val = token / quoted-str-nf
chunk-size = 1*HEXDIG
comment = "(" *( ctext / quoted-cpair / comment ) ")"
connection-token = token
ctext = OWS / %x21-27 ; '!'-'''
 / %x2A-5B ; '*'-'['
 / %x5D-7E ; ']'-'~'
 / obs-text

date1 = day SP month SP year
date2 = day "-" month "-" 2DIGIT
date3 = month SP ( 2DIGIT / ( SP DIGIT ) )
day = 2DIGIT
day-name = %x4D.6F.6E ; Mon
 / %x54.75.65 ; Tue
 / %x57.65.64 ; Wed
 / %x54.68.75 ; Thu
 / %x46.72.69 ; Fri
 / %x53.61.74 ; Sat
 / %x53.75.6E ; Sun
day-name-l = %x4D.6F.6E.64.61.79 ; Monday
 / %x54.75.65.73.64.61.79 ; Tuesday
 / %x57.65.64.6E.65.73.64.61.79 ; Wednesday
 / %x54.68.75.72.73.64.61.79 ; Thursday
 / %x46.72.69.64.61.79 ; Friday
 / %x53.61.74.75.72.64.61.79 ; Saturday
 / %x53.75.6E.64.61.79 ; Sunday

entity-body = <entity-body, defined in [Part3], Section 3.2>
entity-header = <entity-header, defined in [Part3], Section 3.1>
```

```
field-content = *( WSP / VCHAR / obs-text )
field-name = token
field-value = *( field-content / OWS )

general-header = Cache-Control / Connection / Date / Pragma / Trailer
 / Transfer-Encoding / Upgrade / Via / Warning

header-field = field-name ":" OWS [ field-value ] OWS
hour = 2DIGIT
http-URI = "http://" authority path-abempty [ "?" query ]
https-URI = "https://" authority path-abempty [ "?" query ]

last-chunk = 1*"0" *WSP [ chunk-ext ] CRLF

message-body = entity-body /
 <entity-body encoded as per Transfer-Encoding>
minute = 2DIGIT
month = %x4A.61.6E ; Jan
 / %x46.65.62 ; Feb
 / %x4D.61.72 ; Mar
 / %x41.70.72 ; Apr
 / %x4D.61.79 ; May
 / %x4A.75.6E ; Jun
 / %x4A.75.6C ; Jul
 / %x41.75.67 ; Aug
 / %x53.65.70 ; Sep
 / %x4F.63.74 ; Oct
 / %x4E.6F.76 ; Nov
 / %x44.65.63 ; Dec

obs-date = rfc850-date / asctime-date
obs-fold = CRLF
obs-text = %x80-FF

partial-URI = relative-part [ "?" query ]
path-abempty = <path-abempty, defined in [RFC3986], Section 3.3>
path-absolute = <path-absolute, defined in [RFC3986], Section 3.3>
port = <port, defined in [RFC3986], Section 3.2.3>
product = token [ "/" product-version ]
product-version = token
protocol-name = token
protocol-version = token
pseudonym = token

qdtext = OWS / "!" / %x23-5B ; '#'-'['
 / %x5D-7E ; ']'-'~'
 / obs-text
qdtext-nf = WSP / "!" / %x23-5B ; '#'-'['
 / %x5D-7E ; ']'-'~'
```

```
 / obs-text
query = <query, defined in [RFC3986], Section 3.4>
quoted-cpair = "\" ( WSP / VCHAR / obs-text )
quoted-pair = "\" ( WSP / VCHAR / obs-text )
quoted-str-nf = DQUOTE *( qdtext-nf / quoted-pair ) DQUOTE
quoted-string = DQUOTE *( qdtext / quoted-pair ) DQUOTE
qvalue = ( "0" [ "." *3DIGIT ] ) / ( "1" [ "." *3"0" ] )

received-by = ( uri-host [ ":" port ] ) / pseudonym
received-protocol = [ protocol-name "/" ] protocol-version
relative-part = <relative-part, defined in [RFC3986], Section 4.2>
request-header = <request-header, defined in [Part2], Section 3>
request-target = "*" / absolute-URI / ( path-absolute [ "?" query ] )
 / authority
response-header = <response-header, defined in [Part2], Section 5>
rfc1123-date = day-name "," SP date1 SP time-of-day SP GMT
rfc850-date = day-name-l "," SP date2 SP time-of-day SP GMT

second = 2DIGIT
special = "(" / ")" / "<" / ">" / "@" / "," / ";" / ":" / "\" /
 DQUOTE / "/" / "[" / "]" / "?" / "=" / "{" / "}"
start-line = Request-Line / Status-Line

t-codings = "trailers" / ( transfer-extension [ te-params ] )
tchar = "!" / "#" / "$" / "%" / "&" / "'" / "*" / "+" / "-" / "." /
 "^" / "_" / "`" / "|" / "~" / DIGIT / ALPHA
te-ext = OWS ";" OWS token [ "=" word ]
te-params = OWS ";" OWS "q=" qvalue *te-ext
time-of-day = hour ":" minute ":" second
token = 1*tchar
trailer-part = *( entity-header CRLF )
transfer-coding = "chunked" / "compress" / "deflate" / "gzip" /
 transfer-extension
transfer-extension = token *( OWS ";" OWS transfer-parameter )
transfer-parameter = attribute BWS "=" BWS value

uri-host = <host, defined in [RFC3986], Section 3.2.2>

value = word

word = token / quoted-string

year = 4DIGIT

ABNF diagnostics:
```

```
; Chunked-Body defined but not used
; Content-Length defined but not used
; HTTP-message defined but not used
; Host defined but not used
; Request defined but not used
; Response defined but not used
; TE defined but not used
; URI-reference defined but not used
; http-URI defined but not used
; https-URI defined but not used
; partial-URI defined but not used
; special defined but not used
```

---

### Appendix D.  Change Log (to be removed by RFC Editor before publication)     <span>TOC</span>

---

### D.1.  Since RFC2616     <span>TOC</span>

Extracted relevant partitions from [RFC2616] (Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1," June 1999.).

---

### D.2.  Since draft-ietf-httpbis-p1-messaging-00     <span>TOC</span>

Closed issues:

*http://tools.ietf.org/wg/httpbis/trac/ticket/1: "HTTP Version should be case sensitive" (http://purl.org/NET/http-errata#verscase)

*http://tools.ietf.org/wg/httpbis/trac/ticket/2: "'unsafe' characters" (http://purl.org/NET/http-errata#unsafe-uri)

*http://tools.ietf.org/wg/httpbis/trac/ticket/3: "Chunk Size Definition" (http://purl.org/NET/http-errata#chunk-size)

*http://tools.ietf.org/wg/httpbis/trac/ticket/4: "Message Length" (http://purl.org/NET/http-errata#msg-len-chars)

*http://tools.ietf.org/wg/httpbis/trac/ticket/8: "Media Type
 Registrations" (http://purl.org/NET/http-errata#media-reg)

*http://tools.ietf.org/wg/httpbis/trac/ticket/11: "URI includes
 query" (http://purl.org/NET/http-errata#uriquery)

*http://tools.ietf.org/wg/httpbis/trac/ticket/15: "No close on 1xx
 responses" (http://purl.org/NET/http-errata#noclose1xx)

*http://tools.ietf.org/wg/httpbis/trac/ticket/16: "Remove
 'identity' token references" (http://purl.org/NET/http-
 errata#identity)

*http://tools.ietf.org/wg/httpbis/trac/ticket/26: "Import query
 BNF"

*http://tools.ietf.org/wg/httpbis/trac/ticket/31: "qdtext BNF"

*http://tools.ietf.org/wg/httpbis/trac/ticket/35: "Normative and
 Informative references"

*http://tools.ietf.org/wg/httpbis/trac/ticket/42: "RFC2606
 Compliance"

*http://tools.ietf.org/wg/httpbis/trac/ticket/45: "RFC977
 reference"

*http://tools.ietf.org/wg/httpbis/trac/ticket/46: "RFC1700
 references"

*http://tools.ietf.org/wg/httpbis/trac/ticket/47: "inconsistency
 in date format explanation"

*http://tools.ietf.org/wg/httpbis/trac/ticket/48: "Date reference
 typo"

*http://tools.ietf.org/wg/httpbis/trac/ticket/65: "Informative
 references"

*http://tools.ietf.org/wg/httpbis/trac/ticket/66: "ISO-8859-1
 Reference"

*http://tools.ietf.org/wg/httpbis/trac/ticket/86: "Normative up-
 to-date references"

Other changes:

*Update media type registrations to use RFC4288 template.

*Use names of RFC4234 core rules DQUOTE and WSP, fix broken ABNF for chunk-data (work in progress on http://tools.ietf.org/wg/httpbis/trac/ticket/36)

---

### D.3.  Since draft-ietf-httpbis-p1-messaging-01

Closed issues:

*http://tools.ietf.org/wg/httpbis/trac/ticket/19: "Bodies on GET (and other) requests"

*http://tools.ietf.org/wg/httpbis/trac/ticket/55: "Updating to RFC4288"

*http://tools.ietf.org/wg/httpbis/trac/ticket/57: "Status Code and Reason Phrase"

*http://tools.ietf.org/wg/httpbis/trac/ticket/82: "rel_path not used"

Ongoing work on ABNF conversion (http://tools.ietf.org/wg/httpbis/trac/ticket/36):

*Get rid of duplicate BNF rule names ("host" -> "uri-host", "trailer" -> "trailer-part").

*Avoid underscore character in rule names ("http_URL" -> "http-URL", "abs_path" -> "path-absolute").

*Add rules for terms imported from URI spec ("absoluteURI", "authority", "path-absolute", "port", "query", "relativeURI", "host) -- these will have to be updated when switching over to RFC3986.

*Synchronize core rules with RFC5234.

*Get rid of prose rules that span multiple lines.

*Get rid of unused rules LOALPHA and UPALPHA.

*Move "Product Tokens" section (back) into Part 1, as "token" is used in the definition of the Upgrade header.

*Add explicit references to BNF syntax and rules imported from other parts of the specification.

     *Rewrite prose rule "token" in terms of "tchar", rewrite prose
      rule "TEXT".

---

### D.4.  Since draft-ietf-httpbis-p1-messaging-02

Closed issues:

     *http://tools.ietf.org/wg/httpbis/trac/ticket/51: "HTTP-date vs.
      rfc1123-date"

     *http://tools.ietf.org/wg/httpbis/trac/ticket/64: "WS in quoted-
      pair"

Ongoing work on IANA Message Header Registration (http://
tools.ietf.org/wg/httpbis/trac/ticket/40):

     *Reference RFC 3984, and update header registrations for headers
      defined in this document.

Ongoing work on ABNF conversion (http://tools.ietf.org/wg/httpbis/trac/
ticket/36):

     *Replace string literals when the string really is case-sensitive
      (HTTP-Version).

---

### D.5.  Since draft-ietf-httpbis-p1-messaging-03

Closed issues:

     *http://tools.ietf.org/wg/httpbis/trac/ticket/28: "Connection
      closing"

     *http://tools.ietf.org/wg/httpbis/trac/ticket/97: "Move
      registrations and registry information to IANA Considerations"

     *http://tools.ietf.org/wg/httpbis/trac/ticket/120: "need new URL
      for PAD1995 reference"

     *http://tools.ietf.org/wg/httpbis/trac/ticket/127: "IANA
      Considerations: update HTTP URI scheme registration"

     *http://tools.ietf.org/wg/httpbis/trac/ticket/128: "Cite HTTPS URI
      scheme definition"

*[http://tools.ietf.org/wg/httpbis/trac/ticket/129](http://tools.ietf.org/wg/httpbis/trac/ticket/129): "List-type headers vs Set-Cookie"

Ongoing work on ABNF conversion ([http://tools.ietf.org/wg/httpbis/trac/ticket/36](http://tools.ietf.org/wg/httpbis/trac/ticket/36)):

*Replace string literals when the string really is case-sensitive (HTTP-Date).

*Replace HEX by HEXDIG for future consistence with RFC 5234's core rules.

---

### D.6.  Since draft-ietf-httpbis-p1-messaging-04

Closed issues:

*[http://tools.ietf.org/wg/httpbis/trac/ticket/34](http://tools.ietf.org/wg/httpbis/trac/ticket/34): "Out-of-date reference for URIs"

*[http://tools.ietf.org/wg/httpbis/trac/ticket/132](http://tools.ietf.org/wg/httpbis/trac/ticket/132): "RFC 2822 is updated by RFC 5322"

Ongoing work on ABNF conversion ([http://tools.ietf.org/wg/httpbis/trac/ticket/36](http://tools.ietf.org/wg/httpbis/trac/ticket/36)):

*Use "/" instead of "|" for alternatives.

*Get rid of RFC822 dependency; use RFC5234 plus extensions instead.

*Only reference RFC 5234's core rules.

*Introduce new ABNF rules for "bad" whitespace ("BWS"), optional whitespace ("OWS") and required whitespace ("RWS").

*Rewrite ABNFs to spell out whitespace rules, factor out header value format definitions.

---

### D.7.  Since draft-ietf-httpbis-p1-messaging-05

Closed issues:

*[http://tools.ietf.org/wg/httpbis/trac/ticket/30](http://tools.ietf.org/wg/httpbis/trac/ticket/30): "Header LWS"

*[http://tools.ietf.org/wg/httpbis/trac/ticket/52](http://tools.ietf.org/wg/httpbis/trac/ticket/52): "Sort 1.3
    Terminology"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/63](http://tools.ietf.org/wg/httpbis/trac/ticket/63): "RFC2047 encoded
    words"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/74](http://tools.ietf.org/wg/httpbis/trac/ticket/74): "Character
    Encodings in TEXT"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/77](http://tools.ietf.org/wg/httpbis/trac/ticket/77): "Line Folding"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/83](http://tools.ietf.org/wg/httpbis/trac/ticket/83): "OPTIONS * and
    proxies"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/94](http://tools.ietf.org/wg/httpbis/trac/ticket/94): "Reason-Phrase
    BNF"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/111](http://tools.ietf.org/wg/httpbis/trac/ticket/111): "Use of TEXT"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/118](http://tools.ietf.org/wg/httpbis/trac/ticket/118): "Join
    "Differences Between HTTP Entities and RFC 2045 Entities"?"

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/134](http://tools.ietf.org/wg/httpbis/trac/ticket/134): "RFC822
    reference left in discussion of date formats"

Final work on ABNF conversion ([http://tools.ietf.org/wg/httpbis/trac/ticket/36](http://tools.ietf.org/wg/httpbis/trac/ticket/36)):

   *Rewrite definition of list rules, deprecate empty list elements.

   *Add appendix containing collected and expanded ABNF.

Other changes:

   *Rewrite introduction; add mostly new Architecture Section.

   *Move definition of quality values from Part 3 into Part 1; make
    TE request header grammar independent of accept-params (defined
    in Part 3).

---

### D.8.  Since draft-ietf-httpbis-p1-messaging-06                    TOC

Closed issues:

   *[http://tools.ietf.org/wg/httpbis/trac/ticket/161](http://tools.ietf.org/wg/httpbis/trac/ticket/161): "base for
    numeric protocol elements"

* http://tools.ietf.org/wg/httpbis/trac/ticket/162: "comment ABNF"

Partly resolved issues:

* http://tools.ietf.org/wg/httpbis/trac/ticket/88: "205 Bodies"
   (took out language that implied that there may be methods for
   which a request body MUST NOT be included)

* http://tools.ietf.org/wg/httpbis/trac/ticket/163: "editorial
   improvements around HTTP-date"

---

### D.9.  Since draft-ietf-httpbis-p1-messaging-07

Closed issues:

* http://tools.ietf.org/wg/httpbis/trac/ticket/93: "Repeating
   single-value headers"

* http://tools.ietf.org/wg/httpbis/trac/ticket/131: "increase
   connection limit"

* http://tools.ietf.org/wg/httpbis/trac/ticket/157: "IP addresses
   in URLs"

* http://tools.ietf.org/wg/httpbis/trac/ticket/172: "take over HTTP
   Upgrade Token Registry"

* http://tools.ietf.org/wg/httpbis/trac/ticket/173: "CR and LF in
   chunk extension values"

* http://tools.ietf.org/wg/httpbis/trac/ticket/184: "HTTP/0.9
   support"

* http://tools.ietf.org/wg/httpbis/trac/ticket/188: "pick IANA
   policy (RFC5226) for Transfer Coding / Content Coding"

* http://tools.ietf.org/wg/httpbis/trac/ticket/189: "move
   definitions of gzip/deflate/compress to part 1"

* http://tools.ietf.org/wg/httpbis/trac/ticket/194: "disallow
   control characters in quoted-pair"

Partly resolved issues:

* http://tools.ietf.org/wg/httpbis/trac/ticket/148: "update IANA
   requirements wrt Transfer-Coding values" (add the IANA
   Considerations subsection)

## D.10. Since draft-ietf-httpbis-p1-messaging-08

Closed issues:

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/201: "header parsing, treatment of leading and trailing OWS"

Partly resolved issues:

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/60: "Placement of 13.5.1 and 13.5.2"

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/200: "use of term "word" when talking about header structure"

## D.11. Since draft-ietf-httpbis-p1-messaging-09

Closed issues:

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/73: "Clarification of the term 'deflate'"

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/83: "OPTIONS * and proxies"

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/143: "IANA registry for content/transfer encodings"

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/165: "Case-sensitivity of HTTP-date"

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/200: "use of term "word" when talking about header structure"

Partly resolved issues:

*   *http://tools.ietf.org/wg/httpbis/trac/ticket/196: "Term for the requested resource's URI"

**Index**

---

**Authors' Addresses**

| | |
|---|---|
| | Roy T. Fielding (editor) |
| | Day Software |
| | 23 Corporate Plaza DR, Suite 280 |
| | Newport Beach, CA 92660 |
| | USA |
| Phone: | +1-949-706-5300 |
| Fax: | +1-949-706-5305 |
| EMail: | fielding@gbiv.com |
| URI: | http://roy.gbiv.com/ |
| | |
| | Jim Gettys |
| | Alcatel-Lucent Bell Labs |
| | 21 Oak Knoll Road |
| | Carlisle, MA 01741 |
| | USA |
| EMail: | jg@freedesktop.org |
| URI: | http://gettys.wordpress.com/ |
| | |
| | Jeffrey C. Mogul |
| | Hewlett-Packard Company |
| | HP Labs, Large Scale Systems Group |
| | 1501 Page Mill Road, MS 1177 |
| | Palo Alto, CA 94304 |
| | USA |
| EMail: | JeffMogul@acm.org |
| | |
| | Henrik Frystyk Nielsen |
| | Microsoft Corporation |
| | 1 Microsoft Way |

```
             Redmond, WA 98052
             USA
   EMail: henrikn@microsoft.com


             Larry Masinter
             Adobe Systems, Incorporated
             345 Park Ave
             San Jose, CA 95110
             USA
   EMail: LMM@acm.org
     URI: http://larry.masinter.net/


             Paul J. Leach
             Microsoft Corporation
             1 Microsoft Way
             Redmond, WA 98052
   EMail: paulle@microsoft.com


             Tim Berners-Lee
             World Wide Web Consortium
             MIT Computer Science and Artificial Intelligence Laboratory
             The Stata Center, Building 32
             32 Vassar Street
             Cambridge, MA 02139
             USA
   EMail: timbl@w3.org
     URI: http://www.w3.org/People/Berners-Lee/


             Yves Lafon (editor)
             World Wide Web Consortium
             W3C / ERCIM
             2004, rte des Lucioles
             Sophia-Antipolis, AM 06902
             France
   EMail: ylafon@w3.org
     URI: http://www.raubacapeu.net/people/yves/


             Julian F. Reschke (editor)
             greenbytes GmbH
             Hafenweg 16
             Muenster, NW 48155
             Germany
   Phone: +49 251 2807760
     Fax: +49 251 2807761
   EMail: julian.reschke@greenbytes.de
     URI: http://greenbytes.de/tech/webdav/
```