

Click to prove
you're human



[illegible]

Use Requests for Comments (RFCs) started a few years later in a coordinated effort by the Internet Engineering Task Force (IETF) and the World Wide Web Consortium (W3C), with work later moving to the IETF. HTTP/1.6 was finalized and fully documented (as version 1.0) in 1996.[41] It evolved (as version 1.1) in 1997 and then its specifications were updated in 1999, 2014, and 2022.[5] Its security variants, such as HTTPS, is used by more than 85% of web browsers. HTTP/2, published in 2015, updates the expression of HTTP's semantics "on a binary level".[42] As of August 2024,[update] it is supported by 66.2% of websites[118] (35.3% of HTTP/2 + 6 HTTP/2.0 backends compatibility) and supported by almost all web browsers (over 98% of users).[9] It is also supported by major web servers over Transport Layer Security (TLS) using an Application-Layer Protocol Negotiation (ALPN) extension[10] where TLS 1.2 or newer is required.[11][12] HTTP/3, the successor to HTTP/2, was published in 2022.[13] As of February 2024,[update] it is now used on 30.9% of websites[14] and is supported by most web browsers, i.e. (at least partially) supported by 97% of users.[15] HTTP/3 uses QUIC instead of TCP for the underlying transport protocol. Like HTTP/2, it does not obsolete previous major versions of the protocol. Support for HTTP/3 was added to Cloudflare and Google Chrome first.[16][17] and is also enabled in Firefox.[18] HTTP/3 has lower latency for real-world web pages, if enabled on the server, and loads faster than with HTTP/2, in some cases over three times faster than HTTP/1.1 (which is still commonly only enabled).[19] HTTP functions as a request-response protocol in the client-server model. A web browser, for example, may be the client whereas a process, named web server, running on a computer hosting one or more websites may be the server. The client submits an HTTP request message to the server. The server, which provides resources such as HTML files and other content or performs other functions on behalf of the client, returns a response message to the client. The response contains completion status information about the request and may also contain requested content in its message body. A web browser is an example of a user agent (UA). Other types of user agent include the indexing software used by search providers (web crawlers), voice browsers, mobile apps, and other software that accesses, consumes, or displays web content. HTTP is designed to permit intermediate network elements to improve or enable communications between clients and servers. High-traffic websites often benefit from web cache servers that deliver content on behalf of upstream servers to improve response time. Web brokers cache previously accessed web resources and return them whenever possible, to reduce network traffic. HTTP proxy servers at private network boundaries facilitate communication for clients without a globally routable address, by relaying messages with external servers. To allow intermediate HTTP nodes (proxy servers, web caches, etc.) to accomplish their functions, some of the HTTP headers (found in HTTP requests/responses) are managed hop-by-hop whereas other HTTP headers are managed end-to-end (managed only by the source client and by the target web server). HTTP is an application layer protocol designed within the framework of the Internet protocol suite. Its definition presumes an underlying and reliable transport layer protocol.[20] In HTTP/3, the Transmission Control Protocol (TCP) is no longer used, but the older versions are still more used and they most commonly use TCP. They have also been adapted to use unreliable protocols such as the User Datagram Protocol (UDP), which HTTP/3 also (indirectly) always builds on, for example in HTTP/1 and Simple Service Discovery Protocol (SSDP). HTTP resources are identified and located on the network by Uniform Resource Locators (URLs), using the Uniform Resource Identifiers (URIs) schemes http and https. As defined in RFC 3986, URIs are encoded as hyperlinks in HTML documents, so as to form interlinked hypertext documents. In HTTP/1.0 a separate TCP connection to the same server is made for every resource request.[21] In HTTP/1.1 instead a TCP connection can be reused to make multiple resource requests (i.e. of HTML pages, frames, images, scripts, stylesheets, etc.).[22][23] HTTP/1.1 communications therefore experience less latency as the establishment of TCP connections presents considerable overhead, especially under high traffic conditions.[24] HTTP/2 is a revision of previous HTTP/1.1 in order to maintain the same client-server model and the same protocol methods but with these differences in order: to use a compressed binary representation of metadata (HTTP headers) instead of a textual one, so that headers require much less space; to use a single TCP/IP (usually encrypted) connection per accessed server domain instead of 2 to 8 TCP/IP connections; to use one or more bidirectional streams per TCP/IP connection in which HTTP requests and responses are broken down and transmitted in small packets to almost solve the problem of the H2B (head-of-line blocking) (note 1) to add a push capability to allow server application to send data to clients whenever new data is available (without forcing clients to request periodically new data to server by using polling methods).[25] HTTP/2 communications therefore experience much less latency and, in most cases, even higher speeds than HTTP/1.1 communications. HTTP/3 is a revision of previous HTTP/2 in order to use QUIC + UDP transport protocols instead of TCP. Before that version, TCP/IP connections were used; but now, only the IP layer is used (which UDP, like TCP, builds on). This slightly improves the average speed of communications and to avoid the occasional (very rare) problem of TCP connection congestion that can temporarily block or slow down the data flow of all its streams (another form of "head of line blocking"). Tim Berners-Lee The term *hypertext* was coined by Ted Nelson in 1965 in the Xanadu Project, which was in turn inspired by Vannevar Bush's 1930s vision of the microfilm-based information retrieval and management "memex" system described in his 1945 essay "As We May Think". Tim Berners-Lee and his team at CERN are credited with inventing the original HTTP, along with HTML and the associated technology for a web server and a client user interface called web browser. Berners-Lee designed HTTP in order to help with the adoption of his other idea: the "WorldWideWeb" project, which was first proposed in 1989, now known as the World Wide Web. The first web server went live in 1990.[26][27] The protocol used had only one method, namely GET, which would request a page from a server.[28] The response from the server was always an HTML page.[2] Version Year introduced Current status Usage in August 2024[update] Support in August 2024[update] HTTP/0.9 1991 Obsolete 0 100% HTTP/1.0 1996 Obsolete 0 100% HTTP/1.1 1997 Standard 33.8% 100% HTTP/2 2015 Standard 35.3% 66.2% HTTP/3 2022 Standard 30.9% 30.9% In 1991, the first documented official version of HTTP was written as a plain document, less than 700 words long, and this version was named HTTP/0.9, which supported only GET method, allowing clients to only retrieve HTML documents from the server, but not supporting any other file formats or information upload.[2] Since 1992, a new document was written to specify the evolution of the basic protocol toward its next full version. It supported both the simple request method of the 0.9 version and the full GET request that included the client HTTP version. This was the first of the many unofficial HTTP/1.0 drafts that preceded the final work on HTTP/1.0.[3] After having decided that new features of HTTP protocol were required and that they had to be fully documented as official RFCs, in early 1995 the HTTP Working Group (HTTP WG, led by Dave Raggett) was constituted with the aim to standardize and expand the protocol with extended operations, extended negotiation, richer meta-information, tied with a security protocol which became more efficient by adding additional methods and header fields.[29][30] The HTTP WG planned to revise and publish new versions of the protocol as HTTP/1.0 and HTTP/1.1 within 1995, but, because of the many revisions, that timeline lasted much more than one year.[31] The HTTP WG planned also to specify a far future version of HTTP called HTTP-NG (HTTP Next Generation) that would have solved all remaining problems, of previous versions, related to performances, low latency responses, etc. but this work started only a few years later and it was never completed. In May 1996, RFC 1945 was published as a final HTTP/1.0 revision of what had been used in previous 4 years as a pre-standard HTTP/1.0-draft which was already used by many web browsers and web servers. In early 1996 developers started to even include unofficial extensions of the HTTP/1.0 protocol (i.e. keep-alive connections, etc.) into their products by using drafts of the upcoming HTTP/1.1 specifications.[32] Since early 1996, major web browsers and web server developers also started to implement new features specified by pre-standard HTTP/1.1 drafts specifications. End-user adoption of the new versions of browsers and servers was rapid. In March 1996, one web hosting company reported that over 40% of browsers in use on the Internet used the new HTTP/1.1 header "Host" to enable virtual hosting, and that by June 1996, 65% of all browsers accessing their servers were pre-standard HTTP/1.1 compliant.[33] In January 1997, RFC 2068 was officially released as HTTP/1.1 specifications. In June 1999, RFC 2616 was released to include all improvements and updates based on previous (obsolete) HTTP/1.1 specifications. Resuming the old 1995 plan of previous HTTP Working Group, in 1997 an HTTP-NG Working Group was formed to develop a new HTTP protocol named HTTP-NG (HTTP New Generation). A few proposals / drafts were produced for the new protocol to use multiplexing of HTTP transactions inside single TCP/IP connection, but in 1999 the group stopped its activity passing the technical problem to IETF.[34] In 2007, the IETF HTTP Working Group (HTTP WG bis or HTTPbis) was restarted firstly to revise and clarify previous HTTP/1.1 specifications and secondly to write and refine future HTTP/2 specifications (named httpbis)[35][36] In 2009, Google, a private company, announced that it had developed and tested a new HTTP binary protocol named SPDY. The implicit aim was to greatly speed up web traffic (specially between future web browsers and its servers). SPDY was indeed much faster than HTTP/1.1 in many tests and so it was quickly adopted by Chromium and then by other major web browsers.[37] Some of the ideas about multiplexing HTTP streams over a single TCP/IP connection were taken from various sources, including the work of W3C HTTP-NG Working Group. In January-March 2011, HTTP Working Group (HTTPbis) announced the need to start to focus on a new HTTP/2 protocol (while finishing the revision of HTTP/1.1 specifications), maybe taking in consideration ideas and work done for SPDY.[38][39] After a few months about what to do to develop a new version of HTTP, it was decided to derive it from SPDY.[40] In May 2015, HTTP/2 was published as RFC 7540 and quickly adopted by all web browsers already supporting SPDY and more slowly by web servers. In June 2014, the HTTP Working Group released an updated six-part HTTP/1.1 specification obsoleting RFC 2616: RFC 7230, HTTP/1.1: Message Syntax and Routing RFC 7231, HTTP/1.1: Semantics and Content RFC 7232, HTTP/1.1: Conditional Requests RFC 7233, HTTP/1.1: Range Requests RFC 7234, HTTP/1.1: Caching RFC 7235, HTTP/1.1: Authentication In RFC 7230 Appendix-A, HTTP/0.9 was deprecated for servers supporting HTTP/1.1 version (and higher).[41] Since HTTP/0.9 did not support header fields in a request, there is no mechanism for it to support name-based virtual hosts (selection of resource by inspection of the Host header field). Any server that implements name-based virtual hosts ought to disable support for HTTP/0.9. Most requests that appear to be HTTP/0.9 are, in fact, badly constructed HTTP/1.x requests caused by a client failing to properly encode the request-target. Since 2016 many product managers and developers of user agents (browsers, etc.) and web servers have begun planning to gradually deprecate and dismiss support for HTTP/0.9 protocol, mainly for the following reasons:[42] It is so simple that an RFC document was never written (there is only the original document).[42] it has no HTTP headers and lacks many other features that nowadays are required for minimal security reasons; it has not been widespread since 1999..2000 (because of HTTP/1.0 and HTTP/1.1) and is commonly used only by some very old network hardware, i.e. routers, etc. [note 2] In 2020, the first drafts HTTP/3 were published and major web browsers and web servers started to adopt it. On 6 June 2022, IETF standardized HTTP/3 as RFC 9114.[43] In June 2022, a batch of RFCs was published, deprecating many of the previous documents and introducing a few minor changes and a refactoring of HTTP semantics description into a separate document. RFC 9110, HTTP Semantics RFC 9111, HTTP Caching RFC 9112, HTTP/1.1 RFC 9113, HTTP/2 RFC 9114, HTTP/3 (see also the section above) RFC 9204, OPAQ: Field Compression for HTTP/3 RFC 9218, Extensible Prioritization Scheme for HTTP HTTP/1.1 is a stateless application-level protocol and it requires a reliable network transport connection to exchange data between client and server.[20] In HTTP implementations, TCP/IP connections are used using well-known ports (typically port 80 if the connection is unencrypted or port 443 if the connection is encrypted, see also List of TCP and UDP port numbers).[44][45] In HTTP/2, a TCP/IP connection plus multiple protocol channels are used. In HTTP/3, the application transport protocol QUIC over UDP is used. Data is exchanged through a sequence of request-response messages which are exchanged by a session layer transport connection.[20] An HTTP client initially tries to connect to a server establishing a connection (real or virtual). As HTTP/3 server listening on that port accept the connection and then waits for a client's request message. The client sends its HTTP request message. Upon receiving the request the server sends back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[22] Main article: HTTP persistent connection In HTTP/0.9, the TCP/IP connection is always closed after server response has been sent, so it is never persistent. In HTTP/1.0, as stated in RFC 1945, the TCP/IP connection should always be closed by server after a response has been sent.[note 3] In HTTP/1.1 a keep-alive-mechanism was officially introduced so that a connection could be reused for more than one request/response. Such persistent connections reduce request latency perceptibly because the client does not need to re-negotiate the TCP-3-Way-Handshake connection after the first request has been sent. Another positive side effect is that, in general, the connection becomes faster with time due to TCP's slow-start-mechanism. HTTP/1.1 added also HTTP pipelining in order to further reduce lag time when using persistent connections by allowing clients to send multiple requests before waiting for each response. This optimization was never considered really safe because a few web servers and many proxy servers, specially transparent proxy servers placed in Internet / Intranets between clients and servers, did not handle pipelined requests properly (they served only the first request discarding the others, they closed the connection because they saw more data after the first request or some proxies even returned responses out of order etc.). Because of this, only HEAD and some GET requests (i.e. limited to real file requests and so with URLs without query string used as a command, etc.) could be pipelined in a safe and idempotent method. After many years of struggling with the problems introduced by enabling pipelining, this feature was first disabled and then removed from most browsers also because of the announced adoption of HTTP/2. HTTP/2 extended the concept of persistent connections by multiplexing many concurrent requests/responses through a single TCP/IP connection. The request message is sent to the server and the server responds back an HTTP response message, which includes headers(s) plus a body if it is required. The body of this response message is typically the requested resource, although an error message or other information may also be returned. At any time (for many reasons) client or server can close the connection. Closing a connection is usually advertised in advance by using one or more HTTP headers in the last request/response message sent to server or client.[2