


Understanding Web Internals

HTTP

The Definitive Guide



O'REILLY®

*David Gourley & Brian Totty
with Marjorie Sayer, Sailu Reddy & Anshu Aggarwal*

HTTP

The Definitive Guide

David Gourley and Brian Totty
with Marjorie Sayer, Sailu Reddy, and Anshu Aggarwal

O'REILLY®

Beijing • Boston • Farnham • Sebastopol • Tokyo

HTTP: The Definitive Guide

by David Gourley and Brian Totty
with Marjorie Sayer, Sailu Reddy, and Anshu Aggarwal

Copyright © 2002 O'Reilly Media, Inc. All rights reserved.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol,
CA 95472.

O'Reilly Media, Inc. books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (safari.oreilly.com). For more information, contact our corporate/institutional sales department: (800) 998-9938 or corporate@oreilly.com.

Editor: Linda Mui
Production Editor: Rachel Wheeler
Cover Designer: Ellie Volckhausen
Interior Designers: David Futato and Melanie Wang

Printing History:

September 2002: First Edition.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc. *HTTP: The Definitive Guide*, the image of a thirteen-lined ground squirrel, and related trade dress are trademarks of O'Reilly Media, Inc. Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps.

While every precaution has been taken in the preparation of this book, the publisher and authors assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

ISBN-10: 1-56592-509-2
ISBN-13: 978-1-56592-509-0
[LSI]

[03/2019]

Example 1-1. An HTTP transaction using telnet (continued)

```
<HTML>
<HEAD><TITLE>Joe's Tools</TITLE></HEAD>
<BODY>
<H1>Tools Page</H1>
<H2>Hammers</H2>
<P>Joe's Hardware Online has the largest selection of hammers on the earth.</P>
<H2><A NAME=drills></A>Drills</H2>
<P>Joe's Hardware has a complete line of cordless and corded drills, as well as the latest
in plutonium-powered atomic drills, for those big around the house jobs.</P> ...
</BODY>
</HTML>
Connection closed by foreign host.
```

Telnet looks up the hostname and opens a connection to the *www.joes-hardware.com* web server, which is listening on port 80. The three lines after the command are output from Telnet, telling us it has established a connection.

We then type in our basic request command, “GET /tools.html HTTP/1.1”, and send a Host header providing the original hostname, followed by a blank line, asking the server to GET us the resource “/tools.html” from the server *www.joes-hardware.com*. After that, the server responds with a response line, several response headers, a blank line, and finally the body of the HTML document.

Beware that Telnet mimics HTTP clients well but doesn't work well as a server. And automated Telnet scripting is no fun at all. For a more flexible tool, you might want to check out nc (netcat). The nc tool lets you easily manipulate and script UDP- and TCP-based traffic, including HTTP. See <http://netcat.sourceforge.net> for details.

Protocol Versions

There are several versions of the HTTP protocol in use today. HTTP applications need to work hard to robustly handle different variations of the HTTP protocol. The versions in use are:

HTTP/0.9

The 1991 prototype version of HTTP is known as HTTP/0.9. This protocol contains many serious design flaws and should be used only to interoperate with legacy clients. HTTP/0.9 supports only the GET method, and it does not support MIME typing of multimedia content, HTTP headers, or version numbers. HTTP/0.9 was originally defined to fetch simple HTML objects. It was soon replaced with HTTP/1.0.

HTTP/1.0

1.0 was the first version of HTTP that was widely deployed. HTTP/1.0 added version numbers, HTTP headers, additional methods, and multimedia object handling. HTTP/1.0 made it practical to support graphically appealing web

pages and interactive forms, which helped promote the wide-scale adoption of the World Wide Web. This specification was never well specified. It represented a collection of best practices in a time of rapid commercial and academic evolution of the protocol.

HTTP/1.0+

Many popular web clients and servers rapidly added features to HTTP in the mid-1990s to meet the demands of a rapidly expanding, commercially successful World Wide Web. Many of these features, including long-lasting “keep-alive” connections, virtual hosting support, and proxy connection support, were added to HTTP and became unofficial, *de facto* standards. This informal, extended version of HTTP is often referred to as HTTP/1.0+.

HTTP/1.1

HTTP/1.1 focused on correcting architectural flaws in the design of HTTP, specifying semantics, introducing significant performance optimizations, and removing mis-features. HTTP/1.1 also included support for the more sophisticated web applications and deployments that were under way in the late 1990s. HTTP/1.1 is the current version of HTTP.

HTTP-NG (a.k.a. HTTP/2.0)

HTTP-NG is a prototype proposal for an architectural successor to HTTP/1.1 that focuses on significant performance optimizations and a more powerful framework for remote execution of server logic. The HTTP-NG research effort concluded in 1998, and at the time of this writing, there are no plans to advance this proposal as a replacement for HTTP/1.1. See Chapter 10 for more information.

Architectural Components of the Web

In this overview chapter, we’ve focused on how two web applications (web browsers and web servers) send messages back and forth to implement basic transactions. There are many other web applications that you interact with on the Internet. In this section, we’ll outline several other important applications, including:

Proxies

HTTP intermediaries that sit between clients and servers

Caches

HTTP storehouses that keep copies of popular web pages close to clients

Gateways

Special web servers that connect to other applications

Tunnels

Special proxies that blindly forward HTTP communications

Agents

Semi-intelligent web clients that make automated HTTP requests

making optimizations, and your server doesn't get faster than about 500 transactions/sec, make sure you are not experiencing `TIME_WAIT` port exhaustion. You can fix this problem by using more client load-generator machines or making sure the client and server rotate through several virtual IP addresses to add more connection combinations.

Even if you do not suffer port exhaustion problems, be careful about having large numbers of open connections or large numbers of control blocks allocated for connection in wait states. Some operating systems slow down dramatically when there are numerous open connections or control blocks.

HTTP Connection Handling

The first two sections of this chapter provided a fire-hose tour of TCP connections and their performance implications. If you'd like to learn more about TCP networking, check out the resources listed at the end of the chapter.

We're going to switch gears now and get squarely back to HTTP. The rest of this chapter explains the HTTP technology for manipulating and optimizing connections. We'll start with the HTTP Connection header, an often misunderstood but important part of HTTP connection management. Then we'll talk about HTTP's connection optimization techniques.

The Oft-Misunderstood Connection Header

HTTP allows a chain of HTTP intermediaries between the client and the ultimate origin server (proxies, caches, etc.). HTTP messages are forwarded hop by hop from the client, through intermediary devices, to the origin server (or the reverse).

In some cases, two adjacent HTTP applications may want to apply a set of options to their shared connection. The HTTP Connection header field has a comma-separated list of *connection tokens* that specify options for the connection that aren't propagated to other connections. For example, a connection that must be closed after sending the next message can be indicated by `Connection: close`.

The Connection header sometimes is confusing, because it can carry three different types of tokens:

- HTTP header field names, listing headers relevant for only this connection
- Arbitrary token values, describing nonstandard options for this connection
- The value `close`, indicating the persistent connection will be closed when done

If a connection token contains the name of an HTTP header field, that header field contains connection-specific information and must not be forwarded. Any header fields listed in the Connection header must be deleted before the message is forwarded. Placing a hop-by-hop header name in a Connection header is known as

“protecting the header,” because the Connection header protects against accidental forwarding of the local header. An example is shown in Figure 4-9.




Figure 4-9. The Connection header allows the sender to specify connection-specific options

When an HTTP application receives a message with a Connection header, the receiver parses and applies all options requested by the sender. It then deletes the Connection header and all headers listed in the Connection header before forwarding the message to the next hop. In addition, there are a few hop-by-hop headers that might not be listed as values of a Connection header, but must not be proxied. These include Proxy-Authenticate, Proxy-Connection, Transfer-Encoding, and Upgrade. For more about the Connection header, see Appendix C.

Serial Transaction Delays

TCP performance delays can add up if the connections are managed naively. For example, suppose you have a web page with three embedded images. Your browser needs to issue four HTTP transactions to display this page: one for the top-level HTML and three for the embedded images. If each transaction requires a new connection, the connection and slow-start delays can add up (see Figure 4-10).*




Figure 4-10. Four transactions (serial)

* For the purpose of this example, assume all objects are roughly the same size and are hosted from the same server, and that the DNS entry is cached, eliminating the DNS lookup time.

In addition to the real delay imposed by serial loading, there is also a psychological perception of slowness when a single image is loading and nothing is happening on the rest of the page. Users prefer multiple images to load at the same time.*

Another disadvantage of serial loading is that some browsers are unable to display anything onscreen until enough objects are loaded, because they don't know the sizes of the objects until they are loaded, and they may need the size information to decide where to position the objects on the screen. In this situation, the browser may be making good progress loading objects serially, but the user may be faced with a blank white screen, unaware that any progress is being made at all.†

Several current and emerging techniques are available to improve HTTP connection performance. The next several sections discuss four such techniques:

Parallel connections

Concurrent HTTP requests across multiple TCP connections

Persistent connections

Reusing TCP connections to eliminate connect/close delays

Pipelined connections

Concurrent HTTP requests across a shared TCP connection

Multiplexed connections

Interleaving chunks of requests and responses (experimental)

Parallel Connections

As we mentioned previously, a browser could naively process each embedded object serially by completely requesting the original HTML page, then the first embedded object, then the second embedded object, etc. But this is too slow!

HTTP allows clients to open multiple connections and perform multiple HTTP transactions in parallel, as sketched in Figure 4-11. In this example, four embedded images are loaded in parallel, with each transaction getting its own TCP connection.‡

Parallel Connections May Make Pages Load Faster

Composite pages consisting of embedded objects may load faster if they take advantage of the dead time and bandwidth limits of a single connection. The delays can be

* This is true even if loading multiple images at the same time is *slower* than loading images one at a time! Users often perceive multiple-image loading as faster.

† HTML designers can help eliminate this “layout delay” by explicitly adding width and height attributes to HTML tags for embedded objects such as images. Explicitly providing the width and height of the embedded image allows the browser to make graphical layout decisions before it receives the objects from the server.

‡ The embedded components do not all need to be hosted on the same web server, so the parallel connections can be established to multiple servers.




Figure 4-11. Each component of a page involves a separate HTTP transaction

overlapped, and if a single connection does not saturate the client's Internet bandwidth, the unused bandwidth can be allocated to loading additional objects.

Figure 4-12 shows a timeline for parallel connections, which is significantly faster than Figure 4-10. The enclosing HTML page is loaded first, and then the remaining three transactions are processed concurrently, each with their own connection.* Because the images are loaded in parallel, the connection delays are overlapped.




Figure 4-12. Four transactions (parallel)

Parallel Connections Are Not Always Faster

Even though parallel connections may be faster, however, they are not *always* faster. When the client's network bandwidth is scarce (for example, a browser connected to

* There will generally still be a small delay between each connection request due to software overheads, but the connection requests and transfer times are mostly overlapped.

the Internet through a 28.8-Kbps modem), most of the time might be spent just transferring data. In this situation, a single HTTP transaction to a fast server could easily consume all of the available modem bandwidth. If multiple objects are loaded in parallel, each object will just compete for this limited bandwidth, so each object will load proportionally slower, yielding little or no performance advantage.*

Also, a large number of open connections can consume a lot of memory and cause performance problems of their own. Complex web pages may have tens or hundreds of embedded objects. Clients might be able to open hundreds of connections, but few web servers will want to do that, because they often are processing requests for many other users at the same time. A hundred simultaneous users, each opening 100 connections, will put the burden of 10,000 connections on the server. This can cause significant server slowdown. The same situation is true for high-load proxies.

In practice, browsers do use parallel connections, but they limit the total number of parallel connections to a small number (often four). Servers are free to close excessive connections from a particular client.

Parallel Connections May “Feel” Faster

Okay, so parallel connections don’t always make pages load faster. But even if they don’t actually speed up the page transfer, as we said earlier, parallel connections often make users *feel* that the page loads faster, because they can see progress being made as multiple component objects appear onscreen in parallel.† Human beings perceive that web pages load faster if there’s lots of action all over the screen, even if a stopwatch actually shows the aggregate page download time to be slower!

Persistent Connections

Web clients often open connections to the same site. For example, most of the embedded images in a web page often come from the same web site, and a significant number of hyperlinks to other objects often point to the same site. Thus, an application that initiates an HTTP request to a server likely will make more requests to that server in the near future (to fetch the inline images, for example). This property is called *site locality*.

For this reason, HTTP/1.1 (and enhanced versions of HTTP/1.0) allows HTTP devices to keep TCP connections open after transactions complete and to reuse the preexisting connections for future HTTP requests. TCP connections that are kept

* In fact, because of the extra overhead from multiple connections, it’s quite possible that parallel connections could take longer to load the entire page than serial downloads.

† This effect is amplified by the increasing use of progressive images that produce low-resolution approximations of images first and gradually increase the resolution.

open after transactions complete are called *persistent* connections. Nonpersistent connections are closed after each transaction. Persistent connections stay open across transactions, until either the client or the server decides to close them.

By reusing an idle, persistent connection that is already open to the target server, you can avoid the slow connection setup. In addition, the already open connection can avoid the slow-start congestion adaptation phase, allowing faster data transfers.

Persistent Versus Parallel Connections

As we've seen, parallel connections can speed up the transfer of composite pages. But parallel connections have some disadvantages:

- Each transaction opens/closes a new connection, costing time and bandwidth.
- Each new connection has reduced performance because of TCP slow start.
- There is a practical limit on the number of open parallel connections.

Persistent connections offer some advantages over parallel connections. They reduce the delay and overhead of connection establishment, keep the connections in a tuned state, and reduce the potential number of open connections. However, persistent connections need to be managed with care, or you may end up accumulating a large number of idle connections, consuming local resources and resources on remote clients and servers.

Persistent connections can be most effective when used in conjunction with parallel connections. Today, many web applications open a small number of parallel connections, each persistent. There are two types of persistent connections: the older HTTP/1.0+ “keep-alive” connections and the modern HTTP/1.1 “persistent” connections. We'll look at both flavors in the next few sections.

HTTP/1.0+ Keep-Alive Connections

Many HTTP/1.0 browsers and servers were extended (starting around 1996) to support an early, experimental type of persistent connections called *keep-alive connections*. These early persistent connections suffered from some interoperability design problems that were rectified in later revisions of HTTP/1.1, but many clients and servers still use these earlier keep-alive connections.

Some of the performance advantages of keep-alive connections are visible in Figure 4-13, which compares the timeline for four HTTP transactions over serial connections against the same transactions over a single persistent connection. The timeline is compressed because the connect and close overheads are removed.*

* Additionally, the request and response time might also be reduced because of elimination of the slow-start phase. This performance benefit is not depicted in the figure.




Figure 4-13. Four transactions (serial versus persistent)

Keep-Alive Operation

Keep-alive is deprecated and no longer documented in the current HTTP/1.1 specification. However, keep-alive handshaking is still in relatively common use by browsers and servers, so HTTP implementors should be prepared to interoperate with it. We'll take a quick look at keep-alive operation now. Refer to older versions of the HTTP/1.1 specification (such as RFC 2068) for a more complete explanation of keep-alive handshaking.

Clients implementing HTTP/1.0 keep-alive connections can request that a connection be kept open by including the Connection: Keep-Alive request header.

If the server is willing to keep the connection open for the next request, it will respond with the same header in the response (see Figure 4-14). If there is no Connection: keep-alive header in the response, the client assumes that the server does not support keep-alive and that the server will close the connection when the response message is sent back.

Keep-Alive Options

Note that the keep-alive headers are just requests to keep the connection alive. Clients and servers do not need to agree to a keep-alive session if it is requested. They




Figure 4-14. HTTP/1.0 keep-alive transaction header handshake

can close idle keep-alive connections at any time and are free to limit the number of transactions processed on a keep-alive connection.

The keep-alive behavior can be tuned by comma-separated options specified in the Keep-Alive general header:

- The timeout parameter is sent in a Keep-Alive response header. It estimates how long the server is likely to keep the connection alive for. This is not a guarantee.
- The max parameter is sent in a Keep-Alive response header. It estimates how many more HTTP transactions the server is likely to keep the connection alive for. This is not a guarantee.
- The Keep-Alive header also supports arbitrary unprocessed attributes, primarily for diagnostic and debugging purposes. The syntax is *name [= value]*.

The Keep-Alive header is completely optional but is permitted only when Connection: Keep-Alive also is present. Here's an example of a Keep-Alive response header indicating that the server intends to keep the connection open for at most five more transactions, or until it has sat idle for two minutes:

```
Connection: Keep-Alive
Keep-Alive: max=5, timeout=120
```

Keep-Alive Connection Restrictions and Rules

Here are some restrictions and clarifications regarding the use of keep-alive connections:

- Keep-alive does not happen by default in HTTP/1.0. The client must send a Connection: Keep-Alive request header to activate keep-alive connections.
- The Connection: Keep-Alive header must be sent with all messages that want to continue the persistence. If the client does not send a Connection: Keep-Alive header, the server will close the connection after that request.

- Clients can tell if the server will close the connection after the response by detecting the absence of the Connection: Keep-Alive response header.
- The connection can be kept open only if the length of the message's entity body can be determined without sensing a connection close—this means that the entity body must have a correct Content-Length, have a multipart media type, or be encoded with the chunked transfer encoding. Sending the wrong Content-Length back on a keep-alive channel is bad, because the other end of the transaction will not be able to accurately detect the end of one message and the start of another.
- Proxies and gateways must enforce the rules of the Connection header; the proxy or gateway must remove any header fields named in the Connection header, and the Connection header itself, before forwarding or caching the message.
- Formally, keep-alive connections should not be established with a proxy server that isn't guaranteed to support the Connection header, to prevent the problem with dumb proxies described below. This is not always possible in practice.
- Technically, any Connection header fields (including Connection: Keep-Alive) received from an HTTP/1.0 device should be ignored, because they may have been forwarded mistakenly by an older proxy server. In practice, some clients and servers bend this rule, although they run the risk of hanging on older proxies.
- Clients must be prepared to retry requests if the connection closes before they receive the entire response, unless the request could have side effects if repeated.

Keep-Alive and Dumb Proxies

Let's take a closer look at the subtle problem with keep-alive and dumb proxies. A web client's Connection: Keep-Alive header is intended to affect just the single TCP link leaving the client. This is why it is named the "connection" header. If the client is talking to a web server, the client sends a Connection: Keep-Alive header to tell the server it wants keep-alive. The server sends a Connection: Keep-Alive header back if it supports keep-alive and doesn't send it if it doesn't.

The Connection header and blind relays

The problem comes with proxies—in particular, proxies that don't understand the Connection header and don't know that they need to remove the header before proxying it down the chain. Many older or simple proxies act as *blind relays*, tunneling bytes from one connection to another, without specially processing the Connection header.

Imagine a web client talking to a web server through a dumb proxy that is acting as a blind relay. This situation is depicted in Figure 4-15.

Here's what's going on in this figure:

1. In Figure 4-15a, a web client sends a message to the proxy, including the Connection: Keep-Alive header, requesting a keep-alive connection if possible. The client waits for a response to learn if its request for a keep-alive channel was granted.




Figure 4-15. Keep-alive doesn't interoperate with proxies that don't support Connection headers

2. The dumb proxy gets the HTTP request, but it doesn't understand the Connection header (it just treats it as an extension header). The proxy has no idea what keep-alive is, so it passes the message verbatim down the chain to the server (Figure 4-15b). But the Connection header is a hop-by-hop header; it applies to only a single transport link and shouldn't be passed down the chain. Bad things are about to happen.
3. In Figure 4-15b, the relayed HTTP request arrives at the web server. When the web server receives the proxied Connection: Keep-Alive header, it mistakenly concludes that the proxy (which looks like any other client to the server) wants to speak keep-alive! That's fine with the web server—it agrees to speak keep-alive and sends a Connection: Keep-Alive response header back in Figure 4-15c. So, at this point, the web server thinks it is speaking keep-alive with the proxy and will adhere to rules of keep-alive. But the proxy doesn't know the first thing about keep-alive. Uh-oh.
4. In Figure 4-15d, the dumb proxy relays the web server's response message back to the client, passing along the Connection: Keep-Alive header from the web server. The client sees this header and assumes the proxy has agreed to speak keep-alive. So at this point, both the client and server believe they are speaking keep-alive, but the proxy they are talking to doesn't know anything about keep-alive.
5. Because the proxy doesn't know anything about keep-alive, it reflects all the data it receives back to the client and then waits for the origin server to close the connection. But the origin server will not close the connection, because it believes the proxy explicitly asked the server to keep the connection open. So the proxy will hang waiting for the connection to close.
6. When the client gets the response message back in Figure 4-15d, it moves right along to the next request, sending another request to the proxy on the keep-alive connection (see Figure 4-15e). Because the proxy never expects another request

on the same connection, the request is ignored and the browser just spins, making no progress.

7. This miscommunication causes the browser to hang until the client or server times out the connection and closes it.*

Proxies and hop-by-hop headers

To avoid this kind of proxy miscommunication, modern proxies must never proxy the Connection header or any headers whose names appear inside the Connection values. So if a proxy receives a Connection: Keep-Alive header, it shouldn't proxy either the Connection header or any headers named Keep-Alive.

In addition, there are a few hop-by-hop headers that might not be listed as values of a Connection header, but must not be proxied or served as a cache response either. These include Proxy-Authenticate, Proxy-Connection, Transfer-Encoding, and Upgrade. For more information, refer back to “The Oft-Misunderstood Connection Header.”

The Proxy-Connection Hack

Browser and proxy implementors at Netscape proposed a clever workaround to the blind relay problem that didn't require all web applications to support advanced versions of HTTP. The workaround introduced a new header called Proxy-Connection and solved the problem of a single blind relay interposed directly after the client—but not all other situations. Proxy-Connection is implemented by modern browsers when proxies are explicitly configured and is understood by many proxies.

The idea is that dumb proxies get into trouble because they blindly forward hop-by-hop headers such as Connection: Keep-Alive. Hop-by-hop headers are relevant only for that single, particular connection and must not be forwarded. This causes trouble when the forwarded headers are misinterpreted by downstream servers as requests from the proxy itself to control its connection.

In the Netscape workaround, browsers send nonstandard Proxy-Connection extension headers to proxies, instead of officially supported and well-known Connection headers. If the proxy is a blind relay, it relays the nonsense Proxy-Connection header to the web server, which harmlessly ignores the header. But if the proxy is a smart proxy (capable of understanding persistent connection handshaking), it replaces the nonsense Proxy-Connection header with a Connection header, which is then sent to the server, having the desired effect.

Figure 4-16a–d shows how a blind relay harmlessly forwards Proxy-Connection headers to the web server, which ignores the header, causing no keep-alive connection to

* There are many similar scenarios where failures occur due to blind relays and forwarded handshaking.

be established between the client and proxy or the proxy and server. The smart proxy in Figure 4-16e–h understands the Proxy-Connection header as a request to speak keep-alive, and it sends out its own Connection: Keep-Alive headers to establish keep-alive connections.




Figure 4-16. Proxy-Connection header fixes single blind relay

This scheme works around situations where there is only one proxy between the client and server. But if there is a smart proxy on either side of the dumb proxy, the problem will rear its ugly head again, as shown in Figure 4-17.

Furthermore, it is becoming quite common for “invisible” proxies to appear in networks, either as firewalls, intercepting caches, or reverse proxy server accelerators. Because these devices are invisible to the browser, the browser will not send them Proxy-Connection headers. It is critical that transparent web applications implement persistent connections correctly.

HTTP/1.1 Persistent Connections

HTTP/1.1 phased out support for keep-alive connections, replacing them with an improved design called *persistent connections*. The goals of persistent connections are the same as those of keep-alive connections, but the mechanisms behave better.




Figure 4-17. Proxy-Connection still fails for deeper hierarchies of proxies

Unlike HTTP/1.0+ keep-alive connections, HTTP/1.1 persistent connections are active by default. HTTP/1.1 assumes *all* connections are persistent unless otherwise indicated. HTTP/1.1 applications have to explicitly add a `Connection: close` header to a message to indicate that a connection should close after the transaction is complete. This is a significant difference from previous versions of the HTTP protocol, where keep-alive connections were either optional or completely unsupported.

An HTTP/1.1 client assumes an HTTP/1.1 connection will remain open after a response, unless the response contains a `Connection: close` header. However, clients and servers still can close idle connections at any time. Not sending `Connection: close` does not mean that the server promises to keep the connection open forever.

Persistent Connection Restrictions and Rules

Here are the restrictions and clarifications regarding the use of persistent connections:

- After sending a `Connection: close` request header, the client can't send more requests on that connection.
- If a client does not want to send another request on the connection, it should send a `Connection: close` request header in the final request.
- The connection can be kept persistent only if all messages on the connection have a correct, self-defined message length—i.e., the entity bodies must have correct `Content-Lengths` or be encoded with the chunked transfer encoding.

- HTTP/1.1 proxies must manage persistent connections separately with clients and servers—each persistent connection applies to a single transport hop.
- HTTP/1.1 proxy servers should not establish persistent connections with an HTTP/1.0 client (because of the problems of older proxies forwarding Connection headers) unless they know something about the capabilities of the client. This is, in practice, difficult, and many vendors bend this rule.
- Regardless of the values of Connection headers, HTTP/1.1 devices may close the connection at any time, though servers should try not to close in the middle of transmitting a message and should always respond to at least one request before closing.
- HTTP/1.1 applications must be able to recover from asynchronous closes. Clients should retry the requests as long as they don't have side effects that could accumulate.
- Clients must be prepared to retry requests if the connection closes before they receive the entire response, unless the request could have side effects if repeated.
- A single user client should maintain at most two persistent connections to any server or proxy, to prevent the server from being overloaded. Because proxies may need more connections to a server to support concurrent users, a proxy should maintain at most $2N$ connections to any server or parent proxy, if there are N users trying to access the servers.

Pipelined Connections

HTTP/1.1 permits optional *request pipelining* over persistent connections. This is a further performance optimization over keep-alive connections. Multiple requests can be enqueued before the responses arrive. While the first request is streaming across the network to a server on the other side of the globe, the second and third requests can get underway. This can improve performance in high-latency network conditions, by reducing network round trips.

Figure 4-18a-c shows how persistent connections can eliminate TCP connection delays and how pipelined requests (Figure 4-18c) can eliminate transfer latencies.

There are several restrictions for pipelining:

- HTTP clients should not pipeline until they are sure the connection is persistent.
- HTTP responses must be returned in the same order as the requests. HTTP messages are not tagged with sequence numbers, so there is no way to match responses with requests if the responses are received out of order.
- HTTP clients must be prepared for the connection to close at any time and be prepared to redo any pipelined requests that did not finish. If the client opens a persistent connection and immediately issues 10 requests, the server is free to close the connection after processing only, say, 5 requests. The remaining 5




Figure 4-18. Four transactions (pipelined connections)

requests will fail, and the client must be willing to handle these premature closes and reissue the requests.

- HTTP clients should not pipeline requests that have side effects (such as POSTs). In general, on error, pipelining prevents clients from knowing which of a series of pipelined requests were executed by the server. Because nonidempotent requests such as POSTs cannot safely be retried, you run the risk of some methods never being executed in error conditions.

The Mysteries of Connection Close

Connection management—particularly knowing when and how to close connections—is one of the practical black arts of HTTP. This issue is more subtle than many developers first realize, and little has been written on the subject.

“At Will” Disconnection

Any HTTP client, server, or proxy can close a TCP transport connection at any time. The connections normally are closed at the end of a message,* but during error conditions, the connection may be closed in the middle of a header line or in other strange places.

This situation is common with pipelined persistent connections. HTTP applications are free to close persistent connections after any period of time. For example, after a persistent connection has been idle for a while, a server may decide to shut it down.

However, the server can never know for sure that the client on the other end of the line wasn’t about to send data at the same time that the “idle” connection was being shut down by the server. If this happens, the client sees a connection error in the middle of writing its request message.

Content-Length and Truncation

Each HTTP response should have an accurate Content-Length header to describe the size of the response body. Some older HTTP servers omit the Content-Length header or include an erroneous length, depending on a server connection close to signify the actual end of data.

When a client or proxy receives an HTTP response terminating in connection close, and the actual transferred entity length doesn’t match the Content-Length (or there is no Content-Length), the receiver should question the correctness of the length.

If the receiver is a caching proxy, the receiver should not cache the response (to minimize future compounding of a potential error). The proxy should forward the questionable message intact, without attempting to “correct” the Content-Length, to maintain semantic transparency.

Connection Close Tolerance, Retries, and Idempotency

Connections can close at any time, even in non-error conditions. HTTP applications have to be ready to properly handle unexpected closes. If a transport connection closes while the client is performing a transaction, the client should reopen the

* Servers shouldn’t close a connection in the middle of a response unless client or network failure is suspected.

connection and retry one time, unless the transaction has side effects. The situation is worse for pipelined connections. The client can enqueue a large number of requests, but the origin server can close the connection, leaving numerous requests unprocessed and in need of rescheduling.

Side effects are important. When a connection closes after some request data was sent but before the response is returned, the client cannot be 100% sure how much of the transaction actually was invoked by the server. Some transactions, such as GETting a static HTML page, can be repeated again and again without changing anything. Other transactions, such as POSTing an order to an online book store, shouldn't be repeated, or you may risk multiple orders.

A transaction is *idempotent* if it yields the same result regardless of whether it is executed once or many times. Implementors can assume the GET, HEAD, PUT, DELETE, TRACE, and OPTIONS methods share this property.* Clients shouldn't pipeline nonidempotent requests (such as POSTs). Otherwise, a premature termination of the transport connection could lead to indeterminate results. If you want to send a nonidempotent request, you should wait for the response status for the previous request.

Nonidempotent methods or sequences must not be retried automatically, although user agents may offer a human operator the choice of retrying the request. For example, most browsers will offer a dialog box when reloading a cached POST response, asking if you want to post the transaction again.

Graceful Connection Close

TCP connections are bidirectional, as shown in Figure 4-19. Each side of a TCP connection has an input queue and an output queue, for data being read or written. Data placed in the output of one side will eventually show up on the input of the other side.




Figure 4-19. TCP connections are bidirectional

Full and half closes

An application can close either or both of the TCP input and output channels. A `close()` sockets call closes both the input and output channels of a TCP connection.

*Administrators who use GET-based dynamic forms should make sure the forms are idempotent.

This is called a “full close” and is depicted in Figure 4-20a. You can use the `shutdown()` sockets call to close either the input or output channel individually. This is called a “half close” and is depicted in Figure 4-20b.




Figure 4-20. Full and half close

TCP close and reset errors

Simple HTTP applications can use only full closes. But when applications start talking to many other types of HTTP clients, servers, and proxies, and when they start using pipelined persistent connections, it becomes important for them to use half closes to prevent peers from getting unexpected write errors.

In general, closing the output channel of your connection is always safe. The peer on the other side of the connection will be notified that you closed the connection by getting an end-of-stream notification once all the data has been read from its buffer.

Closing the input channel of your connection is riskier, unless you know the other side doesn't plan to send any more data. If the other side sends data to your closed input channel, the operating system will issue a TCP “connection reset by peer” message back to the other side's machine, as shown in Figure 4-21. Most operating systems treat this as a serious error and erase any buffered data the other side has not read yet. This is very bad for pipelined connections.

Say you have sent 10 pipelined requests on a persistent connection, and the responses already have arrived and are sitting in your operating system's buffer (but the application hasn't read them yet). Now say you send request #11, but the server decides you've used this connection long enough, and closes it. Your request #11 will arrive at a closed connection and will reflect a reset back to you. This reset will erase your input buffers.




Figure 4-21. Data arriving at closed connection generates “connection reset by peer” error

When you finally get to reading data, you will get a connection reset by peer error, and the buffered, unread response data will be lost, even though much of it successfully arrived at your machine.

Graceful close

The HTTP specification counsels that when clients or servers want to close a connection unexpectedly, they should “issue a graceful close on the transport connection,” but it doesn’t describe how to do that.

In general, applications implementing graceful closes will first close their output channels and then wait for the peer on the other side of the connection to close *its* output channels. When both sides are done telling each other they won’t be sending any more data (i.e., closing output channels), the connection can be closed fully, with no risk of reset.

Unfortunately, there is no guarantee that the peer implements or checks for half closes. For this reason, applications wanting to close gracefully should half close their output channels and periodically check the status of their input channels (looking for data or for the end of the stream). If the input channel isn’t closed by the peer within some timeout period, the application may force connection close to save resources.

For More Information

This completes our overview of the HTTP plumbing trade. Please refer to the following reference sources for more information about TCP performance and HTTP connection-management facilities.

HTTP Connections

<http://www.ietf.org/rfc/rfc2616.txt>

RFC 2616, “Hypertext Transfer Protocol—HTTP/1.1,” is the official specification for HTTP/1.1; it explains the usage of and HTTP header fields for implementing

parallel, persistent, and pipelined HTTP connections. This document does not cover the proper use of the underlying TCP connections.

<http://www.ietf.org/rfc/rfc2068.txt>

RFC 2068 is the 1997 version of the HTTP/1.1 protocol. It contains explanation of the HTTP/1.0+ Keep-Alive connections that is missing from RFC 2616.

<http://www.ics.uci.edu/pub/ietf/http/draft-ietf-http-connection-00.txt>

This expired Internet draft, "HTTP Connection Management," has some good discussion of issues facing HTTP connection management.

HTTP Performance Issues

<http://www.w3.org/Protocols/HTTP/Performance/>

This W3C web page, entitled "HTTP Performance Overview," contains a few papers and tools related to HTTP performance and connection management.

<http://www.w3.org/Protocols/HTTP/1.0/HTTPPerformance.html>

This short memo by Simon Spero, "Analysis of HTTP Performance Problems," is one of the earliest (1994) assessments of HTTP connection performance. The memo gives some early performance measurements of the effect of connection setup, slow start, and lack of connection sharing.

<ftp://gatekeeper.dec.com/pub/DEC/WRL/research-reports/WRL-TR-95.4.pdf>

"The Case for Persistent-Connection HTTP."

http://www.isi.edu/lam/publications/phttp_tcp_interactions/paper.html

"Performance Interactions Between P-HTTP and TCP Implementations."

<http://www.sun.com/sun-on-net/performance/tcp.slowstart.html>

"TCP Slow Start Tuning for Solaris" is a web page from Sun Microsystems that talks about some of the practical implications of TCP slow start. It's a useful read, even if you are working with different operating systems.

TCP/IP

The following three books by W. Richard Stevens are excellent, detailed engineering texts on TCP/IP. These are extremely useful for anyone using TCP:

TCP Illustrated, Volume I: The Protocols

W. Richard Stevens, Addison Wesley

UNIX Network Programming, Volume 1: Networking APIs

W. Richard Stevens, Prentice-Hall

UNIX Network Programming, Volume 2: The Implementation

W. Richard Stevens, Prentice-Hall

The following papers and specifications describe TCP/IP and features that affect its performance. Some of these specifications are over 20 years old and, given the worldwide success of TCP/IP, probably can be classified as historical treasures:

<http://www.acm.org/sigcomm/ccr/archive/2001/jan01/ccr-200101-mogul.pdf>

In “Rethinking the TCP Nagle Algorithm,” Jeff Mogul and Greg Minshall present a modern perspective on Nagle’s algorithm, outline what applications should and should not use the algorithm, and propose several modifications.

<http://www.ietf.org/rfc/rfc2001.txt>

RFC 2001, “TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms,” defines the TCP slow-start algorithm.

<http://www.ietf.org/rfc/rfc1122.txt>

RFC 1122, “Requirements for Internet Hosts—Communication Layers,” discusses TCP acknowledgment and delayed acknowledgments.

<http://www.ietf.org/rfc/rfc896.txt>

RFC 896, “Congestion Control in IP/TCP Internetworks,” was released by John Nagle in 1984. It describes the need for TCP congestion control and introduces what is now called “Nagle’s algorithm.”

<http://www.ietf.org/rfc/rfc0813.txt>

RFC 813, “Window and Acknowledgement Strategy in TCP,” is a historical (1982) specification that describes TCP window and acknowledgment implementation strategies and provides an early description of the delayed acknowledgment technique.

<http://www.ietf.org/rfc/rfc0793.txt>

RFC 793, “Transmission Control Protocol,” is Jon Postel’s classic 1981 definition of the TCP protocol.