




Figure 17-2. If servers vary on specific request headers, caches must match those request headers in addition to the regular content-negotiation headers before sending back cached responses

If a server's Vary header looked like this, the huge number of different User-Agent and Cookie values could generate many variants:

Vary: User-Agent, Cookie

A cache would have to store each document version corresponding to each variant. When the cache does a lookup, it first does content matching with the content-negotiation headers, then matches the request's variant with cached variants. If there is no match, the cache fetches the document from the origin server.

Transcoding

We have discussed in some detail the mechanism by which clients and servers can choose between a set of documents for a URL and send the one that best matches the

client's needs. These mechanisms rely on the presence of documents that match the client's needs—whether they match the needs perfectly or not so well.

What happens, however, when a server does not have a document that matches the client's needs at all? The server may have to respond with an error, but theoretically, the server may be able to transform one of its existing documents into something that the client can use. This option is called *transcoding*.

Table 17-4 lists some hypothetical transcodings.

Table 17-4. Hypothetical transcodings

Before	After
HTML document	WML document
High-resolution image	Low-resolution image
Image in 64K colors	Black-and-white image
Complex page with frames	Simple text page without frames or images
HTML page with Java applets	HTML page without Java applets
Page with ads	Page with ads removed

There are three categories of transcoding: format conversion, information synthesis, and content injection.

Format Conversion

Format conversion is the transformation of data from one format to another to make it viewable by a client. A wireless device seeking to access a document typically viewed by a desktop client may be able to do so with an HTML-to-WML conversion. A client accessing a web page over a slow link that is not very interested in high-resolution images may be able to view an image-rich page more easily if the images are reduced in size and resolution by converting them from color to black and white and shrinking them.

Format conversion is driven by the content-negotiation headers listed in Table 17-2, although it may also be driven by the User-Agent header. Note that content transformation or transcoding is different from content encoding or transfer encoding, in that the latter two typically are used for more efficient or safe transport of content, whereas the former is used to make content viewable on the access device.

Information Synthesis

The extraction of key pieces of information from a document—known as *information synthesis*—can be a useful transcoding process. A simple example of this is the generation of an outline of a document based on section headings, or the removal of advertisements and logos from a page.

More sophisticated technologies that categorize pages based on keywords in content also are useful in summarizing the essence of a document. This technology often is used by automatic web page-classification systems, such as web-page directories at portal sites.

Content Injection

The two categories of transcodings described so far typically reduce the amount of content in web documents, but there is another category of transformations that increases the amount of content: *content-injection* transcodings. Examples of content-injection transcodings are automatic ad generators and user-tracking systems.

Imagine the appeal (and offence) of an ad-insertion transcoder that automatically adds advertisements to each HTML page as it goes by. Transcoding of this type has to be dynamic—it must be done on the fly in order to be effective in adding ads that currently are relevant or somehow have been targeted for a particular user. User-tracking systems also can be built to add content to pages dynamically, for the purpose of collecting statistics about how the page is viewed and how clients surf the Web.

Transcoding Versus Static Pregeneration

An alternative to transcodings is to build different copies of web pages at the web server—for example, one with HTML, one with WML, one with high-resolution images, one with low-resolution images, one with multimedia content, and one without. This, however, is not a very practical technique, for many reasons: any small change in a page requires multiple pages to be modified, more space is necessary to store all the different versions of each page, and it's harder to catalog pages and program web servers to serve the right ones. Some transcodings, such as ad insertion (especially targeted ad insertion), cannot be done statically—the ad inserted will depend upon the user requesting the page.

An on-the-fly transformation of a single root page can be an easier solution than static pregeneration. It can come, however, at the cost of increased latency in serving the content. Some of this computation can, however, be done by a third party, thereby off-loading the computation from the web server—the transformation can be done by an external agent at a proxy or cache. Figure 17-3 illustrates transcoding at a proxy cache.

Next Steps

The story of content negotiation does not end with the Accept and Content headers, for a couple of reasons:

- Content negotiation in HTTP incurs some performance limits. Searching through many variants for appropriate content, or trying to “guess” the best match, can




Figure 17-3. Content transformation or transcoding at a proxy cache

be costly. Are there ways to streamline and focus the content-negotiation protocol? RFCs 2295 and 2296 attempt to address this question for transparent HTTP content negotiation.

- HTTP is not the only protocol that needs to do content negotiation. Streaming media and fax are two other examples where client and server need to discuss the best answer to the client's request. Can a general content-negotiation protocol be developed on top of TCP/IP application protocols? The Content Negotiation Working Group was formed to tackle this question. The group is now closed, but it contributed several RFCs. See the next section for a link to the group's web site.

For More Information

The following Internet drafts and online documentation can give you more details about content negotiation:

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

RFC 2616, "Hypertext Transfer Protocol—HTTP/1.1," is the official specification for HTTP/1.1, the current version of the HTTP protocol. The specification is a well-written, well-organized, detailed reference for HTTP, but it isn't ideal for readers who want to learn the underlying concepts and motivations of HTTP or the differences between theory and practice. We hope that this book fills in the underlying concepts, so you can make better use of the specification.

<http://search.ietf.org/rfc/rfc2295.txt>

RFC 2295, “Transparent Content Negotiation in HTTP,” is a memo describing a transparent content-negotiation protocol on top of HTTP. The status of this memo remains experimental.

<http://search.ietf.org/rfc/rfc2296.txt>

RFC 2296, “HTTP Remote Variant Selection Algorithm—RVSA 1.0,” is a memo describing an algorithm for the transparent selection of the “best” content for a particular HTTP request. The status of this memo remains experimental.

<http://search.ietf.org/rfc/rfc2936.txt>

RFC 2936, “HTTP MIME Type Handler Detection,” is a memo describing an approach for determining the actual MIME type handlers that a browser supports. This approach can help if the Accept header is not specific enough.

<http://www.imc.org/ietf-medfree/index.htm>

This is a link to the Content Negotiation (CONNEG) Working Group, which looked into transparent content negotiation for HTTP, fax, and print. This group is now closed.

Content Publishing and Distribution

Part V talks all about the technology for publishing and disseminating web content:

- Chapter 18, *Web Hosting*, discusses the ways people deploy servers in modern web hosting environments, HTTP support for virtual web hosting, and how to replicate content across geographically distant servers.
- Chapter 19, *Publishing Systems*, discusses the technologies for creating web content and installing it onto web servers.
- Chapter 20, *Redirection and Load Balancing*, surveys the tools and techniques for distributing incoming web traffic among a collection of servers.
- Chapter 21, *Logging and Usage Tracking*, covers log formats and common questions.

Web Hosting

When you place resources on a public web server, you make them available to the Internet community. These resources can be as simple as text files or images, or as complicated as real-time driving maps or e-commerce shopping gateways. It's critical that this rich variety of resources, owned by different organizations, can be conveniently published to web sites and placed on web servers that offer good performance at a fair price.

The collective duties of storing, brokering, and administering content resources is called *web hosting*. Hosting is one of the primary functions of a web server. You need a server to hold, serve, log access to, and administer your content. If you don't want to manage the required hardware and software yourself, you need a hosting service, or *hoster*. Hosters rent you serving and web-site administration services and provide various degrees of security, reporting, and ease of use. Hosters typically pool web sites on heavy-duty web servers for cost-efficiency, reliability, and performance.

This chapter explains some of the most important features of web hosting services and how they interact with HTTP applications. In particular, this chapter covers:

- How different web sites can be “virtually hosted” on the same server, and how this affects HTTP
- How to make web sites more reliable under heavy traffic
- How to make web sites load faster

Hosting Services

In the early days of the World Wide Web, individual organizations purchased their own computer hardware, built their own computer rooms, acquired their own network connections, and managed their own web server software.

As the Web quickly became mainstream, everyone wanted a web site, but few people had the skills or time to build air-conditioned server rooms, register domain

names, or purchase network bandwidth. To save the day, many new businesses emerged, offering professionally managed web hosting services. Many levels of service are available, from physical facilities management (providing space, air conditioning, and wiring) to full-service web hosting, where all the customer does is provide the content.

This chapter focuses on what the hosting web server provides. Much of what makes a web site work—as well as, for example, its ability to support different languages and its ability to do secure e-commerce transactions—depends on what capabilities the hosting web server supports.

A Simple Example: Dedicated Hosting

Suppose that Joe's Hardware Online and Mary's Antique Auction both want fairly high-volume web sites. Irene's ISP has racks and racks full of identical, high-performance web servers that it can lease to Joe and Mary, instead of having Joe and Mary purchase their own servers and maintain the server software.

In Figure 18-1, both Joe and Mary sign up for the *dedicated web hosting* service offered by Irene's ISP. Joe leases a dedicated web server that is purchased and maintained by Irene's ISP. Mary gets a different dedicated server from Irene's ISP. Irene's ISP gets to buy server hardware in volume and can select hardware that is reliable, time-tested, and low-cost. If either Joe's Hardware Online or Mary's Antique Auction grows in popularity, Irene's ISP can offer Joe or Mary additional servers immediately.




Figure 18-1. Outsourced dedicated hosting

In this example, browsers send HTTP requests for *www.joes-hardware.com* to the IP address of Joe's server and requests for *www.marys-antiques.com* to the (different) IP address of Mary's server.

Virtual Hosting

Many folks want to have a web presence but don't have high-traffic web sites. For these people, providing a dedicated web server may be a waste, because they're paying many hundreds of dollars a month to lease a server that is mostly idle!

Many web hosters offer lower-cost web hosting services by sharing one computer between several customers. This is called *shared hosting* or *virtual hosting*. Each web site appears to be hosted by a different server, but they really are hosted on the same physical server. From the end user's perspective, virtually hosted web sites should be indistinguishable from sites hosted on separate dedicated servers.

For cost efficiency, space, and management reasons, a virtual hosting company wants to host tens, hundreds, or thousands of web sites on the same server—but this does not necessarily mean that 1,000 web sites are served from only one PC. Hosters can create banks of replicated servers (called *server farms*) and spread the load across the farm of servers. Because each server in the farm is a clone of the others, and hosts many virtual web sites, administration is much easier. (We'll talk more about server farms in Chapter 20.)

When Joe and Mary started their businesses, they might have chosen virtual hosting to save money until their traffic levels made a dedicated server worthwhile (see Figure 18-2).




Figure 18-2. Outsourced virtual hosting

Virtual Server Request Lacks Host Information

Unfortunately, there is a design flaw in HTTP/1.0 that makes virtual hosters pull their hair out. The HTTP/1.0 specification didn't give any means for shared web servers to identify which of the virtual web sites they're hosting is being accessed.

Recall that HTTP/1.0 requests send only the path component of the URL in the request message. If you try to get `http://www.joes-hardware.com/index.html`, the browser connects to the server `www.joes-hardware.com`, but the HTTP/1.0 request says “GET /index.html”, with no further mention of the hostname. If the server is virtually hosting multiple sites, this isn’t enough information to figure out what virtual web site is being accessed. For example, in Figure 18-3:

- If client A tries to access `http://www.joes-hardware.com/index.html`, the request “GET /index.html” will be sent to the shared web server.
- If client B tries to access `http://www.marys-antiques.com/index.html`, the identical request “GET /index.html” will be sent to the shared web server.




Figure 18-3. HTTP/1.0 server requests don’t contain hostname information

As far as the web server is concerned, there is not enough information to determine which web site is being accessed! The two requests look the same, even though they are for totally different documents (from different web sites). The problem is that the web site host information has been stripped from the request.

As we saw in Chapter 6, HTTP surrogates (reverse proxies) and intercepting proxies also need site-specifying information.

Making Virtual Hosting Work

The missing host information was an oversight in the original HTTP specification, which mistakenly assumed that each web server would host exactly one web site. HTTP’s designers didn’t provide support for virtually hosted, shared servers. For this reason, the hostname information in the URL was viewed as redundant and stripped away; only the path component was required to be sent.

Because the early specifications did not make provisions for virtual hosting, web hosters needed to develop workarounds and conventions to support shared virtual hosting. The problem could have been solved simply by requiring all HTTP request

messages to send the full URL instead of just the path component. HTTP/1.1 does require servers to handle full URLs in the request lines of HTTP messages, but it will be a long time before all legacy applications are upgraded to this specification. In the meantime, four techniques have emerged:

Virtual hosting by URL path

Adding a special path component to the URL so the server can determine the site.

Virtual hosting by port number

Assigning a different port number to each site, so requests are handled by separate instances of the web server.

Virtual hosting by IP address

Dedicating different IP addresses for different virtual sites and binding all the IP addresses to a single machine. This allows the web server to identify the site name by IP address.

Virtual hosting by Host header

Many web hosts pressured the HTTP designers to solve this problem. Enhanced versions of HTTP/1.0 and the official version of HTTP/1.1 define a Host request header that carries the site name. The web server can identify the virtual site from the Host header.

Let's take a closer look at each technique.

Virtual hosting by URL path

You can use brute force to isolate virtual sites on a shared server by assigning them different URL paths. For example, you could give each logical web site a special path prefix:

- Joe's Hardware store could be *http://www.joes-hardware.com/joe/index.html*.
- Mary's Antiques store could be *http://www.marys-antiques.com/mary/index.html*.

When the requests arrive at the server, the hostname information is not present in the request, but the server can tell them apart based on the path:

- The request for Joe's hardware is "GET /joe/index.html".
- The request for Mary's antiques is "GET /mary/index.html".

This is not a good solution. The "/joe" and "/mary" prefixes are redundant and confusing (we already mentioned "joe" in the hostname). Worse, the common convention of specifying *http://www.joes-hardware.com* or *http://www.joes-hardware.com/index.html* for the home page won't work.

In general, URL-based virtual hosting is a poor solution and seldom is used.

Virtual hosting by port number

Instead of changing the pathname, Joe and Mary could each be assigned a different port number on the web server. Instead of port 80, for example, Joe could get 82 and

Mary could have 83. But this solution has the same problem: an end user would expect to find the resources without having to specify a nonstandard port in the URL.

Virtual hosting by IP address

A much better approach (in common use) is virtual IP addressing. Here, each virtual web site gets one or more unique IP addresses. The IP addresses for all of the virtual web sites are attached to the same shared server. The server can look up the destination IP address of the HTTP connection and use that to determine what web site the client thinks it is connected to.

Say a hoster assigned the IP address 209.172.34.3 to *www.joes-hardware.com*, assigned 209.172.34.4 to *www.marys-antiques.com*, and tied both IP addresses to the same physical server machine. The web server could then use the destination IP address to identify which virtual site is being requested, as shown in Figure 18-4:

- Client A fetches *http://www.joes-hardware.com/index.html*.
- Client A finds the IP address for *www.joes-hardware.com*, getting 209.172.34.3.
- Client A opens a TCP connection to the shared web server at 209.172.34.3.
- Client A sends the request “GET /index.html HTTP/1.0”.
- Before the web server serves a response, it notes the actual destination IP address (209.172.34.3), determines that this is a virtual IP address for Joe’s web site, and fulfills the request from the */joe* subdirectory. The page */joe/index.html* is returned.




Figure 18-4. Virtual IP hosting

Similarly, if client B asks for *http://www.marys-antiques.com/index.html*:

- Client B finds the IP address for *www.marys-antiques.com*, getting 209.172.34.4.
- Client B opens a TCP connection to the web server at 209.172.34.4.
- Client B sends the request “GET /index.html HTTP/1.0”.
- The web server determines that 209.172.34.4 is Mary’s web site and fulfills the request from the */mary* subdirectory, returning the document */mary/index.html*.

Virtual IP hosting works, but it causes some difficulties, especially for large hosters:

- Computer systems usually have a limit on how many virtual IP addresses can be bound to a machine. Hosters that want hundreds or thousands of virtual sites to be hosted on a shared server may be out of luck.
- IP addresses are a scarce commodity. Hosters with many virtual sites might not be able to obtain enough virtual IP addresses for the hosted web sites.
- The IP address shortage is made worse when hosters replicate their servers for additional capacity. Different virtual IP addresses may be needed on each replicated server, depending on the load-balancing architecture, so the number of IP addresses needed can multiply by the number of replicated servers.

Despite the address consumption problems with virtual IP hosting, it is used widely.

Virtual hosting by Host header

To avoid excessive address consumption and virtual IP limits, we'd like to share the same IP address among virtual sites, but still be able to tell the sites apart. But as we've seen, because most browsers send just the path component of the URL to servers, the critical virtual hostname information is lost.

To solve this problem, browser and server implementors extended HTTP to provide the original hostname to servers. But browsers couldn't just send a full URL, because that would break many servers that expected to receive only a path component. Instead, the hostname (and port) is passed in a Host extension header in all requests.

In Figure 18-5, client A and client B both send Host headers that carry the original hostname being accessed. When the server gets the request for `/index.html`, it can use the Host header to decide which resources to use.




Figure 18-5. Host headers distinguish virtual host requests

Host headers were first introduced with HTTP/1.0+, a vendor-extended superset of HTTP/1.0. Host headers are required for HTTP/1.1 compliance. Host headers are supported by most modern browsers and servers, but there are still a few clients and servers (and robots) that don't support them.

HTTP/1.1 Host Headers

The Host header is an HTTP/1.1 request header, defined in RFC 2068. Virtual servers are so common that most HTTP clients, even if they are not HTTP/1.1-compliant, implement the Host header.

Syntax and usage

The Host header specifies the Internet host and port number for the resource being requested, as obtained from the original URL:

```
Host = "Host" ":" host [ ":" port ]
```

In particular:

- If the Host header does not contain a port, the default port for the scheme is assumed.
- If the URL contains an IP address, the Host header should contain the same address.
- If the URL contains a hostname, the Host header must contain the same name.
- If the URL contains a hostname, the Host header should *not* contain the IP address equivalent to the URL's hostname, because this will break virtually hosted servers, which layer multiple virtual sites over a single IP address.
- If the URL contains a hostname, the Host header should not contain another alias for this hostname, because this also will break virtually hosted servers.
- If the client is using an explicit proxy server, the client must include the name and port of the *origin* server in the Host header, *not* the proxy server. In the past, several web clients had bugs where the outgoing Host header was set to the hostname of the proxy, when the client's proxy setting was enabled. This incorrect behavior causes proxies and origin servers to misbehave.
- Web clients must include a Host header field in all request messages.
- Web proxies must add Host headers to request messages before forwarding them.
- HTTP/1.1 web servers must respond with a 400 status code to any HTTP/1.1 request message that lacks a Host header field.

Here is a sample HTTP request message used to fetch the home page of *www.joes-hardware.com*, along with the required Host header field:

```
GET http://www.joes-hardware.com/index.html HTTP/1.0
Connection: Keep-Alive
User-Agent: Mozilla/4.51 [en] (X11; U; IRIX 6.2 IP22)
```

```
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, image/png, */*
Accept-Encoding: gzip
Accept-Language: en
Host: www.joes-hardware.com
```

Missing Host headers

A small percentage of old browsers in use do not send Host headers. If a virtual hosting server is using Host headers to determine which web site to serve, and no Host header is present, it probably will either direct the user to a default web page (such as the web page of the ISP) or return an error page suggesting that the user upgrade her browser.

Interpreting Host headers

An origin server that isn't virtually hosted, and doesn't allow resources to differ by the requested host, may ignore the Host header field value. But any origin server that does differentiate resources based on the host must use the following rules for determining the requested resource on an HTTP/1.1 request:

1. If the URL in the HTTP request message is absolute (i.e., contains a scheme and host component), the value in the Host header is ignored in favor of the URL.
2. If the URL in the HTTP request message doesn't have a host, and the request contains a Host header, the value of the host/port is obtained from the Host header.
3. If no valid host can be determined through Steps 1 or 2, a 400 Bad Response response is returned to the client.

Host headers and proxies

Some browser versions send incorrect Host headers, especially when configured to use proxies. For example, when configured to use a proxy, some older versions of Apple and PointCast clients mistakenly sent the name of the proxy instead of the origin server in the Host header.

Making Web Sites Reliable

There are several times during which web sites commonly break:

- Server downtime
- Traffic spikes: suddenly everyone wants to see a particular news broadcast or rush to a sale. Sudden spikes can overload a web server, slowing it down or stopping it completely.
- Network outages or losses

This section presents some ways of anticipating and dealing with these common problems.

Mirrored Server Farms

A server farm is a bank of identically configured web servers that can cover for each other. The content on each server in the farm can be mirrored, so that if one has a problem, another can fill in.

Often, mirrored servers follow a hierarchical relationship. One server might act as the “content authority”—the server that contains the original content (perhaps a server to which the content authors post). This server is called the *master origin server*. The mirrored servers that receive content from the master origin server are called *replica origin servers*. One simple way to deploy a server farm is to use a network switch to distribute requests to the servers. The IP address for each of the web sites hosted on the servers is the IP address of the switch.

In the mirrored server farm shown in Figure 18-6, the master origin server is responsible for sending content to the replica origin servers. To the outside world, the IP address for this content is the IP address of the switch. The switch is responsible for sending requests to the servers.




Figure 18-6. Mirrored server farm

Mirrored web servers can contain copies of the exact same content at different locations. Figure 18-7 illustrates four mirrored servers, with a master server in Chicago and replicas in New York, Miami, and Little Rock. The master server serves clients in the Chicago area and also has the job of propagating its content to the replica servers.

In the Figure 18-7 scenario, there are a couple of ways that client requests would be directed to a particular server:

HTTP redirection

The URL for the content could resolve to the IP address of the master server, which could then send redirects to replica servers.




Figure 18-7. Dispersed mirrored servers

DNS redirection

The URL for the content could resolve to four IP addresses, and the DNS server could choose the IP address that it sends to clients.

See Chapter 20 for more details.

Content Distribution Networks

A *content distribution network* (CDN) is simply a network whose purpose is the distribution of specific content. The nodes of the network can be web servers, surrogates, or proxy caches.

Surrogate Caches in CDNs

Surrogate caches can be used in place of replica origin servers in Figures 18-6 and 18-7. Surrogates, also known as reverse proxies, receive server requests for content just as mirrored web servers do. They receive server requests on behalf of a specific set of origin servers (this is possible because of the way IP addresses for content are advertised; there usually is a working relationship between origin server and surrogate, and surrogates expect to receive requests aimed at specific origin servers).

The difference between a surrogate and a mirrored server is that surrogates typically are demand-driven. They do not store entire copies of the origin server content; they store whatever content their clients request. The way content is distributed in their caches depends on the requests that they receive; the origin server does not have the responsibility to update their content. For easy access to “hot” content (content that is in high demand), some surrogates have “prefetching” features that enable them to pull content in advance of user requests.

An added complexity in CDNs with surrogates is the possibility of cache hierarchies.

Proxy Caches in CDNs

Proxy caches also can be deployed in configurations similar to those in Figures 18-6 and 18-7. Unlike surrogates, traditional proxy caches can receive requests aimed at any web servers (there need not be any working relationship or IP address agreement between a proxy cache and an origin server). As with surrogates, however, proxy cache content typically is demand-driven and is not expected to be an exact duplicate of the origin server content. Some proxy caches also can be preloaded with hot content.

Demand-driven proxy caches can be deployed in other kinds of configurations—in particular, interception configurations, where a layer-2 or -3 device (switch or router) intercepts web traffic and sends it to a proxy cache (see Figure 18-8).




Figure 18-8. Client requests intercepted by a switch and sent to a proxy

An interception configuration depends on being able to set up the network between clients and servers so that all of the appropriate HTTP requests are physically channeled to the cache. (See Chapter 20). The content is distributed in the cache according to the requests it receives.

Making Web Sites Fast

Many of the technologies mentioned in the previous section also help web sites load faster. Server farms and distributed proxy caches or surrogate servers distribute network traffic, avoiding congestion. Distributing the content brings it closer to end users, so that the travel time from server to client is lower. The key to speed of resource access is how requests and responses are directed from client to server and back across the Internet. See Chapter 20 for details on redirection methods.

Another approach to speeding up web sites is encoding the content for fast transportation. This can mean, for example, compressing the content, assuming that the receiving client can uncompress it. See Chapter 15 for details.

For More Information

See Part III, *Identification, Authorization, and Security*, for details on how to make web sites secure. The following Internet drafts and documentation can give you more details about web hosting and content distribution:

<http://www.ietf.org/rfc/rfc3040.txt>

RFC 3040, “Internet Web Replication and Caching Taxonomy,” is a reference for the vocabulary of web replication and caching applications.

<http://search.ietf.org/internet-drafts/draft-ietf-cdi-request-routing-reqs-00.txt>

“Request-Routing Requirements for Content Internetworking.”

Apache: The Definitive Guide

Ben Laurie and Peter Laurie, O’Reilly & Associates, Inc. This book describes how to run the open source Apache web server.

Publishing Systems

How do you create web pages and get them onto a web server? In the dark ages of the Web (let's say, 1995), you might have hand-crafted your HTML in a text editor and manually uploaded the content to the web server using FTP. This procedure was painful, difficult to coordinate with coworkers, and not particularly secure.

Modern-day publishing tools make it much more convenient to create, publish, and manage web content. Today, you can interactively edit web content as you'll see it on the screen and publish that content to servers with a single click, while being notified of any files that have changed.

Many of the tools that support remote publishing of content use extensions to the HTTP protocol. In this chapter, we explain two important technologies for web-content publishing based on HTTP: FrontPage and DAV.

FrontPage Server Extensions for Publishing Support

FrontPage (commonly referred to as FP) is a versatile web authoring and publishing toolkit provided by Microsoft Corp. The original idea for FrontPage (FrontPage 1.0) was conceived in 1994, at Vermeer Technologies, Inc., and was dubbed the first product to combine web site management and creation into a single, unified tool. Microsoft purchased Vermeer and shipped FrontPage 1.1 in 1996. The latest version, FrontPage Version 2002, is the sixth version in the line and a core part of the Microsoft Office suite.

FrontPage Server Extensions

As part of the “publish anywhere” strategy, Microsoft released a set of server-side software called FrontPage Server Extensions (FPSE). These server-side components integrate with the web server and provide the necessary translation between the web site and the client running FrontPage (and other clients that support these extensions).

Our primary interest lies in the publishing protocol between the FP clients and FPSE. This protocol provides an example of designing extensions to the core services available in HTTP without changing HTTP semantics.

The FrontPage publishing protocol implements an RPC layer on top of the HTTP POST request. This allows the FrontPage client to send commands to the server to update documents on the web site, perform searches, collaborate amongst the web authors, etc. Figure 19-1 gives an overview of the communication.




Figure 19-1. FrontPage publishing architecture

The web server sees POST requests addressed to the FPSE (implemented as a set of CGI programs, in the case of a non-Microsoft IIS server) and directs those requests accordingly. As long as intervening firewalls and proxy servers are configured to allow the POST method, FrontPage can continue communicating with the server.

FrontPage Vocabulary

Before we dive deeper into the RPC layer defined by FPSE, it may help to establish the common vocabulary:

Virtual server

One of the multiple web sites running on the same server, each with a unique domain name and IP address. In essence, a virtual server allows a single web server to host multiple web sites, each of which appears to a browser as being hosted by its own web server. A web server that supports virtual servers is called a *multi-hosting* web server. A machine that is configured with multiple IP addresses is called a *multi-homed* server (for more details, please refer to “Virtual Hosting” in Chapter 18).

Root web

The default, top-level content directory of a web server, or, in a multi-hosting environment, the top-level content directory of a virtual web server. To access the root web, it is enough to specify the URL of the server without specifying a page name. There can be only one root web per web server.

Subweb

A named subdirectory of the root web or another subweb that is a complete FPSE extended web. A subweb can be a complete independent entity with the ability to specify its own administration and authoring permissions. In addition, subwebs may provide scoping for methods such as searches.

The FrontPage RPC Protocol

The FrontPage client and FPSE communicate using a proprietary RPC protocol. This protocol is layered on top of HTTP POST by embedding the RPC methods and their associated variables in the body of the POST request.

To start the process, the client needs to determine the location and the name of the target programs on the server (the part of the FPSE package that can execute the POST request). It then issues a special GET request (see Figure 19-2).




Figure 19-2. Initial request

When the file is returned, the FrontPage client reads the response and finds the values associated with `FPShtmlScriptUrl`, `FPAuthorScriptUrl`, and `FPAdminScriptUrl`. Typically, this may look like:

```
FPShtmlScriptUrl="_vti_bin/_vti_rpc/shtml.dll"
FPAuthorScriptUrl="_vti_bin/_vti_aut/author.dll"
FPAdminScriptUrl="_vti_bin/_vti_adm/admin.dll"
```

`FPShtmlScriptUrl` tells the client where to POST requests for “browse time” commands (e.g., getting the version of FPSE) to be executed.

`FPAuthorScriptUrl` tells the client where to POST requests for “authoring time” commands to be executed. Similarly, `FPAdminScriptUrl` tells FrontPage where to POST requests for administrative actions.

Now that we know where the various programs are located, we are ready to send a request.

Request

The body of the POST request contains the RPC command, in the form of “method=<command>” and any required parameters. For example, consider the RPC message requesting a list of documents, as follows:

```
POST /_vti_bin/_vti_aut/author.dll HTTP/1.1
Date: Sat, 12 Aug 2000 20:32:54 GMT
User-Agent: MSFrontPage/4.0
.....

<BODY>
method=list+documents%3a4%2e0%2e2%2e3717&service%5fname=&listHiddenDocs=false&listExp
lorerDocs=false&listRecurse=false&listFiles=true&listFolders=true&listLinkInfo=true&l
istIncludeParent=true&listDerived=false
&listBorders=false&listChildWebs=true&initialUrl=&folderList=%5b%3bTW%7c12+Aug+2000+2
0%3a33%3a04+%2d0000%5d
```

The body of the POST command contains the RPC command being sent to the FPSE. As with CGI programs, the spaces in the method are encoded as plus sign (+) characters. All other nonalphanumeric characters in the method are encoded using %XX format, where the XX stands for the ASCII representation of the character. Using this notation, a more readable version of the body would look like the following:

```
method=list+documents:4.0.1.3717
&service_name=
&listHiddenDocs=false
&listExplorerDocs=false
.....
```

Some of the elements listed are:

service_name

The URL of the web site on which the method should act. Must be an existing folder or one level below an existing folder.

listHiddenDocs

Shows the hidden documents in a web if its value is “true”. The “hidden” documents are designated by URLs with path components starting with “_”.

listExploreDocs

If the value is “true”, lists the task lists.

Response

Most RPC protocol methods have return values. Most common return values are for successful methods and errors. Some methods also have a third subsection, “Sample Return Code.” FrontPage properly interprets the codes to provide accurate feedback to the user.

Continuing with our example, the FPSE processes the “list+documents” request and returns the necessary information. A sample response follows:

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Sat, 12 Aug 2000 22:49:50 GMT
Content-type: application/x-vermeer-rpc
X-FrontPage-User-Name: IUSER_MINSTAR

<html><head><title>RPC packet</title></head>
<body>
<p>method=list documents: 4.0.2.3717
<p>document_list=
<ul>
  <li>document_name=help.gif
</ul>
```

As you can see from the response, a formatted list of documents available on the web server is returned to the FP client. You can find the complete list of commands and responses at the Microsoft web site.

FrontPage Security Model

Any publishing system directly accessing web server content needs to be very conscious of the security implications of its actions. For the most part, FPSE depends on the web server to provide the security.

The FPSE security model defines three kinds of users: administrators, authors, and browsers, with administrators having complete control. All permissions are cumulative; i.e., all administrators may author and browse the FrontPage web. Similarly, all authors have browsing permissions.

The list of administrators, authors, and browsers is defined for a given FPSE extended web. All of the subwebs may inherit the permissions from the root web or set their own. For non-IIS web servers, all the FPSE programs are required to be stored in directories marked “executable” (the same restriction as for any other CGI program). *Fpsrvadm*, the FrontPage server administrator utility, may be used for this purpose. On IIS servers, the integrated Windows security model prevails.

On non-IIS servers, web server access-control mechanisms specify the users who are allowed to access a given program. On Apache and NCSA web servers, the file is named *.htaccess*; on Netscape servers, it is named *.nsconfig*. The access file associates users, groups, and IP addresses with various levels of permissions: GET (read), POST (execute), etc. For example, for a user to be an author on an Apache web server, the *.htaccess* file should permit that user to POST to *author.exe*. These access-specification files often are defined on a per-directory basis, providing greater flexibility in defining the permissions.

On IIS servers, the permissions are checked against the ACLs for a given root or sub-root. When IIS gets a request, it first logs on and impersonates the user, then sends the request to one of the three extension dynamic link libraries (DLLs). The DLL checks the impersonation credentials against the ACL defined for the destination folder. If the check is successful, the requested operation is executed by the extension DLL. Otherwise, a “permission denied” message is sent back to the client. Given the tight integration of Windows security with IIS, the User Manager may be used to define fine-grained control.

In spite of this elaborate security model, enabling FPSE has gained notoriety as a nontrivial security risk. In most cases, this is due to sloppy practices adopted by web site administrators. However, the earlier versions of FPSE did have severe security loopholes and thus contributed to the general perception of security risk. This problem also was exacerbated by the arcane practices needed to fully implement a tight security model.

WebDAV and Collaborative Authoring

Web Distributed Authoring and Versioning (WebDAV) adds an extra dimension to web publishing—collaboration. Currently, the most common practice of collaboration is decidedly low-tech: predominantly email, sometimes combined with distributed fileshares. This practice has proven to be very inconvenient and error-prone, with little or no control over the process. Consider an example of launching a multinational, multilingual web site for an automobile manufacturer. It’s easy to see the need for a robust system with secure, reliable publishing primitives, along with collaboration primitives such as locking and versioning.

WebDAV (published as RFC 2518) is focused on extending HTTP to provide a suitable platform for collaborative authoring. It currently is an IETF effort with support from various vendors, including Adobe, Apple, IBM, Microsoft, Netscape, Novell, Oracle, and Xerox.

WebDAV Methods

WebDAV defines a set of new HTTP methods and modifies the operational scope of a few other HTTP methods. The new methods added by WebDAV are:

PROPFIND

Retrieves the properties of a resource.

PROPPATCH

Sets one or more properties on one or many resources.

MKCOL

Creates collections.

COPY

Copies a resource or a collection of resources from a given source to a given destination. The destination need not be on the same machine.

MOVE

Moves a resource or a collection of resources from a given source to a given destination. The destination need not be on the same machine.

LOCK

Locks a resource or multiple resources.

UNLOCK

Unlocks a previously locked resource.

HTTP methods modified by WebDAV are DELETE, PUT, and OPTIONS. Both the new and the modified methods are discussed in detail later in this chapter.

WebDAV and XML

WebDAV's methods generally require a great deal of information to be associated with both requests and responses. HTTP usually communicates this information in message headers. However, transporting necessary information in headers alone imposes some limitations, including the difficulties of selective application of header information to multiple resources in a request, to represent hierarchy, etc.

To solve this problem, WebDAV embraces the Extensible Markup Language (XML), a meta-markup language that provides a format for describing structured data. XML provides WebDAV with:

- A method of formatting instructions describing how data is to be handled
- A method of formatting complex responses from the server
- A method of communicating customized information about the collections and resources handled
- A flexible vehicle for the data itself
- A robust solution for most of the internationalization issues

Traditionally, the schema definition for XML documents is kept in a Document Type Definition (DTD) file that is referenced within the XML document itself. Therefore, when trying to interpret an XML document, the DOCTYPE definition entity gives the name of the DTD file associated with the XML document in question.

WebDAV defines an explicit XML namespace, "DAV:". Without going into many details, an XML namespace is a collection of names of elements or attributes. The namespace qualifies the embedded names uniquely across the domain, thus avoiding any name collisions.

The complete XML schema is defined in the WebDAV specification, RFC 2518. The presence of a predefined schema allows the parsing software to make assumptions on the XML schema without having to read in DTD files and interpret them correctly.

WebDAV Headers

WebDAV does introduce several HTTP headers to augment the functionality of the new methods. This section provides a brief overview; see RFC 2518 for more information. The new headers are:

DAV

Used to communicate the WebDAV capabilities of the server. All resources supported by WebDAV are required to return this header in the response to the OPTIONS request. See “The OPTIONS method” for more details.

```
DAV = "DAV" ":" "1" [", " "2"] [", " 1#extend]
```

Depth

The crucial element for extending WebDAV to grouped resources with multiple levels of hierarchy (for more detailed explanation about collections, please refer to “Collections and Namespace Management”).

```
Depth = "Depth" ":" ("0" | "1" | "infinity")
```

Let’s look at a simple example. Consider a directory *DIR_A* with files *file_1.html* and *file_2.html*. If a method uses Depth: 0, the method applies to the *DIR_A* directory alone, and Depth: 1 applies to the *DIR_A* directory and its files, *file_1.html* and *file_2.html*.

The Depth header modifies many WebDAV-defined methods. Some of the methods that use the Depth header are LOCK, COPY, and MOVE.

Destination

Defined to assist the COPY or MOVE methods in identifying the destination URI.

```
Destination = "Destination" ":" absoluteURI
```

If

The only defined state token is a lock token (see “The LOCK Method”). The If header defines a set of conditionals; if they all evaluate to false, the request will fail. Methods such as COPY and PUT conditionalize the applicability by specifying preconditions in the If header. In practice, the most common precondition to be satisfied is the prior acquisition of a lock.

```
If = "If" ":" (1*No-tag-list | 1*Tagged-list)
No-tag-list = List
Tagged-list = Resource 1*List
Resource = Coded-URL
List = "(" 1*([ "Not" ](State-token | "[" entity-tag "]")) ")"
State-token = Coded-URL
Coded-URL = "<" absoluteURI ">"
```

Lock-Token

Used by the UNLOCK method to specify the lock that needs to be removed. A response to a LOCK method also has a Lock-Token header, carrying the necessary information about the lock taken.

```
Lock-Token = "Lock-Token" ":" Coded-URL
```

Overwrite

Used by the COPY and MOVE methods to designate whether the destination should be overwritten. See the discussion of the COPY and MOVE methods later in this chapter for more details.

```
Overwrite = "Overwrite" ":" ("T" | "F")
```

Timeout

A request header used by a client to specify a desired lock timeout value. For more information, refer to the section “Lock refreshes and the Timeout header.”

```
Timeout = "Timeout" ":" 1#TimeType  
TimeType = ("Second-" DAVTimeOutVal | "Infinite" | Other)  
DAVTimeOutVal = 1*digit  
Other = "Extend" field-value
```

Now that we have sketched the intent and implementation of WebDAV, let's look more closely at the functions provided.

WebDAV Locking and Overwrite Prevention

By definition, collaboration requires more than one person working on a given document. The inherent problem associated with collaboration is illustrated in Figure 19-3.




Figure 19-3. Lost update problem

In this example, authors A and B are jointly writing a specification. A and B independently make a set of changes to the document. A pushes the updated document to the repository, and at a later point, B posts her own version of the document into the repository. Unfortunately, because B never knew about A's changes, she never merged her version with A's version, resulting in A's work being lost.

To ameliorate the problem, WebDAV supports the concept of locking. Locking alone will not fully solve the problem. Versioning and messaging support are needed to complete the solution.

WebDAV supports two types of locks:

- Exclusive write locking of a resource or a collection
- Shared write locking of a resource or a collection

An *exclusive* write lock guarantees write privileges only to the lock owner. This type of locking completely eliminates potential conflicts. A *shared* write lock allows a group of people to work on a given document. This type of locking works well in an environment where all the authors are aware of each other's activities. WebDAV provides a property discovery mechanism, via PROPFIND, to determine the support for locking and the types of locks supported.

WebDAV has two new methods to support locking: LOCK and UNLOCK.

To accomplish locking, there needs to be a mechanism for identifying the author. WebDAV requires digest authentication (discussed in Chapter 13).

When a lock is granted, the server returns a token that is unique across the domain to the client. The specification refers to this as the opaque lock token URI scheme. When the client subsequently wants to perform a write, it connects to the server and completes the digest authentication sequence. Once the authentication is complete, the WebDAV client presents the lock token, along with the PUT request. Thus, the combination of the correct user and the lock token is required to complete the write.

The LOCK Method

A powerful feature of WebDAV is its ability to lock multiple resources with a single LOCK request. WebDAV locking does not require the client to stay connected to the server.

For example, here's a simple LOCK request:

```
LOCK /ch-publish.fm HTTP/1.1
Host: minstar
Content-Type: text/xml
User-Agent: Mozilla/4.0 (compatible; MSIE 5.0; Windows NT)
Content-Length: 201

<?xml version="1.0"?>
<a:lockinfo xmlns:a="DAV:">
  <a:lockscope><a:exclusive/></a:lockscope>
  <a:locktype><a:write/></a:locktype>
  <a:owner><a:href>AuthorA</a:href></a:owner>
</a:lockinfo>
```

The XML being submitted has the <lockinfo> element as its base element. Within the <lockinfo> structure, there are three subelements:

<locktype>

Indicates the type of lock. Currently there is only one, “write.”

<lockscope>

Indicates whether this is an exclusive lock or a shared lock.

<owner>

Field is set with the person who holds the current lock.

Here’s a successful response to our LOCK request:

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Fri, 10 May 2002 20:56:18 GMT
Content-Type: text/xml
Content-Length: 419

<?xml version="1.0"?>
<a:prop xmlns:a="DAV:">
<a:lockdiscovery><a:activelock>
<a:locktype><a:write/></a:locktype>
<a:lockscope><a:exclusive/></a:lockscope>
<a:owner xmlns:a="DAV:"><a:href>AutherA</a:href></a:owner>
<a:locktoken><a:href>opaquelocktoken:*****</a:href></a:locktoken>
<a:depth>0</a:depth>
<a:timeout>Second-180</a:timeout>
</a:activelock></a:lockdiscovery>
</a:prop>
```

The <lockdiscovery> element acts as a container for information about the lock. Embedded in the <lockdiscovery> element is an <activelock> subelement that holds the information sent with the request (<locktype>, <lockscope>, and <owner>). In addition, <activelock> has the following subelements:

<locktoken>

Uniquely identifies the lock in a URI scheme called opaquelocktoken. Given the stateless nature of HTTP, this token is used to identify the ownership of the lock in future requests.

<depth>

Mirrors the value of the Depth header.

<timeout>

Indicates the timeout associated with the lock. In the above response (Figure 19-3), the timeout value is 180 seconds.

The opaquelocktoken scheme

The opaquelocktoken scheme is designed to provide a unique token across all resources for all times. To guarantee uniqueness, the WebDAV specification mandates the use of the universal unique identifier (UUID) mechanism, as described in ISO-11578.

When it comes to actual implementation, there is some leeway. The server has the choice of generating a UUID for each LOCK request, or generating a single UUID and maintaining the uniqueness by appending extra characters at the end. For performance considerations, the latter choice is better. However, if the server chooses to implement the latter choice, it is required to guarantee that none of the added extensions will ever be reused.

The <lockdiscovery> XML element

The <lockdiscovery> XML element provides a mechanism for active lock discovery. If others try to lock the file while a lock is in place, they will receive a <lockdiscovery> XML element that indicates the current owner. The <lockdiscovery> element lists all outstanding locks along with their properties.

Lock refreshes and the Timeout header

To refresh a lock, a client needs to resubmit a lock request with the lock token in the If header. The timeout value returned may be different from the earlier timeout values.

Instead of accepting the timeout value given by the server, a client may indicate the timeout value required in the LOCK request. This is done through the Timeout header. The syntax of the Timeout header allows the client to specify a few options in a comma-separated list. For example:

```
Timeout : Infinite, Second-86400
```

The server is not obligated to honor either of the options. However, it is required to provide the lock expiration time in the <timeout> XML element. In all cases, lock timeout is only a guideline and is not necessarily binding on the server. The administrator may do a manual reset, or some other extraordinary event may cause the server to reset the lock. The clients should avoid taking lengthy locks.

In spite of these primitives, we may not completely solve the “lost update problem” illustrated in Figure 19-3. To completely solve it, a cooperative event system with a versioning control is needed.

The UNLOCK Method

The UNLOCK method removes a lock on a resource, as follows:

```
UNLOCK /ch-publish.fm HTTP/1.1
Host: minstar.inktomi.com
User-Agent: Mozilla/4.0 (compatible; MSIE 5.0; Windows NT)
Lock-Token:
opaquelocktoken:*****
```

```
HTTP/1.1 204 OK
Server: Microsoft-IIS/5.0
Date: Fri, 10 May 2002 20:56:18 GMT
```

As with most resource management requests, WebDAV has two requirements for UNLOCK to succeed: prior completion of a successful digest authentication sequence, and matching the lock token that is sent in the Lock-Token header.

If the unlock is successful, a 204 No Content status code is returned to client. Table 19-1 summarizes the possible status codes with the LOCK and UNLOCK methods.

Table 19-1. Status codes for LOCK and UNLOCK methods

Status code	Defined by	Method	Effect
200 OK	HTTP	LOCK	Indicates successful locking.
201 Created	HTTP	LOCK	Indicates that a lock on a nonexistent resource succeeded by creating the resource.
204 No Content	HTTP	UNLOCK	Indicates successful unlocking.
207 Multi-Status	WebDAV	LOCK	The request was for locking multiple resources. Not all status codes returned were the same. Hence, they are all encapsulated in a 207 response.
403 Forbidden	HTTP	LOCK	Indicates that the client does not have permission to lock the resource.
412 Precondition Failed	HTTP	LOCK	Either the XML sent with the LOCK command indicated a condition to be satisfied and the server failed to complete the required condition, or the lock token could not be enforced.
422 Unprocessable Property	WebDAV	LOCK	Inapplicable semantics—an example may be specifying a non-zero Depth for a resource that is not a collection.
423 Locked	WebDAV	LOCK	Already locked.
424 Failed Dependency	WebDAV	UNLOCK	UNLOCK specifies other actions and their success as a condition for the unlocking. This error is returned if the dependency fails to complete.

Properties and META Data

Properties describe information about the resource, including the author's name, modification date, content rating, etc. META tags in HTML do provide a mechanism to embed this information as part of the content; however, many resources (such as any binary data) have no capability for embedding META data.

A distributed collaborative system such as WebDAV adds more complexity to the property requirement. For example, consider an author property: when a document gets edited, this property needs to be updated to reflect the new authors. WebDAV terms such dynamically modifiable properties “live” properties. The more permanent, static properties, such as Content-Type, are termed “dead” properties.

To support discovery and modification of properties, WebDAV extends HTTP to include two new methods, PROPFIND and PROPPATCH. Examples and corresponding XML elements are described in the following sections.

The PROPFIND Method

The PROPFIND (property find) method is used for retrieving the properties of a given file or a group of files (also known as a “collection”). PROPFIND supports three types of operations:

- Request all properties and their values.
- Request a selected set of properties and values.
- Request all property names.

Here’s the scenario where all the properties and their values are requested:

```
PROPFIND /ch-publish.fm HTTP/1.1
Host: minstar.inktomi.com
User-Agent: Mozilla/4.0 (compatible; MSIE 5.0; Windows NT)
Depth: 0
Cache-Control: no-cache
Connection: Keep-Alive
Content-Length: 0
```

The <propfind> request element specifies the properties to be returned from a PROPFIND method. The following list summarizes a few XML elements that are used with PROPFIND requests:

<allprop>

Requires all property names and values to be returned. To request all properties and their values, a WebDAV client may either send an <allprop> XML subelement as part of the <propfind> element, or submit a request with no body.

<propname>

Specifies the set of property names to be returned.

<prop>

A subelement of the <propfind> element. Specifies a specific property whose value is to be returned. For example: “<a:prop> <a:owner />..... </a:prop>”.

Here’s a response to a sample PROPFIND request:

```
HTTP/1.1 207 Multi-Status
Server: Microsoft-IIS/5.0
.....

<?xml version="1.0"?>
<a:multistatusxmlns:b="urn:uuid:*****/" xmlns:c="xml:" xmlns:a="DAV:">
  <a:response>
    <a:href>http://minstar/ch-publish.fm </a:href>
    <a:propstat>
      <a:status>HTTP/1.1 200OK</a:status>
      <a:prop>
        <a:getcontentlength b:dt="int">1155</a:getcontentlength>
        .....
        .....
```

```

    <a:ishidden b:dt="boolean">0</a:ishidden>
    <a:iscollection b:dt="boolean">0</a:iscollection>
  </a:prop>
</a:propstat>
</a:response></a:multistatus>

```

In this example, the server responds with a 207 Multi-Status code. WebDAV uses the 207 response for PROPFIND and a few other WebDAV methods that act simultaneously on multiple resources and potentially have different responses for each resource.

A few XML elements in the response need to be defined:

<multistatus>

A container for multiple responses.

<href>

Identifies the resource's URI.

<status>

Contains the HTTP status code for the particular request.

<propstat>

Groups one <status> element and one <prop> element. The <prop> element may contain one or more property name/value pairs for the given resource.

In the sample response listed above, the response is for one URI, *http://minstar/ch-publish.fm*. The <propstat> element embeds one <status> element and one <prop> element. For this URI, the server returned a 200 OK response, as defined by the <status> element. The <prop> element has several subelements; only some are listed in the example.

One instant application of PROPFIND is the support for directory listing. Given the expressability of a PROPFIND request, one single call can retrieve the entire hierarchy of the collection with all the properties of individual entities.

The PROPPATCH Method

The PROPPATCH method provides an atomic mechanism to set or remove multiple properties on a given resource. The atomicity will guarantee that either all of the requests are successful or none of them made it.

The base XML element for the PROPPATCH method is <propertyupdate>. It acts as a container for all the properties that need updating. The XML elements <set> and <remove> are used to specify the operation:

<set>

Specifies the property values to be set. The <set> contains one or more <prop> subelements, which in turn contains the name/value pairs of the properties to be set for the resource. If the property already exists, the value is replaced.

<remove>

Specifies the properties that are to be removed. Unlike with <set>, only the names of the properties are listed in the <prop> container.

This trivial example sets and removes the “owner” property:

```
<d:propertyupdate xmlns:d="DAV:" xmlns:o="http://name-space/scheme/">
  <d:set>
    <d:prop>
      <o:owner>Author A</o:owner>
    </d:prop>
  </d:set>

  <d:remove>
    <d:prop>
      <o:owner/>
    </d:prop>
  </d:remove>
</d:propertyupdate>
```

The response to PROPPATCH requests is very similar to that for PROPFIND requests. For more information, refer to RFC 2518.

Table 19-2 summarizes the status codes for the PROPFIND and PROPPATCH methods.

Table 19-2. Status codes for PROPFIND and PROPPATCH methods

Status code	Defined by	Methods	Effect
200 OK	HTTP	PROPFIND, PROPPATCH	Command success.
207 Multi-Status	WEBDAV	PROPFIND, PROPPATCH	When acting on one or more resources (or a collection), the status for each object is encapsulated into one 207 response. This is a typical success response.
401 Unauthorized	HTTP	PROPPATCH	Requires authorization to complete the property modification operation.
403 Forbidden	HTTP	PROPFIND, PROPPATCH	For PROPFIND, the client is not allowed to access the property. For PROPPATCH, the client may not change the property.
404 Not Found	HTTP	PROPFIND	No such property.
409 Conflict	HTTP	PROPPATCH	Conflict of update semantics—for example, trying to update a read-only property.
423 Locked	WebDAV	PROPPATCH	Destination resource is locked and there is no lock token or the lock token does not match.
507 Insufficient Storage	WebDAV	PROPPATCH	Not enough space for registering the modified property.

Collections and Namespace Management

A collection refers to a logical or physical grouping of resources in a predefined hierarchy. A classic example of a collection is a directory. Like directories in a filesystem,

collections act as containers of other resources, including other collections (equivalent to directories on the filesystem).

WebDAV uses the XML namespace mechanism. Unlike traditional namespaces, XML namespace partitions allow for precise structural control while preventing any namespace collisions.

WebDAV provides five methods for manipulating the namespace: DELETE, MKCOL, COPY, MOVE, and PROPFIND. PROPFIND was discussed previously in this chapter, but let's talk about the other methods.

The MKCOL Method

The MKCOL method allows clients to create a collection at the indicated URL on the server. At first sight, it may seem rather redundant to define an entire new method just for creating a collection. Overlaying on top of a PUT or POST method seems like a perfect alternative. The designers of the WebDAV protocol did consider these alternatives and still chose to define a new method. Some of the reasons behind that decision are:

- To have a PUT or a POST create a collection, the client needs to send some extra “semantic glue” along with the request. While this certainly is feasible, defining an ad hoc protocol may become tedious and error-prone.
- Most of the access-control mechanisms are based on the type of methods—only a few are allowed to create and delete resources in the repository. If we overload other methods, these access-control mechanisms will not work.

For example, a request might be:

```
MKCOL /publishing HTTP/1.1
Host: minstar
Content-Length: 0
Connection: Keep-Alive
```

And the response might be:

```
HTTP/1.1 201 Created
Server: Microsoft-IIS/5.0
Date: Fri, 10 May 2002 23:20:36 GMT
Location: http://minstar/publishing/
Content-Length: 0
```

Let us examine a few pathological cases:

- Suppose the collection already exists. If a MKCOL /colA request is made and colA already exists (i.e., namespace conflict), the request will fail with a 405 Method Not Allowed status code.
- If there are no write permissions, the MKCOL request will fail with a 403 Forbidden status code.

- If a request such as MKCOL /colA/colB is made and colA does not exist, the request will fail with a 409 Conflict status code.

Once the file or collection is created, we can delete it with the DELETE method.

The DELETE Method

We already saw the DELETE method in Chapter 3. WebDAV extends the semantics to cover collections.

If we need to delete a directory, the Depth header is needed. If the Depth header is not specified, the DELETE method assumes the Depth header to be set to infinity—that is, all the files in the directory and any subdirectories thereof are deleted. The response also has a Content-Location header identifying the collection that just got deleted. The request might read:

```
DELETE /publishing HTTP/1.0
Host: minstar
```

And the response might read:

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Tue, 14 May 2002 16:41:44 GMT
Content-Location: http://minstar/publishing/
Content-Type: text/xml
Content-Length: 0
```

When removing collections, there always is a chance that a file in the collection is locked by someone else and can't be deleted. In such a case, the collection itself can't be deleted, and the server replies with a 207 Multi-Status status code. The request might read:

```
DELETE /publishing HTTP/1.0
Host: minstar
```

And the response might read:

```
HTTP/1.1 207 Multi-Status
Server: Microsoft-IIS/5.0
Content-Location: http://minstar/publishing/
.....
<?xml version="1.0"?>
<a:multistatus xmlns:a="DAV:">
  <a:response>
    <a:href>http://minstar/index3/ch-publish.fm</a:href>
    <a:status> HTTP/1.1 423 Locked </a:status>
  </a:response>
</a:multistatus>
```

In this transaction, the <status> XML element contains the status code 403 Locked, indicating that the resource *ch-publish.fm* is locked by another user.

The COPY and MOVE Methods

As with MKCOL, there are alternatives to defining new methods for COPY and MOVE operations. One such alternative for the COPY method is to do a GET request on the source, thus downloading the resource, and then to upload it back to the server with a PUT request. A similar scenario could be envisioned for MOVE (with the additional DELETE operation). However, this process does not scale well—consider all the issues involved in managing a COPY or MOVE operation on a multilevel collection.

Both the COPY and MOVE methods use the request URL as the source and the contents of the Destination HTTP header as the target. The MOVE method performs some additional work beyond that of the COPY method: it copies the source URL to the destination, checks the integrity of the newly created URI, and then deletes the source. The request might read:

```
{COPY,MOVE} /publishing HTTP/1.1
Destination: http://minstar/pub-new
Depth: infinity
Overwrite: T
Host: minstar
```

And the response might read:

```
HTTP/1.1 201 Created
Server: Microsoft-IIS/5.0
Date: Wed, 15 May 2002 18:29:53 GMT
Location: http://minstar.inktomi.com/pub-new/
Content-Type: text/xml
Content-Length: 0
```

When acting on a collection, the behavior of COPY or MOVE is affected by the Depth header. In the absence of the Depth header, infinity is assumed (i.e., by default, the entire structure of the source directory will be copied or moved). If the Depth is set to zero, the method is applied just to the resource. If we are doing a copy or a move of a collection, only a collection with properties identical to those of the source is created at the destination—no internal members of the collection are copied or moved.

For obvious reasons, only a Depth value of infinity is allowed with the MOVE method.

Overwrite header effect

The COPY and MOVE methods also may use the Overwrite header. The Overwrite header can be set to either T or F. If it's set to T and the destination exists, a DELETE with a Depth value of infinity is performed on the destination resource before a COPY or MOVE operation. If the Overwrite flag is set to F and the destination resource exists, the operation will fail.

COPY/MOVE of properties

When a collection or an element is copied, all of its properties are copied by default. However, a request may contain an optional XML body that supplies additional information for the operation. You can specify that all properties must be copied successfully for the operation to succeed, or define which properties must be copied for the operation to succeed.

A couple of pathological cases to consider are:

- Suppose COPY or MOVE is applied to the output of a CGI program or other script that generates content. To preserve the semantics, if a file generated by a CGI script is to be copied or moved, WebDAV provides “src” and “link” XML elements that point to the location of the program that generated the page.
- The COPY and MOVE methods may not be able to completely duplicate all of the live properties. For example, consider a CGI program. If it is copied away from the *cgi-bin* directory, it may no longer be executed. The current specification of WebDAV makes COPY and MOVE a “best effort” solution, copying all the static properties and the appropriate live properties.

Locked resources and COPY/MOVE

If a resource currently is locked, both COPY and MOVE are prohibited from moving or duplicating the lock at the destination. In both cases, if the destination is to be created under an existing collection with its own lock, the duplicated or moved resource is added to the lock. Consider the following example:

```
COPY /publishing HTTP/1.1
Destination: http://minstar/archived/publishing-old
```

Let’s assume that */publishing* and */archived* already are under two different locks, lock1 and lock2. When the COPY operation completes, */publishing* continues to be under the scope of lock1, while, by virtue of moving into a collection that’s already locked by lock2, *publishing-old* gets added to lock2. If the operation was a MOVE, just *publishing-old* gets added to lock2.

Table 19-3 lists most of the possible status codes for the MKCOL, DELETE, COPY, and MOVE methods.

Table 19-3. Status codes for the MKCOL, DELETE, COPY, and MOVE methods

Status code	Defined by	Methods	Effect
102 Processing	WebDAV	MOVE, COPY	If the request takes longer than 20 seconds, the server sends this status code to keep clients from timing out. This usually is seen with a COPY or MOVE of a large collection.
201 Created	HTTP	MKCOL, COPY, MOVE	For MKCOL, a collection has been created. For COPY and MOVE, a resource/collection was copied or moved successfully.

Table 19-3. Status codes for the MKCOL, DELETE, COPY, and MOVE methods (continued)

Status code	Defined by	Methods	Effect
204 No Content	HTTP	DELETE, COPY, MOVE	For DELETE, a standard success response. For COPY and MOVE, the resource was copied over successfully or moved to replace an existing entity.
207 Multi-Status	WebDAV	MKCOL, COPY, MOVE	For MKCOL, a typical success response. For COPY and MOVE, if an error is associated with a resource other than the request URI, the server returns a 207 response with the XML body detailing the error.
403 Forbidden	HTTP	MKCOL, COPY, MOVE	For MKCOL, the server does not allow creation of a collection at the specified location. For COPY and MOVE, the source and destination are the same.
409 Conflict	HTTP	MKCOL, COPY, MOVE	In all cases, the methods are trying to create a collection or a resource when an intermediate collection does not exist—for example, trying to create colA/colB when colA does not exist.
412 Precondition Failed	HTTP	COPY, MOVE	Either the Overwrite header is set to F and the destination exists, or the XML body specifies a certain requirement (such as keeping the “liveness” property) and the COPY or MOVE methods are not able to retain the property.
415 Unsupported Media Type	HTTP	MKCOL	The server does not support or understand the creation of the request entity type.
422 Unprocessable Entity	WebDAV	MKCOL	The server does not understand the XML body sent with the request.
423 Locked	WebDAV	DELETE, COPY, MOVE	The source or the destination resource is locked, or the lock token supplied with the method does not match.
502 Bad Gateway	HTTP	COPY, MOVE	The destination is on a different server and permissions are missing.
507 Insufficient Storage	WebDAV	MKCOL COPY	There is not enough free space to create the resource.

Enhanced HTTP/1.1 Methods

WebDAV modifies the semantics of the HTTP methods DELETE, PUT, and OPTIONS. Semantics for the GET and HEAD methods remain unchanged. Operations performed by POST always are defined by the specific server implementation, and WebDAV does not modify any of the POST semantics. We already covered the DELETE method, in “Collections and Namespace Management.” We’ll discuss the PUT and OPTIONS methods here.

The PUT method

Though PUT is not defined by WebDAV, it is the only way for an author to transport the content to a shared site. We discussed the general functionality of PUT in Chapter 3. WebDAV modifies its behavior to support locking.

Consider the following example:

```
PUT /ch-publish.fm HTTP/1.1
Accept: */*
If:<http://minstar/index.htm>(<opaquelocktoken:*****>)
User-Agent: DAV Client (C)
Host: minstar.inktomi.com
Connection: Keep-Alive
Cache-Control: no-cache
Content-Length: 1155
```

To support locking, WebDAV adds an If header to the PUT request. In the above transaction, the semantics of the If header state that if the lock token specified with the If header matches the lock on the resource (in this case, *ch-publish.fm*), the PUT operation should be performed. The If header also is used with a few other methods, such as PROPPATCH, DELETE, MOVE, LOCK, UNLOCK, etc.

The OPTIONS method

We discussed OPTIONS in Chapter 3. This usually is the first request a WebDAV-enabled client makes. Using the OPTIONS method, the client tries to establish the capability of the WebDAV server. Consider a transaction in which the request reads:

```
OPTIONS /ch-publish.fm HTTP/1.1
Accept: */*
Host: minstar.inktomi.com
```

And the response reads:

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
MS-Author-Via: DAV
DASL: <DAV:sql>
DAV: 1, 2
Public: OPTIONS, TRACE, GET, HEAD, DELETE, PUT, POST, COPY, MOVE, MKCOL,PROPFIND,
PROPPATCH, LOCK, UNLOCK, SEARCH
Allow: OPTIONS, TRACE, GET, HEAD, DELETE, PUT, COPY, MOVE, PROPFIND,PROPPATCH,
SEARCH, LOCK, UNLOCK
```

There are several interesting headers in the response to the OPTIONS method. A slightly out-of-order examination follows:

- The DAV header carries the information about DAV compliance classes. There are two classes of compliance:

Class 1 compliance

Requires the server to comply with all MUST requirements in all sections of RFC 2518. If the resource complies only at the Class 1 level, it will send 1 with the DAV header.

Class 2 compliance

Meets all the Class 1 requirements and adds support for the LOCK method. Along with LOCK, Class 2 compliance requires support for the Timeout and

Lock-Token headers and the <supportedlock> and <lockdiscovery> XML elements. A value of 2 in the DAV header indicates Class 2 compliance.

In the above example, the DAV header indicates both Class 1 and Class 2 compliance.

- The Public header lists all methods supported by this particular server.
- The Allow header usually contains a subset of the Public header methods. It lists only those methods that are allowed on this particular resource (*ch-publish.fm*).
- The DASL header provides the type of query grammar used in the SEARCH method. In this case, it is sql. More details about the DASL header are provided at <http://www.webdav.org>.

Version Management in WebDAV

It may be ironic, given the “V” in “DAV,” but versioning is a feature that did not make the first cut. In a multi-author, collaborative environment, version management is critical. In fact, to completely fix the lost update problem (illustrated in Figure 19-3), locking and versioning are essential. Some of the common features associated with versioning are the ability to store and access previous document versions and the ability to manage the change history and any associated annotations detailing the changes.

Versioning was added to WebDAV in RFC 3253.

Future of WebDAV

WebDAV is well supported today. Working implementations of clients include IE 5.x and above, Windows Explorer, and Microsoft Office. On the server side, implementations include IIS5.x and above, Apache with mod_dav, and many others. Both Windows XP and Mac OS 10.x provide support for WebDAV out of the box; thus, any applications written to run on these operating systems are WebDAV-enabled natively.

For More Information

For more information, refer to:

<http://officeupdate.microsoft.com/frontpage/wpp/serk/>

Microsoft FrontPage 2000 Server Extensions Resource Kit.

<http://www.ietf.org/rfc/rfc2518.txt?number=2518>

“HTTP Extensions for Distributed Authoring—WEBDAV,” by Y. Goland, J. Whitehead, A. Faizi, S. Carter, and D. Jensen.

<http://www.ietf.org/rfc/rfc3253.txt?number=3253>

“Versioning Extensions to WebDAV,” by G. Clemm, J. Amsden, T. Ellison, C. Kaler, and J. Whitehead.

http://www.ics.uci.edu/pub/ietf/webdav/intro/webdav_intro.pdf

“WEBDAV: IETF Standard for Collaborative Authoring on the Web,” by J. Whitehead and M. Wiggins.

http://www.ics.uci.edu/~ejw/http-future/whitehead/http_pos_paper.html

“Lessons from WebDAV for the Next Generation Web Infrastructure,” by J. Whitehead.

<http://www.microsoft.com/msj/0699/dav/davtop.htm>

“Distributed Authoring and Versioning Extensions for HTTP Enable Team Authoring,” by L. Braginski and M. Powell.

<http://www.webdav.org/dasl/protocol/draft-dasl-protocol-00.html>

“DAV Searching & Locating,” by S. Reddy, D. Lowry, S. Reddy, R. Henderson, J. Davis, and A. Babich.

Redirection and Load Balancing

HTTP does not walk the Web alone. The data in an HTTP message is governed by many protocols on its journey. HTTP cares only about the endpoints of the journey—the sender and the receiver—but in a world with mirrored servers, web proxies, and caches, the destination of an HTTP message is not necessarily straightforward.

This chapter is about redirection technologies—network tools, techniques, and protocols that determine the final destination of an HTTP message. Redirection technologies usually determine whether the message ends up at a proxy, a cache, or a particular web server in a server farm. Redirection technologies may send your messages to places a client didn't explicitly request.

In this chapter, we'll take a look at the following redirection techniques, how they work, and what their load-balancing capabilities are (if any):

- HTTP redirection
- DNS redirection
- Anycast routing
- Policy routing
- IP MAC forwarding
- IP address forwarding
- The Web Cache Coordination Protocol (WCCP)
- The Intercache Communication Protocol (ICP)
- The Hyper Text Caching Protocol (HTCP)
- The Network Element Control Protocol (NECP)
- The Cache Array Routing Protocol (CARP)
- The Web Proxy Autodiscovery Protocol (WPAD)

Why Redirect?

Redirection is a fact of life in the modern Web because HTTP applications always want to do three things:

- Perform HTTP transactions reliably
- Minimize delay
- Conserve network bandwidth

For these reasons, web content often is distributed in multiple locations. This is done for reliability, so that if one location fails, another is available; it is done to lower response times, because if clients can access a nearer resource, they receive their requested content faster; and it's done to lower network congestion, by spreading out target servers. You can think of redirection as a set of techniques that help to find the “best” distributed content.

The subject of load balancing is included because redirection and load balancing coexist. Most redirection deployments include some form of load balancing; that is, they are capable of spreading incoming message load among a set of servers. Conversely, any form of load balancing involves redirection, because incoming messages must somehow be somehow among the servers sharing the load.

Where to Redirect

Servers, proxies, caches, and gateways all appear to clients as servers, in the sense that a client sends them an HTTP request, and they process it. Many redirection techniques work for servers, proxies, caches, and gateways because of their common, server-like traits. Other redirection techniques are specially designed for a particular class of endpoint and are not generally applicable. We'll see general techniques and specialized techniques in later sections of this chapter.

Web servers handle requests on a per-IP basis. Distributing requests to duplicate servers means that each request for a specific URL should be sent to an optimal web server (the one nearest to the client, or the least-loaded one, or some other optimization). Redirecting to a server is like sending all drivers in search of gasoline to the nearest gas station.

Proxies tend to handle requests on a per-protocol basis. Ideally, all HTTP traffic in the neighborhood of a proxy should go through the proxy. For instance, if a proxy cache is near various clients, all requests ideally will flow through the proxy cache, because the cache will store popular documents and serve them directly, avoiding longer and more expensive trips to the origin servers. Redirecting to a proxy is like siphoning off traffic on a main access road (no matter where it is headed) to a local shortcut.

Overview of Redirection Protocols

The goal of redirection is to send HTTP messages to available web servers as quickly as possible. The direction that an HTTP message takes on its way through the Internet is affected by the HTTP applications and routing devices it passes from, through, and toward. For example:

- The browser application that creates the client's message could be configured to send it to a proxy server.
- DNS resolvers choose the IP address that is used for addressing the message. This IP address can be different for different clients in different geographical locations.
- As the message passes through networks, it is divided into addressed packets; switches and routers examine the TCP/IP addressing on the packets and make decisions about routing the packets on that basis.
- Web servers can bounce requests back to different web servers with HTTP redirects.

Browser configuration, DNS, TCP/IP routing, and HTTP all provide mechanisms for redirecting messages. Notice that some methods, such as browser configuration, make sense only for redirecting traffic to proxies, while others, such as DNS redirection, can be used to send traffic to any server.

Table 20-1 summarizes the redirection methods used to redirect messages to servers, each of which is discussed later in this chapter.

Table 20-1. General redirection methods

Mechanism	How it works	Basis for rerouting	Limitations
HTTP redirection	Initial HTTP request goes to a first web server that chooses a "best" web server to serve the content. The first web server sends the client an HTTP redirect to the chosen server. The client resends the request to the chosen server.	Many options, from round-robin load balancing, to minimizing latency, to choosing the shortest path.	Can be slow—every transaction involves the extra redirect step. Also, the first server must be able to handle the request load.
DNS redirection	DNS server decides which IP address, among several, to return for the host-name in the URL.	Many options, from round-robin load balancing, to minimizing latency, to choosing the shortest path.	Need to configure DNS server.
Anycast addressing	Several servers use the same IP address. Each server masquerades as a backbone router. The other routers send packets addressed to the shared IP to the nearest server (believing they are sending packets to the nearest router).	Routers use built-in shortest-path routing capabilities.	Need to own/configure routers. Risks address conflicts. Established TCP connections can break if routing changes and packets associated with a connection get sent to different servers.

Table 20-1. General redirection methods (continued)

Mechanism	How it works	Basis for rerouting	Limitations
IP MAC forwarding	A network element such as a switch or router reads a packet's destination address; if the packet should be redirected, the switch gives the packet the destination MAC address of a server or proxy.	Save bandwidth and improve QOS. Load balance.	Server or proxy must be one hop away.
IP address forwarding	Layer-4 switch evaluates a packet's destination port and changes the IP address of a redirect packet to that of a proxy or mirrored server.	Save bandwidth and improve QOS. Load balance.	IP address of the client can be lost to the server/proxy.

Table 20-2 summarizes the redirection methods used to redirect messages to proxy servers.

Table 20-2. Proxy and cache redirection techniques

Mechanism	How it works	Basis for rerouting	Limitations
Explicit browser configuration	Web browser is configured to send HTTP messages to a nearby proxy, usually a cache. The configuration can be done by the end user or by a service that manages the browser.	Save bandwidth and improve QOS. Load balance.	Depends on ability to configure the browser.
Proxy auto-configuration (PAC)	Web browser retrieves a PAC file from a configuration server. The PAC file tells the browser what proxy to use for each URL.	Save bandwidth and improve QOS. Load balance.	Browser must be configured to query the configuration server.
Web Proxy Autodiscovery Protocol (WPAD)	Web browser asks a configuration server for the URL of a PAC file. Unlike PAC alone, the browser does not have to be configured with a specific configuration server.	The configuration server bases the URL on information in client HTTP request headers. Load balance.	Only a few browsers support WPAD.
Web Cache Coordination Protocol (WCCP)	Router evaluates a packet's destination address and encapsulates redirect packets with the IP address of a proxy or mirrored server. Works with many existing routers. Packet can be encapsulated, so the client's IP address is not lost.	Save bandwidth and improve QOS. Load balance.	Must use routers that support WCCP. Some topological limitations.
Internet Cache Protocol (ICP)	A proxy cache can query a group of sibling caches for requested content. Also supports cache hierarchies.	Obtaining content from a sibling or parent cache is faster than applying to the origin server.	False cache hits can arise because only the URL is used to request content.
Cache Array Routing Protocol (CARP)	A proxy cache hashing protocol. Allows a cache to forward a request to a parent cache. Unlike with ICP, the content on the caches is disjoint, and the group of caches acts as a single large cache.	Obtaining content from a nearby peer cache is faster than applying to the origin server.	CARP cannot support sibling relationships. All CARP clients must agree on the configuration; otherwise, different clients will send the same URI to different parents, reducing hit ratios.

Table 20-2. Proxy and cache redirection techniques (continued)

Mechanism	How it works	Basis for rerouting	Limitations
Hyper Text Caching Protocol (HTCP)	Participating proxy caches can query a group of sibling caches for requested content. Supports HTTP 1.0 and 1.1 headers to fine-tune cache queries.	Obtaining content from a sibling or parent cache is faster than applying to the origin server.	

General Redirection Methods

In this section, we will delve deeper into the various redirection methods that are commonly used for both servers and proxies. These techniques can be used to redirect traffic to a different (presumably more optimal) server or to vector traffic through a proxy. Specifically, we'll cover HTTP redirection, DNS redirection, any-cast addressing, IP MAC forwarding, and IP address forwarding.

HTTP Redirection

Web servers can send short redirect messages back to clients, telling them to try someplace else. Some web sites use HTTP redirection as a simple form of load balancing; the server that handles the redirect (the redirecting server) finds the least-loaded content server available and redirects the browser to that server. For widely distributed web sites, determining the “best” available server gets more complicated, taking into account not only the servers’ load but the Internet distance between the browser and the server. One advantage of HTTP redirection over some other forms of redirection is that the redirecting server knows the client’s IP address; in theory, it may be able to make a more informed choice.

Here’s how HTTP redirection works. In Figure 20-1a, Alice sends a request to *www.joes-hardware.com*:

```
GET /hammers.html HTTP/1.0
Host: www.joes-hardware.com
User-Agent: Mozilla/4.51 [en] (X11; U; IRIX 6.2 IP22)
```

In Figure 20-1b, instead of sending back a web page body with HTTP status code 200, the server sends back a redirect message with status code 302:

```
HTTP/1.0 302 Redirect
Server: Stronghold/2.4.2 Apache/1.3.6
Location: http://161.58.228.45/hammers.html
```

Now, in Figure 20-1c, the browser resends the request using the redirected URL, this time to host 161.58.228.45:

```
GET /hammers.html HTTP/1.0
Host: 161.58.228.45
User-Agent: Mozilla/4.51 [en] (X11; U; IRIX 6.2 IP22)
```

Another client could get redirected to a different server. In Figure 20-1d–f, Bob’s request gets redirected to 161.58.228.46.




Figure 20-1. HTTP redirection

HTTP redirection can vector requests across servers, but it has several disadvantages:

- A significant amount of processing power is required from the original server to determine which server to redirect to. Sometimes almost as much server horsepower is required to issue the redirect as would be to serve up the page itself.
- User delays are increased, because two round trips are required to access pages.
- If the redirecting server is broken, the site will be broken.

Because of these weaknesses, HTTP redirection usually is used in combination with some of the other redirection technique.

DNS Redirection

Every time a client tries to access Joe's Hardware's web site, the domain name `www.joes-hardware.com` must be resolved to an IP address. The DNS resolver may be the client's own operating system, a DNS server in the client's network, or a more remote DNS server. DNS allows several IP addresses to be associated to a single domain, and DNS resolvers can be configured or programmed to return varying IP addresses. The basis on which the resolver returns the IP address can run from the simple (round robin) to the complex (such as checking the load on several servers and returning the IP address of the least-loaded server).

In Figure 20-2, Joe runs four servers for *www.joes-hardware.com*. The DNS server has to decide which of four IP addresses to return for *www.joes-hardware.com*. The easiest DNS decision algorithm is a simple round robin.




Figure 20-2. DNS-based redirection

For a run-through of the DNS resolution process, see the DNS reference listed at the end of this chapter.

DNS round robin

One of the most common redirection techniques also is one of the simplest. DNS round robin uses a feature of DNS hostname resolution to balance load across a farm of web servers. It is a pure load-balancing strategy, and it does not take into account any factors about the location of the client relative to the server or the current stress on the server.

Let's look at what CNN.com really does. In early May of 2000, we used the *nslookup* Unix tool to find the IP addresses associated with CNN.com. Example 20-1 shows the results.*

Example 20-1. IP addresses for *www.cnn.com*

```
% nslookup www.cnn.com
Name:      cnn.com
```

* DNS results as of May 7, 2000 and resolved from Northern California. The particular values likely will change over time, and some DNS systems return different values based on client location.

Example 20-1. IP addresses for www.cnn.com (continued)

```
Addresses: 207.25.71.5, 207.25.71.6, 207.25.71.7, 207.25.71.8
           207.25.71.9, 207.25.71.12, 207.25.71.20, 207.25.71.22, 207.25.71.23
           207.25.71.24, 207.25.71.25, 207.25.71.26, 207.25.71.27, 207.25.71.28
           207.25.71.29, 207.25.71.30, 207.25.71.82, 207.25.71.199, 207.25.71.245
           207.25.71.246
Aliases:  www.cnn.com
```

The web site *www.cnn.com* actually is a farm of 20 distinct IP addresses! Each IP address might typically translate to a different physical server.

Multiple addresses and round-robin address rotation

Most DNS clients just use the first address of the multi-address set. To balance load, most DNS servers rotate the addresses each time a lookup is done. This address rotation often is called *DNS round robin*.

For example, three consecutive DNS lookups of *www.cnn.com* might return rotated lists of IP addresses like those shown in Example 20-2.

Example 20-2. Rotating DNS address lists

```
% nslookup www.cnn.com
Name:      cnn.com
Addresses: 207.25.71.5, 207.25.71.6, 207.25.71.7, 207.25.71.8
           207.25.71.9, 207.25.71.12, 207.25.71.20, 207.25.71.22, 207.25.71.23
           207.25.71.24, 207.25.71.25, 207.25.71.26, 207.25.71.27, 207.25.71.28
           207.25.71.29, 207.25.71.30, 207.25.71.82, 207.25.71.199, 207.25.71.245
           207.25.71.246

% nslookup www.cnn.com
Name:      cnn.com
Addresses: 207.25.71.6, 207.25.71.7, 207.25.71.8, 207.25.71.9
           207.25.71.12, 207.25.71.20, 207.25.71.22, 207.25.71.23, 207.25.71.24
           207.25.71.25, 207.25.71.26, 207.25.71.27, 207.25.71.28, 207.25.71.29
           207.25.71.30, 207.25.71.82, 207.25.71.199, 207.25.71.245, 207.25.71.246
           207.25.71.5

% nslookup www.cnn.com
Name:      cnn.com
Addresses: 207.25.71.7, 207.25.71.8, 207.25.71.9, 207.25.71.12
           207.25.71.20, 207.25.71.22, 207.25.71.23, 207.25.71.24, 207.25.71.25
           207.25.71.26, 207.25.71.27, 207.25.71.28, 207.25.71.29, 207.25.71.30
           207.25.71.82, 207.25.71.199, 207.25.71.245, 207.25.71.246, 207.25.71.5
           207.25.71.6
```

In Example 20-2:

- The first address of the first DNS lookup is 207.25.71.5.
- The first address of the second DNS lookup is 207.25.71.6.
- The first address of the third DNS lookup is 207.25.71.7.

DNS round robin for load balancing

Because most DNS clients just use the first address, the DNS rotation serves to balance load among servers. If DNS did not rotate the addresses, most clients would always send load to the first client.

Figure 20-3 shows how DNS round-robin rotation acts to balance load:

- When Alice tries to connect to *www.cnn.com*, she looks up the IP address using DNS and gets back 207.25.71.5 as the first IP address. Alice connects to the web server 207.25.71.5 in Figure 20-3c.
- When Bob subsequently tries to connect to *www.cnn.com*, he also looks up the IP address using DNS, but he gets back a different result because the address list has been rotated one position, based on Alice's previous request. Bob gets back 207.25.71.6 as the first IP address, and he connects to this server in Figure 20-3f.




Figure 20-3. DNS round robin load balances across servers in a server farm

The impact of DNS caching

DNS address rotation spreads the load around, because each DNS lookup to a server gets a different ordering of server addresses. However, this load balancing isn't perfect, because the results of the DNS lookup may be memorized and reused by applications, operating systems, and some primitive child DNS servers. Many web browsers

perform a DNS lookup for a host but then use the same address over and over again, to eliminate the cost of DNS lookups and because some servers prefer to keep talking to the same client. Furthermore, many operating systems perform the DNS lookup automatically, and cache the result, but don't rotate the addresses. Consequently, DNS round robin generally doesn't balance the load of a single client—one client typically will be stuck to one server for a long period of time.

But, even though DNS doesn't deal out the transactions of a single client across server replicas, it does a decent job of spreading the aggregate load of multiple clients. As long as there is a modestly large number of clients with similar demand, the load will be relatively well distributed across servers.

Other DNS-based redirection algorithms

We've already discussed how DNS rotates address lists with each request. However, some enhanced DNS servers use other techniques for choosing the order of the addresses:

Load-balancing algorithms

Some DNS servers keep track of the load on the web servers and place the least-loaded web servers at the front of the list.

Proximity-routing algorithms

DNS servers can attempt to direct users to nearby web servers, when the farm of web servers is geographically dispersed.

Fault-masking algorithms

DNS servers can monitor the health of the network and route requests away from service interruptions or other faults.

Typically, the DNS server that runs sophisticated server-tracking algorithms is an authoritative server that is under the control of the content provider (see Figure 20-4).

Several distributed hosting services use this DNS redirection model. One drawback of the model for services that look for nearby servers is that the only information that the authoritative DNS server uses to make its decision is the IP address of the local DNS server, not the IP address of the client.

Anycast Addressing

In anycast addressing, several geographically dispersed web servers have the exact same IP address and rely on the "shortest-path" routing capabilities of backbone routers to send client requests to the server nearest to the client. One way this method can work is for each web server to advertise itself as a router to a neighboring backbone router. The web server talks to its neighboring backbone router using a router communication protocol. When the backbone router receives packets aimed at the anycast address, it looks (as it usually would) for the nearest "router" that




Figure 20-4. DNS request involving authoritative server

accepts that IP address. Because the server will have advertised itself as a router for that address, the backbone router will send the server the packet.

In Figure 20-5, three servers front the same IP address, `10.10.10.1`. The Los Angeles (LA) server advertises this address to the LA router, the New York (NY) server advertises the same address to the NY router, and so on. The servers communicate with the routers using a router protocol. The routers automatically route client requests aimed at `10.10.10.1` to the nearest server that advertises the address. In Figure 20-5, a request for the IP address `10.10.10.1` will be routed to server 3.




Figure 20-5. Distributed anycast addressing

Anycast addressing is still an experimental technique. For distributed anycast to work, the servers must “speak router language” and the routers must be able to handle possible address conflicts, because Internet addressing basically assumes one server for one address. (If done improperly, this can lead to serious problems known as “route leaks.”) Distributed anycast is an emerging technology and might be a solution for content providers who control their own backbone networks.

IP MAC Forwarding

In Ethernet networks, HTTP messages are sent in the form of addressed data packets. Each packet has a layer-4 address, consisting of the source and destination IP address and TCP port numbers; this is the address to which layer 4–aware devices pay attention. Each packet also has a layer-2 address, the Media Access Control (MAC) address, to which layer-2 devices (commonly switches and hubs) pay attention. The job of layer-2 devices is to receive packets with particular incoming MAC addresses and forward them to particular outgoing MAC addresses.

In Figure 20-6, for example, the switch is programmed to send all traffic from MAC address “MAC3” to MAC address “MAC4.”




Figure 20-6. Layer-2 switch sending client requests to a gateway

A layer 4–aware switch is able to examine the layer-4 addressing (IP addresses and TCP port numbers) and make routing decisions based on this information. For example, a layer-4 switch could send all port 80–destined web traffic to a proxy. In Figure 20-7, the switch is programmed to send all port 80 traffic from MAC3 to MAC6 (a proxy cache). All other MAC3 traffic goes to MAC5.

Typically, if the requested HTTP content is in the cache and is fresh, the proxy cache serves it; otherwise, the proxy cache sends an HTTP request to the origin server for the content, on the client’s behalf. The switch sends port 80 requests from the proxy (MAC6) to the Internet gateway (MAC5).

Layer-4 switches that support MAC forwarding usually can forward requests to several proxy caches and balance the load among them. Likewise, HTTP traffic also can be forwarded to alternate HTTP servers.




Figure 20-7. MAC forwarding using a layer-4 switch

Because MAC address forwarding is point-to-point only, the server or proxy has to be located one hop away from the switch.

IP Address Forwarding

In IP address forwarding, a switch or other layer 4-aware device examines TCP/IP addressing on incoming packets and routes packets accordingly by changing the destination IP address, instead of the destination MAC address. An advantage over MAC forwarding is that the destination server need not be one hop away; it just needs to be located upstream from the switch, and the usual layer-3 end-to-end Internet routing gets the packet to the right place. This type of forwarding also is called Network Address Translation (NAT).

There is a catch, however: routing symmetry. The switch that accepts the incoming TCP connection from the client is managing that connection; the switch must send the response back to the client on that TCP connection. Therefore, any response from the destination server or proxy must return to the switch (see Figure 20-8).




Figure 20-8. A switch doing IP forwarding to a caching proxy or mirrored web server

Two ways to control the return path of the response are:

- Change the source IP address of the packet to the IP address of the switch. That way, regardless of the network configuration between the switch and server, the response packet goes to the switch. This is called *full NAT*, where the IP forwarding device translates both destination and source IP addresses. Figure 20-9 shows the effect of full NAT on a TCP/IP datagram. The consequence is that the client IP address is unknown to the web server, which might want it for authentication or billing purposes, for example.
- If the source IP address remains the client's IP address, make sure (from a hardware perspective) that no routes exist directly from server to client (bypassing the switch). This sometimes is called *half NAT*. The advantage here is that the server obtains the client IP address, but the disadvantage is the requirement of some control of the entire network between client and server.




Figure 20-9. Full NAT of a TCP/IP datagram

Network Element Control Protocol

The Network Element Control Protocol (NECP) allows network elements (NEs)—devices such as routers and switches that forward IP packets—to talk with server elements (SEs)—devices such as web servers and proxy caches that serve application layer requests. NECP does not explicitly support load balancing; it only offers a way for an SE to send an NE load-balancing information so that the NE can load balance as it sees fit. Like WCCP, NECP offers several ways to forward packets: MAC forwarding, GRE encapsulation, and NAT.

NECP supports the idea of exceptions. The SE can decide that it cannot service particular source IP addresses, and send those addresses to the NE. The NE can then forward requests from those IP addresses to the origin server.

Messages

The NECP messages are described in Table 20-3.

Table 20-3. NECP messages

Message	Who sends it	Meaning
NECP_NOOP		No operation—do nothing.
NECP_INIT	SE	SE initiates communication with NE. SE sends this message to NE after opening TCP connection with NE. SE must know which NE port to connect to.
NECP_INIT_ACK	NE	Acknowledges NECP_INIT.
NECP_KEEPALIVE	NE or SE	Asks if peer is alive.
NECP_KEEPALIVE_ACK	NE or SE	Answers keep-alive message.
NECP_START	SE	SE says “I am here and ready to accept network traffic.” Can specify a port.
NECP_START_ACK	NE	Acknowledges NECP_START.
NECP_STOP	SE	SE tells NE “stop sending me traffic.”
NECP_STOP_ACK	NE	NE acknowledges stop.
NECP_EXCEPTION_ADD	SE	SE says to add one or more exceptions to NE’s list. Exceptions can be based on source IP, destination IP, protocol (above IP), or port.
NECP_EXCEPTION_ADD_ACK	NE	Confirms EXCEPTION_ADD.
NECP_EXCEPTION_DEL	SE	Asks NE to delete one or more exceptions from its list.
NECP_EXCEPTION_DEL_ACK	NE	Confirms EXCEPTION_DEL.
NECP_EXCEPTION_RESET	SE	Asks NE to delete entire exception list.
NECP_EXCEPTION_RESET_ACK	NE	Confirms EXCEPTION_RESET.
NECP_EXCEPTION_QUERY	SE	Queries NE’s entire exception list.
NECP_EXCEPTION_RESP	NE	Responds to exception query.

Proxy Redirection Methods

So far, we have talked about general redirection methods. Content also may need to be accessed through various proxies (potentially for security reasons), or there might be a proxy cache in the network that a client should take advantage of (because it likely will be much faster to retrieve the cached content than it would be to go directly to the origin server).

But how do clients such as web browsers know to go to a proxy? There are three ways to determine this: by explicit browser configuration, by dynamic automatic configuration, and by transparent interception. We will discuss these three techniques in this section.

A proxy can, in turn, redirect client requests to a different proxy. For example, a proxy cache that does not have the content in its cache may choose to redirect the client to another cache. As this results in the response coming from a location different from the one from which the client requested the resource, we also will discuss several protocols used for peer proxy-cache redirection: the Internet Cache Protocol (ICP), the Cache Array Routing Protocol (CARP), and the Hyper Text Caching Protocol (HTCP).

Explicit Browser Configuration

Most browsers can be configured to contact a proxy server for content—there is a pull-down menu where the user can enter the proxy’s name or IP address and port number. The browser then contacts the proxy for all requests. Rather than relying on users to correctly configure their browsers to use proxies, some service providers require users to download preconfigured browsers. These browsers know the address of the proxy to contact.

Explicit browser configuration has two main disadvantages:

- Browsers configured to use proxies do not contact the origin server even if the proxy is not responding. If the proxy is down or if the browser is incorrectly configured, the user experiences connectivity problems.
- It is difficult to make changes in network architecture and propagate those changes to all end users. If a service provider wants to add more proxies or take some out of service, browser users have to change their proxy settings.

Proxy Auto-configuration

Explicit configuration of browsers to contact specific proxies can restrict changes in network architecture, because it depends on users to intervene and reconfigure their browsers. An automatic configuration methodology that allows browsers to dynamically configure themselves to contact the correct proxy server solves this problem. Such a methodology exists; it is called the Proxy Auto-configuration (PAC) protocol. PAC was defined by Netscape and is supported by the Netscape Navigator and Microsoft Internet Explorer browsers.

The basic idea behind PAC is to have browsers retrieve a special file, called the PAC file, which specifies the proxy to contact for each URL. The browser must be configured to contact a specific server for the PAC file. The browser then fetches the PAC file every time it is restarted.

The PAC file is a JavaScript file, which must define the function:

```
function FindProxyForURL(url, host)
```

Browsers call this function for every requested URL, as follows:

```
return_value = FindProxyForURL(url_of_request, host_in_url);
```

where the return value is a string specifying where the browser should request this URL. The return value can be a list of the names of proxies to contact (for example, “PROXY proxy1.domain.com; PROXY proxy2.domain.com”) or the string “DIRECT”, which means that the browser should go directly to the origin server, bypassing any proxies.

The sequence of operations that illustrate the request for and response to a browser’s request for the PAC file are illustrated in Figure 20-10. In this example, the server

sends back a PAC file with a JavaScript program. The JavaScript program has a function called “FindProxyForURL” that tells the browser to contact the origin server directly if the host in the requested URL is in the “netscape.com” domain, and to go to “proxy1.joes-cache.com” for all other requests. The browser calls this function for each URL it requests and connects according to the results returned by the function.




Figure 20-10. Proxy auto-configuration

The PAC protocol is quite powerful: the JavaScript program can ask the browser to choose a proxy based on any of a number of parameters related to the hostname, such as the DNS address and subnet, and even the day of week or time of day. PAC allows browsers automatically to contact the right proxy with changes in network architecture, as long as the PAC file is updated at the server to reflect changes to the proxy locations. The main drawback with PAC is that the browser must be configured to know which server to fetch the PAC file from, so it is not a completely automatic configuration system. WPAD, discussed in the next section, addresses this problem.

PAC, like preconfigured browsers, is used by some major ISPs today.

Web Proxy Autodiscovery Protocol

The Web Proxy Autodiscovery Protocol (WPAD) aims to provide a way for web browsers to find and use nearby proxies, without requiring the end user to manually

configure a proxy setting and without relying on transparent traffic interception. The general problem of defining a web proxy autodiscovery protocol is complicated by the existence of many discovery protocols to choose from and the differences in proxy-use configurations in different browsers.

This section contains an abbreviated and slightly reorganized version of the WPAD Internet draft. The draft currently is being developed as part of the Web Intermediaries Working Group of the IETF.

PAC file autodiscovery

WPAD enables HTTP clients to locate a PAC file and use the PAC file to discover the name of an appropriate proxy server. WPAD does not directly determine the name of the proxy server, because that would circumvent the additional capabilities provided by PAC files (load balancing, request routing to an array of servers, automated failover to backup proxy servers, and so on).

As shown in Figure 20-11, the WPAD protocol discovers a PAC file URL, also known as a *configuration URL* (CURL). The PAC file executes a JavaScript program that returns the address of an appropriate proxy server.




Figure 20-11. WPAD determines the PAC URL, which determines the proxy server

An HTTP client that implements the WPAD protocol:

- Uses WPAD to find the PAC file CURL
- Fetches the PAC file (a.k.a. configuration file, or CFILE) corresponding to the CURL
- Executes the PAC file to determine the proxy server
- Sends HTTP requests to the proxy server returned by the PAC file

WPAD algorithm

WPAD uses a series of resource-discovery techniques to determine the proper PAC file CURL. Multiple discovery techniques are specified, because not all organizations

can use all techniques. WPAD clients attempt each technique, one by one, until they succeed in obtaining a CURL.

The current WPAD specification defines the following techniques, in order:

- DHCP (Dynamic Host Discovery Protocol)
- SLP (Service Location Protocol)
- DNS well-known hostnames
- DNS SRV records
- DNS service URLs in TXT records

Of these five mechanisms, only the DHCP and DNS well-known hostname techniques are required for WPAD clients. We present more details in subsequent sections.

The WPAD client sends a series of resource-discovery requests, using the discovery mechanisms mentioned above, in order. Clients attempt only mechanisms that they support. Whenever a discovery attempt succeeds, the client uses the information obtained to construct a PAC CURL.

If a PAC file is retrieved successfully at that CURL, the process completes. If not, the client resumes where it left off in the predefined series of resource-discovery requests. If, after trying all discovery mechanisms, no PAC file is retrieved, the WPAD protocol fails and the client is configured to use no proxy server.

The client tries DHCP first, followed by SLP. If no PAC file is retrieved, the client moves on to the DNS-based mechanisms.

The client cycles through the DNS SRV, well-known hostnames, and DNS TXT record methods multiple times. Each time, the DNS query QNAME is made less and less specific. In this manner, the client can locate the most specific configuration information possible, but still can fall back on less specific information. Every DNS lookup has the QNAME prefixed with “wpad” to indicate the resource type being requested.

Consider a client with hostname *johns-desktop.development.foo.com*. This is the sequence of discovery attempts a complete WPAD client would perform:

- DHCP
- SLP
- DNS A lookup on “QNAME=wpad.development.foo.com”
- DNS SRV lookup on “QNAME=wpad.development.foo.com”
- DNS TXT lookup on “QNAME=wpad.development.foo.com”
- DNS A lookup on “QNAME=wpad.foo.com”
- DNS SRV lookup on “QNAME=wpad.foo.com”
- DNS TXT lookup on “QNAME=wpad.foo.com”

Refer to the WPAD specification to get detailed pseudocode that addresses the entire sequence of operations. The following sections discuss the two required mechanisms, DHCP and DNS A lookup. For more details about the remainder of the CURL discovery methods, refer to the WPAD specification.

CURL discovery using DHCP

For this mechanism to work, the CURLs must be stored on DHCP servers that WPAD clients can query. The WPAD client obtains the CURL by sending a DHCP query to a DHCP server. The CURL is contained in DHCP option code 252 (if the DHCP server is configured with this information). All WPAD client implementations are required to support DHCP. The DHCP protocol is detailed in RFC 2131. See RFC 2132 for a list of existing DHCP options.

If the WPAD client already has conducted DHCP queries during its initialization, the DHCP server might already have supplied that value. If the value is not available through a client OS API, the client sends a DHCPINFORM message to query the DHCP server to obtain the value.

The DHCP option code 252 for WPAD is of type STRING and is of arbitrary size. This string contains a URL that points to an appropriate PAC file. For example:

```
"http://server.domain/proxyconfig.pac"
```

DNS A record lookup

For this mechanism to work, the IP addresses of suitable proxy servers must be stored on DNS servers that the WPAD clients can query. The WPAD client obtains the CURL by sending an A record lookup to a DNS server. The result of a successful lookup contains an IP address for an appropriate proxy server.

WPAD client implementations are required to support this mechanism. This should be straightforward, as only basic DNS lookup of A records is required. See RFC 2219 for a description of using well-known DNS aliases for resource discovery. For WPAD, the specification uses “well known alias” of “wpad” for web proxy autodiscovery.

The client performs the following DNS lookup:

```
QNAME=wpad.TGTDOM., QCLASS=IN, QTYPE=A
```

A successful lookup contains an IP address from which the WPAD client constructs the CURL.

Retrieving the PAC file

Once a candidate CURL is created, the WPAD client usually makes a GET request to the CURL. When making requests, WPAD clients are required to send Accept headers with appropriate CFILEREAD format information that they are capable of handling. For example:

```
Accept: application/x-ns-proxy-autoconfig
```

In addition, if the CURL results in a redirect, the clients are required to follow the redirect to its final destination.

When to execute WPAD

The web proxy autodiscovery process is required to occur at least as frequently as one of the following:

- Upon startup of the web client—WPAD is performed only for the start of the first instance. Subsequent instances inherit the settings.
- Whenever there is an indication from the networking stack that the IP address of the client host has changed.

A web client can use either option, depending on what makes sense in its environment. In addition, the client must attempt a discovery cycle upon expiration of a previously downloaded PAC file in accordance with HTTP expiration. It's important that the client obey the timeouts and rerun the WPAD process when the PAC file expires.

Optionally, the client also may implement rerunning the WPAD process on failure of the currently configured proxy if the PAC file does not provide an alternative.

Whenever the client decides to invalidate the current PAC file, it must rerun the entire WPAD protocol to ensure it discovers the currently correct CURL. Specifically, there is no provision in the protocol to do an If-Modified-Since conditional fetch of the PAC file.

A number of network round trips might be required during the WPAD protocol broadcast and/or multicast communications. The WPAD protocol should not be invoked at a more frequent rate than specified above (such as per-URL retrieval).

WPAD spoofing

The IE 5 implementation of WPAD enabled web clients to detect proxy settings automatically, without user intervention. The algorithm used by WPAD prepends the hostname “wpad” to the fully qualified domain name and progressively removes subdomains until it either finds a WPAD server answering the hostname or reaches the third-level domain. For instance, web clients in the domain *a.b.microsoft.com* would query *wpad.a.b.microsoft*, *wpad.b.microsoft.com*, then *wpad.microsoft.com*.

This exposed a security hole, because in international usage (and certain other configurations), the third-level domain may not be trusted. A malicious user could set up a WPAD server and serve proxy configuration commands of her choice. Subsequent versions of IE (5.01 and later) rectified the problem.

Timeouts

WPAD goes through multiple levels of discovery, and clients must make sure that each phase is time-bound. When possible, limiting each phase to 10 seconds is

considered reasonable, but implementors may choose a different value that is more appropriate to their network properties. For example, a device implementation, operating over a wireless network, might use a much larger timeout to account for low bandwidth or high latency.

Administrator considerations

Administrators should configure at least one of the DHCP or DNS A record lookup methods in their environments, as those are the only two that all compatible clients are required to implement. Beyond that, configuring to support mechanisms earlier in the search order will improve client startup time.

One of the major motivations for this protocol structure was to support client location of nearby proxy servers. In many environments, there are several proxy servers (workgroup, corporate gateway, ISP, backbone).

There are a number of possible points at which “nearness” decisions can be made in the WPAD framework:

- DHCP servers for different subnets can return different answers. They also can base decisions on the client `ciaddr` field or the client identifier option.
- DNS servers can be configured to return different SRV/A/TXT resource records (RRs) for different domain suffixes (for example, QNAMEs *wpad.marketing.bigcorp.com* and *wpad.development.bigcorp.com*).
- The web server handling the CURL request can make decisions based on the User-Agent header, Accept header, client IP address/subnet/hostname, topological distribution of nearby proxy servers, etc. This can occur inside a CGI executable created to handle the CURL. As mentioned earlier, it even can be a proxy server handling the CURL requests and making these decisions.
- The PAC file may be expressive enough to select from a set of alternatives at runtime on the client. CARP is based on this premise for an array of caches. It is not inconceivable that the PAC file could compute some network distance or fitness metrics to a set of candidate proxy servers and then select the “closest” or “most responsive” server.

Cache Redirection Methods

We’ve discussed techniques to redirect traffic to general servers and specialized techniques to vector traffic to proxies and gateways. This final section will explain some of the more sophisticated redirection techniques used for caching proxy servers. These techniques are more complex than the previously discussed protocols because they try to be reliable, high-performance, and content-aware—dispatching requests to locations likely to have particular pieces of content.

WCCP Redirection

Cisco Systems developed the Web Cache Coordination Protocol (WCCP) to enable routers to redirect web traffic to proxy caches. WCCP governs communication between routers and caches so that routers can verify caches (make sure they are up and running), load balance among caches, and send specific types of traffic to specific caches. WCCP Version 2 (WCCP2) is an open protocol. We'll discuss WCCP2 here.

How WCCP redirection works

Here's a brief overview of how WCCP redirection works for HTTP (WCCP redirects other protocols similarly):

- Start with a network containing WCCP-enabled routers and caches that can communicate with one another.
- A set of routers and their target caches form a WCCP service group. The configuration of the service group specifies what traffic is sent where, how traffic is sent, and how load should be balanced among the caches in the service group.
- If the service group is configured to redirect HTTP traffic, routers in the service group send HTTP requests to caches in the service group.
- When an HTTP request arrives at a router in the service group, the router chooses one of the caches in the service group to serve the request (based on either a hash on the request's IP address or a mask/value set pairing scheme).
- The router sends the request packets to the cache, either by encapsulating the packets with the cache's IP address or by IP MAC forwarding.
- If the cache cannot serve the request, the packets are returned to the router for normal forwarding.
- The members of the service group exchange heartbeat messages with one another, continually verifying one another's availability.

WCCP2 messages

There are four WCCP2 messages, described in Table 20-4.

Table 20-4. WCCP2 messages

Message name	Who sends it	Information carried
WCCP2_HERE_I_AM	Cache to router	These messages tell routers that caches are available to receive traffic. The messages contain all of the cache's service group information. As soon as a cache joins a service group, it sends these messages to all routers in the group. These messages negotiate with routers sending WCCP2_I_SEE_YOU messages.
WCCP2_I_SEE_YOU	Router to cache	These messages respond to WCCP2_HERE_I_AM messages. They are used to negotiate the packet forwarding method, assignment method (who is the designated cache), packet return method, and security.

Table 20-4. WCCP2 messages (continued)

Message name	Who sends it	Information carried
WCCP2_REDIRECT_ASSIGN	Designated cache to router	These messages make assignments for load balancing; they send bucket information for hash table load balancing or mask/value set pair information for mask/value load balancing.
WCCP2_REMOVAL_QUERY	Router to cache that has not sent WCCP2_HERE_I_AM messages for $2.5 \times \text{HERE_I_AM_T}$ seconds	If a router does not receive WCCP2_HERE_I_AM messages regularly, the router sends this message to see if the cache should be removed from the service group. The proper response from a cache is three identical WCCP2_HERE_I_AM messages, separated by $\text{HERE_I_AM_T}/10$ seconds.

The WCCP2_HERE_I_AM message format is:

WCCP Message Header
 Security Info Component
 Service Info Component
 Web-cache Identity Info Component
 Web-cache View Info Component
 Capability Info Component (optional)
 Command Extension Component (optional)

The WCCP2_I_SEE_YOU message format is:

WCCP Message Header
 Security Info Component
 Service Info Component
 Router Identity Info Component
 Router View Info Component
 Capability Info Component (optional)
 Command Extension Component (optional)

The WCCP2_REDIRECT_ASSIGN message format is:

WCCP Message Header
 Security Info Component
 Service Info Component
 Assignment Info Component, or Alternate Assignment Component

The WCCP2_REMOVAL_QUERY message format is:

WCCP Message Header
 Security Info Component
 Service Info Component
 Router Query Info Component

Message components

Each WCCP2 message consists of a header and components. The WCCP header information contains the message type (Here I Am, I See You, Assignment, or Removal Query), WCCP version, and message length (not including the length of the header).

The components each begin with a four-octet header describing the component type and length. The component length does not include the length of the component header. The message components are described in Table 20-5.

Table 20-5. WCCP2 message components

Component	Description
Security Info	Contains the security option and security implementation. The security option can be: WCCP2_NO_SECURITY (0) WCCP2_MD5_SECURITY (1) If the option is no security, the security implementation field does not exist. If the option is MD5, the security implementation field is a 16-octet field containing the message checksum and Service Group password. The password can be no more than eight octets.
Service Info	Describes the service group. The service type ID can have two values: WCCP2_SERVICE_STANDARD (0) WCCP2_SERVICE_DYNAMIC (1) If the service type is standard, the service is a well-known service, defined entirely by service ID. HTTP is an example of a well-known service. If the service type is dynamic, the following settings define the service: priority, protocol, service flags (which determine hashing), and port.
Router Identity Info	Contains the router IP address and ID, and lists (by IP address) all of the web caches with which the router intends to communicate.
Web Cache Identity Info	Contains the web cache IP address and redirection hash table mapping.
Router View Info	Contains the router's view of the service group (identities of the routers and caches).
Web Cache View Info	Contains the web cache's view of the service group.
Assignment Info	Shows the assignment of a web cache to a particular hashing bucket.
Router Query Info	Contains the router's IP address, address of the web cache being queried, and ID of the last router in the service group that received a Here I Am message from the web cache.
Capabilities Info	Used by routers to advertise supported packet forwarding, load balancing, and packet return methods; used by web caches to let routers know what method the web cache prefers.
Alternate Assignment	Contains hash table assignment information for load balancing.
Assignment Map	Contains mask/value set elements for service group.
Command Extension	Used by web caches to tell routers they are shutting down; used by routers to acknowledge a cache shutdown.

Service groups

A *service group* consists of a set of WCCP-enabled routers and caches that exchange WCCP messages. The routers send web traffic to the caches in the service group. The configuration of the service group determines how traffic is distributed to caches in the service group. The routers and caches exchange service group configuration information in Here I Am and I See You messages.

GRE packet encapsulation

Routers that support WCCP redirect HTTP packets to a particular server by encapsulating them with the server's IP address. The packet encapsulation also contains an IP header proto field that indicates Generic Router Encapsulation (GRE). The existence of the proto field tells the receiving proxy that it has an encapsulated packet.

Because the packet is encapsulated, the client IP address is not lost. Figure 20-12 illustrates GRE packet encapsulation.




Figure 20-12. How a WCCP router changes an HTTP packet's destination IP address

WCCP load balancing

In addition to routing, WCCP routers can balance load among several receiving servers. WCCP routers and their receiving servers exchange *heartbeat messages* to let one another know they are up and running. If a particular receiving server stops sending heartbeat messages, the WCCP router sends request traffic directly to the Internet, instead of redirecting it to that node. When the node returns to service, the WCCP router begins receiving heartbeat messages again and resumes sending request traffic to the node.

Internet Cache Protocol

The Internet Cache Protocol (ICP) allows caches to look for content hits in sibling caches. If a cache does not have the content requested in an HTTP message, it can find out if the content is in a nearby sibling cache and, if so, retrieve the content from there, hopefully avoiding a more costly query to an origin server. ICP can be thought of as a cache clustering protocol. It is a redirection protocol in the sense that the final destination of an HTTP request message can be determined by a series of ICP queries.

ICP is an object discovery protocol. It asks nearby caches, all at the same time, if any of them have a particular URL in their caches. The nearby caches send back a short message saying “HIT” if they have that URL or “MISS” if they don’t. The cache is then free to open an HTTP connection to a neighbor cache that has the object.

ICP is simple and lightweight. ICP messages are 32-bit packed structures in network byte order, making them easy to parse. They are carried in UDP datagrams for

efficiency. UDP is an unreliable Internet protocol, which means that the data can get destroyed in transit, so programs that speak ICP need to have timeouts to detect lost datagrams.

Here is a brief description of the parts of an ICP message:

Opcode

The opcode is an 8-bit value that describes the meaning of the ICP message. Basic opcodes are ICP_OP_QUERY request messages and ICP_OP_HIT and ICP_OP_MISS response messages.

Version

The 8-bit version number describes the version number of the ICP protocol. The version of ICP used by Squid, documented in Internet RFC 2186, is Version 2.

Message length

The total size in bytes of the ICP message. Because there are only 16 bits, the ICP message size cannot be larger than 16,383 bytes. URLs usually are shorter than 16 KB; if they're longer than that, many web applications will not process them.

Request number

ICP-enabled caches use the request number to keep track of multiple simultaneous requests and replies. An ICP reply message always must contain the same request number as the ICP request message that triggered the reply.

Options

The 32-bit ICP options field is a bit vector containing flags that modify ICP behavior. ICPv2 defines two flags, both of which modify ICP_OP_QUERY requests. The ICP_FLAG_HIT_OBJ flag enables and disables the return of document data in ICP responses. The ICP_FLAG_SRC_RTT flag requests an estimate of the round-trip time to the origin server, as measured by a sibling cache.

Option data

The 32-bit option data is reserved for optional features. ICPv2 uses the low 16 bits of the option data to hold an optional round-trip time estimate from the sibling to the origin server.

Sender host address

A historic field carrying the 32-bit IP address of the message sender; not used in practice.

Payload

The contents of the payload vary depending on the message type. For ICP_OP_QUERY, the payload is a 4-byte original requester host address followed by a NUL-terminated URL. For ICP_OP_HIT_OBJ, the payload is a NUL-terminated URL followed by a 16 bit object size, followed by the object data.

For more information about ICP, refer to informational RFCs 2186 and 2187. Excellent ICP and peering references also are available from the U.S. National Laboratory for Applied Network Research (<http://www.nlanr.net/Squid/>).

Cache Array Routing Protocol

Proxy servers greatly reduce traffic to the Internet by intercepting requests from individual users and serving cached copies of the requested web objects. However, as the number of users grows, a high volume of traffic can overload the proxy servers themselves.

One solution to this problem is to use multiple proxy servers to distribute the load to a collection of servers. The Cache Array Routing Protocol (CARP) is a standard proposed by Microsoft Corporation and Netscape Communication Corporation to administer a collection of proxy servers such that an array of proxy servers appears to clients as one logical cache.

CARP is an alternative to ICP. Both CARP and ICP allow administrators to improve performance by using multiple proxy servers. This section discusses how CARP differs from ICP, the advantages and disadvantages of using CARP over ICP, and the technical details of how the CARP protocol is implemented.

Upon a cache miss in ICP, the proxy server queries neighboring caches using an ICP message format to determine the availability of the web object. The neighboring caches respond with either a “HIT” or a “MISS,” and the requesting proxy server uses these responses to select the most appropriate location from which to retrieve the object. If the ICP proxy servers were arranged in a hierarchical fashion, a miss would be elevated to the parent. Figure 20-13 diagrammatically shows how hits and misses are resolved using ICP.




Figure 20-13. ICP queries

Note that each of the proxy servers, connected together using the ICP protocol, is a standalone cache server with redundant mirrors of content, meaning that duplicate entries of web objects across proxy servers is possible. In contrast, the collection of servers connected using CARP operates as a single, large server with each component server containing only a fraction of the total cached documents. By applying a hash function to the URL of a web object, CARP maps web objects to a specific proxy server. Because each web object has a unique home, we can determine the location of the object by a single lookup, rather than polling each of the proxy servers configured in the collection. Figure 20-14 summarizes the CARP approach.




Figure 20-14. CARP redirection

Although Figure 20-14 shows the caching proxy as being the intermediary between clients and proxy servers that distributes the load to the various proxy servers, it is possible for this function to be served by the clients themselves. Commercial browsers such as Internet Explorer and Netscape Navigator can be configured to compute the hash function in the form of a plug-in that determines the proxy server to which the request should be sent.

Deterministic resolution of the proxy server in CARP means that it isn't necessary to send queries to all the neighbors, which means that this method requires fewer inter-cache messages to be sent out. As more proxy servers are added to the configuration, the collective cache system will scale fairly well. However, a disadvantage of CARP is that if one of the proxy servers becomes unavailable, the hash function needs to be modified to reflect this change, and the contents of the proxy servers must be reshuffled across the existing proxy servers. This can be expensive if the proxy server crashes often. In contrast, redundant content in ICP proxy servers

means that reshuffling is not required. Another potential problem is that, because CARP is a new protocol, existing proxy servers running only the ICP protocol may not be included readily in a CARP collection.

Having described the difference between CARP and ICP, let us now describe CARP in a little more detail. The CARP redirection method involves the following tasks:

- Keep a table of participating proxy servers. These proxy servers are polled periodically to see which ones are still active.
- For each participating proxy server, compute a hash function. The value returned by the hash function takes into account the amount of load this proxy can handle.
- Define a separate hash function that returns a number based on the URL of the requested web object.
- Take the sum of the hash function of the URL and the hash function of the proxy servers to get an array of numbers. The maximum value of these numbers determines the proxy server to use for the URL. Because the computed values are deterministic, subsequent requests for the same web object will be forwarded to the same proxy server.

These four chores can either be carried out on the browser, in a plug-in, or be computed on an intermediate server.

For each collection of proxy servers, create a table listing all of the servers in the collection. Each entry in the table should contain information about load factors, time-to-live (TTL) countdown values, and global parameters such as how often members should be polled. The load factor indicates how much load that machine can handle, which depends on the CPU speed and hard drive capacity of that machine. The table can be maintained remotely via an RPC interface. Once the fields in the tables have been updated by RPC, they can be made available or published to downstream clients and proxies. This publication is done in HTTP, allowing any client or proxy server to consume the table information without introducing another inter-proxy protocol. Clients and proxy servers simply use a well-known URL to retrieve the table.

The hash function used must ensure that the web objects are statistically distributed across the participating proxy servers. The load factor of the proxy server should be used to determine the statistic probability of a web object being assigned to that proxy.

In summary, the CARP protocol allows a group of proxy servers to be viewed as single collective cache, instead of a group of cooperating but separate caches (as in ICP). A deterministic request resolution path finds the home of a specific web object within a single hop. This eliminates the inter-proxy traffic that often is generated to

find the web object in a group of proxy servers in ICP. CARP also avoids duplicate copies of web objects being stored on different proxy servers, which has the advantage that the cache system collectively has a larger capacity for storing web objects but also has the disadvantage that a failure in any one proxy requires reshuffling some of the cache contents to existing proxies.

Hyper Text Caching Protocol

Earlier, we discussed ICP, a protocol that allows proxy caches to query siblings about the presence of documents. ICP, however, was designed with HTTP/0.9 in mind and therefore allows caches to send just the URL when querying a sibling about the presence of a resource. Versions 1.0 and 1.1 of HTTP introduced many new request headers that, along with the URL, are used to make decisions about document matching, so simply sending the URL in a request may not result in accurate responses.

The Hyper Text Caching Protocol (HTCP) reduces the probability of false hits by allowing siblings to query each other for the presence of documents using the URL and all of the request and response headers. Further, HTCP allows sibling caches to monitor and request the addition and deletion of selected documents in each other's caches and to make changes in the caching policies of each other's cached documents.

Figure 20-13, which illustrates an ICP transaction, also can be used to illustrate an HTCP transaction—HTCP is just another object discovery protocol. If a nearby cache has the document, the requesting cache can open an HTTP connection to the cache to get a copy of the document. The difference between an ICP and an HTCP transaction is in the level of detail in the requests and responses.

The structure of HTCP messages is illustrated in Figure 20-15. The Header portion includes the message length and message versions. The Data portion starts with the data length and includes opcodes, response codes, and some flags and IDs, and it terminates with the actual data. An optional Authentication section may follow the Data section.

Details of the message fields are as follows:

Header

The Header section consists of a 32-bit message length, an 8-bit major protocol version, and an 8-bit minor protocol version. The message length includes all of the header, data, and authentication sizes.

Data

The Data section contains the HTCP message and has the structure illustrated in Figure 20-15. The data components are described in Table 20-6.




Figure 20-15. HTCP message format

Table 20-6. HTCP data components

Component	Description
Data length	A 16-bit value of the number of bytes in the Data section including the length of the Length field itself.
Opcode	The 4-bit operation code for the HTCP transaction. The full list of opcodes is provided in Table 20-7.
Response code	A 4-bit key indicating the success or failure of the transaction. The possible values are: 0—Authentication was not used, but is needed 1—Authentication was used, but is not satisfactory 2—Unimplemented opcode 3—Major version not supported 4—Minor version not supported 5—Inappropriate, disallowed, or undesirable opcode
F1	F1 is overloaded—if the message is a request, F1 is a 1-bit flag set by the requestor indicating that it needs a response (F1=1); if the message is a response, F1 is a 1-bit flag indicating whether the response is to be interpreted as a response to the overall message (F1=1) or just as a response to the Opcode data fields (F1=0).
RR	A 1-bit flag indicating that the message is a request (RR=0) or a response (RR=1).
Transaction ID	A 32-bit value that, combined with the requestor's network address, uniquely identifies the HTCP transaction.
Opcode data	Opcode data is opcode-dependent. See Table 20-7.

Table 20-7 lists the HTCP opcodes and their corresponding data types.

Table 20-7. HTCP opcodes

Opcode	Value	Description	Response codes	Opcode data
NOP	0	Essentially a “ping” operation.	Always 0	None
TST	1		0 if entity is present, 1 if entity is not present	Contains the URL and request headers in the request and just response headers in the response
MON	2		0 if accepted, 1 if refused	
SET	3	The SET message allows caches to request changes in caching policies. See Table 20-9 for a list of the headers that can be used in SET messages.	0 if accepted, 1 if ignored	
CLR	4		0 if I had it, but it’s now gone; 1 if I had it, but I am keeping it; and 2 if I didn’t have it	

HTCP Authentication

The authentication portion of the HTCP message is optional. Its structure is illustrated in Figure 20-15, and its components are described in Table 20-8.

Table 20-8. HTCP authentication components

Component	Description
Auth length	The 16-bit number of bytes in the Authentication section of the message, including the length of the Length field itself.
Sig time	A 32-bit number representing the number of seconds since 00:00:00 Jan 1, 1970 GMT at the time that the signature is generated.
Sig expire	A 32-bit number representing the number of seconds since 00:00:00 Jan 1, 1970 GMT when the signature will expire.
Key name	A string that specifies the name of the shared secret. The Key section has two parts: the 16-bit length in bytes of the string that follows, followed by the stream of uninterrupted bytes of the string.
Signature	The HMAC-MD5 digest with a B value of 64 (representing the source and destination IP addresses and ports), the major and minor HTCP versions of the message, the Sig time and Sig expires values, the full HTCP data, and the key. The Signature also has two parts: the 16-bit length in bytes of the string, followed by the string.

Setting Caching Policies

The SET message allows caches to request changes in the caching policies of cached documents. The headers that can be used in SET messages are described in Table 20-9.

Table 20-9. List of Cache headers for modifying caching policies

Header	Description
Cache-Vary	The requestor has learned that the content varies on a set of headers different from the set in the response Vary header. This header overrides the response Vary header.
Cache-Location	The list of proxy caches that also may have copies of this object.
Cache-Policy	The requestor has learned the caching policies for this object in more detail than is specified in the response headers. Possible values are: "no-cache," meaning that the response is not cacheable but may be shareable among simultaneous requestors; "no-share," meaning that the object is not shareable; and "no-cache-cookie," meaning that the content may change as a result of cookies and caching therefore is not advised.
Cache-Flags	The requestor has modified the object's caching policies and the object may have to be treated specially and not necessarily in accordance with the object's actual policies.
Cache-Expiry	The actual expiration time for the document as learned by the requestor.
Cache-MD5	The requestor-computed MD5 checksum of the object, which may be different from the value in the Content-MD5 header, or may be supplied because the object does not have a Content-MD5 header.
Cache-to-Origin	The requestor-measured round-trip time to an origin server. The format of the values in this header is <i><origin server name or ip> <average round-trip time in seconds> <number of samples> <number of router hops between requestor and origin server></i> .

By allowing request and response headers to be sent in query messages to sibling caches, HTCP can decrease the false-hit rate in cache queries. By further allowing sibling caches to exchange policy information with each other, HTCP can improve sibling caches' ability to cooperate with each other.

For More Information

For more information, consult the following references:

DNS and Bind

Cricket Liu, Paul Albitz, and Mike Loukides, O'Reilly & Associates, Inc.

<http://www.wrec.org/Drafts/draft-cooper-webi-wpad-00.txt>

"Web Proxy Auto-Discovery Protocol."

<http://home.netscape.com/eng/mozilla/2.0/relnotes/demo/proxy-live.html>

"Navigator Proxy Auto-Config File Format."

<http://www.ietf.org/rfc/rfc2186.txt>

IETF RFC 2186, "Inter-cache Communication Protocol (ICP) Version 2," by D. Wessels and K. Claffy.

<http://icp.ircache.net/carp.txt>

"Cache Array Routing Protocol v1.0."

<http://www.ietf.org/rfc/rfc2756.txt>

IETF RFC 2756, "Hyper Text Caching Protocol (HTCP/0.0)," by P. Vixie and D. Wessels.

*<http://www.ietf.org/internet-drafts/draft-wilson-wrec-wccp-v2-00.txt>
[draft-wilson-wrec-wccp-v2-01.txt](http://www.ietf.org/internet-drafts/draft-wilson-wrec-wccp-v2-01.txt), “Web Cache Communication Protocol V2.0,”
by M. Cieslak, D. Forster, G. Tiwana, and R. Wilson.*

*<http://www.ietf.org/rfc/rfc2131.txt?number=2131>
“Dynamic Host Configuration Protocol.”*

*<http://www.ietf.org/rfc/rfc2132.txt?number=2132>
“DHCP Options and BOOTP Vendor Extensions.”*

*<http://www.ietf.org/rfc/rfc2608.txt?number=2608>
“Service Location Protocol, Version 2.”*

*<http://www.ietf.org/rfc/rfc2219.txt?number=2219>
“Use of DNS Aliases for Network Services.”*

Logging and Usage Tracking

Almost all servers and proxies log summaries of the HTTP transactions they process. This is done for a variety of reasons: usage tracking, security, billing, error detection, and so on. In this chapter, we take a brief tour of logging, examining what information about HTTP transactions typically is logged and what some of the common log formats contain.

What to Log?

For the most part, logging is done for two reasons: to look for problems on the server or proxy (e.g., which requests are failing), and to generate statistics about how web sites are accessed. Statistics are useful for marketing, billing, and capacity planning (for instance, determining the need for additional servers or bandwidth).

You could log all of the headers in an HTTP transaction, but for servers and proxies that process millions of transactions per day, the sheer bulk of all of that data quickly would get out of hand. You also would end up logging a lot of information that you don't really care about and may never even look at.

Typically, just the basics of a transaction are logged. A few examples of commonly logged fields are:

- HTTP method
- HTTP version of client and server
- URL of the requested resource
- HTTP status code of the response
- Size of the request and response messages (including any entity bodies)
- Timestamp of when the transaction occurred
- Referer and User-Agent header values

The HTTP method and URL tell what the request was trying to do—for example, GETting a resource or POSTing an order form. The URL can be used to track popularity of pages on the web site.

The version strings give hints about the client and server, which are useful in debugging strange or unexpected interactions between clients and servers. For example, if requests are failing at a higher-than-expected rate, the version information may point to a new release of a browser that is unable to interact with the server.

The HTTP status code tells what happened to the request: whether it was successful, the authorization attempt failed, the resource was found, etc. (See “Status Codes” in Chapter 3 for a list of HTTP status codes.)

The size of the request/response and the timestamp are used mainly for accounting purposes; i.e., to track how many bytes flowed into, out of, or through the application. The timestamp also can be used to correlate observed problems with the requests that were being made at the time.

Log Formats

Several log formats have become standard, and we’ll discuss some of the most common formats in this section. Most commercial and open source HTTP applications support logging in one or more of these common formats. Many of these applications also support the ability of administrators to configure log formats and create their own custom formats.

One of the main benefits of supporting (for applications) and using (for administrators) these more standard formats rests in the ability to leverage the tools that have been built to process and generate basic statistics from these logs. Many open source and commercial packages exist to crunch logs for reporting purposes, and by utilizing standard formats, applications and their administrators can plug into these resources.

Common Log Format

One of the most common log formats in use today is called, appropriately, the Common Log Format. Originally defined by NCSA, many servers use this log format as a default. Most commercial and open source servers can be configured to use this format, and many commercial and freeware tools exist to help parse common log files. Table 21-1 lists, in order, the fields of the Common Log Format.

Table 21-1. Common Log Format fields

Field	Description
remotehost	The hostname or IP address of the requestor’s machine (IP if the server was not configured to perform reverse DNS or cannot look up the requestor’s hostname)
username	If an <i>ident</i> lookup was performed, the requestor’s authenticated username ^a

Table 21-1. Common Log Format fields (continued)

Field	Description
auth-username	If authentication was performed, the username with which the requestor authenticated
timestamp	The date and time of the request
request-line	The exact text of the HTTP request line, "GET /index.html HTTP/1.1"
response-code	The HTTP status code that was returned in the response
response-size	The Content-Length of the response entity—if no entity was returned in the response, a zero is logged

^a RFC 931 describes the *ident* lookup used in this authentication. The *ident* protocol was discussed in Chapter 5.

Example 21-1 lists a few examples of Common Log Format entries.

Example 21-1. Common Log Format

```
209.1.32.44 - - [03/Oct/1999:14:16:00 -0400] "GET / HTTP/1.0" 200 1024
http-guide.com - dg [03/Oct/1999:14:16:32 -0400] "GET / HTTP/1.0" 200 477
http-guide.com - dg [03/Oct/1999:14:16:32 -0400] "GET /foo HTTP/1.0" 404 0
```

In these examples, the fields are assigned as follows:

Field	Entry 1	Entry 2	Entry 2
remotehost	209.1.32.44	http-guide.com	http-guide.com
username	<empty>	<empty>	<empty>
auth-username	<empty>	dg	dg
timestamp	03/Oct/1999:14:16:00 -0400	03/Oct/1999:14:16:32 -0400	03/Oct/1999:14:16:32 -0400
request-line	GET / HTTP/1.0	GET / HTTP/1.0	GET /foo HTTP/1.0
response-code	200	200	404
response-size	1024	477	0

Note that the *remotehost* field can be either a hostname, as in *http-guide.com*, or an IP address, such as 209.1.32.44.

The dashes in the second (username) and third (auth-username) fields indicate that the fields are empty. This indicates that either an *ident* lookup did not occur (second field empty) or authentication was not performed (third field empty).

Combined Log Format

Another commonly used log format is the Combined Log Format. This format is supported by servers such as Apache. The Combined Log Format is very similar to the Common Log Format; in fact, it mirrors it exactly, with the addition of two fields (listed in Table 21-2). The User-Agent field is useful in noting which HTTP client applications are making the logged requests, while the Referer field provides more detail about where the requestor found this URL.

Table 21-2. Additional Combined Log Format fields

Field	Description
Referer	The contents of the Referer HTTP header
User-Agent	The contents of the User-Agent HTTP header

Example 21-2 gives an example of a Combined Log Format entry.

Example 21-2. Combined Log Format

```
209.1.32.44 - - [03/Oct/1999:14:16:00 -0400] "GET / HTTP/1.0" 200 1024 "http://www.joes-  
hardware.com/" "5.0: Mozilla/4.0 (compatible; MSIE 5.0; Windows 98)"
```

In Example 21-2, the Referer and User-Agent fields are assigned as follows:

Field	Value
Referer	http://www.joes-hardware.com/
User-Agent	5.0: Mozilla/4.0 (compatible; MSIE 5.0; Windows 98)

The first seven fields of the example Combined Log Format entry in Example 21-2 are exactly as they would be in the Common Log Format (see the first entry in Example 21-1). The two new fields, Referer and User-Agent, are tacked onto the end of the log entry.

Netscape Extended Log Format

When Netscape entered into the commercial HTTP application space, it defined for its servers many log formats that have been adopted by other HTTP application developers. Netscape's formats derive from the NCSA Common Log Format, but they extend that format to incorporate fields relevant to HTTP applications such as proxies and web caches.

The first seven fields in the Netscape Extended Log Format are identical to those in the Common Log Format (see Table 21-1). Table 21-3 lists, in order, the new fields that the Netscape Extended Log Format introduces.

Table 21-3. Additional Netscape Extended Log Format fields

Field	Description
proxy-response-code	If the transaction went through a proxy, the HTTP response code from the server to the proxy
proxy-response-size	If the transaction went through a proxy, the Content-Length of the server's response entity sent to the proxy
client-request-size	The Content-Length of any body or entity in the client's request to the proxy
proxy-request-size	If the transaction went through a proxy, the Content-Length of any body or entity in the proxy's request to the server
client-request-hdr-size	The length, in bytes, of the client's request headers

Table 21-3. Additional Netscape Extended Log Format fields (continued)

Field	Description
proxy-response-hdr-size	If the transaction went through a proxy, the length, in bytes, of the proxy's response headers that were sent to the requestor
proxy-request-hdr-size	If the transaction went through a proxy, the length, in bytes, of the proxy's request headers that were sent to the server
server-response-hdr-size	The length, in bytes, of the server's response headers
proxy-timestamp	If the transaction went through a proxy, the elapsed time for the request and response to travel through the proxy, in seconds

Example 21-3 gives an example of a Netscape Extended Log Format entry.

Example 21-3. Netscape Extended Log Format

```
209.1.32.44 - - [03/Oct/1999:14:16:00-0400] "GET / HTTP/1.0" 200 1024 200 1024 0 0 215 260
279 254 3
```

In this example, the extended fields are assigned as follows:

Field	Value
proxy-response-code	200
proxy-response-size	1024
client-request-size	0
proxy-request-size	0
client-request-hdr-size	215
proxy-response-hdr-size	260
proxy-request-hdr-size	279
server-response-hdr-size	254
proxy-timestamp	3

The first seven fields of the example Netscape Extended Log Format entry in Example 21-3 mirror the entries in the Common Log Format example (see the first entry in Example 21-1).

Netscape Extended 2 Log Format

Another Netscape log format, the Netscape Extended 2 Log Format, takes the Extended Log Format and adds further information relevant to HTTP proxy and web caching applications. These extra fields help paint a better picture of the interactions between an HTTP client and an HTTP proxy application.

The Netscape Extended 2 Log Format derives from the Netscape Extended Log Format, and its initial fields are identical to those listed in Table 21-3 (it also extends the Common Log Format fields listed in Table 21-1).

Table 21-4 lists, in order, the additional fields of the Netscape Extended 2 Log Format.

Table 21-4. Additional Netscape Extended 2 Log Format fields

Field	Description
route	The route that the proxy used to make the request for the client (see Table 21-5)
client-finish-status-code	The client finish status code; specifies whether the client request to the proxy completed successfully (FIN) or was interrupted (INTR)
proxy-finish-status-code	The proxy finish status code; specifies whether the proxy request to the server completed successfully (FIN) or was interrupted (INTR)
cache-result-code	The cache result code; tells how the cache responded to the request ^a

^a Table 21-7 lists the Netscape cache result codes.

Example 21-4 gives an example of a Netscape Extended 2 Log Format entry.

Example 21-4. Netscape Extended 2 Log Format

```
209.1.32.44 - - [03/Oct/1999:14:16:00-0400] "GET / HTTP/1.0" 200 1024 200 1024 0 0 215 260
279 254 3 DIRECT FIN FIN WRITTEN
```

The extended fields in this example are assigned as follows:

Field	Value
route	DIRECT
client-finish-status-code	FIN
proxy-finish-status-code	FIN
cache-result-code	WRITTEN

The first 16 fields in the Netscape Extended 2 Log Format entry in Example 21-4 mirror the entries in the Netscape Extended Log Format example (see Example 21-3).

Table 21-5 lists the valid Netscape route codes.

Table 21-5. Netscape route codes

Value	Description
DIRECT	The resource was fetched directly from the server.
PROXY(host:port)	The resource was fetched through the proxy "host."
SOCKS(socks:port)	The resource was fetched through the SOCKS server "host."

Table 21-6 lists the valid Netscape finish codes.

Table 21-6. Netscape finish status codes

Value	Description
-	The request never even started.
FIN	The request was completed successfully.

Table 21-6. Netscape finish status codes (continued)

Value	Description
INTR	The request was interrupted by the client or ended by a proxy/server.
TIMEOUT	The request was timed out by the proxy/server.

Table 21-7 lists the valid Netscape cache codes.*

Table 21-7. Netscape cache codes

Code	Description
-	The resource was uncacheable.
WRITTEN	The resource was written into the cache.
REFRESHED	The resource was cached and it was refreshed.
NO-CHECK	The cached resource was returned; no freshness check was done.
UP-TO-DATE	The cached resource was returned; a freshness check was done.
HOST-NOT-AVAILABLE	The cached resource was returned; no freshness check was done because the remote server was not available.
CL-MISMATCH	The resource was not written to the cache; the write was aborted because the Content-Length did not match the resource size.
ERROR	The resource was not written to the cache due to some error; for example, a timeout occurred or the client aborted the transaction.

Netscape applications, like many other HTTP applications, have other log formats too, including a Flexible Log Format and a means for administrators to output custom log fields. These formats allow administrators greater control and the ability to customize their logs by choosing which parts of the HTTP transaction (headers, status, sizes, etc.) to report in their logs.

The ability for administrators to configure custom formats was added because it is difficult to predict what information administrators will be interested in getting from their logs. Many other proxies and servers also have the ability to emit custom logs.

Squid Proxy Log Format

The Squid proxy cache (<http://www.squid-cache.org>) is a venerable part of the Web. Its roots trace back to one of the early web proxy cache projects (<ftp://ftp.cs.colorado.edu/pub/techreports/schwartz/Harvest.Conf.ps.Z>). Squid is an open source project that has been extended and enhanced by the open source community over the years. Many tools have been written to help administer the Squid application, including tools to help process, audit, and mine its logs. Many subsequent proxy caches adopted the Squid format for their own logs so that they could leverage these tools.

* Chapter 7 discusses HTTP caching in detail.

The format of a Squid log entry is fairly simple. Its fields are summarized in Table 21-8.

Table 21-8. Squid Log Format fields

Field	Description
timestamp	The timestamp when the request arrived, in seconds since January 1, 1970 GMT.
time-elapsed	The elapsed time for request and response to travel through the proxy, in milliseconds.
host-ip	The IP address of the client's (requestor's) host machine.
result-code/status	The result field is a Squid-ism that tells what action the proxy took during this request ^a ; the code field is the HTTP response code that the proxy sent to the client.
size	The length of the proxy's response to the client, including HTTP response headers and body, in bytes.
method	The HTTP method of the client's request.
url	The URL in the client's request. ^b
rfc931-ident ^c	The client's authenticated username. ^d
hierarchy/from	Like the route field in Netscape formats, the hierarchy field tells what route the proxy used to make the request for the client. ^e The from field tells the name of the server that the proxy used to make the request.
content-type	The Content-Type of the proxy response entity.

^a Table 21-9 lists the various result codes and their meanings.

^b Recall from Chapter 2 that proxies often log the entire requested URL, so if a username and password component are in the URL, a proxy can inadvertently record this information.

^c The rfc931-ident, hierarchy/from, and content-type fields were added in Squid 1.1. Previous versions did not have these fields.

^d RFC 931 describes the *ident* lookup used in this authentication.

^e <http://squid.nlanr.net/Doc/FAQ/FAQ-6.html#ss6.6> lists all of the valid Squid hierarchy codes.

Example 21-5 gives an example of a Squid Log Format entry.

Example 21-5. Squid Log Format

```
99823414 3001 209.1.32.44 TCP_MISS/200 4087 GET http://www.joes-hardware.com - DIRECT/
proxy.com text/html
```

The fields are assigned as follows:

Field	Value
timestamp	99823414
time-elapsed	3001
host-ip	209.1.32.44
action-code	TCP_MISS
status	200
size	4087
method	GET
URL	http://www.joes-hardware.com

Field	Value
RFC 931 ident	-
hierarchy	DIRECT ^a
from	proxy.com
content-type	text/html

^a The DIRECT Squid hierarchy value is the same as the DIRECT route value in Netscape log formats.

Table 21-9 lists the various Squid result codes.*

Table 21-9. Squid result codes

Action	Description
TCP_HIT	A valid copy of the resource was served out of the cache.
TCP_MISS	The resource was not in the cache.
TCP_REFRESH_HIT	The resource was in the cache but needed to be checked for freshness. The proxy revalidated the resource with the server and found that the in-cache copy was indeed still fresh.
TCP_REF_FAIL_HIT	The resource was in the cache but needed to be checked for freshness. However, the revalidation failed (perhaps the proxy could not connect to the server), so the “stale” resource was returned.
TCP_REFRESH_MISS	The resource was in the cache but needed to be checked for freshness. Upon checking with the server, the proxy learned that the resource in the cache was out of date and received a new version.
TCP_CLIENT_REFRESH_MISS	The requestor sent a Pragma: no-cache or similar Cache-Control directive, so the proxy was forced to fetch the resource.
TCP_IMS_HIT	The requestor issued a conditional request, which was validated against the cached copy of the resource.
TCP_SWAPFAIL_MISS	The proxy thought the resource was in the cache but for some reason could not access it.
TCP_NEGATIVE_HIT	A cached response was returned, but the response was a negatively cached response. Squid supports the notion of caching errors for resources—for example, caching a 404 Not Found response—so if multiple requests go through the proxy-cache for an invalid resource, the error is served from the proxy cache.
TCP_MEM_HIT	A valid copy of the resource was served out of the cache, and the resource was in the proxy cache’s memory (as opposed to having to access the disk to retrieve the cached resource).
TCP_DENIED	The request for this resource was denied, probably because the requestor does not have permission to make requests for this resource.
TCP_OFFLINE_HIT	The requested resource was retrieved from the cache during its <i>offline</i> mode. Resources are not validated when Squid (or another proxy using this format) is in offline mode.
UDP_*	The UDP_* codes indicate that requests were received through the UDP interface to the proxy. HTTP normally uses the TCP transport protocol, so these requests are not using the HTTP protocol. ^a

* Several of these action codes deal more with the internals of the Squid proxy cache, so not all of them are used by other proxies that implement the Squid Log Format.

Table 21-9. Squid result codes (continued)

Action	Description
UDP_HIT	A valid copy of the resource was served out of the cache.
UDP_MISS	The resource was not in the cache.
UDP_DENIED	The request for this resource was denied, probably because the requestor does not have permission to make requests for this resource.
UDP_INVALID	The request that the proxy received was invalid.
UDP_MISS_NOFETCH	Used by Squid during specific operation modes or in the cache of frequent failures. A cache miss was returned and the resource was not fetched.
NONE	Logged sometimes with errors.
TCP_CLIENT_REFRESH	See TCP_CLIENT_REFRESH_MISS.
TCP_SWAPFAIL	See TCP_SWAPFAIL_MISS.
UDP_RELOADING	See UDP_MISS_NOFETCH.

^a Squid has its own protocol for making these requests: ICP. This protocol is used for cache-to-cache requests. See <http://www.squid-cache.org> for more information.

Hit Metering

Origin servers often keep detailed logs for billing purposes. Content providers need to know how often URLs are accessed, advertisers want to know how often their ads are shown, and web authors want to know how popular their content is. Logging works well for tracking these things when clients visit web servers directly.

However, caches stand between clients and servers and prevent many accesses from reaching servers (the very purpose of caches).^{*} Because caches handle many HTTP requests and satisfy them without visiting the origin server, the server has no record that a client accessed its content, creating omissions in log files.

Missing log data makes content providers resort to *cache busting* for their most important pages. Cache busting refers to a content producer intentionally making certain content uncacheable, so all requests for this content must go to the origin server.[†] This allows the origin server to log the access. Defeating caching might yield better logs, but it slows down requests and increases load on the origin server and network.

Because proxy caches (and some clients) keep their own logs, if servers could get access to these logs—or at least have a crude way to determine how often their content is served by a proxy cache—cache busting could be avoided. The proposed Hit Metering protocol, an extension to HTTP, suggests a solution to this problem. The Hit Metering protocol requires caches to periodically report cache access statistics to origin servers.

^{*} Recall that virtually every browser has a cache.

[†] Chapter 7 describes how HTTP responses can be marked as uncacheable.

RFC 2227 defines the Hit Metering protocol in detail. This section provides a brief tour of the proposal.

Overview

The Hit Metering protocol defines an extension to HTTP that provides a few basic facilities that caches and servers can implement to share access information and to regulate how many times cached resources can be used.

Hit Metering is, by design, not a complete solution to the problem caches pose for logging access, but it does provide a basic means for obtaining metrics that servers want to track. The Hit Metering protocol has not been widely implemented or deployed (and may never be). That said, a cooperative scheme like Hit Metering holds some promise of providing accurate access statistics while retaining caching performance gains. Hopefully, that will be motivation to implement the Hit Metering protocol instead of marking content uncacheable.

The Meter Header

The Hit Metering extension proposes the addition of a new header, Meter, that caches and servers can use to pass to each other directives about usage and reporting, much like the Cache-Control header allows caching directives to be exchanged.

Table 21-10 defines the various directives and who can pass them in the Meter header.

Table 21-10. Hit Metering directives

Directive	Abbreviation	Who	Description
will-report-and-limit	w	Cache	The cache is capable of reporting usage and obeying any usage limits the server specifies.
wont-report	x	Cache	The cache is able to obey usage limits but won't report usage.
wont-limit	y	Cache	The cache is able to report usage but won't limit usage.
count	c	Cache	The reporting directive, specified as "uses/reuses" integers—for example, "count=2/4". ^a
max-uses	u	Server	Allows the server to specify the maximum number of times a response can be used by a cache—for example, "max-uses=100".
max-reuses	r	Server	Allows the server to specify the maximum number of times a response can be reused by a cache—for example, "max-reuses=100".
do-report	d	Server	The server requires proxies to send usage reports.
dont-report	e	Server	The server does not want usage reports.
timeout	t	Server	Allows the server to specify a timeout on the metering of a resource. The cache should send a report at or before the specified timeout, plus or minus 1 minute. The timeout is specified in minutes—for example, "timeout=60".
wont-ask	n	Server	The server does not want any metering information.

^a Hit Metering defines a *use* as satisfying a request with the response, whereas a *reuse* is revalidating a client request.

Figure 21-1 shows an example of Hit Metering in action. The first part of the transaction is just a normal HTTP transaction between a client and proxy cache, but in the proxy request, note the insertion of the Meter header and the response from the server. Here, the proxy is informing the server that it is capable of doing Hit Metering, and the server in turn is asking the proxy to report its hit counts.




Figure 21-1. Hit Metering example

The request completes as it normally would, from the client's perspective, and the proxy begins tracking hits to that resource on behalf of the server. Later, the proxy tries to revalidate the resource with the server. The proxy embeds the metered information it has been tracking in the conditional request to the server.

A Word on Privacy

Because logging really is an administrative function that servers and proxies perform, the whole operation is transparent to users. Often, they may not even be aware that their HTTP transactions are being logged—in fact, many users probably do not even know that they are using the HTTP protocol when accessing content on the Web.

Web application developers and administrators need to be aware of the implications of tracking a user's HTTP transactions. Much can be gleaned about a user based on the information he retrieves. This information obviously can be put to bad use—discrimination, harassment, blackmail, etc. Web servers and proxies that log must be vigilant in protecting the privacy of their end users.

Sometimes, such as in work environments, tracking a user's usage to make sure he is not goofing off may be appropriate, but administrators also should make public the fact that people's transactions are being monitored.

In short, logging is a very useful tool for the administrator and developer—just be aware of the privacy infringements that logs can have without the permission or knowledge of the users whose actions are being logged.

For More Information

For more information on logging, refer to:

<http://httpd.apache.org/docs/logs.html>

“Apache HTTP Server: Log Files.” Apache HTTP Server Project web site.

<http://www.squid-cache.org/Doc/FAQ/FAQ-6.html>

“Squid Log Files.” Squid Proxy Cache web site.

<http://www.w3.org/Daemon/User/Config/Logging.html#common-logfile-format>

“Logging Control in W3C httpd.”

<http://www.w3.org/TR/WD-logfile.html>

“Extended Log File Format.”

<http://www.ietf.org/rfc/rfc2227.txt>

RFC 2227, “Simple Hit-Metering and Usage-Limiting for HTTP,” by J. Mogul and P. Leach.

Appendixes

This collection of appendixes contains useful reference tables, background information, and tutorials on a variety of topics relevant to HTTP architecture and implementation:

- Appendix A, *URI Schemes*
- Appendix B, *HTTP Status Codes*
- Appendix C, *HTTP Header Reference*
- Appendix D, *MIME Types*
- Appendix E, *Base-64 Encoding*
- Appendix F, *Digest Authentication*
- Appendix G, *Language Tags*
- Appendix H, *MIME Charset Registry*

APPENDIX A

URI Schemes

Many URI schemes have been defined, but few are in common use. Generally speaking, those URI schemes with associated RFCs are in more common use, though there are a few schemes that have been developed by leading software corporations (notably Netscape and Microsoft), but not formalized, that also are in wide use.

The W3C maintains a list of URI schemes, which you can view at:

<http://www.w3.org/Addressing/schemes.html>

The IANA also maintains a list of URL schemes, at:

<http://www.iana.org/assignments/uri-schemes>

Table A-1 informally describes some of the schemes that have been proposed and those that are in active use. Note that many of the approximately 90 schemes in the table are not widely used, and many are extinct.

Table A-1. URI schemes from the W3C registry

Scheme	Description	RFCs
about	Netscape scheme to explore aspects of the browser. For example: about by itself is the same as choosing "About Communicator" from the Navigator Help menu, about:cache displays disk-cache statistics, and about:plugins displays information about configured plug-ins. Other browsers, such as Microsoft Internet Explorer, also use this scheme.	2244
acap	Application Configuration Access Protocol.	
afp	For file-sharing services using the Apple Filing Protocol (AFP) protocol, defined as part of the expired IETF <i>draft-ietf-svrlac-afp-service-01.txt</i> .	
afs	Reserved for future use by the Andrew File System.	
callto	Initiates a Microsoft NetMeeting conference session, such as: <i>callto:ws3.joes-hardware.com/joe@joes-hardware.com</i>	
chttp	The CHTTP caching protocol defined by Real Networks. RealPlayer does not cache all items streamed by HTTP. Instead, you designate files to cache by using <i>chttp://</i> instead of <i>http://</i> in the file's URL. When RealPlayer reads a CHTTP URL in a SMIL file, it first checks its disk cache for the file. If the file isn't present, it requests the file through HTTP, storing the file in its cache.	

Table A-1. URI schemes from the W3C registry (continued)

Scheme	Description	RFCs
cid	The use of [MIME] within email to convey web pages and their associated images requires a URL scheme to permit the HTML to refer to the images or other data included in the message. The Content-ID URL, "cid:", serves that purpose.	2392 2111
clsid	Allows Microsoft OLE/COM (Component Object Model) classes to be referenced. Used to insert active objects into web pages.	
data	Allows inclusion of small, constant data items as "immediate" data. This URL encodes the text/plain string "A brief note": <i>data:A%20brief%20note</i>	2397
date	Proposal for scheme to support dates, as in <i>date:1999-03-04T20:42:08</i> .	
dav	To ensure correct interoperation based on this specification, the IANA must reserve the URI namespaces starting with "DAV:" and with "opaquelocktoken:" for use by this specification, its revisions, and related WebDAV specifications.	2518
dns	Used by REBOL software. See http://www.rebol.com/users/valurl.html .	
eid	The external ID (eid) scheme provides a mechanism by which the local application can reference data that has been obtained by other, non-URL scheme means. The scheme is intended to provide a general escape mechanism to allow access to information for applications that are too specialized to justify their own schemes. There is some controversy about this URI. See http://www.ics.uci.edu/pub/ietf/uri/draft-finseth-url-00.txt .	
fax	The "fax" scheme describes a connection to a terminal that can handle telefaxes (facsimile machines).	2806
file	Designates files accessible on a particular host computer. A hostname can be included, but the scheme is unusual in that it does not specify an Internet protocol or access method for such files; as such, its utility in network protocols between hosts is limited.	1738
finger	The finger URL has the form: <i>finger://host[:port][/<request>]</i> The <request> must conform with the RFC 1288 request format. See http://www.ics.uci.edu/pub/ietf/uri/draft-ietf-uri-url-finger-03.txt .	
freenet	URLs for information in the Freenet distributed information system. See http://freenet.sourceforge.net .	
ftp	File Transfer Protocol scheme.	1738
gopher	The archaic gopher protocol.	1738
gsm-sms	URLs for the GSM mobile phone short message service.	
h323, h324	Multimedia conferencing URI schemes. See http://www.ics.uci.edu/pub/ietf/uri/draft-cordell-sg16-conv-url-00.txt .	
hdl	The Handle System is a comprehensive system for assigning, managing, and resolving persistent identifiers, known as "handles," for digital objects and other resources on the Internet. Handles can be used as URNs. See http://www.handle.net .	
hnews	HNEWS is an HTTP-tunneling variant of the NNTP news protocol. The syntax of hnews URLs is designed to be compatible with the current common usage of the news URL scheme. See http://www.ics.uci.edu/pub/ietf/uri/draft-stockwell-hnews-url-00.txt .	

Table A-1. URI schemes from the W3C registry (continued)

Scheme	Description	RFCs
http	The HTTP protocol. Read this book for more information.	2616
https	HTTP over SSL. See http://sitesearch.netscape.com/eng/ssl3/draft302.txt .	
iioploc	CORBA extensions. The Interoperable Name Service defines one URL-format object reference, iioploc, that can be typed into a program to reach defined services at remote locations, including the Naming Service. For example, this iioploc identifier: <i>iioploc://www.omg.org/NameService</i> would resolve to the CORBA Naming Service running on the machine whose IP address corresponded to the domain name <i>www.omg.org</i> . See http://www.omg.org .	2192
ilu	The Inter-Language Unification (ILU) system is a multilingual object interface system. The object interfaces provided by ILU hide implementation distinctions between different languages, different address spaces, and different operating system types. ILU can be used to build multilingual object-oriented libraries ("class libraries") with well-specified, language-independent interfaces. It also can be used to implement distributed systems. See ftp://parcftp.parc.xerox.com/pub/ilu/ilu.html .	
imap	The IMAP URL scheme is used to designate IMAP servers, mailboxes, messages, MIME bodies [MIME], and search programs on Internet hosts accessible using the IMAP protocol.	2192
IOR	CORBA interoperable object reference. See http://www.omg.org .	
irc	The irc URL scheme is used to refer to either Internet Relay Chat (IRC) servers or individual entities (channels or people) on IRC servers. See http://www.w3.org/Addressing/draft-mirashi-url-irc-01.txt .	2192
isbn	Proposed scheme for ISBN book references. See http://lists.w3.org/Archives/Public/www-talk/1991NovDec/0008.html .	
java	Identifies Java classes.	2192
javascript	The Netscape browser processes javascript URLs, evaluates the expression after the colon (:), if there is one, and loads a page containing the string value of the expression, unless it is undefined.	
jdbc	Used in the Java SQL API.	2255
ldap	Allows Internet clients direct access to the LDAP protocol.	
lid	The Local Identifier (lid:) scheme. See <i>draft-blackletter-lid-00</i> .	2255
lifn	A Location-Independent File Name (LIFN) for the Bulk File Distribution distributed storage system developed at UTK.	
livescript	Old name for JavaScript.	2255
lrq	See h323.	
mailto	The mailto URL scheme is used to designate the Internet mailing address of an individual or service.	2368
mailserver	Old proposal from 1994–1995 to let an entire message be encoded in a URL, so that (for example) the URL can automatically send email to a mail server for subscribing to a mailing list.	

Table A-1. URI schemes from the W3C registry (continued)

Scheme	Description	RFCs
md5	MD5 is a cryptographic checksum.	
mid	The mid scheme uses (a part of) the message-id of an email message to refer to a specific message.	2392 2111
mocha	See javascript.	
modem	The modem scheme describes a connection to a terminal that can handle incoming data calls.	2806
mms, mmst, mmsu	Scheme for Microsoft Media Server (MMS) to stream Active Streaming Format (ASF) files. To force UDP transport, use the mmsu scheme. To force TCP transport, use mmst.	
news	The news URL scheme is used to refer to either news groups or individual articles of USENET news. A news URL takes one of two forms: <i>news:<newsgroup-name></i> or <i>news:<message-id></i> .	1738 1036
nfs	Used to refer to files and directories on NFS servers.	2224
nnntp	An alternative method of referencing news articles, useful for specifying news articles from NNTP servers. An nnntp URL looks like: <i>nntp://<host>:<port>/<newsgroup-name>/<article-num></i> Note that while nnntp URLs specify a unique location for the article resource, most NNTP servers currently on the Internet are configured to allow access only from local clients, and thus nnntp URLs do not designate globally accessible resources. Hence, the news form of URL is preferred as a way of identifying news articles.	1738 977
opaquelocktoken	A WebDAV lock token, represented as a URI, that identifies a particular lock. A lock token is returned by every successful LOCK operation in the lockdiscovery property in the response body and also can be found through lock discovery on a resource. See RFC 2518.	
path	The path scheme defines a uniformly hierarchical namespace where a path URN is a sequence of components and an optional opaque string. See <i>http://www.hypernews.org/~liberte/www/path.html</i> .	
phone	Used in "URLs for Telephony"; replaced with tel: in RFC 2806.	
pop	The POP URL designates a POP email server, and optionally a port number, authentication mechanism, authentication ID, and/or authorization ID.	2384
pnm	Real Networks's streaming protocol.	
pop3	The POP3 URL scheme allows a URL to specify a POP3 server, allowing other protocols to use a general "URL to be used for mail access" in place of an explicit reference to POP3. Defined in expired <i>draft-earhart-url-pop3-00.txt</i> .	
printer	Abstract URLs for use with the Service Location standard. See <i>draft-ietf-srvloc-printer-scheme-02.txt</i> .	
prospero	Names resources to be accessed via the Prospero Directory Service.	1738
res	Microsoft scheme that specifies a resource to be obtained from a module. Consists of a string or numerical resource type, and a string or numerical ID.	
rtsp	Real-time streaming protocol that is the basis for Real Networks's modern streaming control protocols.	2326
rvp	URLs for the RVP rendezvous protocol, used to notify the arrival of users on a computer network. See <i>draft-calsyn-rvp-01</i> .	

Table A-1. URI schemes from the W3C registry (continued)

Scheme	Description	RFCs
rwhois	RWhois is an Internet directory access protocol, defined in RFC 1714 and RFC 2167. The RWhois URL gives clients direct access to rwhois. See http://www.rwhois.net/rwhois/docs/ .	
rx	An architecture to allow remote graphical applications to display data inside web pages. See http://www.w3.org/People/daniel/papers/mobgui/ .	
sdp	Session Description Protocol (SDP) URLs. See RFC 2327.	
service	The service scheme is used to provide access information for arbitrary network services. These URLs provide an extensible framework for client-based network software to obtain configuration information required to make use of network services.	2609
sip	The sip* family of schemes are used to establish multimedia conferences using the Session Initiation Protocol (SIP).	2543
shttp	S-HTTP is a superset of HTTP designed to secure HTTP connections and provide a wide variety of mechanisms to provide for confidentiality, authentication, and integrity. It has not been widely deployed, and it has mostly been supplanted with HTTPS SSL-encrypted HTTP. See http://www.homeport.org/~adam/shttp.html .	
snews	SSL-encrypted news.	
STANF	Old proposal for stable network filenames. Related to URNs. See http://web3.w3.org/Addressing/#STANF .	
t120	See h323.	
tel	URL to place a call using the telephone network.	2806
telephone	Used in previous drafts of tel.	
telnet	Designates interactive services that may be accessed by the Telnet protocol. A telnet URL takes the form: <i>telnet://<user>:<password>@<host>:<port>/</i>	1738
tip	Supports TIP atomic Internet transactions.	2371 2372
tn3270	Reserved, as per ftp://ftp.isi.edu/in-notes/iana/assignments/url-schemes .	
tv	The TV URL names a particular television broadcast channel.	2838
uuid	Universally unique identifiers (UUIDs) contain no information about location. They also are known as globally unique identifiers (GUIDs). They are persistent over time, like URNs, and consist of a 128-bit unique ID. UUID URLs are useful in situations where a unique identifier is required that cannot or should not be tied to a particular physical root namespace (such as a DNS name). See <i>draft-kindel-uuid-uri-00.txt</i> .	
urn	Persistent, location-independent, URNs.	2141
vemmi	Allows versatile multimedia interface (VEMMI) client software and VEMMI terminals to connect to VEMMI-compliant services. VEMMI is an international standard for online multimedia services.	2122
videotex	Allows videotex client software or terminals to connect to videotex services compliant with the ITU-T and ETSI videotex standards. See http://www.ics.uci.edu/pub/ietf/uri/draft-mavrakis-videotex-url-spec-01.txt .	

Table A-1. URI schemes from the W3C registry (continued)

Scheme	Description	RFCs
view-source	Netscape Navigator source viewers. These view-source URLs display HTML that was generated with JavaScript.	
wais	The wide area information service—an early form of search engine.	1738
whois++	URLs for the WHOIS++ simple Internet directory protocol. <i>See http://martinh.net/wip/whois-url.txt.</i>	1835
whodp	The Widely Hosted Object Data Protocol (WhoDP) exists to communicate the current location and state of large numbers of dynamic, relocatable objects. A WhoDP program “subscribes” to locate and receive information about an object and “publishes” to control the location and visible state of an object. <i>See draft-mohr-whodp-00.txt.</i>	
z39.50r, z39.50s	Z39.50 session and retrieval URLs. Z39.50 is an information retrieval protocol that does not fit neatly into a retrieval model designed primarily around the stateless fetch of data. Instead, it models a general user inquiry as a session-oriented, multi-step task, any step of which may be suspended temporarily while the server requests additional parameters from the client before continuing.	2056

HTTP Status Codes

This appendix is a quick reference of HTTP status codes and their meanings.

Status Code Classifications

HTTP status codes are segmented into five classes, shown in Table B-1.

Table B-1. Status code classifications

Overall range	Defined range	Category
100–199	100–101	Informational
200–299	200–206	Successful
300–399	300–305	Redirection
400–499	400–415	Client error
500–599	500–505	Server error

Status Codes

Table B-2 is a quick reference for all the status codes defined in the HTTP/1.1 specification, providing a brief summary of each. “Status Codes” in Chapter 3 goes into more detailed descriptions of these status codes and their uses.

Table B-2. Status codes

Status code	Reason phrase	Meaning
100	Continue	An initial part of the request was received, and the client should continue.
101	Switching Protocols	The server is changing protocols, as specified by the client, to one listed in the Upgrade header.
200	OK	The request is okay.
201	Created	The resource was created (for requests that create server objects).

Table B-2. Status codes (continued)

Status code	Reason phrase	Meaning
202	Accepted	The request was accepted, but the server has not yet performed any action with it.
203	Non-Authoritative Information	The transaction was okay, except the information contained in the entity headers was not from the origin server, but from a copy of the resource.
204	No Content	The response message contains headers and a status line, but no entity body.
205	Reset Content	Another code primarily for browsers; basically means that the browser should clear any HTML form elements on the current page.
206	Partial Content	A partial request was successful.
300	Multiple Choices	A client has requested a URL that actually refers to multiple resources. This code is returned along with a list of options; the user can then select which one he wants.
301	Moved Permanently	The requested URL has been moved. The response should contain a Location URL indicating where the resource now resides.
302	Found	Like the 301 status code, but the move is temporary. The client should use the URL given in the Location header to locate the resource temporarily.
303	See Other	Tells the client that the resource should be fetched using a different URL. This new URL is in the Location header of the response message.
304	Not Modified	Clients can make their requests conditional by the request headers they include. This code indicates that the resource has not changed.
305	Use Proxy	The resource must be accessed through a proxy, the location of the proxy is given in the Location header.
306	(Unused)	This status code currently is not used.
307	Temporary Redirect	Like the 301 status code; however, the client should use the URL given in the Location header to locate the resource temporarily.
400	Bad Request	Tells the client that it sent a malformed request.
401	Unauthorized	Returned along with appropriate headers that ask the client to authenticate itself before it can gain access to the resource.
402	Payment Required	Currently this status code is not used, but it has been set aside for future use.
403	Forbidden	The request was refused by the server.
404	Not Found	The server cannot find the requested URL.
405	Method Not Allowed	A request was made with a method that is not supported for the requested URL. The Allow header should be included in the response to tell the client what methods are allowed on the requested resource.
406	Not Acceptable	Clients can specify parameters about what types of entities they are willing to accept. This code is used when the server has no resource matching the URL that is acceptable for the client.
407	Proxy Authentication Required	Like the 401 status code, but used for proxy servers that require authentication for a resource.

Table B-2. Status codes (continued)

Status code	Reason phrase	Meaning
408	Request Timeout	If a client takes too long to complete its request, a server can send back this status code and close down the connection.
409	Conflict	The request is causing some conflict on a resource.
410	Gone	Like the 404 status code, except that the server once held the resource.
411	Length Required	Servers use this code when they require a Content-Length header in the request message. The server will not accept requests for the resource without the Content-Length header.
412	Precondition Failed	If a client makes a conditional request and one of the conditions fails, this response code is returned.
413	Request Entity Too Large	The client sent an entity body that is larger than the server can or wants to process.
414	Request URI Too Long	The client sent a request with a request URL that is larger than what the server can or wants to process.
415	Unsupported Media Type	The client sent an entity of a content type that the server does not understand or support.
416	Requested Range Not Satisfiable	The request message requested a range of a given resource, and that range either was invalid or could not be met.
417	Expectation Failed	The request contained an expectation in the Expect request header that could not be satisfied by the server.
500	Internal Server Error	The server encountered an error that prevented it from servicing the request.
501	Not Implemented	The client made a request that is beyond the server's capabilities.
502	Bad Gateway	A server acting as a proxy or gateway encountered a bogus response from the next link in the request response chain.
503	Service Unavailable	The server cannot currently service the request but will be able to in the future.
504	Gateway Timeout	Similar to the 408 status code, except that the response is coming from a gateway or proxy that has timed out waiting for a response to its request from another server.
505	HTTP Version Not Supported	The server received a request in a version of the protocol that it can't or won't support.

HTTP Header Reference

It's almost amusing to remember that the first version of HTTP, 0.9, had no headers. While this certainly had its down sides, its fun to marvel in its simplistic elegance.

Well, back to reality. Today there are a horde of HTTP headers, many part of the specification and still others that are extensions to it. This appendix provides some background on these official and extension headers. It also acts as an index for the various headers in this book, pointing out where their concepts and features are discussed in the running text. Most of these headers are simple up-front; it's the interactions with each other and other features of HTTP where things get hairy. This appendix provides a bit of background for the headers listed and directs you to the sections of the book where they are discussed at length.

The headers listed in this appendix are drawn from the HTTP specifications, related documents, and our own experience poking around with HTTP messages and the various servers and clients on the Internet.

This list is far from exhaustive. There are many other extension headers floating around on the Web, not to mention those potentially used in private intranets. Nonetheless, we have attempted to make this list as complete as possible. See RFC 2616 for the current version of the HTTP/1.1 specification and a list of official headers and their specification descriptions.

Accept

The Accept header is used by clients to let servers know what media types are *acceptable*. The value of the Accept header field is a list of media types that the client can use. For instance, your web browser cannot display every type of multimedia object on the Web. By including an Accept header in your requests, your browser can save you from downloading a video or other type of object that you can't use.

The Accept header field also may include a list of quality values (q values) that tell the server which media type is preferred, in case the server has multiple versions of the media type. See Chapter 17 for a complete discussion of content negotiation and q values.

Type	Request header
Notes	“*” is a special value that is used to wildcard media types. For example, “*/*” represents all types, and “image/*” represents all image types.
Examples	Accept: text/*, image/* Accept: text/*, image/gif, image/jpeg;q=1

Accept-Charset

The Accept-Charset header is used by clients to tell servers what character sets are acceptable or preferred. The value of this request header is a list of character sets and possibly quality values for the listed character sets. The quality values let the server know which character set is preferred, in case the server has the document in multiple acceptable character sets. See Chapter 17 for a complete discussion of content negotiation and q values.

Type	Request header
Notes	As with the Accept header, “*” is a special character. If present, it represents all character sets, except those that also are mentioned explicitly in the value. If it’s not present, any charset not in the value field has a default q value of zero, with the exception of the iso-latin-1 charset, which gets a default of 1.
Basic Syntax	Accept-Charset: 1# ((charset “*”) [“;” “q” “=” qvalue])
Example	Accept-Charset: iso-latin-1

Accept-Encoding

The Accept-Encoding header is used by clients to tell servers what encodings are acceptable. If the content the server is holding is encoded (perhaps compressed), this request header lets the server know whether the client will accept it. Chapter 17 contains a complete description of the Accept-Encoding header.

Type	Request header
Basic Syntax	Accept-Encoding: 1# ((content-coding “*”) [“;” “q” “=” qvalue])
Examples*	Accept-Encoding: Accept-Encoding: gzip Accept-Encoding: compress;q=0.5, gzip;q=1

* The empty Accept-Encoding example is not a typo. It refers to the identity encoding—that is, the unencoded content. If the Accept-Encoding header is present and empty, only the unencoded content is acceptable.

Accept-Language

The Accept-Language request header functions like the other Accept headers, allowing clients to inform the server about what languages (e.g., the natural language for content) are acceptable or preferred. Chapter 17 contains a complete description of the Accept-Language header.

Type	Request header
Basic Syntax	Accept-Language: 1# (language-range [";" "q" "=" qvalue]) language-range = ((1*8ALPHA * ("-" 1*8ALPHA)) "**")
Examples	Accept-Language: en Accept-Language: en;q=0.7, en-gb;q=0.5

Accept-Ranges

The Accept-Ranges header differs from the other Accept headers—it is a response header used by servers to tell clients whether they accept requests for ranges of a resource. The value of this header tells what type of ranges, if any, the server accepts for a given resource.

A client can attempt to make a range request on a resource without having received this header. If the server does not support range requests for that resource, it can respond with an appropriate status code* and the Accept-Ranges value “none”. Servers might want to send the “none” value for normal requests to discourage clients from making range requests in the future.

Chapter 17 contains a complete description of the Accept-Ranges header.

Type	Response header
Basic Syntax	Accept-Ranges: 1# range-unit none
Examples	Accept-Ranges: none Accept-Ranges: bytes

Age

The Age header tells the receiver how old a response is. It is the sender’s best guess as to how long ago the response was generated by or revalidated with the origin server. The value of the header is the sender’s guess, a delta in seconds. See Chapter 7 for more on the Age header.

Type	Response header
Notes	HTTP/1.1 caches must include an Age header in every response they send.

* For example, status code 416 (see “400–499: Client Error Status Codes” in Chapter 3).

Basic Syntax Age: delta-seconds

Example Age: 60

Allow

The Allow header is used to inform clients what HTTP methods are supported on a particular resource.

Type Response header

Notes An HTTP/1.1 server sending a 405 Method Not Allowed response must include an Allow header.*

Basic Syntax Allow: #Method

Example Allow: GET, HEAD

Authorization

The Authorization header is sent by a client to authenticate itself with a server. A client will include this header in its request after receiving a 401 Authentication Required response from a server. The value of this header depends on the authentication scheme in use. See Chapter 14 for a detailed discussion of the Authorization header.

Type Response header

Basic Syntax Authorization: authentication-scheme #authentication-param

Example Authorization: Basic YnJpYW4tdG90dHk6T3ch

Cache-Control

The Cache-Control header is used to pass information about how an object can be cached. This header is one of the more complex headers introduced in HTTP/1.1. Its value is a caching directive, giving caches special instructions about an object's cacheability.

In Chapter 7, we discuss caching in general as well as the specific details about this header.

Type General header

Example Cache-Control: no-cache

* See "Status Codes" in Chapter 3 for more on the 405 status code.

Client-ip

The Client-ip header is an extension header used by some older clients and some proxies to transmit the IP address of the machine on which the client is running.

Type	Extension request header
Notes	Implementors should be aware that the information provided in the value of this header is not secure.
Basic Syntax	Client-ip: ip-address
Example	Client-ip: 209.1.33.49

Connection

The Connection header is a somewhat overloaded header that can lead to a bit of confusion. This header was used in HTTP/1.0 clients that were extended with keep-alive connections for control information.* In HTTP/1.1, the older semantics are mostly recognized, but the header has taken on a new function.

In HTTP/1.1, the Connection header's value is a list of tokens that correspond to header names. Applications receiving an HTTP/1.1 message with a Connection header are supposed to parse the list and remove any of the headers in the message that are in the Connection header list. This is mainly for proxies, allowing a server or other proxy to specify hop-by-hop headers that should not be passed along.

One special token value is "close". This token means that the connection is going to be closed after the response is completed. HTTP/1.1 applications that do not support persistent connections need to insert the Connection header with the "close" token in all requests and responses.

Type	General header
Notes	While RFC 2616 does not specifically mention keep-alive as a connection token, some browsers (including those sending HTTP/1.1 as their versions) use it in making requests.
Basic Syntax	Connection: 1# (connection-token)
Examples	Connection: close

* See Chapter 4 for more on keep-alive and persistent connections.

Content-Base

The Content-Base header provides a way for a server to specify a base URL for resolving URLs found in the entity body of a response.* The value of the Content-Base header is an absolute URL that can be used to resolve relative URLs found inside the entity.

Type	Entity header
Notes	This header is not defined in RFC 2616; it was previously defined in RFC 2068, an earlier draft of the HTTP/1.1 specification, and has since been removed from the official specification.
Basic Syntax	Content-Base: absoluteURL
Example	Content-Base: http://www.joes-hardware.com/

Content-Encoding

The Content-Encoding header is used to specify whether any encodings have been performed on the object. By encoding the content, a server can compress it before sending the response. The value of the Content-Encoding header tells the client what type or types of encoding have been performed on the object. With that information, the client can then decode the message.

Sometimes more than one encoding is applied to an entity, in which case the encodings must be listed in the order in which they were performed.

Type	Entity header
Basic Syntax	Content-Encoding: 1# content-coding
Examples	Content-Encoding: gzip Content-Encoding: compress, gzip

Content-Language

The Content-Language header tells the client the natural language that should be understood in order to understand the object. For instance, a document written in French would have a Content-Language value indicating French. If this header is not present in the response, the object is intended for all audiences. Multiple languages in the header's value indicate that the object is suitable for audiences of each language listed.

One caveat about this header is that the header's value may just represent the natural language of the intended audience of this object, not all or any of the languages contained

* See Chapter 2 for more on base URLs.

in the object. Also, this header is not limited to text or written data objects; images, video, and other media types can be tagged with their intended audiences' natural languages.

Type	Entity header
Basic Syntax	Content-Language: 1# language-tag
Examples	Content-Language: en Content-Language: en, fr

Content-Length

The Content-Length header gives the length or size of the entity body. If the header is in a response message to a HEAD HTTP request, the value of the header indicates the size that the entity body would have been had it been sent.

Type	Entity header
Basic Syntax	Content-Length: 1*DIGIT
Example	Content-Length: 2417

Content-Location

The Content-Location header is included in an HTTP message to give the URL corresponding to the entity in the message. For objects that may have multiple URLs, a response message can include a Content-Location header indicating the URL of the object used to generate the response. The Content-Location can be different from the requested URL. This generally is used by servers that are directing or redirecting a client to a new URL.

If the URL is relative, it should be interpreted relative to the Content-Base header. If the Content-Base header is not present, the URL used in the request should be used.

Type	Entity header
Basic Syntax	Content-Location: (absoluteURL relativeURL)
Example	Content-Location: http://www.joes-hardware.com/index.html

Content-MD5

The Content-MD5 header is used by servers to provide a message-integrity check for the message body. Only an origin server or requesting client should insert a Content-MD5

header in the message. The value of the header is an MD5 digest* of the (potentially encoded) message body.

The value of this header allows for an end-to-end check on the data, useful for detecting unintentional modifications to the data in transit. It is not intended to be used for security purposes.

RFC 1864 defines this header in more detail.

Type	Entity header
Notes	The MD5 digest value is a base-64 (see Appendix E) or 128-bit MD5 digest, as defined in RFC 1864.
Basic Syntax	Content-MD5: md5-digest
Example	Content-MD5: Q2h1Y2sgSW51ZwDIAXR5IQ==

Content-Range

The Content-Range header is sent as the result of a request that transmitted a range of a document. It provides the location (range) within the original entity that this entity represents. It also gives the length of the entire entity.

If an “*” is present in the value instead of the length of the entire entity, this means that the length was not known when the response was sent.

See Chapter 15 for more on the Content-Range header.

Type	Entity header
Notes	Servers responding with the 206 Partial Content response code must not include a Content-Range header with an “*” as the length.
Example	Content-Range: bytes 500-999 / 5400

Content-Type

The Content-Type header tells the media type of the object in the message.

Type	Entity header
Basic Syntax	Content-Type: media-type
Example	Content-Type: text/html; charset=iso-latin-1

* The MD5 digest is defined in RFC 1864.

Cookie

The Cookie header is an extension header used for client identification and tracking. Chapter 11 talks about the Cookie header and its use in detail (also see “Set-Cookie”).

Type Extension request header

Example Cookie: ink=IUOK164y59BC708378908CFF89OE5573998A115

Cookie2

The Cookie2 header is an extension header used for client identification and tracking. Cookie2 is used to identify what version of cookies a requestor understands. It is defined in greater detail in RFC 2965.

Chapter 11 talks about the Cookie2 header and its use in detail.

Type Extension request header

Example Cookie2: \$version="1"

Date

The Date header gives the date and time at which the message was created. This header is required in servers’ responses because the time and date at which the server believes the message was created can be used by caches in evaluating the freshness of a response. For clients, this header is completely optional, although it’s good form to include it.

Type General header

Basic Syntax Date: HTTP-date

Examples Date: Tue, 3 Oct 1997 02:15:31 GMT

HTTP has a few specific date formats. This one is defined in RFC 822 and is the preferred format for HTTP/1.1 messages. However, in earlier specifications of HTTP, the date format was not spelled out as well, so server and client implementors have used other formats, which need to be supported for the sake of legacy. You will run into date formats like the one specified in RFC 850, as well as dates in the format produced by the *asctime()* system call. Here they are for the date represented above:

 Date: Tuesday, 03-Oct-97 02:15:31 GMT *RFC 850 format*
 Date: Tue Oct 3 02:15:31 1997 *asctime() format*

The *asctime()* format is looked down on because it is in local time and it does not specify its time zone (e.g., GMT). In general, the date header should be in GMT; however, robust applications should handle dates that either do not specify the time zone or include Date values in non-GMT time.

ETag

The ETag header provides the *entity tag* for the entity contained in the message. An entity tag is basically a way of identifying a resource.

Entity tags and their relationship to resources are discussed in detail in Chapter 15.

Type	Entity header
Basic Syntax	ETag: entity-tag
Examples	ETag: "11e92a-457b-31345aa" ETag: W/"11e92a-457b-3134b5aa"

Expect

The Expect header is used by clients to let servers know that they expect certain behavior. This header currently is closely tied to the response code 100 Continue (see “100–199: Informational Status Codes” in Chapter 3).

If a server does not understand the Expect header’s value, it should respond with a status code of 417 Expectation Failed.

Type	Request header
Basic Syntax	Expect: 1# ("100-continue" expectation-extension)
Example	Expect: 100-continue

Expires

The Expires header gives a date and time at which the response is no longer valid. This allows clients such as your browser to cache a copy and not have to ask the server if it is still valid until after this time has expired.

Chapter 7 discusses how the Expires header is used—in particular, how it relates to caches and having to revalidate responses with the origin server.

Type	Entity header
Basic Syntax	Expires: HTTP-date
Example	Expires: Thu, 03 Oct 1997 17:15:00 GMT

From

The From header says who the request is coming from. The format is just a valid Internet email address (specified in RFC 1123) for the user of the client.

There are potential privacy issues with using/populating this header. Client implementors should be careful to inform their users and give them a choice before including this header in a request message. Given the potential for abuse by people collecting email addresses for unsolicited mail messages, woe to the implementor who broadcasts this header unannounced and has to answer to angry users.

Type Request header

Basic Syntax From: mailbox

Example From: slurp@inktomi.com

Host

The Host header is used by clients to provide the server with the Internet hostname and port number of the machine from which the client wants to make a request. The hostname and port are those from the URL the client was requesting.

The Host header allows servers to differentiate different relative URLs based on the hostname, giving the server the ability to host several different hostnames on the same machine (i.e., the same IP address).

Type Request header

Notes HTTP/1.1 clients must include a Host header in all requests. All HTTP/1.1 servers must respond with the 400 Bad Request status code to HTTP/1.1 clients that do not provide a Host header.

Basic Syntax Host: host [":" port]

Example Host: www.hotbot.com:80

If-Modified-Since

The If-Modified-Since request header is used to make conditional requests. A client can use the GET method to request a resource from a server, having the response hinge on whether the resource has been modified since the client last requested it.

If the object has not been modified, the server will respond with a 304 Not Modified response, instead of with the resource. If the object has been modified, the server will respond as if it was a non-conditional GET request. Chapter 7 discusses conditional requests in detail.

Type Request header

Basic Syntax If-Modified-Since: HTTP-date

Example If-Modified-Since: Thu, 03 Oct 1997 17:15:00 GMT

If-Match

Like the If-Modified-Since header, the If-Match header can be used to make a request conditional. Instead of a date, the If-Match request uses an entity tag. The server compares the entity tag in the If-Match header with the current entity tag of the resource and returns the object if the tags match.

The server should use the If-Match value of "*" to match any entity tag it has for a resource; "*" will always match, unless the server no longer has the resource.

This header is useful for updating resources that a client or cache already has. The resource is returned only if it has changed—that is, if the previously requested object's entity tag does not match the entity tag of the current version on the server. Chapter 7 discusses conditional requests in detail.

Type	Request header
Basic Syntax	If-Match: ("*" 1# entity-tag)
Example	If-Match: "11e92a-457b-31345aa"

If-None-Match

The If-None-Match header, like all the If headers, can be used to make a request conditional. The client supplies the server with a list of entity tags, and the server compares those tags against the entity tags it has for the resource, returning the resource only if none match.

This allows a cache to update resources only if they have changed. Using the If-None-Match header, a cache can use a single request to both invalidate the entities it has and receive the new entity in the response. Chapter 7 discusses conditional requests in detail.

Type	Request header
Basic Syntax	If-None-Match: ("*" 1# entity-tag)
Example	If-None-Match: "11e92a-457b-31345aa"

If-Range

The If-Range header, like all the If headers, can be used to make a request conditional. It is used when an application has a copy of a range of a resource, to revalidate the range or get the complete resource if the range is no longer valid. Chapter 7 discusses conditional requests in detail.

Type	Request header
Basic Syntax	If-Range: (HTTP-date entity-tag)

Examples If-Range: Tue, 3 Oct 1997 02:15:31 GMT
 If-Range: "11e92a-457b-3134b5aa"

If-Unmodified-Since

The If-Unmodified-Since header is the twin of the If-Modified-Since header. Including it in a request makes the request conditional. The server should look at the date value of the header and return the object only if it has not been modified since the date provided. Chapter 7 discusses conditional requests in detail.

Type Request header

Basic Syntax If-Unmodified-Since: HTTP-date

Example If-Unmodified-Since: Thu, 03 Oct 1997 17:15:00 GMT

Last-Modified

The Last-Modified header tries to provide information about the last time this entity was changed. This could mean a lot of things. For example, resources typically are files on a server, so the Last-Modified value could be the last-modified time provided by the server's filesystem. On the other hand, for dynamically created resources such as those created by scripts, the Last-Modified value could be the time the response was created.

Servers need to be careful that the Last-Modified time is not in the future. HTTP/1.1 servers should reset the Last-Modified time if it is later than the value that would be sent in the Date header.

Type Entity header

Basic Syntax Last-Modified: HTTP-date

Example Last-Modified: Thu, 03 Oct 1997 17:15:00 GMT

Location

The Location header is used by servers to direct clients to the location of a resource that either was moved since the client last requested it or was created in response to the request.

Type Response header

Basic Syntax Location: absoluteURL

Example Location: http://www.hotbot.com

Max-Forwards

This header is used only with the TRACE method, to limit the number of proxies or other intermediaries that a request goes through. Its value is an integer. Each application that receives a TRACE request with this header should decrement the value before it forwards the request along.

If the value is zero when the application receives the request, it should send back a 200 OK response to the request, with an entity body containing the original request. If the Max-Forwards header is missing from a TRACE request, assume that there is no maximum number of forwards.

For other HTTP methods, this header should be ignored. See “Methods” in Chapter 3 for more on the TRACE method.

Type	Request header
Basic Syntax	Max-Forwards: 1*DIGIT
Example	Max-Forwards: 5

MIME-Version

MIME is HTTP’s cousin. While they are radically different, some HTTP servers do construct messages that are valid under the MIME specification. When this is the case, the MIME-Version header can be supplied by the server.

This header has never been part of the official specification, although it is mentioned in the HTTP/1.0 specification. Many older servers send messages with this header, however, those messages often are not valid MIME messages, making this header both confusing and impossible to trust.

Type	Extension general header
Basic Syntax	MIME-Version: DIGIT "." DIGIT
Example	MIME-Version: 1.0

Pragma

The Pragma header is used to pass directions along with the message. These directions could be almost anything, but often they are used to control caching behavior. Proxies and gateways must not remove the Pragma header, because it could be intended for all applications that receive the message.

The most common form of Pragma, Pragma: no-cache, is a request header that forces caches to request or revalidate the document from the origin server even when a fresh copy is available in the cache. It is sent by browsers when users click on the Reload/Refresh button. Many servers send Pragma: no-cache as a response header (as an equivalent to

Cache-Control: no-cache), but despite its common use, this behavior is technically undefined. Not all applications support Pragma response headers.

Chapter 7 discusses the Pragma header and how it is used by HTTP/1.0 applications to control caches.

Type	Request header
Basic Syntax	Pragma: 1# pragma-directive*
Example	Pragma: no-cache

Proxy-Authenticate

The Proxy-Authenticate header functions like the WWW-Authenticate header. It is used by proxies to challenge an application sending a request to authenticate itself. The full details of this challenge/response, and other security mechanisms of HTTP, are discussed in detail in Chapter 14.

If an HTTP/1.1 proxy server is sending a 407 Proxy Authentication Required response, it must include the Proxy-Authenticate header.

Proxies and gateways must be careful in interpreting all the Proxy headers. They generally are hop-by-hop headers, applying only to the current connection. For instance, the Proxy-Authenticate header requests authentication for the current connection.

Type	Response header
Basic Syntax	Proxy-Authenticate: challenge
Example	Proxy-Authenticate: Basic realm="Super Secret Corporate Financial Documents"

Proxy-Authorization

The Proxy-Authorization header functions like the Authorization header. It is used by client applications to respond to Proxy-Authenticate challenges. See Chapter 14 for more on how the challenge/response security mechanism works.

Type	Request header
Basic Syntax	Proxy-Authorization: credentials
Example	Proxy-Authorization: Basic YnJpYW4tdG90dHk6T3ch

* The only specification-defined Pragma directive is “no-cache”; however, you may run into other Pragma headers that have been defined as extensions to the specification.

Proxy-Connection

The Proxy-Connection header was meant to have similar semantics to the HTTP/1.0 Connection header. It was to be used between clients and proxies to specify options about the connections (chiefly keep-alive connections).^{*} It is not a standard header and is viewed as an ad hoc header by the standards committee. However, it is widely used by browsers and proxies.

Browser implementors created the Proxy-Connection header to solve the problem of a client sending an HTTP/1.0 Connection header that gets blindly forwarded by a dumb proxy. A server receiving the blindly forwarded Connection header could confuse the capabilities of the client connection with those of the proxy connection.

The Proxy-Connection header is sent instead of the Connection header when the client knows that it is going through a proxy. Because servers don't recognize the Proxy-Connection header, they ignore it, allowing dumb proxies that blindly forward the header to do so without causing harm.

The problem with this solution occurs if there is more than one proxy in the path of the client to the server. If the first one blindly forwards the header to the second, which understands it, the second proxy can suffer from the same confusion the server did with the Connection header.

This is the problem that the HTTP working group had with this solution—they saw it as a hack that solved the case of a single proxy, but not the bigger problem. Nonetheless, it does handle some of the more common cases, and because older versions of both Netscape Navigator and Microsoft Internet Explorer implement it, proxy implementors need to deal with it. See Chapter 4 for more information.

Type	General header
Basic Syntax	Proxy-Connection: 1# (connection-token)
Example	Proxy-Connection: close

Public

The Public header allows a server to tell a client what methods it supports. These methods can be used in future requests by the client. Proxies need to be careful when they receive a response from a server with the Public header. The header indicates the capabilities of the server, not the proxy, so the proxy needs to either edit the list of methods in the header or remove the header before it sends the response to the client.

Type	Response header
-------------	-----------------

^{*} See Chapter 4 for more on keep-alive and persistent connections.

Notes This header is not defined in RFC 2616. It was previously defined in RFC 2068, an earlier draft of the HTTP/1.1 specification, but it has since been removed from the official specification.

Basic Syntax Public: 1# HTTP-method

Example Public: OPTIONS, GET, HEAD, TRACE, POST

Range

The Range header is used in requests for parts or ranges of an entity. Its value indicates the range of the entity that is included in the message.

Requests for ranges of a document allow for more efficient requests of large objects (by requesting them in segments) or for recovery from failed transfers (allowing a client to request the range of the resource that did not make it). Range requests and the headers that make the requests possible are discussed in detail in Chapter 15.

Type Entity header

Example Range: bytes=500-1500

Referer

The Referer header is inserted into client requests to let the server know where the client got the URL from. This is a voluntary effort, for the server's benefit; it allows the server to better log the requests or perform other tasks. The misspelling of "Referer" hearkens back to the early days of HTTP, to the frustration of English-speaking copyeditors throughout the world.

What your browser does is fairly simple. If you get home page A and click on a link to go to home page B, your browser will insert a Referer header in the request with value A. Referer headers are inserted by your browser only when you click on links; requests for URLs you type in yourself will not contain a Referer header.

Because some pages are private, there are some privacy concerns with this header. While some of this is unwarranted paranoia, this header does allow web servers and their administrators to see where you came from, potentially allowing them to better track your surfing. As a result, the HTTP/1.1 specification recommends that application writers allow the user to decide whether this header is transmitted.

Type Request header

Basic Syntax Referer: (absoluteURL | relativeURL)

Example Referer: http://www.inktomi.com/index.html

Retry-After

Servers can use the Retry-After header to tell a client when to retry its request for a resource. It is used with the 503 Service Unavailable status code to give the client a specific date and time (or number of seconds) at which it should retry its request.

A server can also use this header when it is redirecting clients to resources, giving the client a time to wait before making a request on the resource to which it is redirected.* This can be very useful to servers that are creating dynamic resources, allowing the server to redirect the client to the newly created resource but giving time for the resource to be created.

Type	Response header
Basic Syntax	Retry-After: (HTTP-date delta-seconds)
Examples	Retry-After: Tue, 3 Oct 1997 02:15:31 GMT Retry-After: 120

Server

The Server header is akin to the User-Agent header; it provides a way for servers to identify themselves to clients. Its value is the server name and an optional comment about the server.

Because the Server header identifies the server product and can contain additional comments about the product, its format is somewhat free-form. If you are writing software that depends on how a server identifies itself, you should experiment with the server software to see what it sends back, because these tokens vary from product to product and release to release.

As with the User-Agent header, don't be surprised if an older proxy or gateway inserts what amounts to a Via header in the Server header itself.

Type	Response header
Basic Syntax	Server: 1* (product comment)
Examples	Server: Microsoft-Internet-Information-Server/1.0 Server: websitepro/1.1f (s/n wpo-07d0) Server: apache/1.2b6 via proxy gateway CERN-HTTPD/3.0 libwww/2.13

Set-Cookie

The Set-Cookie header is the partner to the Cookie header; in Chapter 11, we discuss the use of this header in detail.

* See "Redirection status codes and reason phrases" in Chapter 3 for more on server redirect responses.

Type	Extension response header
Basic Syntax	Set-Cookie: command
Examples	Set-Cookie: lastorder=00183; path=/orders Set-Cookie: private_id=519; secure

Set-Cookie2

The Set-Cookie2 header is an extension of the Set-Cookie header; in Chapter 11, we discuss the use of this header in detail.

Type	Extension response header
Basic Syntax	Set-Cookie2: command
Examples	Set-Cookie2: ID="29046"; Domain=".joes-hardware.com" Set-Cookie2: color=blue

TE

The poorly named TE header functions like the Accept-Encoding header, but for transfer encodings (it could have been named Accept-Transfer-Encoding, but it wasn't). The TE header also can be used to indicate whether a client can handle headers in the trailer of a response that has been through the chunked encoding. See Chapter 15 for more on the TE header, chunked encoding, and trailers.

Type	Request header
Notes	If the value is empty, only the chunked transfer encoding is acceptable. The special token “trailers” indicates that trailer headers are acceptable in a chunked response.
Basic Syntax	TE: # (transfer-codings) transfer-codings= "trailers" (transfer-extension [accept-params])
Examples	TE: TE: chunked

Trailer

The Trailer header is used to indicate which headers are present in the trailer of a message. Chapter 15 discusses chunked encodings and trailers in detail.

Type	General header
-------------	----------------

Basic Syntax	Trailer: 1#field-name
Example	Trailer: Content-Length

Title

The Title header is a non-specification header that is supposed to give the title of the entity. This header was part of an early HTTP/1.0 extension and was used primarily for HTML pages, which have clear title markers that servers can use. Because many, if not most, media types on the Web do not have such an easy way to extract a title, this header has limited usefulness. As a result, it never made it into the official specification, though some older servers on the Net still send it faithfully.

Type	Response header
Notes	The Title header is not defined in RFC 2616. It was originally defined in the HTTP/1.0 draft definition (http://www.w3.org/Protocols/HTTP/HTTP2.html) but has since been removed from the official specification.
Basic Syntax	Title: document-title
Example	Title: CNN Interactive

Transfer-Encoding

If some encoding had to be performed to transfer the HTTP message body safely, the message will contain the Transfer-Encoding header. Its value is a list of the encodings that were performed on the message body. If multiple encodings were performed, they are listed in order.

The Transfer-Encoding header differs from the Content-Encoding header because the transfer encoding is an encoding that was performed by a server or other intermediary application to transfer the message.

Transfer encodings are discussed in Chapter 15.

Type	General header
Basic Syntax	Transfer-Encoding: 1# transfer-coding
Example	Transfer-Encoding: chunked

UA-(CPU, Disp, OS, Color, Pixels)

These User-Agent headers are nonstandard and no longer common. They provide information about the client machine that could allow for better content selection by a server. For

instance, if a server knew that a user's machine had only an 8-bit color display, the server could select images that were optimized for that type of display.

With any header that gives information about the client that otherwise would be unavailable, there are some security concerns (see Chapter 14 for more information).

Type	Extension request headers	
Notes	These headers are not defined in RFC 2616, and their use is frowned upon.	
Basic Syntax	"UA" "-" ("CPU" "Disp" "OS" "Color" "Pixels") ":" machine-value machine-value = (cpu screensize os-name display-color-depth)	
Examples	UA-CPU: x86	<i>CPU of client's machine</i>
	UA-Disp: 640, 480, 8	<i>Size and color depth of client's display</i>
	UA-OS: Windows 95	<i>Operating system of client machine</i>
	UA-Color: color8	<i>Color depth of client's display</i>
	UA-Pixels: 640x480	<i>Size of client's display</i>

Upgrade

The Upgrade header provides the sender of a message with a means of broadcasting the desire to use another, perhaps completely different, protocol. For instance, an HTTP/1.1 client could send an HTTP/1.0 request to a server and include an Upgrade header with the value "HTTP/1.1", allowing the client to test the waters and see whether the server speaks HTTP/1.1.

If the server is capable, it can send an appropriate response letting the client know that it is okay to use the new protocol. This provides an efficient way to move to other protocols. Most servers currently are only HTTP/1.0-compliant, and this strategy allows a client to avoid confusing a server with too many HTTP/1.1 headers until it determines whether the server is indeed capable of speaking HTTP/1.1.

When a server sends a 101 Switching Protocols response, it must include this header.

Type	General header
Basic Syntax	Upgrade: 1# protocol
Example	Upgrade: HTTP/2.0

User-Agent

The User-Agent header is used by client applications to identify themselves, much like the Server header for servers. Its value is the product name and possibly a comment describing the client application.

This header's format is somewhat free-form. Its value varies from client product to product and release to release. This header sometimes even contains information about the machine on which the client is running.

As with the Server header, don't be surprised if older proxy or gateway applications insert what amounts to a Via header in the User-Agent header itself.

Type; Request header

Basic Syntax User-Agent: 1* (product | comment)

Example User-Agent: Mozilla/4.0 (compatible; MSIE 5.5; Windows NT 5.0)

Vary

The Vary header is used by servers to inform clients what headers from a client's request will be used in server-side negotiation.* Its value is a list of headers that the server looks at to determine what to send the client as a response.

An example of this would be a server that sends special HTML pages based on your web browser's features. A server sending these special pages for a URL would include a Vary header that indicated that it looked at the User-Agent header of the request to determine what to send as a response.

The Vary header also is used by caching proxies; see Chapter 7 for more on how the Vary header relates to cached HTTP responses.

Type Response header

Basic Syntax Vary: ("*" | 1# field-name)

Example Vary: User-Agent

Via

The Via header is used to trace messages as they pass through proxies and gateways. It is an informational header that can be used to see what applications are handling requests and responses.

When a message passes through an HTTP application on its way to a client or a server, that application can use the Via header to tag the message as having gone *via* it. This is an HTTP/1.1 header; many older applications insert a Via-like string in the User-Agent or Server headers of requests and responses.

If the message passes through multiple in-between applications, each one should tack on its Via string. The Via header must be inserted by HTTP/1.1 proxies and gateways.

* See Chapter 17 for more on content negotiation.

Type	General header
Basic Syntax*	Via: 1# (received-protocol received-by [comment])
Example	Via: 1.1 joes-hardware.com (Joes-Server/1.0) The above says that the message passed through the Joes Server Version 1.0 software running on the machine <i>joes-hardware.com</i>. Joe's Server was speaking HTTP 1.1. The Via header should be formatted like this: HTTP-Version machine-hostname (Application-Name-Version)

Warning

The Warning header is used to give a little more information about what happened during a request. It provides the server with a way to send additional information that is not in the status code or reason phrase. Several warning codes are defined in the HTTP/1.1 specification:

101 Response Is Stale

When a response message is known to be stale—for instance, if the origin server is unavailable for revalidation—this warning must be included.

111 Revalidation Failed

If a cache attempts to revalidate a response with an origin server and the revalidation fails because the cache cannot reach the origin server, this warning must be included in the response to the client.

112 Disconnected Operation

An informative warning; should be used if a cache's connectivity to the network is removed.

113 Heuristic Expiration

Caches must include this warning if their freshness heuristic is greater than 24 hours and they are returning a response with an age greater than 24 hours.

199 Miscellaneous Warning

Systems receiving this warning must not take any automated response; the message may and probably should contain a body with additional information for the user.

214 Transformation Applied

Must be added by any intermediate application, such as a proxy, if the application performs any transformation that changes the content encoding of the response.

299 Miscellaneous Persistent Warning

Systems receiving this warning must not take any automated reaction; the error may contain a body with more information for the user.

* See the HTTP/1.1 specification for the complete Via header syntax.

Type	Response header
Basic Syntax	Warning: 1# warning-value
Example	Warning: 113

WWW-Authenticate

The WWW-Authenticate header is used in 401 Unauthorized responses to issue a challenge authentication scheme to the client. Chapter 14 discusses the WWW-Authenticate header and its use in HTTP's basic challenge/response authentication system.

Type	Response header
Basic Syntax	WWW-Authenticate: 1# challenge
Example	WWW-Authenticate: Basic realm="Your Private Travel Profile"

X-Cache

The X headers are all extension headers. The X-Cache header is used by Squid to inform a client whether a resource is available.

Type	Extension response header
Example	X-Cache: HIT

X-Forwarded-For

This header is used by many proxy servers (e.g., Squid) to note whom a request has been forwarded for. Like the Client-ip header mentioned earlier, this request header notes the address from which the request originates.

Type	Extension request header
Basic Syntax	X-Forwarded-For: addr
Example	X-Forwarded-For: 64.95.76.161

X-Pad

This header is used to overcome a bug related to response header length in some browsers; it pads the response message headers with extra bytes to work around the bug.

Type Extension general header

Basic Syntax X-Pad: pad-text

Example X-Pad: bogosity

X-Serial-Number

The X-Serial-Number header is an extension header. It was used by some older HTTP applications to insert the serial number of the licensed software in the HTTP message.

Its use has pretty much died out, but it is listed here as an example of the X headers that are out there.

Type Extension general header

Basic Syntax X-Serial-Number: serialno

Example X-Serial-Number: 010014056

APPENDIX D

MIME Types

MIME media types (MIME types, for short) are standardized names that describe the contents of a message entity body (e.g., text/html, image/jpeg). This appendix explains how MIME types work, how to register new ones, and where to go for more information.

In addition, this appendix contains 10 convenient tables, detailing hundreds of MIME types, gathered from many sources around the globe. This may be the most detailed tabular listing of MIME types ever compiled. We hope these tables are useful to you.

In this appendix, we will:

- Outline the primary reference material, in “Background.”
- Explain the structure of MIME types, in “MIME Type Structure.”
- Show you how to register MIME types, in “MIME Type IANA Registration.”
- Make it easier for you to look up MIME types.

The following MIME type tables are included in this appendix:

- application/*—Table D-3
- audio/*—Table D-4
- chemical/*—Table D-5
- image/*—Table D-6
- message/*—Table D-7
- model/*—Table D-8
- multipart/*—Table D-9
- text/*—Table D-10
- video/*—Table D-11
- Other—Table D-12

Background

MIME types originally were developed for multimedia email (MIME stands for Multipurpose Internet Mail Extensions), but they have been reused for HTTP and several other protocols that need to describe the format and purpose of data objects.

MIME is defined by five primary documents:

RFC 2045, "MIME: Format of Internet Message Bodies"

Describes the overall MIME message structure, and introduces the Content-Type header, borrowed by HTTP

RFC 2046, "MIME: Media Types"

Introduces MIME types and their structure

RFC 2047, "MIME: Message Header Extensions for Non-ASCII Text"

Defines ways to include non-ASCII characters in headers

RFC 2048, "MIME: Registration Procedures"

Defines how to register MIME values with the Internet Assigned Numbers Authority (IANA)

RFC 2049, "MIME: Conformance Criteria and Examples"

Details rules for compliance, and provides examples

For the purposes of HTTP, we are most interested in RFC 2046 (Media Types) and RFC 2048 (Registration Procedures).

MIME Type Structure

Each MIME media type consists of a type, a subtype, and a list of optional parameters. The type and subtype are separated by a slash, and the optional parameters begin with a semicolon, if they are present. In HTTP, MIME media types are widely used in Content-Type and Accept headers. Here are a few examples:

```
Content-Type: video/quicktime
Content-Type: text/html; charset="iso-8859-6"
Content-Type: multipart/mixed; boundary=gcOp4Jq0M2Yt08j34c0p
Accept: image/gif
```

Discrete Types

MIME types can directly describe the object type, or they can describe collections or packages of other object types. If a MIME type describes an object type directly, it is a *discrete type*. These include text files, videos, and application-specific file formats.

Composite Types

If a MIME type describes a collection or encapsulation of other content, the MIME type is called a *composite type*. A composite type describes the format of the enclosing

package. When the enclosing package is opened, each enclosed object will have its own type.

Multipart Types

Multipart media types are composite types. A multipart object consists of multiple component types. Here's an example of multipart/mixed content, where each component has its own MIME type:

```
Content-Type: multipart/mixed; boundary=unique-boundary-1

--unique-boundary-1
Content-type: text/plain; charset=US-ASCII

Hi there, I'm some boring ASCII text...

--unique-boundary-1
Content-Type: multipart/parallel; boundary=unique-boundary-2

--unique-boundary-2
Content-Type: audio/basic

    ... 8000 Hz single-channel mu-law-format
    audio data goes here ...

--unique-boundary-2
Content-Type: image/jpeg

    ... image data goes here ...

--unique-boundary-2--

--unique-boundary-1
Content-type: text/enriched

This is <bold><italic>enriched.</italic></bold>
<smaller>as defined in RFC 1896</smaller>

Isn't it <bigger><bigger>cool?</bigger></bigger>

--unique-boundary-1
Content-Type: message/rfc822

From: (mailbox in US-ASCII)
To: (address in US-ASCII)
Subject: (subject in US-ASCII)
Content-Type: Text/plain; charset=ISO-8859-1
Content-Transfer-Encoding: Quoted-printable

    ... Additional text in ISO-8859-1 goes here ...

--unique-boundary-1--
```

Syntax

As we stated earlier, MIME types consist of a primary type, a subtype, and an optional list of parameters.

The primary type can be a predefined type, an IETF-defined extension token, or an experimental token (beginning with “x-”). Some common primary types are described in Table D-1.

Table D-1. Common primary MIME types

Type	Description
application	Application-specific content format (discrete type)
audio	Audio format (discrete type)
chemical	Chemical data set (discrete IETF extension type)
image	Image format (discrete type)
message	Message format (composite type)
model	3-D model format (discrete IETF extension type)
multipart	Collection of multiple objects (composite type)
text	Text format (discrete type)
video	Video movie format (discrete type)

Subtypes can be primary types (as in “text/text”), IANA-registered subtypes, or experimental extension tokens (beginning with “x-”).

Types and subtypes are made up of a subset of US-ASCII characters. Spaces and certain reserved grouping and punctuation characters, called “tspecials,” are control characters and are forbidden from type and subtype names.

The grammar from RFC 2046 is shown below:

```
TYPE := "application" | "audio" | "image" | "message" | "multipart" |
      "text" | "video" | IETF-TOKEN | X-TOKEN
SUBTYPE := IANA-SUBTOKEN | IETF-TOKEN | X-TOKEN

IETF-TOKEN := <extension token with RFC and registered with IANA>
IANA-SUBTOKEN := <extension token registered with IANA>
X-TOKEN := <"X-" or "x-" prefix, followed by any token>

PARAMETER := TOKEN "=" VALUE
VALUE := TOKEN / QUOTED-STRING
TOKEN := 1*<any (US-ASCII) CHAR except SPACE, CTLs, or TSPECIALS>
TSPECIALS := "(" | ")" | "<" | ">" | "@" |
            "," | ";" | ":" | "\" | "<"> |
            "/" | "[" | "]" | "?" | "="
```

MIME Type IANA Registration

The MIME media type registration process is described in RFC 2048. The goal of the registration process is to make it easy to register new media types but also to provide some sanity checking to make sure the new types are well thought out.

Registration Trees

MIME type tokens are split into four classes, called “registration trees,” each with its own registration rules. The four trees—IETF, vendor, personal, and experimental—are described in Table D-2.

Table D-2. Four MIME media type registration trees

Registration tree	Example	Description
IETF	text/html (HTML text)	The IETF tree is intended for types that are of general significance to the Internet community. New IETF tree media types require approval by the Internet Engineering Steering Group (IESG) and an accompanying standards-track RFC. IETF tree types have no periods (.) in tokens.
Vendor (vnd.)	image/vnd.fpx (Kodak FlashPix image)	The vendor tree is intended for media types used by commercially available products. Public review of new vendor types is encouraged but not required. Vendor tree types begin with “vnd.”.
Personal/Vanity (prs.)	image/prs.btif (internal check-management format used by Nations Bank)	Private, personal, or vanity media types can be registered in the personal tree. These media types will not be distributed commercially. Personal tree types begin with “prs.”.
Experimental (x- or x.)	application/x-tar (Unix tar archive)	The experimental tree is for unregistered or experimental media types. Because it’s relatively simple to register a new vendor or personal media type, software should not be distributed widely using x- types. Experimental tree types begin with “x.” or “x-”.

Registration Process

Read RFC 2048 carefully for the details of MIME media type registration.

The basic registration process is not a formal standards process; it’s just an administrative procedure intended to sanity check new types with the community, and record them in a registry, without much delay. The process follows the following steps:

1. Present the media type to the community for review.

Send a proposed media type registration to the ietf-types@iana.org mailing list for a two-week review period. The public posting solicits feedback about the choice of name, interoperability, and security implications. The “x-” prefix specified in RFC 2045 can be used until registration is complete.

2. IESG approval (for IETF tree only).

If the media type is being registered in the IETF tree, it must be submitted to the IESG for approval and must have an accompanying standards-track RFC.

3. IANA registration.

As soon as the media type meets the approval requirements, the author can submit the registration request to the IANA, using the email template in Example D-1 and mailing the information to *ietf-types@iana.org*. The IANA will register the media type and make the media type application available to the community at <http://www.isi.edu/in-notes/iana/assignments/media-types/>.

Registration Rules

The IANA will register media types in the IETF tree only in response to a communication from the IESG stating that a given registration has been approved.

Vendor and personal types will be registered by the IANA automatically and without any formal review as long as the following minimal conditions are met:

1. Media types must function as actual media formats. Types that act like transfer encodings or character sets may not be registered as media types.
2. All media types must have proper type and subtype names. All type names must be defined by standards-track RFCs. All subtype names must be unique, must conform to the MIME grammar for such names, and must contain the proper tree prefixes.
3. Personal tree types must provide a format specification or a pointer to one.
4. Any security considerations given must not be obviously bogus. Everyone who is developing Internet software needs to do his part to prevent security holes.

Registration Template

The actual IANA registration is done via email. You complete a registration form using the template shown in Example D-1, and mail it to *ietf-types@iana.org*.*

Example D-1. IANA MIME registration email template

To: *ietf-types@iana.org*
Subject: Registration of MIME media type XXX/YYY

MIME media type name:

* The lightly structured nature of the form makes the submitted information fine for human consumption but difficult for machine processing. This is one reason why it is difficult to find a readable, well-organized summary of MIME types, and the reason we created the tables that end this appendix.

Example D-1. IANA MIME registration email template (continued)

MIME subtype name:

Required parameters:

Optional parameters:

Encoding considerations:

Security considerations:

Interoperability considerations:

Published specification:

Applications which use this media type:

Additional information:

 Magic number(s):

 File extension(s):

 Macintosh File Type Code(s):

Person & email address to contact for further information:

Intended usage:

(One of COMMON, LIMITED USE or OBSOLETE)

Author/Change controller:

(Any other information that the author deems interesting may be added below this line.)

MIME Media Type Registry

The submitted forms are accessible from the IANA web site (<http://www.iana.org>). At the time of writing, the actual database of MIME media types is stored on an ISI web server, at <http://www.isi.edu/in-notes/iana/assignments/media-types/>.

The media types are stored in a directory tree, structured by primary type and subtype, with one leaf file per media type. Each file contains the email submission. Unfortunately, each person completes the registration template slightly differently, so the quality and format of information varies across submissions. (In the tables in this appendix, we tried to fill in the holes omitted by registrants.)

MIME Type Tables

This section summarizes hundreds of MIME types in 10 tables. Each table lists the MIME media types within a particular primary type (image, text, etc.).

The information is gathered from many sources, including the IANA media type registry, the Apache *mime.types* file, and assorted Internet web pages. We spent several days refining the data, plugging holes, and including descriptive summaries from cross-references to make the data more useful.

This may well be the most detailed tabular listing of MIME types ever compiled. We hope you find it handy!

application/*

Table D-3 describes many of the application-specific MIME media types.

Table D-3. “Application” MIME types

MIME type	Description	Extension	Contact and reference
application/activemessage	Supports the Active Mail groupware system.		“Active Mail: A Framework for Integrated Groupware Applications” in <i>Readings in Groupware and Computer-Supported Cooperative Work</i> , Ronald M. Baecker, ed., Morgan Kaufmann, ISBN 1558602410
application/andrew-inset	Supports the creation of multimedia content with the Andrew toolkit.	ez	<i>Multimedia Applications Development with the Andrew Toolkit</i> , Nathaniel S. Borenstein, Prentice Hall, ASIN 0130366331 nsb@bellcore.com
application/applefile	Permits MIME-based transmission of data with Apple/Macintosh-specific information, while allowing general access to nonspecific user data.		RFC 1740
application/atomicmail	ATOMICMAIL was an experimental research project at Bellcore, designed for including programs in electronic mail messages that are executed when mail is read. ATOMICMAIL is rapidly becoming obsolete in favor of safe-tcl.		“ATOMICMAIL Language Reference Manual,” Nathaniel S. Borenstein, Bellcore Technical Memorandum TM ARH-018429
application/batch-SMTP	Defines a MIME content type suitable for tunneling an ESMTP mail transaction through any MIME-capable transport.		RFC 2442
application/beep+xml	Supports the interaction protocol called BEEP. BEEP permits simultaneous and independent exchanges of MIME messages between peers, where the messages usually are XML-structured text.		RFC 3080

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/cals-1840	Supports MIME email exchanges of U.S. Department of Defense digital data that was previously exchanged by tapem, as defined by MIL-STD-1840.		RFC 1895
application/commonground	Common Ground is an electronic document exchange and distribution program that lets users create documents that anyone can view, search, and print, without requiring that they have the creating applications or fonts on their systems.		Nick Gault No Hands Software ngault@nohands.com
application/cybercash	Supports credit card payment through the CyberCash protocol. When a user starts payment, a message is sent by the merchant to the customer as the body of a message of MIME type application/cybercash.		RFC 1898
application/dca-rft	IBM Document Content Architecture.		“IBM Document Content Architecture/Revisable Form Text Reference,” document number SC23-0758-1, International Business Machines
application/dec-dx	DEC Document Transfer Format.		“Digital Document Transmission (DX) Technical Notebook,” document number EJ29141-86, Digital Equipment Corporation
application/dvcs	Supports the protocols used by a Data Validation and Certification Server (DVCS), which acts as a trusted third party in a public-key security infrastructure.		RFC 3029
application/EDI-Consent	Supports bilateral trading via electronic data interchange (EDI), using nonstandard specifications.		http://www.isi.edu/in-notes/iana/assignments/media-types/application/EDI-Consent
application/EDI-X12	Supports bilateral trading via electronic data interchange (EDI), using the ASC X12 EDI specifications.		http://www.isi.edu/in-notes/iana/assignments/media-types/application/EDI-X12
application/EDIFACT	Supports bilateral trading via electronic data interchange (EDI), using the EDIFACT specifications.		http://www.isi.edu/in-notes/iana/assignments/media-types/application/EDIFACT
application/eshop	Unknown.		Steve Katz System Architecture Shop steve_katz@eshop.com
application/font-tdpfr	Defines a Portable Font Resource (PFR) that contains a set of glyph shapes, each associated with a character code.		RFC 3073

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/http	Used to enclose a pipeline of one or more HTTP request or response messages (not intermixed).		RFC 2616
application/hyperstudio	Supports transfer of HyperStudio educational hypermedia files.	stk	http://www.hyperstudio.com
application/iges	A commonly used format for CAD model interchange.		“ANS/US PRO/IPO-100” U.S. Product Data Association 2722 Merrilee Drive, Suite 200 Fairfax, VA 22031-4499
application/index application/index.cmd application/index.obj application/index.response application/index.vnd	Support the Common Indexing Protocol (CIP). CIP is an evolution of the Whois++ directory service, used to pass indexing information from server to server in order to redirect and replicate queries through a distributed database system.		RFC 2652, and RFCs 2651, 1913, and 1914
application/iotp	Supports Internet Open Trading Protocol (IOTP) messages over HTTP.		RFC 2935
application/ipp	Supports Internet Printing Protocol (IPP) over HTTP.		RFC 2910
application/mac-binhex40	Encodes a string of 8-bit bytes into a string of 7-bit bytes, which is safer for some applications (though not quite as safe as the 6-bit base-64 encoding).	hqx	RFC 1341
application/mac-compactpro	From Apache <i>mime.types</i> .	cpt	
application/macwriteii	Claris MacWrite II.		
application/marc	MARC objects are Machine-Readable Cataloging records—standards for the representation and communication of bibliographic and related information.	mrc	RFC 2220
application/mathematica application/mathematica-old	Supports Mathematica and Math-Reader numerical analysis software.	nb, ma, mb	<i>The Mathematica Book</i> , Stephen Wolfram, Cambridge University Press, ISBN 0521643147
application/msword	Microsoft Word MIME type.	doc	
application/news-message-id			RFCs 822 (message IDs), 1036 (application to news), and 977 (NNTP)
application/news-transmission	Allows transmission of news articles by email or other transport.		RFC 1036
application/ocsp-request	Supports the Online Certificate Status Protocol (OCSP), which provides a way to check on the validity of a digital certificate without requiring local certificate revocation lists.	orq	RFC 2560

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/ocsp-response	Same as above.	ors	RFC 2560
application/octet-stream	Unclassified binary data.	bin, dms, lha, lzh, exe, class	RFC 1341
application/oda	Used for information encoded according to the Office Document Architecture (ODA) standards, using the Office Document Interchange Format (ODIF) representation format. The Content-Type line also should specify an attribute/value pair that indicates the document application profile (DAP), as in: Content-Type: application/oda; profile=Q112	oda	RFC 1341 ISO 8613; “Information Processing: Text and Office System; Office Document Architecture (ODA) and Interchange Format (ODIF),” Part 1-8, 1989
application/parityfec	Forward error correction parity encoding for RTP data streams.		RFC 3009
application/pdf	Adobe PDF files.	pdf	See <i>Portable Document Format Reference Manual</i> , Adobe Systems, Inc., Addison Wesley, ISBN 0201626284
application/pgp-encrypted	PGP encrypted data.		RFC 2015
application/pgp-keys	PGP public-key blocks.		RFC 2015
application/pgp-signature	PGP cryptographic signature.		RFC 2015
application/pkcs10	Public Key Crypto System #10—the application/pkcs10 body type <i>must</i> be used to transfer a PKCS #10 certification request.	p10	RFC 2311
application/pkcs7-mime	Public Key Crypto System #7—this type is used to carry PKCS #7 objects of several types including envelopedData and signedData.	p7m	RFC 2311
application/pkcs7-signature	Public Key Crypto System #7—this type always contains a single PKCS #7 object of type signedData.	p7s	RFC 2311
application/pkix-cert	Transports X.509 certificates.	cer	RFC 2585
application/pkix-crl	Transports X.509 certificate revocation lists.	crl	RFC 2585
application/pkixcmp	Message format used by X.509 Public Key Infrastructure Certificate Management Protocols.	pki	RFC 2510
application/postscript	An Adobe PostScript graphics file (program).	ai, ps, eps	RFC 2046
application/prs.alvestrand.titrax-sheet	“TimeTracker” program by Harald T. Alvestrand.		http://domen.uninett.no/~hta/titrax/

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/prs.cww	CU-Writer for Windows.	cw, cww	Dr. Somchai Prasitjutrakul somchaip@chulkn.car.chula.ac.th
application/prs.nprend	Unknown.	rnd, rct	John M. Doggett jdoggett@tiac.net
application/remote-printing	Contains meta information used when remote printing, for the printer cover sheet.		RFC 1486 Marshall T. Rose mrose@dbc.mtview.ca.us
application/riscos	Acorn RISC OS binaries.		<i>RISC OS Programmer's Reference Manuals</i> , Acorn Computers, Ltd., ISBN1852501103
application/sdp	SDP is intended for describing live multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation.		RFC 2327 Henning Schulzrinne hgs@cs.columbia.edu
application/set-payment-initiation	Supports the SET secure electronic transaction payment protocol.		http://www.visa.com
application/set-payment-initiation			http://www.mastercard.com
application/set-registration-initiation			
application/sgml-open-catalog	Intended for use with systems that support the SGML Open TR9401:1995 “Entity Management” specification.		SGML Open 910 Beaver Grade Road, #3008 Coraopolis, PA 15109 info@sgmlopen.org
application/sieve	Sieve mail filtering script.		RFC 3028
application/slate	The BBN/Slate document format is published as part of the standard documentation set distributed with the BBN/Slate product.		BBN/Slate Product Mgr BBN Systems and Technologies 10 Moulton Street Cambridge, MA 02138
application/smil	The Synchronized Multimedia Integration Language (SMIL) integrates a set of independent multimedia objects into a synchronized multimedia presentation.	smi, smil	http://www.w3.org/AudioVideo/
application/tve-trigger	Supports embedded URLs in enhanced television receivers.		“SMPTe: Declarative Data Essence, Content Level 1,” produced by the Society of Motion Picture and Television Engineers http://www.smpete.org
application/vemmi	Enhanced videotex standard.		RFC 2122
application/vnd.3M.Post-it-Notes	Used by the “Post-it® Notes for Internet Designers” Internet control/plug-in.	pwn	http://www.3M.com/psnotes/

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.accpac. simply.aso	Simply Accounting v7.0 and higher. Files of this type conform to Open Financial Exchange v1.02 specifications.	aso	http://www.ofx.net
application/vnd.accpac. simply.imp	Used by Simply Accounting v7.0 and higher, to import its own data.	imp	http://www.ofx.net
application/vnd.acucobol	ACUCOBOL-GT Runtime.		Dovid Lubin dovid@acucobol.com
application/vnd.aether.imp	Supports airtime-efficient Instant Message communications between an Instant Messaging service, such as AOL Instant Messenger, Yahoo! Messenger, or MSN Messenger, and a special set of Instant Messaging client software on a wireless device.		Wireless Instant Messaging Protocol (IMP) specification available from Aether Systems by license
application/vnd.anser-web- certificate-issue-initiation	Trigger for web browsers to launch the ANSER-WEB Terminal Client.	cii	Hiroyoshi Mori mori@mm.rd.nttdata.co.jp
application/vnd.anser-web- funds-transfer-initiation	Same as above.	fti	Same as above
application/vnd.audiograph	AudioGraph.	aep	Horia Cristian H.C.Slusanschi@massey.ac.nz
application/vnd.bmi	BMI graphics format by CADAM Systems.	bmi	Tadashi Gotoh tgotoh@cadamsystems.co.jp
application/vnd. businessobjects	BusinessObjects 4.0 and higher.	rep	
application/vnd.canon-cpdl	Supports Canon, Inc. office imaging products.		Shin Muto shinmuto@pure.cpdn.canon.co.jp
application/vnd.canon-lips			
application/vnd.claymore	Claymore.exe.	cla	Ray Simpson ray@cnation.com
application/vnd.commerce- battelle	Supports a generic mechanism for delimiting smart card-based information, for digital commerce, identification, authentication, and exchange of smart card-based card holder information.	ica, icf, icd, icc, ic0, ic1, ic2, ic3, ic4, ic5, ic6, ic7, ic8	David C. Applebaum applebau@131.167.52.15
application/vnd. commonspace	Allows for proper transmission of CommonSpace™ documents via MIME-based processes. CommonSpace is published by Sixth Floor Media, part of the Houghton-Mifflin Company.	csp, cst	Ravinder Chandhok chandhok@within.com
application/vnd.contact.cmsg	Used for CONTACT software’s CIM DATABASE.	cdbcmsg	Frank Patz fp@contact.de http://www.contact.de

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.cosmocaller	Allows for files containing connection parameters to be downloaded from web sites, invokes the CosmoCaller application to interpret the parameters, and initiates connections with the CosmoCallACD server.	cmc	Steve Dellutri sdellutri@cosmocom.com
application/vnd.ctc-posml	Continuum Technology’s PosML.	pml	Bayard Kohlhepp bayardk@ctcexchange.com http://www.cups.org
application/vnd.cups-postscript	Supports Common UNIX Printing System (CUPS) servers and clients.		
application/vnd.cups-raster			
application/vnd.cups-raw			
application/vnd.cybank	Proprietary data type for Cybank data.		Nor Helmee B. Abd. Halim helmee@cybank.net http://www.cybank.net
application/vnd.dna	DNA is intended to easily Web-enable any 32-bit Windows application.	dna	Meredith Searcy msearcy@newmoon.com
application/vnd.dpggraph	Used by DPGraph 2000 and MathWare Cyclone.	dpg, mwc, dpgraph	David Parker http://www.davidparker.com
application/vnd.dxr	Digital Xpress Reports by PSI Technologies.	dxr	Michael Duffy miked@psiaustin.com http://www.sevencs.com
application/vnd.ecdis-update	Supports ECDIS applications.	mag	Thomas Olsson thomas@vinga.se
application/vnd.ecowin.chart	EcoWin.		
application/vnd.ecowin.filerequest			
application/vnd.ecowin.fileupdate			
application/vnd.ecowin.series			
application/vnd.ecowin.seriesrequest			
application/vnd.ecowin.seriesupdate			
application/vnd.enliven	Supports delivery of Enliven interactive multimedia.	nml	Paul Santinelli psantinelli@narrative.com
application/vnd.epson.esf	Proprietary content for Seiko Epson QUASS Stream Player.	esf	Shoji Hoshina Hoshina.Shoji@exc.epson.co.jp
application/vnd.epson.msf	Proprietary content for Seiko Epson QUASS Stream Player.	msf	Same as above
application/vnd.epson.quickanime	Proprietary content for Seiko Epson QuickAnime Player.	qam	Yu Gu guyu@rd.oda.epson.co.jp
application/vnd.epson.salt	Proprietary content for Seiko Epson SimpleAnimeLite Player.	slt	Yasuhito Nagatomo naga@rd.oda.epson.co.jp
application/vnd.epson.ssf	Proprietary content for Seiko Epson QUASS Stream Player.	ssf	Shoji Hoshina Hoshina.Shoji@exc.epson.co.jp

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.ericsson.quickcall	Phone Doubler Quick Call.	qcall, qca	Paul Tidwell paul.tidwell@ericsson.com http://www.ericsson.com
application/vnd.eudora.data	Eudora Version 4.3 and later.		Pete Resnick presnick@qualcomm.com
application/vnd.fdf	Adobe Forms Data Format.		“Forms Data Format,” Technical Note 5173, Adobe Systems
application/vnd.ffsns	Used for application communication with FirstFloor’s Smart Delivery.		Mary Holstege holstege@firstfloor.com
application/vnd.FloGraphIt	NpGraphIt.	gph	
application/vnd.framemaker	Adobe FrameMaker files.	fm, mif, book	http://www.adobe.com
application/vnd.fsc.weblaunch	Supports Friendly Software Corporation’s golf simulation software.	fsc	Derek Smith derek@friendlysoftware.com
application/vnd.fujitsu.oasys	Supports Fujitsu’s OASYS software.	oas	Nobukazu Togashi togashi@ai.cs.fujitsu.co.jp
application/vnd.fujitsu.oasys2	Supports Fujitsu’s OASYS V2 software.	oa2	Same as above
application/vnd.fujitsu.oasys3	Support’s Fujitsu’s OASYS V5 software.	oa3	Seiji Okudaira okudaira@candy.paso.fujitsu.co.jp
application/vnd.fujitsu.oasysgp	Supports Fujitsu’s OASYS GraphPro software.	fg5	Masahiko Sugimoto sugimoto@sz.sel.fujitsu.co.jp
application/vnd.fujitsu.oasysprs	Support’s Fujitsu’s OASYS Presentation software.	bh2	Masumi Ogita ogita@oa.tfl.fujitsu.co.jp
application/vnd.fujixerox.ddd	Supports Fuji Xerox’s EDMICS 2000 and DocuFile.	ddd	Masanori Onda Masanori.Onda@fujixerox.co.jp
application/vnd.fujixerox.docuworks	Supports Fuji Xerox’s DocuWorks Desk and DocuWorks Viewer software.	xdw	Yasuo Taguchi yasuo.taguchi@fujixerox.co.jp
application/vnd.fujixerox.docuworks.binder	Supports Fuji Xerox’s DocuWorks Desk and DocuWorks Viewer software.	xbd	Same as above.
application/vnd.fut-misnet	Unknown.		Jaan Pruulmann jaan@fut.ee
application/vnd.grafeq	Lets users of GrafEq exchange GrafEq documents through the Web and email.	gqf, gqs	http://www.peda.com
application/vnd.groove-account	Groove is a peer-to-peer communication system implementing a virtual space for small group interaction.	gac	Todd Joseph todd_joseph@groove.net
application/vnd.groove-identity-message	Same as above.	gim	Same as above
application/vnd.groove-injector	Same as above.	grv	Same as above
application/vnd.groove-tool-message	Same as above.	gtm	Same as above

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.groove-tool-template	Same as above.	tpl	Same as above
application/vnd.groove-vcard	Same as above.	vcg	Same as above
application/vnd.hhe.lesson-player	Supports the LessonPlayer and PresentationEditor software.	les	Randy Jones Harcourt E-Learning randy_jones@archipelago.com
application/vnd.hp-HPGL	HPGL files.		<i>The HP-GL/2 and HP RTL Reference Guide</i> , Addison Wesley, ISBN 0201310147
application/vnd.hp-hpid	Supports Hewlett-Packard’s Instant Delivery Software.	hpi, hpid	http://www.instant-delivery.com
application/vnd.hp-hps	Supports Hewlett-Packard’s Web-PrintSmart software.	hps	http://www.hp.com/go/webprintsmart_mime_type_specs/
application/vnd.hp-PCL application/vnd.hp-PCLXL	PCL printer files.	pcl	“PCL-PJL Technical Reference Manual Documentation Package,” HP Part No. 5012-0330
application/vnd.httpphone	HTTPPhone asynchronous voice over IP system.		Franck LeFevre franck@k1info.com
application/vnd.hzn-3d-crossword	Used to encode crossword puzzles by Horizon, A Glimpse of Tomorrow.	x3d	James Minnis james_minnis@glimpse-of-tomorrow.com
application/vnd.ibm.afplinedata	Print Services Facility (PSF), AFP Conversion and Indexing Facility (ACIF).		Roger Buis buis@us.ibm.com
application/vnd.ibm.Minipay	Minipay authentication and payment software.	mpy	Amir Herzberg amirh@vnet.ibm.com
application/vnd.ibm.modcap	Mixed Object Document Content.	list3820, listaftp, afp, pseg3820	Reinhard Hohensee rhohensee@vnet.ibm.com “Mixed Object Document Content Architecture Reference,” IBM publication SC31-6802
application/vnd.informix-visionary	Informix Visionary.	vis	Christopher Gales christopher.gales@informix.com
application/vnd.intercon.formnet	Supports Intercon Associates FormNet software.	xpw, xpx	Thomas A. Gurak assoc@intercon.roc.servtech.com
application/vnd.intertrust.digibox application/vnd.intertrust.nncp	Supports InterTrust architecture for secure electronic commerce and digital rights management.		InterTrust Technologies 460 Oakmead Parkway Sunnyvale, CA 94086 USA info@intertrust.com http://www.intertrust.com
application/vnd.intu.qbo	Intended for use only with Quick-Books 6.0 (Canada).	qbo	Greg Scratchley greg_scratchley@intuit.com Format of these files discussed in the Open Financial Exchange specs, available from http://www.ofx.net

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.intu.qfx	Intended for use only with Quicken 99 and following versions.	qfx	Same as above
application/vnd.is-xpr	Express by Infoseek.	xpr	Satish Natarajan satish@infoseek.com
application/vnd.japannet-directory-service	Supports Mitsubishi Electric’s Japan-Net security, authentication, and payment software.		Jun Yoshitake yositake@iss.isl.melco.co.jp
application/vnd.japannet-jpnstore-wakeup			
application/vnd.japannet-payment-wakeup			
application/vnd.japannet-registration			
application/vnd.japannet-registration-wakeup			
application/vnd.japannet-setstore-wakeup			
application/vnd.japannet-verification			
application/vnd.japannet-verification-wakeup			
application/vnd.koan	Supports the automatic playback of Koan music files over the Internet, by helper applications such as SSEYO Koan Netscape Plugin.	skp, skd, skm, skt	Peter Cole pcole@sseyod.demon.co.uk
application/vnd.lotus-1-2-3	Lotus 1-2-3 and Lotus approach.	123, wk1, wk3, wk4	Paul Wattenberger Paul_Wattenberger@lotus.com
application/vnd.lotus-approach	Lotus Approach.	apr, vew	Same as above
application/vnd.lotus-freelance	Lotus Freelance.	prz, pre	Same as above
application/vnd.lotus-notes	Lotus Notes.	nsf, ntf, ndl, ns4, ns3, ns2, nsh, nsq	Michael Laramie laramiem@btv.ibm.com
application/vnd.lotus-organizer	Lotus Organizer.	or3, or2, org	Paul Wattenberger Paul_Wattenberger@lotus.com
application/vnd.lotus-screencam	Lotus ScreenCam.	scm	Same as above
application/vnd.lotus-wordpro	Lotus Word Pro.	lwp, sam	Same as above
application/vnd.mcd	Micro CADAM CAD software.	mcd	Tadashi Gotoh tgotoh@cadamsystems.co.jp http://www.cadamsystems.co.jp

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd. mediastation.cdkey	Supports Media Station’s CDKey remote CDRom communications protocol.	cdkey	Henry Flurry henryf@mediastation.com
application/vnd.meridian- slingshot	Slingshot by Meridian Data.		Eric Wedel Meridian Data, Inc. 5615 Scotts Valley Drive Scotts Valley, CA 95066 ewedel@meridian-data.com
application/vnd.mif	FrameMaker interchange format.	mif	ftp://ftp.frame.com/pub/techsup/techinfo/dos/mif4.zip Mike Wexler Adobe Systems, Inc 333 W. San Carlos St. San Jose, CA 95110 USA mwexler@adobe.com
application/vnd.minisoft- hp3000-save	NetMail 3000 save format.		Minisoft, Inc. support@minisoft.com ftp://ftp.3k.com/DOC/ms92-save-format.txt
application/vnd.mitsubishi. misty-guard.trustweb	Supports Mitsubishi Electric’s Trustweb software.		Manabu Tanaka mtana@iss.isl.melco.co.jp
application/vnd.Mobius.DAF	Supports Mobius Management Systems software.	daf	Celso Rodriguez crodrigu@mobius.com Greg Chrzczon gchrzco@mobius.com
application/vnd.Mobius.DIS	Same as above.	dis	Same as above
application/vnd.Mobius.MBK	Same as above.	mbk	Same as above
application/vnd.Mobius.MQY	Same as above.	mqy	Same as above
application/vnd.Mobius.MSL	Same as above.	msl	Same as above
application/vnd.Mobius.PLC	Same as above.	plc	Same as above
application/vnd.Mobius.TXF	Same as above.	txf	Same as above
application/vnd.motorola. flexsuite	FLEXsuite™ is a collection of wireless messaging protocols. This type is used by the network gateways of wireless messaging service providers as well as wireless OSs and applications.		Mark Patton Motorola Personal Networks Group fmp014@email.mot.com FLEXsuite™ specification available from Motorola under appropriate licensing agreement
application/vnd.motorola. flexsuite.adsi	FLEXsuite™ is a collection of wireless messaging protocols. This type provides a wireless-friendly format for enabling various data-encryption solutions.		Same as above

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.motorola.flexsuite.fis	FLEXsuite™ is a collection of wireless messaging protocols. This type is a wireless-friendly format for the efficient delivery of structured information (e.g., news, stocks, weather) to a wireless device.		Same as above
application/vnd.motorola.flexsuite.gotap	FLEXsuite™ is a collection of wireless messaging protocols. This type provides a common wireless-friendly format for the programming of wireless device attributes via over-the-air messages.		Same as above
application/vnd.motorola.flexsuite.kmr	FLEXsuite™ is a collection of wireless messaging protocols. This type provides a wireless-friendly format for encryption key management.		Same as above
application/vnd.motorola.flexsuite.ttc	FLEXsuite™ is a collection of wireless messaging protocols. This type supports a wireless-friendly format for the efficient delivery of text using token text compression.		Same as above
application/vnd.motorola.flexsuite.wem	FLEXsuite™ is a collection of wireless messaging protocols. This type provides a wireless-friendly format for the communication of Internet email to wireless devices.		Same as above
application/vnd.mozilla.xul+xml	Supports the Mozilla Internet application suite.	xul	Dan Rosen2 dr@netscape.com
application/vnd.ms-artgalry	Supports Microsoft’s Art Gallery.	cil	deansl@microsoft.com
application/vnd.ms-asf	ASF is a multimedia file format whose contents are designed to be streamed across a network to support distributed multimedia applications. ASF content may include any combination of any media type (e.g., audio, video, images, URLs, HTML content, MIDI, 2-D and 3-D modeling, scripts, and objects of various types).	asf	Eric Fleischman ericf@microsoft.com http://www.microsoft.com/mind/0997/netshow/netshow.asp
application/vnd.ms-excel	Microsoft Excel spreadsheet.	xls	Sukvinder S. Gill sukvg@microsoft.com
application/vnd.ms-lrm	Microsoft proprietary.	lrm	Eric Ledoux ericle@microsoft.com
application/vnd.ms-powerpoint	Microsoft PowerPoint presentation.	ppt	Sukvinder S. Gill sukvg@microsoft.com
application/vnd.ms-project	Microsoft Project file.	mpp	Same as above

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.ms-tnef	Identifies an attachment that in general would be processable only by a MAPI-aware application. This type is an encapsulated format of rich MAPI properties, such as Rich Text and Icon information, that may otherwise be degraded by the messaging transport.		Same as above
application/vnd.ms-works	Microsoft Works software.		Same as above
application/vnd.mseq	MSEQ is a compact multimedia format suitable for wireless devices.	mseq	Gwenaél Le Bodic Gwenaél.le_bodic@alcatel.fr http://www.3gpp.org
application/vnd.msimg	Used by applications implementing the msimg protocol, which requests signatures from mobile devices.		Malte Borchering Malte.Borchering@brokat.com
application/vnd.music-niff	NIFF music files.		Cindy Grande 72723.1272@compuserve.com ftp://blackbox.cartah.washington.edu/pub/NIFF/NIFF6A.TXT
application/vnd.musician	MUSICIAN scoring language/encoding conceived and developed by Renaissance Corporation.	mus	Robert G. Adams gadams@renaissance.com
application/vnd.netfpx	Intended for dynamic retrieval of multiresolution image information, as used by Hewlett-Packard Company Imaging for Internet.	fpx	Andy Mutz andy_mutz@hp.com
application/vnd.noblenet-directory	Supports the NobleNet Directory software, purchased by RogueWave.	nnd	http://www.noblenet.com
application/vnd.noblenet-sealer	Supports the NobleNet Sealer software, purchased by RogueWave.	nns	http://www.noblenet.com
application/vnd.noblenet-web	Supports the NobleNet Web software, purchased by RogueWave.	nnw	http://www.noblenet.com
application/vnd.novadigm.EDM	Supports Novadigm’s RADIA and EDM products.	edm	Phil Burgard pburgard@novadigm.com
application/vnd.novadigm.EDX	Same as above.	edx	Same as above
application/vnd.novadigm.EXT	Same as above.	ext	Same as above
application/vnd.osa.netdeploy	Supports the Open Software Associates netDeploy application deployment software.	ndc	Steve Klos stevek@osa.com http://www.osa.com
application/vnd.palm	Used by PalmOS system software and applications—this new type, “application/vnd.palm,” replaces the old type “application/x-pilot.”	prc, pdb, pqa, oprc	Gavin Peacock gpeacock@palm.com

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.pg.format	Proprietary Proctor & Gamble Standard Reporting System.	str	April Gandert TN152 Procter & Gamble Way Cincinnati, Ohio 45202 (513) 983-4249
application/vnd.pg.osasli	Proprietary Proctor & Gamble Standard Reporting System.	ei6	Same as above
application/vnd. powerbuilder6 application/vnd. powerbuilder6-s application/vnd. powerbuilder7 application/vnd. powerbuilder7-s application/vnd. powerbuilder75 application/vnd. powerbuilder75-s	Used only by Sybase PowerBuilder release 6, 7, and 7.5 runtime environments, nonsecure and secure.	pbd	Reed Shilts reed.shilts@sybase.com
application/vnd. previewsystems.box	Preview Systems ZipLock/VBox product.	box, vbox	Roman Smolgovsky romans@previewsystems.com http://www.previewsystems.com
application/vnd.publishare-delta-tree	Used by Capella Computers' PubliShare runtime environment.	qps	Oren Ben-Kiki publishare-delta-tree@capella.co.il
application/vnd.rapid	Emultek's rapid packaged applications.	zrp	Itay Szekeley etay@emultek.co.il
application/vnd.s3sms	Integrates the transfer mechanisms of the Sonera SmartTrust products into the Internet infrastructure.		Lauri Tarkkala Lauri.Tarkkala@sonera.com http://www.smarttrust.com
application/vnd.seemail	Supports the transmission of SeeMail files. SeeMail is an application that captures video and sound and uses bitwise compression to compress and archive the two pieces into one file.	see	Steven Webb steve@wynde.com http://www.realmediainc.com
application/vnd.shana. informed.formdata	Shana e-forms data formats.	ifm	Guy Selzler Shana Corporation gselzler@shana.com
application/vnd.shana. informed.formtemp	Shana e-forms data formats.	itp	Same as above
application/vnd.shana. informed.interchange	Shana e-forms data formats.	iif, iif1	Same as above
application/vnd.shana. informed.package	Shana e-forms data formats.	ipk, ipkg	Same as above
application/vnd.street-stream	Proprietary to Street Technologies.		Glenn Levitt Street Technologies streetd1@ix.netcom.com

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.svd	Dataware Electronics SVD files.		Scott Becker dataware@compumedia.com
application/vnd.swiftview-ics	Supports SwiftView®.		Randy Prakken tech@ndg.com http://www.ndg.com/svm.htm
application/vnd.triscape.mxs	Supports Triscape Map Explorer.	mxs	Steven Simonoff scs@triscape.com
application/vnd.trueapp	True BASIC files.	tra	J. Scott Hepler scott@truebasic.com
application/vnd.truedoc	Proprietary to Bitstream, Inc.		Brad Chase brad_chase@bitstream.com
application/vnd.ufdl	UWI's UFDL files.	ufdl, ufd, frm	Dave Manning dmanning@uwi.com http://www.uwi.com/
application/vnd.uplanet.alert	Formats used by Unwired Planet (now Openwave) UP browser micro-browser for mobile devices.		iana-registrar@uplanet.com http://www.openwave.com
application/vnd.uplanet.alert-wbxml			
application/vnd.uplanet.bearer-choi-wbxml			
application/vnd.uplanet.bearer-choice			
application/vnd.uplanet.cacheop			
application/vnd.uplanet.cacheop-wbxml			
application/vnd.uplanet.channel			
application/vnd.uplanet.channel-wbxml			
application/vnd.uplanet.list			
application/vnd.uplanet.list-wbxml			
application/vnd.uplanet.listcmd			
application/vnd.uplanet.listcmd-wbxml			
application/vnd.uplanet.signal			
application/vnd.vcx	VirtualCatalog.	vcx	Taisuke Sugimoto sugimototi@noanet.nttdata.co.jp
application/vnd.vectorworks	VectorWorks graphics files.	mcd	Paul C. Pharr pharr@diehlgraphsoft.com
application/vnd.vidsoft.vidconference	VidConference format.	vsc	Robert Hess hess@vidsoft.de
application/vnd.visio	Visio files.	vsd, vst, vsw, vss	Troy Sandal troys@visio.com

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/vnd.vividence.scriptfile	Vividence files.	vsf, vtd, vd	Mark Risher markr@vividence.com
application/vnd.wap.sic	WAP Service Indication format.	sic, wbxml	WAP Forum Ltd http://www.wapforum.org
application/vnd.wap.slc	WAP Service Loading format. Anything that conforms to the Service Loading specification, available at http://www.wapforum.org .	slc, wbxml	Same as above
application/vnd.wap.wbxml	WAP WBXML binary XML format for wireless devices.	wbxml	Same as above “WAP Binary XML Content Format—WBXML version 1.1”
application/vnd.wap.wmlc	WAP WML format for wireless devices.	wmlc, wbxml	Same as above
application/vnd.wap.wmlscriptc	WAP WMLScript format.	wmlsc	Same as above
application/vnd.webturbo	WebTurbo format.	wtb	Yaser Rehem Sapient Corporation yrehem@sapient.com
application/vnd.wrq-hp3000-labelled	Supports HP3000 formats.		support@wrq.com support@3k.com
application/vnd.wt.stf	Supports Worldtalk software.	stf	Bill Wohler wohler@worldtalk.com
application/vnd.xara	Xara files are saved by CorelXARA, an object-oriented vector graphics package written by Xara Limited (and marketed by Corel).	xar	David Matthewman david@xara.com http://www.xara.com
application/vnd.xfdl	UWI’s XFDL files.	xfdl, xfd, frm	Dave Manning dmanning@uwi.com http://www.uwi.com
application/vnd.yellowriver-custom-menu	Supports the Yellow River Custom-Menu plug-in, which provides customized browser drop-down menus.	cmp	yellowriversw@yahoo.com
application/whoispp-query	Defines Whois++ protocol queries within MIME.		RFC 2957
application/whoispp-response	Defines Whois++ protocol responses within MIME.		RFC 2958
application/wita	Wang Information Transfer Architecture.		Document number 715-0050A, Wang Laboratories campbell@redsox.bsw.com
application/wordperfect5.1	WordPerfect documents.		
application/x400-bp	Carries any X.400 body part for which there is no registered IANA mapping.		RFC 1494

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/x-bcpio	Old-style binary CPIO archives.	bcpio	http://www.cdlink.com
application/x-cdlink	Allows integration of CD-ROM media within web pages.	vcd	
application/x-chess-pgm	From Apache <i>mime.types</i> .	pgn	
application/x-compress	Binary data from Unix compress.	z	
application/x-cpio	CPIO archive file.	cpio	
application/x-csh	CSH scripts.	csh	
application/x-director	Macromedia director files.	dcr, dir, dxr	
application/x-dvi	TeX DVI files.	dvi	
application/x-futuresplash	From Apache <i>mime.types</i> .	spl	
application/x-gtar	GNU tar archives.	gtar	
application/x-gzip	GZIP compressed data.	gz	
application/x-hdf	From Apache <i>mime.types</i> .	hdf	
application/x-javascript	JavaScript files.	js	
application/x-koan	Supports the automatic playback of Koan music files over the Internet, by helper applications such as SSEYO Koan Netscape Plugin.	skp, skd, skt, skm	
application/x-latex	LaTeX files.	latex	
application/x-netcdf	NETCDF files.	nc, cdf	
application/x-sh	SH scripts.	sh	
application/x-shar	SHAR archives.	shar	
application/x-shockwave-flash	Macromedia Flash files.	swf	
application/x-stuffit	Stuffit archives.	sit	
application/x-sv4cpio	Unix SysV R4 CPIO archives.	sv4cpio	
application/x-sv4crc	Unix SysV R4 CPIO w/CRC archives.	sv4crc	
application/x-tar	TAR archives.	tar	
application/x-tcl	TCL scripts.	tcl	
application/x-tex	TeX files.	tex	
application/x-texinfo	TeX info files.	texinfo, texi	
application/x-troff	TROFF files.	t, tr, roff	
application/x-troff-man	TROFF Unix manpages.	man	
application/x-troff-me	TROFF+me files.	me	
application/x-troff-ms	TROFF+ms files	ms	

Table D-3. “Application” MIME types (continued)

MIME type	Description	Extension	Contact and reference
application/x-ustar	The extended tar interchange format.	ustar	See the IEEE 1003.1(1990) specifications
application/x-wais-source	WAIS source structure.	src	RFC 2376
application/xml	Extensible Markup Language format file (use text/xml if you want the file treated as plain text by browsers, etc.).	xml, dtd	
application/zip	PKWARE zip archives.	zip	

audio/*

Table D-4 summarizes audio content types.

Table D-4. “Audio” MIME types

MIME type	Description	Extension	Contact and reference
audio/32kadtcm	8 kHz ADPCM audio encoding.	au, snd	RFC 2421
audio/basic	Audio encoded with 8-kHz monaural 8-bit ISDN u-law PCM.		RFC 1341
audio/G.722.1	G.722.1 compresses 50Hz–7kHz audio signals into 24 kbit/s or 32 kbit/s. It may be used for speech, music, and other types of audio.		RFC 3047
audio/L16	Audio/L16 is based on L16, described in RFC 1890. L16 denotes uncompressed audio data, using 16-bit signed representation.	mid, midi, kar	RFC 2586
audio/MP4A-LATM	MPEG-4 audio.		RFC 3016
audio/midi	MIDI music files.		RFC 3003
audio/mpeg	MPEG encoded audio files.	mpga, mp2, mp3	
audio/parityfec	Parity-based forward error correction for RTP audio.		
audio/prs.sid	Commodore 64 SID audio files.	sid, psid	http://www.geocities.com/SiliconValley/Lakes/5147/sidplay/docs.html#fileformats
audio/telephone-event	Logical telephone event.		RFC 2833
audio/tone	Telephonic sound pattern.		RFC 2833
audio/vnd.cns.anp1	Supports voice and unified messaging application features available on the Access NP network services platform from Comverse Network Systems.		Ann McLaughlin Comverse Network Systems amclaughlin@comversens.com

Table D-4. “Audio” MIME types (continued)

MIME type	Description	Extension	Contact and reference
audio/vnd.cns.inf1	Supports voice and unified messaging application features available on the TRILOGUE Infinity network services platform from Comverse Network Systems.		Same as above
audio/vnd.digital-winds	Digital Winds music is never-ending, reproducible, and interactive MIDI music in very small packages (<3K).	eol	Armands Strazds armands.strazds@medienhaus-bremen.de
audio/vnd.everad.plj	Proprietary EverAD audio encoding.	plj	Tomer Weisberg tomer@everad.com
audio/vnd.lucent.voice	Voice messaging including Lucent Technologies’ Intuity™ AUDIX® Multimedia Messaging System and the Lucent Voice Player.	lvp	Frederick Block rickblock@lucent.com http://www.lucent.com/lvp/
audio/vnd.nortel.vbk	Proprietary Nortel Networks Voice Block audio encoding.	vbk	Glenn Parsons Glenn.Parsons@NortelNetworks.com
audio/vnd.nuera.ecelp4800	Proprietary Nuera Communications audio and speech encoding, available in Nuera voice-over-IP gateways, terminals, application servers, and as a media service for various host platforms and OSs.	ecelp4800	Michael Fox mfox@nuera.com
audio/vnd.nuera.ecelp7470	Same as above.	ecelp7470	Same as above
audio/vnd.nuera.ecelp9600	Same as above.	ecelp9600	Same as above
audio/vnd.octel.sbc	Variable-rate encoding averaging 18 kbps used for voice messaging in Lucent Technologies’ Sierra™, Overture™, and IMA™ platforms.		Jeff Bouis jbouis@lucent.com
audio/vnd.qcelp	Qualcomm audio encoding.	qcp	Andy Dejaco adejaco@qualcomm.com
audio/vnd.rhetorex.32kadpcm	32-kbps Rhetorex™ ADPCM audio encoding used in voice messaging products such as Lucent Technologies’ CallPerformer™, Unified Messenger™, and other products.		Jeff Bouis jbouis@lucent.com
audio/vnd.vmx.cvsd	Audio encoding used in voice messaging products including Lucent Technologies’ Overture200™, Overture 300™, and VMX 300™ product lines.		Same as above
audio/x-aiff	AIFF audio file format.	aif, aiff, aifc	
audio/x-pn-realaudio	RealAudio metafile format by Real Networks (formerly Progressive Networks).	ram, rm	
audio/x-pn-realaudio-plugin	From Apache <i>mime.types</i> .	rpm	

Table D-4. “Audio” MIME types (continued)

MIME type	Description	Extension	Contact and reference
audio/x-realaudio	RealAudio audio format by Real Networks (formerly Progressive Networks).	ra	
audio/x-wav	WAV audio files.	wav	

chemical/*

Much of the information in Table D-5 was obtained courtesy of the “Chemical MIME Home Page” (<http://www.ch.ic.ac.uk/chemime/>).

Table D-5. “Chemical” MIME types

MIME type	Description	Extension	Contact and reference
chemical/x-alchemy	Alchemy format	alc	http://www.camsoft.com
chemical/x-cache-csf		csf	
chemical/x-cactvs-binary	CACTVS binary format	cbin	http://cactvs.cit.nih.gov
chemical/x-cactvs-ascii	CACTVS ASCII format	cascii	http://cactvs.cit.nih.gov
chemical/x-cactvs-table	CACTVS table format	ctab	http://cactvs.cit.nih.gov
chemical/x-cdx	ChemDraw eXchange file	cdx	http://www.camsoft.com
chemical/x-cerius	MSI Cerius II format	cer	http://www.msi.com
chemical/x-chemdraw	ChemDraw file	chm	http://www.camsoft.com
chemical/x-cif	Crystallographic Interchange Format	cif	http://www.bernstein-plus-sons.com/software/rasmol/ http://ndbserver.rutgers.edu/NDB/mmcif/examples/index.html
chemical/x-mmCIF	MacroMolecular CIF	mcif	Same as above
chemical/x-chem3d	Chem3D format	c3d	http://www.camsoft.com
chemical/x-cmdf	CrystalMaker Data Format	cmdf	http://www.crystallmaker.co.uk
chemical/x-compass	Compass program of the Takahashi	cpa	
chemical/x-crossfire	Crossfire file	bsd	
chemical/x-cml	Chemical Markup Language	cml	http://www.xml-cml.org
chemical/x-csml	Chemical Style Markup Language	csml, csm	http://www.mdli.com
chemical/x-ctx	Gasteiger group CTX file format	ctx	
chemical/x-cxf		cxf	
chemical/x-daylight-smiles	Smiles format	smi	http://www.daylight.com/dayhtml/smiles/index.html
chemical/x-embl-dl-nucleotide	EMBL nucleotide format	emb	http://mercury.ebi.ac.uk
chemical/x-galactic-spc	SPC format for spectral and chromatographic data	spc	http://www.galactic.com/galactic/Data/spcvue.htm

Table D-5. “Chemical” MIME types (continued)

MIME type	Description	Extension	Contact and reference
chemical/x-gamess-input	GAMESS Input format	inp, gam	http://www.msg.ameslab.gov/GAMESS/Graphics/MacMolPlt.shtml
chemical/x-gaussian-input	Gaussian Input format	gau	http://www.mdli.com
chemical/x-gaussian-checkpoint	Gaussian Checkpoint format	fch, fchk	http://products.camsoft.com
chemical/x-gaussian-cube	Gaussian Cube (Wavefunction) format	cub	http://www.mdli.com
chemical/x-gcg8-sequence		gcg	
chemical/x-genbank	ToGenBank format	gen	
chemical/x-isostar	IsoStar Library of intermolecular interactions	istr, ist	http://www.ccdc.cam.ac.uk
chemical/x-jcamp-dx	JCAMP Spectroscopic Data Exchange format	jdx, dx	http://www.mdli.com
chemical/x-jjc-review-surface	Re_View3 Orbital Contour files	rv3	http://www.brunel.ac.uk/depts/chem/ch241s/re_view/rv3.htm
chemical/x-jjc-review-xyz	Re_View3 Animation files	xyb	http://www.brunel.ac.uk/depts/chem/ch241s/re_view/rv3.htm
chemical/x-jjc-review-vib	Re_View3 Vibration files	rv2, vib	http://www.brunel.ac.uk/depts/chem/ch241s/re_view/rv3.htm
chemical/x-kinemage	Kinetic (Protein Structure) Images	kin	http://www.faseb.org/protein/kinemages/MageSoftware.html
chemical/x-macmolecule	MacMolecule file format	mcm	
chemical/x-macromodel-input	MacroModel Molecular Mechanics	mmd, mmod	http://www.columbia.edu/cu/chemistry/
chemical/x-mdl-molfile	MDL Molfile	mol	http://www.mdli.com
chemical/x-mdl-rdfile	Reaction data file	rd	http://www.mdli.com
chemical/x-mdl-rxnfile	MDL Reaction format	rxn	http://www.mdli.com
chemical/x-mdl-sdfile	MDL Structure data file	sd	http://www.mdli.com
chemical/x-mdl-tgf	MDL Transportable Graphics Format	tgf	http://www.mdli.com
chemical/x-mif		mif	
chemical/x-mol2	Portable representation of a SYBYL molecule	mol2	http://www.tripos.com
chemical/x-molconn-Z	Molconn-Z format	b	http://www.eslc.vabiotech.com/molconn/molconnz.html
chemical/x-mopac-input	MOPAC Input format	mop	http://www.mdli.com
chemical/x-mopac-graph	MOPAC Graph format	gpt	http://products.camsoft.com
chemical/x-ncbi-asn1		asn (old form)	
chemical/x-ncbi-asn1-binary		val	
chemical/x-pdb	Protein DataBank pdb	pdb	http://www.mdli.com

Table D-5. “Chemical” MIME types (continued)

MIME type	Description	Extension	Contact and reference
chemical/x-swissprot	SWISS-PROT protein sequence database	sw	http://www.expasy.ch/spdbv/text/download.htm
chemical/x-vamas-iso14976	Versailles Agreement on Materials and Standards	vms	http://www.acolyte.co.uk/JISO/
chemical/x-vmd	Visual Molecular Dynamics	vmd	http://www.ks.uiuc.edu/Research/vmd/
chemical/x-xtel	Xtelplot file format	xtel	http://www.recipnet.indiana.edu/graphics/xtelplot/xtelplot.htm
chemical/x-xyz	Co-ordinate Animation format	xyz	http://www.mdli.com

image/*

Table D-6 summarizes some of the image types commonly exchanged by email and HTTP.

Table D-6. “Image” MIME types

MIME type	Description	Extension	Contact and reference
image/bmp	Windows BMP image format.	bmp	
image/cgm	Computer Graphics Metafile (CGM) is an International Standard for the portable storage and transfer of 2-D illustrations.		Alan Francis A.H.Francis@open.ac.uk See ISO 8632:1992, IS 8632:1992 Amendment 1 (1994), and IS 8632:1992 Amendment 2 (1995)
image/g3fax	G3 Facsimile byte streams.		RFC 1494
image/gif	CompuServe GIF images.	gif	RFC 1341
image/ief		ief	RFC 1314
image/jpeg	JPEG images.	jpeg, jpg, jpe, jfif	JPEG Draft Standard ISO 10918-1 CD
image/naplps	North American Presentation Layer Protocol Syntax (NAPLPS) images.		ANSI X3.110-1983 CSA T500-1983
image/png	Portable Network Graphics (PNG) images.	png	Internet draft <i>draft-boutell-png-spec-04.txt</i> , “Png (Portable Network Graphics) Specification Version 1.0”
image/prs.btif	Format used by Nations Bank for BTIF image viewing of checks and other applications.	btif, btif	Arthur Rubin arthurr@crt.com
image/prs.pti	PTI encoded images.	pti	Juern Laun juern.laun@gmx.de http://server.hvzgymn.wn.schule-bw.de/pti/
image/tiff	TIFF images.	tiff, tif	RFC 2302

Table D-6. “Image” MIME types (continued)

MIME type	Description	Extension	Contact and reference
image/vnd.cns.inf2	Supports application features available on the TRILOGUE Infinity network services platform from Comverse Network Systems.		Ann McLaughlin Comverse Network Systems amclaughlin@comversens.com
image/vnd.dxf	DXF vector CAD files.	dxf	
image/vnd.fastbidsheet	A FastBid Sheet contains a raster or vector image that represents an engineering or architectural drawing.	fbs	Scott Becker scottb@bxwa.com
image/vnd.fpx	Kodak FlashPix images.	fpx	Chris Wing format_change_request@kodak.com http://www.kodak.com
image/vnd.fst	Image format from FAST Search and Transfer.	fst	Arild Fuldseth Arild.Fuldseth@fast.no
image/vnd.fujixerox.edmics-mmr	Fuji Xerox EDMICS MMR image format.	mmr	Masanori Onda Masanori.Onda@fujixerox.co.jp
image/vnd.fujixerox.edmics-rlc	Fuji Xerox EDMICS RLC image format.	rlc	Same as above
image/vnd.mix	MIX files contain binary data in streams that are used to represent images and related information. They are used by Microsoft PhotDraw and PictureIt software.		Saveen Reddy2 saveenr@microsoft.com
image/vnd.net-fpx	Kodak FlashPix images.		Chris Wing format_change_request@kodak.com http://www.kodak.com
image/vnd.wap.wbmp	From Apache <i>mime.types</i> .	wbmp	
image/vnd.xiff	Extended Image Format used by Pagis software.	xif	Steve Martin smartin@xis.xerox.com
image/x-cmu-raster	From Apache <i>mime.types</i> .	ras	
image/x-portable-anymap	PBM generic images.	pnm	Jeff Poskanzer http://www.acme.com/software/pbmplus/
image/x-portable-bitmap	PBM bitmap images.	pbm	Same as above
image/x-portable-graymap	PBM grayscale images.	pgm	Same as above
image/x-portable-pixmap	PBM color images.	ppm	Same as above
image/x-rgb	Silicon Graphics’s RGB images.	rgb	
image/x-bitmap	X-Window System bitmap images.	xbm	
image/x-pixmap	X-Window System color images.	xpm	
image/x-windowdump	X-Window System screen capture images.	xwd	

message/*

Messages are composite types used to communicate data objects (through email, HTTP, or other transport protocols). Table D-7 describes the common MIME message types.

Table D-7. “Message” MIME types

MIME type	Description	Extension	Contact and reference
message/delivery-status			
message/disposition-notification			RFC 2298
message/external-body			RFC 1341
message/http			RFC 2616
message/news	Defines a way to transmit news articles via email for human reading—message/rfc822 is not sufficient because news headers have semantics beyond those defined by RFC 822.		RFC 1036
message/partial	Permits the fragmented transmission of bodies that are thought to be too large to be sent directly by email.		RFC 1341
message/rfc822	A complete email message.		RFC 1341
message/s-http	Secure HTTP messages, an alternative to HTTP over SSL.		RFC 2660

model/*

The model MIME type is an IETF-registered extension type. It represents mathematical models of physical worlds, for computer-aided design, and 3-D graphics. Table D-8 describes some of the model formats.

Table D-8. “Model” MIME types

MIME type	Description	Extension	Contact and reference
model/iges	The Initial Graphics Exchange Specification (IGES) defines a neutral data format that allows for the digital exchange of information between computer-aided design (CAD) systems.	igs, iges	RFC 2077
model/mesh		msh, mesh, silo	RFC 2077
model/vnd.dwf	DWF CAD files.	dwf	Jason Pratt jason.pratt@autodesk.com
model/vnd.flatland.3dml	Supports 3DML models supported by Flatland products.	3dml, 3dm	Michael Powers pow@flatland.com http://www.flatland.com

Table D-8. “Model” MIME types (continued)

MIME type	Description	Extension	Contact and reference
model/vnd.gdl model/vnd.gs-gdl	The Geometric Description Language (GDL) is a parametric object definition language for ArchiCAD by Graphisoft.	gdl, gsm, win, dor, lmp, rsm, msm, ism	Attila Babits ababits@graphisoft.hu http://www.graphisoft.com
model/vnd.gtw	Gen-Trix models.	gtw	Yutaka Ozaki yutaka_ozaki@gen.co.jp
model/vnd.mts	MTS model format by Virtue.	mts	Boris Rabinovitch boris@virtue3d.com
model/vnd.parasolid.transmit.binary	Binary Parasolid modeling file.	x_b	http://www.ugsolutions.com/products/parasolid/
model/vnd.parasolid.transmit.text	Text Parasolid modeling file.	x_t	http://www.ugsolutions.com/products/parasolid/
model/vnd.vtu	VTU model format by Virtue.	vtu	Boris Rabinovitch boris@virtue3d.com
model/vrml	Virtual Reality Markup Language format files.	wrl, vrml	RFC 2077

multipart/*

Multipart MIME types are composite objects that contain other objects. The sub-type describes the implementation of the multipart packaging and how to process the components. Multipart media types are summarized in Table D-9.

Table D-9. “Multipart” MIME types

MIME type	Description	Extension	Contact and reference
multipart/alternative	The content consists of a list of alternative representations, each with its own Content-Type. The client can select the best supported component.		RFC 1341
multipart/appledouble	Apple Macintosh files contain “resource forks” and other desktop data that describes the actual file contents. This multipart content sends the Apple metadata in one part and the actual content in another part.		http://www.isi.edu/in-notes/iana/assignments/media-types/multipart/appledouble
multipart/byteranges	When an HTTP message includes the content of multiple ranges, these are transmitted in a “multipart/byteranges” object. This media type includes two or more parts, separated by MIME boundaries, each with its own Content-Type and Content-Range fields.		RFC 2068

Table D-9. “Multipart” MIME types (continued)

MIME type	Description	Extension	Contact and reference
multipart/digest	Contains a collection of individual email messages, in an easy-to-read form.		RFC 1341
multipart/encrypted	Uses two parts to support cryptographically encrypted content. The first part contains the control information necessary to decrypt the data in the second body part and is labeled according to the value of the protocol parameter. The second part contains the encrypted data in type application/octet-stream.		RFC 1847
multipart/form-data	Used to bundle up a set of values as the result of a user filling out a form.		RFC 2388
multipart/header-set	Separates user data from arbitrary descriptive metadata.		http://www.isi.edu/in-notes/iana/assignments/media-types/multipart/header-set
multipart/mixed	A collection of objects.		RFC 1341
multipart/parallel	Syntactically identical to multipart/mixed, but all of the parts are intended to be presented simultaneously, on systems capable of doing so.		RFC 1341
multipart/related	Intended for compound objects consisting of several interrelated body parts. The relationships between the body parts distinguish them from other object types. These relationships often are represented by links internal to the object’s components that reference the other components.		RFC 2387
multipart/report	Defines a general container type for electronic mail reports of any kind.		RFC 1892
multipart/signed	Uses two parts to support cryptographically signed content. The first part is the content, including its MIME headers. The second part contains the information necessary to verify the digital signature.		RFC 1847
multipart/voice-message	Provides a mechanism for packaging a voice message into one container that is tagged as VPIM v2–compliant.		RFCs 2421 and 2423

text/*

Text media types contain characters and potential formatting information. Table D-10 summarizes text MIME types.

Table D-10. “Text” MIME types

MIME type	Description	Extension	Contact and reference
text/calendar	Supports the iCalendar calendaring and scheduling standard.		RFC 2445
text/css	Cascading Style Sheets.	css	RFC 2318
text/directory	Holds record data from a directory database, such as LDAP.		RFC 2425
text/enriched	Simple formatted text, supporting fonts, colors, and spacing. SGML-like tags are used to begin and end formatting.		RFC 1896
text/html	HTML file.	html, htm	RFC 2854
text/parityfec	Forward error correction for text streamed in an RTP stream.		RFC 3009
text/plain	Plain old text.	asc, txt	
text/prs.lines.tag	Supports tagged forms, as used for email registration.	tag, dsc	John Lines john@paladin.demon.co.uk http://www.paladin.demon.co.uk/tag-types/
text/rfc822-headers	Used to bundle a set of email headers, such as when sending mail failure reports.		RFC 1892
text/richtext	Older form of enriched text. See text/enriched.	rtx	RFC 1341
text/rtf	The Rich Text Format (RTF) is a method of encoding formatted text and graphics for transfer between applications. The format is widely supported by word-processing applications on the MS-DOS, Windows, OS/2, and Macintosh platforms.	rtf	
text/sgml	SGML markup files.	sgml, sgm	RFC 1874
text/t140	Supports standardized T.140 text, as used in synchronized RTP multimedia.		RFC 2793
text/tab-separated-values	TSV is a popular method of data interchange among databases and spreadsheets and word processors. It consists of a set of lines, with fields separated by tab characters.	tsv	http://www.isi.edu/in-notes/iana/assignments/media-types/text/tab-separated-values
text/uri-list	Simple, commented lists of URLs and URNs used by URN resolvers, and any other applications that need to communicate bulk URI lists.	uris, uri	RFC 2483

Table D-10. “Text” MIME types (continued)

MIME type	Description	Extension	Contact and reference
text/vnd.abc	ABC files are a human-readable format for musical scores.	abc	http://www.gre.ac.uk/~c.walshaw/abc/ http://home1.swipnet.se/~w-11382/abcbnf.htm
text/vnd.curl	Provides a set of content definition languages interpreted by the CURL runtime plug-in.	curl	Tim Hodge thodge@curl.com
text/vnd.DMClientScript	CommonDM Client Script files are used as hyperlinks to non-http sites (such as BYOND, IRC, or telnet) accessed by the Dream Seeker client application.	dms	Dan Bradley dan@dantom.com http://www.byond.com/code/ref/
text/vnd.fly	Fly is a text preprocessor that uses a simple syntax to create an interface between databases and web pages.	fly	John-Mark Gurney jmg@flyidea.com http://www.flyidea.com
text/vnd.fmi.flexstor	For use in the SUVDAMA and UVRAPPF projects.	flx	http://www.ozone.fmi.fi/SUVDAMA/ http://www.ozone.fmi.fi/UVRAPPF/
text/vnd.in3d.3dml	For In3D Player.	3dml, 3dm	Michael Powers powers@insideout.net
text/vnd.in3d.spot	For In3D Player.	spot, spo	Same as above
text/vnd.IPTC.NewsML	NewsML format specified by the International Press Telecommunications Council (IPTC).	xml	David Allen m_director_iptc@dial.pipex.com http://www.iptc.org
text/vnd.IPTC.NITF	NITF format specified by the IPTC.	xml	Same as above http://www.nitf.org
text/vnd.latex-z	Supports LaTeX documents containing Z notation. Z notation (pronounced “zed”), is based on Zermelo-Fraenkel set theory and first order predicate logic, and it is useful for describing computer systems.		http://www.comlab.ox.ac.uk/archive/z/
text/vnd.motorola.reflex	Provides a common method for submitting simple text messages from ReFLEX™ wireless devices.		Mark Patton fmp014@email.mot.com Part of the FLEXsuite™ of Enabling Protocols specification available from Motorola under the licensing agreement
text/vnd.ms-mediapackage	This type is intended to be handled by the Microsoft application programs MStore.exe and 7 storDB.exe.	mpf	Jan Nelson jann@microsoft.com

Table D-10. “Text” MIME types (continued)

MIME type	Description	Extension	Contact and reference
text/vnd.wap.si	Service Indication (SI) objects contain a message describing an event and a URI describing where to load the corresponding service.	si, xml	WAP Forum Ltd http://www.wapforum.org
text/vnd.wap.sl	The Service Loading (SL) content type provides a means to convey a URI to a user agent in a mobile client. The client itself automatically loads the content indicated by that URI and executes it in the addressed user agent without user intervention when appropriate.	sl, xml	Same as above
text/vnd.wap.wml	Wireless Markup Language (WML) is a markup language, based on XML, that defines content and user interface for narrow-band devices, including cellular phones and pagers.	wml	Same as above
text/vnd.wap.wmlscript	WMLScript is an evolution of JavaScript for wireless devices.	wmls	Same as above
text/x-setext	From Apache <i>mime.types</i> .	etx	
text/xml	Extensible Markup Language format file (use application/xml if you want the browser to save to file when downloaded).	xml	RFC 2376

video/*

Table D-11 lists some popular video movie formats. Note that some video formats are classified as application types.

Table D-11. “Video” MIME types

MIME type	Description	Extension	Contact and reference
video/MP4V-ES	MPEG-4 video payload, as carried by RTP.		RFC 3016
video/mpeg	Video encoded per the ISO 11172 CD MPEG standard.	mpeg, mpg, mpe	RFC 1341
video/parityfec	Forward error correcting video format for data carried through RTP streams.		RFC 3009
video/pointer	Transporting pointer position information for presentations.		RFC 2862
video/quicktime	Apple Quicktime video format.	qt, mov	http://www.apple.com
video/vnd.fvt	Video format from FAST Search & Transfer.	fvt	Arild Fuldseth Arild.Fuldseth@fast.no

Table D-11. “Video” MIME types (continued)

MIME type	Description	Extension	Contact and reference
video/vnd.motorola.video video/vnd.motorola.videop	Proprietary formats used by products from Motorola ISG.		Tom McGinty Motorola ISG tmcginty@dma.isg.mot
video/vnd.mpegurl	This media type consists of a series of URLs of MPEG Video files.	mxu	Heiko Recktenwald uzs106@uni-bonn.de “Power and Responsibility: Conversations with Contributors,” Guy van Belle, et al., LMJ 9 (1999), 127–133, 129 (MIT Press)
video/vnd.nokia.interleaved-multimedia	Used in Nokia 9210 Communicator video player and related tools.	nim	Petteri Kangaslampi petteri.kangaslampi@nokia.com
video/x-msvideo	Microsoft AVI movies.	avi	http://www.microsoft.com
video/x-sgi-movie	Silicon Graphics’s movie format.	movie	http://www.sgi.com

Experimental Types

The set of primary types supports most content types. Table D-12 lists one experimental type, for conferencing software, that is configured in some web servers.

Table D-12. Extension MIME types

MIME type	Description	Extension	Contact and reference
x-conference/x-cooltalk	Collaboration tool from Netscape	ice	

Base-64 Encoding

Base-64 encoding is used by HTTP, for basic and digest authentication, and by several HTTP extensions. This appendix explains base-64 encoding and provides conversion tables and pointers to Perl software to help you correctly use base-64 encoding in HTTP software.

Base-64 Encoding Makes Binary Data Safe

The base-64 encoding converts a series of arbitrary bytes into a longer sequence of common text characters that are all legal header field values. Base-64 encoding lets us take user input or binary data, pack it into a safe format, and ship it as HTTP header field values without fear of them containing colons, newlines, or binary values that would break HTTP parsers.

Base-64 encoding was developed as part of the MIME multimedia electronic mail standard, so MIME could transport rich text and arbitrary binary data between different legacy email gateways.* Base-64 encoding is similar in spirit, but more efficient in space, to the uuencode and BinHex standards for textifying binary data. Section 6.8 of MIME RFC 2045 details the base-64 algorithm.

Eight Bits to Six Bits

Base-64 encoding takes a sequence of 8-bit bytes, breaks the sequence into 6-bit pieces, and assigns each 6-bit piece to one of 64 characters comprising the base-64 alphabet. The 64 possible output characters are common and safe to place in HTTP header fields. The 64 characters include upper- and lowercase letters, numbers, +,

* Some mail gateways would silently strip many “non-printing” characters with ASCII values between 0 and 31. Other programs would interpret some bytes as flow control characters or other special control characters, or convert carriage returns to line feeds and the like. Some programs would experience fatal errors upon receiving international characters with a value above 127 because the software was not “8-bit clean.”

and /. The special character = also is used. The base-64 alphabet is shown in Table E-1.

Note that because the base-64 encoding uses 8-bit characters to represent 6 bits of information, base 64–encoded strings are about 33% larger than the original values.

Table E-1. Base-64 alphabet

0	A	8	I	16	Q	24	Y	32	g	40	o	48	w	56	4
1	B	9	J	17	R	25	Z	33	h	41	p	49	x	57	5
2	C	10	K	18	S	26	a	34	i	42	q	50	y	58	6
3	D	11	L	19	T	27	b	35	j	43	r	51	z	59	7
4	E	12	M	20	U	28	c	36	k	44	s	52	0	60	8
5	F	13	N	21	V	29	d	37	l	45	t	53	1	61	9
6	G	14	O	22	W	30	e	38	m	46	u	54	2	62	+
7	H	15	P	23	X	31	f	39	n	47	v	55	3	63	/

Figure E-1 shows a simple example of base-64 encoding. Here, the three-character input value “Ow!” is base 64–encoded, resulting in the four-character base 64–encoded value “T3ch”. It works like this:

1. The string “Ow!” is broken into 3 8-bit bytes (0x4F, 0x77, 0x21).
2. The 3 bytes create the 24-bit binary value 010011110111011100100001.
3. These bits are segmented into the 6-bit sequences 010011, 110111, 01110, 100001.
4. Each of these 6-bit values represents a number from 0 to 63, corresponding to one of 64 characters in the base-64 alphabet. The resulting base 64–encoded string is the 4-character string “T3ch”, which can then be sent across the wire as “safe” 8-bit characters, because only the most portable characters are used (letters, numbers, etc.).




Figure E-1. Base-64 encoding example

Base-64 Padding

Base-64 encoding takes a sequence of 8-bit bytes and segments the bit stream into 6-bit chunks. It is unlikely that the sequence of bits will divide evenly into 6-bit pieces. When the bit sequence does not divide evenly into 6-bit pieces, the bit sequence is padded with zero bits at the end to make the length of the bit sequence a multiple of 24 (the least common multiple of 6 and 8 bits).

When encoding the padded bit string, any group of 6 bits that is completely padding (containing no bits from the original data) is represented by a special 65th symbol: “=”. If a group of 6 bits is partially padded, the padding bits are set to zero.

Table E-2 shows examples of padding. The initial input string “a:a” is 3 bytes long, or 24 bits. 24 is a multiple of 6 and 8, so no padding is required. The resulting base 64-encoded string is “YTph”.

Table E-2. Base-64 padding examples

Input data	Binary sequence (padding noted as “x”)	Encoded data
a:a	011000 010011 101001 100001	YTph
a:aa	011000 010011 101001 100001 011000 01xxxx xxxxxx xxxxxx	YTphYQ==
a:aaa	011000 010011 101001 100001 011000 010110 0001xx xxxxxx	YTphYWE=
a:aaaa	011000 010011 101001 100001 011000 010110 000101 100001	YTphYWZh

However, when another character is added, the input string grows to 32 bits long. The next smallest multiple of 6 and 8 is 48 bits, so 16 bits of padding are added. The first 4 bits of padding are mixed with data bits. The resulting 6-bit group, 01xxxx, is treated as 010000, 16 decimal, or base-64 encoding Q. The remaining two 6-bit groups are all padding and are represented by “=”.

Perl Implementation

MIME::Base64 is a Perl module for base-64 encoding and decoding. You can read about this module at <http://www.perldoc.com/perl5.6.1/lib/MIME/Base64.html>.

You can encode and decode strings using the MIME::Base64 `encode_base64` and `decode_base64` methods:

```
use MIME::Base64;

$encoded = encode_base64('Aladdin:open sesame');
$decoded = decode_base64($encoded);
```

For More Information

For more information on base-64 encoding, see:

<http://www.ietf.org/rfc/rfc2045.txt>

Section 6.8 of RFC 2045, “MIME Part 1: Format of Internet Message Bodies,” provides an official specification of base-64 encoding.

<http://www.perldoc.com/perl5.6.1/lib/MIME/Base64.html>

This web site contains documentation for the MIME::Base64 Perl module that provides encoding and decoding of base-64 strings.

Digest Authentication

This appendix contains supporting data and source code for implementing HTTP digest authentication facilities.

Digest WWW-Authenticate Directives

WWW-Authenticate directives are described in Table F-1, paraphrased from the descriptions in RFC 2617. As always, refer to the official specifications for the most up-to-date details.

Table F-1. Digest WWW-Authenticate header directives (from RFC 2617)

Directive	Description
realm	A string to be displayed to users so they know which username and password to use. This string should contain at least the name of the host performing the authentication and might additionally indicate the collection of users who might have access. An example might be "registered_users@gotham.news.com".
nonce	A server-specified data string that should be uniquely generated each time a 401 response is made. It is recommended that this string be base-64 or hexadecimal data. Specifically, because the string is passed in the header lines as a quoted string, the double-quote character is not allowed. The contents of the nonce are implementation-dependent. The quality of the implementation depends on a good choice. A nonce might, for example, be constructed as the base-64 encoding of: <pre>time-stamp H(time-stamp ":" ETag ":" private-key)</pre> where <i>time-stamp</i> is a server-generated time or other nonrepeating value, <i>ETag</i> is the value of the HTTP ETag header associated with the requested entity, and <i>private-key</i> is data known only to the server. With a nonce of this form, a server would recalculate the hash portion after receiving the client Authentication header and reject the request if it did not match the nonce from that header or if the time-stamp value is not recent enough. In this way, the server can limit the time of the nonce's validity. The inclusion of the ETag prevents a replay request for an updated version of the resource. (Note: including the IP address of the client in the nonce appears to offer the server the ability to limit the reuse of the nonce to the same client that originally got it. However, that would break proxy farms, where requests from a single user often go through different proxies in the farm. Also, IP address spoofing is not that hard.) An implementation might choose not to accept a previously used nonce or a previously used digest, to protect against replay attacks, or it might choose to use one-time nonces or digests for POST or PUT requests and time-stamps for GET requests.

Table F-1. Digest WWW-Authenticate header directives (from RFC 2617) (continued)

Directive	Description
domain	A quoted, space-separated list of URIs (as specified in RFC 2396, "Uniform Resource Identifiers: Generic Syntax") that define the protection space. If a URI is an <code>abs_path</code>, it is relative to the canonical root URL of the server being accessed. An absolute URI in this list may refer to a different server than the one being accessed. The client can use this list to determine the set of URIs for which the same authentication information may be sent: any URI that has a URI in this list as a prefix (after both have been made absolute) may be assumed to be in the same protection space. If this directive is omitted or its value is empty, the client should assume that the protection space consists of all URIs on the responding server. This directive is not meaningful in Proxy-Authenticate headers, for which the protection space is always the entire proxy; if present, it should be ignored.
opaque	A string of data, specified by the server, that should be returned by the client unchanged in the Authorization header of subsequent requests with URIs in the same protection space. It is recommended that this string be base-64 or hexadecimal data.
stale	A flag indicating that the previous request from the client was rejected because the nonce value was stale. If stale is TRUE (case-insensitive), the client may want to retry the request with a new encrypted response, without reprompting the user for a new username and password. The server should set stale to TRUE only if it receives a request for which the nonce is invalid but has a valid digest (indicating that the client knows the correct username/password). If stale is FALSE, or anything other than TRUE, or the stale directive is not present, the username and/or password are invalid, and new values must be obtained.
algorithm	A string indicating a pair of algorithms used to produce the digest and a checksum. If this is not present, it is assumed to be "MD5". If the algorithm is not understood, the challenge should be ignored (and a different one used, if there is more than one). In this document, the string obtained by applying the digest algorithm to the data "data" with secret "secret" will be denoted by "KD(secret, data)", and the string obtained by applying the checksum algorithm to the data "data" will be denoted "H(data)". The notation "unq(X)" means the value of the quoted string "X" without the surrounding quotes. For the MD5 and MD5-sess algorithms: $H(data) = MD5(data)$ $HD(secret, data) = H(concat(secret, ":", data))$ I.e., the digest is the MD5 of the secret concatenated with a colon concatenated with the data. The MD5-sess algorithm is intended to allow efficient third-party authentication servers.
qop	This directive is optional but is made so only for backward compatibility with RFC 2069 [6]; it should be used by all implementations compliant with this version of the digest scheme. If present, it is a quoted string of one or more tokens indicating the "quality of protection" values supported by the server. The value "auth" indicates authentication; the value "auth-int" indicates authentication with integrity protection. Unrecognized options must be ignored.
<extension>	This directive allows for future extensions. Any unrecognized directives must be ignored.

Digest Authorization Directives

Each of the Authorization directives is described in Table F-2, paraphrased from the descriptions in RFC 2617. Refer to the official specifications for the most up-to-date details.

Table F-2. Digest Authorization header directives (from RFC 2617)

Directive	Description
username	The user's name in the specified realm.
realm	The realm passed to the client in the WWW-Authenticate header.
nonce	The same nonce passed to the client in the WWW-Authenticate header.
uri	The URI from the request URI of the request line; duplicated because proxies are allowed to change the request line in transit, and we may need the original URI for proper digest verification calculations.
response	This is the actual digest—the whole point of digest authentication! The response is a string of 32 hexadecimal digits, computed by a negotiated digest algorithm, which proves that the user knows the password.
algorithm	A string indicating a pair of algorithms used to produce the digest and a checksum. If this is not present, it is assumed to be "MD5".
opaque	A string of data, specified by the server in a WWW-Authenticate header, that should be returned by the client unchanged in the Authorization header of subsequent requests with URIs in the same protection space.
cnonce	This must be specified if a qop directive is sent and must not be specified if the server did not send a qop directive in the WWW-Authenticate header field. The cnonce value is an opaque quoted string value provided by the client and used by both client and server to avoid chosen plaintext attacks, to provide mutual authentication, and to provide some message-integrity protection. See the descriptions of the response-digest and request-digest calculations later in this appendix.
qop	Indicates what "quality of protection" the client has applied to the message. If present, its value must be one of the alternatives the server indicated it supports in the WWW-Authenticate header. These values affect the computation of the request digest. This is a single token, not a quoted list of alternatives, as in WWW-Authenticate. This directive is optional, to preserve backward compatibility with a minimal implementation of RFC 2069, but it should be used if the server indicated that qop is supported by providing a qop directive in the WWW-Authenticate header field.
nc	This must be specified if a qop directive is sent and must not be specified if the server did not send a qop directive in the WWW-Authenticate header field. The value is the hexadecimal count of the number of requests (including the current request) that the client has sent with the nonce value in this request. For example, in the first request sent in response to a given nonce value, the client sends nc="00000001". The purpose of this directive is to allow the server to detect request replays by maintaining its own copy of this count—if the same nc value is seen twice, the request is a replay.
<extension>	This directive allows for future extensions. Any unrecognized directive must be ignored.

Digest Authentication-Info Directives

Each of the Authentication-Info directives is described in Table F-3, paraphrased from the descriptions in RFC 2617. Refer to the official specifications for the most up-to-date details.

Table F-3. Digest Authentication-Info header directives (from RFC 2617)

Directive	Description
nextnonce	The value of the nextnonce directive is the nonce the server wants the client to use for a future authentication response. The server may send the Authentication-Info header with a nextnonce field as a means of implementing one-time or otherwise changing nonces. If the nextnonce field is present the client should use it when constructing the Authorization header for its next request. Failure of the client to do so may result in a reauthentication request from the server with "stale=TRUE". Server implementations should carefully consider the performance implications of the use of this mechanism; pipelined requests will not be possible if every response includes a nextnonce directive that must be used on the next request received by the server. Consideration should be given to the performance versus security trade-offs of allowing an old nonce value to be used for a limited time to permit request pipelining. Use of the nonce count can retain most of the security advantages of a new server nonce without the deleterious effects on pipelining.
qop	Indicates the "quality of protection" options applied to the response by the server. The value "auth" indicates authentication; the value "auth-int" indicates authentication with integrity protection. The server should use the same value for the qop directive in the response as was sent by the client in the corresponding request.
rspauth	The optional response digest in the "response auth" directive supports mutual authentication—the server proves that it knows the user's secret, and, with qop="auth-int", it also provides limited integrity protection of the response. The "response-digest" value is calculated as for the "request-digest" in the Authorization header, except that if qop="auth" or qop is not specified in the Authorization header for the request, A2 is: A2 = ":" digest-uri-value and if qop="auth-int", A2 is: A2 = ":" digest-uri-value ":" H(entity-body) where <i>digest-uri-value</i> is the value of the uri directive on the Authorization header in the request. The cnonce and nc values must be the same as the ones in the client request to which this message is a response. The rspauth directive must be present if qop="auth" or qop="auth-int" is specified.
cnonce	The cnonce value must be the same as the one in the client request to which this message is a response. The cnonce directive must be present if qop="auth" or qop="auth-int" is specified.
nc	The nc value must be the same as the one in the client request to which this message is a response. The nc directive must be present if qop="auth" or qop="auth-int" is specified.
<extension>	This directive allows for future extensions. Any unrecognized directive must be ignored.

Reference Code

The following code implements the calculations of H(A1), H(A2), request-digest, and response-digest, from RFC 2617. It uses the MD5 implementation from RFC 1321.

File "digcalc.h"

```
#define HASHLEN 16
typedef char HASH[HASHLEN];
#define HASHHEXLEN 32
typedef char HASHHEX[HASHHEXLEN+1];
#define IN
```

```

#define OUT
/* calculate H(A1) as per HTTP Digest spec */
void DigestCalcHA1(
    IN char * pszAlg,
    IN char * pszUserName,
    IN char * pszRealm,
    IN char * pszPassword,
    IN char * pszNonce,
    IN char * pszCNonce,
    OUT HASHHEX SessionKey
);

/* calculate request-digest/response-digest as per HTTP Digest spec */
void DigestCalcResponse(
    IN HASHHEX HA1,          /* H(A1) */
    IN char * pszNonce,      /* nonce from server */
    IN char * pszNonceCount, /* 8 hex digits */
    IN char * pszCNonce,     /* client nonce */
    IN char * pszQop,        /* qop-value: "", "auth", "auth-int" */
    IN char * pszMethod,     /* method from the request */
    IN char * pszDigestUri,  /* requested URL */
    IN HASHHEX HEntity,     /* H(entity body) if qop="auth-int" */
    OUT HASHHEX Response    /* request-digest or response-digest */
);

```

File “digcalc.c”

```

#include <global.h>
#include <md5.h>
#include <string.h>
#include "digcalc.h"

void CvtHex(
    IN HASH Bin,
    OUT HASHHEX Hex
)
{
    unsigned short i;
    unsigned char j;
    for (i = 0; i < HASHLEN; i++) {
        j = (Bin[i] >> 4) & 0xf;
        if (j <= 9)
            Hex[i*2] = (j + '0');
        else
            Hex[i*2] = (j + 'a' - 10);
        j = Bin[i] & 0xf;
        if (j <= 9)
            Hex[i*2+1] = (j + '0');
        else
            Hex[i*2+1] = (j + 'a' - 10);
    };
    Hex[HASHHEXLEN] = '\0';
};

```

```

/* calculate H(A1) as per spec */
void DigestCalcHA1(
    IN char * pszAlg,
    IN char * pszUserName,
    IN char * pszRealm,
    IN char * pszPassword,
    IN char * pszNonce,
    IN char * pszCNonce,
    OUT HASHHEX SessionKey
)
{
    MD5_CTX Md5Ctx;
    HASH HA1;
    MD5Init(&Md5Ctx);
    MD5Update(&Md5Ctx, pszUserName, strlen(pszUserName));
    MD5Update(&Md5Ctx, ":", 1);
    MD5Update(&Md5Ctx, pszRealm, strlen(pszRealm));
    MD5Update(&Md5Ctx, ":", 1);
    MD5Update(&Md5Ctx, pszPassword, strlen(pszPassword));
    MD5Final(HA1, &Md5Ctx);
    if (strcmp(pszAlg, "md5-sess") == 0) {
        MD5Init(&Md5Ctx);
        MD5Update(&Md5Ctx, HA1, HASHLEN);
        MD5Update(&Md5Ctx, ":", 1);
        MD5Update(&Md5Ctx, pszNonce, strlen(pszNonce));
        MD5Update(&Md5Ctx, ":", 1);
        MD5Update(&Md5Ctx, pszCNonce, strlen(pszCNonce));
        MD5Final(HA1, &Md5Ctx);
    };
    CvtHex(HA1, SessionKey);
};

/* calculate request-digest/response-digest as per HTTP Digest spec */
void DigestCalcResponse(
    IN HASHHEX HA1,           /* H(A1) */
    IN char * pszNonce,       /* nonce from server */
    IN char * pszNonceCount,  /* 8 hex digits */
    IN char * pszCNonce,      /* client nonce */
    IN char * pszQop,         /* qop-value: "", "auth", "auth-int" */
    IN char * pszMethod,      /* method from the request */
    IN char * pszDigestUri,   /* requested URL */
    IN HASHHEX HEntity,       /* H(entity body) if qop="auth-int" */
    OUT HASHHEX Response     /* request-digest or response-digest */
)
{
    MD5_CTX Md5Ctx;
    HASH HA2;
    HASH RespHash;
    HASHHEX HA2Hex;
    // calculate H(A2)
    MD5Init(&Md5Ctx);
    MD5Update(&Md5Ctx, pszMethod, strlen(pszMethod));
    MD5Update(&Md5Ctx, ":", 1);
    MD5Update(&Md5Ctx, pszDigestUri, strlen(pszDigestUri));
    if (strcmp(pszQop, "auth-int") == 0) {

```

```

        MD5Update(&Md5Ctx, ":", 1);
        MD5Update(&Md5Ctx, HEntity, HASHHEXLEN);
    };
    MD5Final(HA2, &Md5Ctx);
    CvtHex(HA2, HA2Hex);
    // calculate response
    MD5Init(&Md5Ctx);
    MD5Update(&Md5Ctx, HA1, HASHHEXLEN);
    MD5Update(&Md5Ctx, ":", 1);
    MD5Update(&Md5Ctx, pszNonce, strlen(pszNonce));
    MD5Update(&Md5Ctx, ":", 1);
    if (*pszQop) {
        MD5Update(&Md5Ctx, pszNonceCount, strlen(pszNonceCount));
        MD5Update(&Md5Ctx, ":", 1);
        MD5Update(&Md5Ctx, pszCNonce, strlen(pszCNonce));
        MD5Update(&Md5Ctx, ":", 1);
        MD5Update(&Md5Ctx, pszQop, strlen(pszQop));
        MD5Update(&Md5Ctx, ":", 1);
    };
    MD5Update(&Md5Ctx, HA2Hex, HASHHEXLEN);
    MD5Final(RespHash, &Md5Ctx);
    CvtHex(RespHash, Response);
};

```

File “digtest.c”

```

#include <stdio.h>
#include "digcalc.h"

void main(int argc, char ** argv) {
    char * pszNonce = "dcd98b7102dd2f0e8b11d0f600bfb0c093";
    char * pszCNonce = "0a4f113b";
    char * pszUser = "Mufasa";
    char * pszRealm = "testrealm@host.com";
    char * pszPass = "Circle Of Life";
    char * pszAlg = "md5";
    char szNonceCount[9] = "00000001";
    char * pszMethod = "GET";
    char * pszQop = "auth";
    char * pszURI = "/dir/index.html";
    HASHHEX HA1;
    HASHHEX HA2 = "";
    HASHHEX Response;
    DigestCalcHA1(pszAlg, pszUser, pszRealm, pszPass,
        pszNonce, pszCNonce, HA1);
    DigestCalcResponse(HA1, pszNonce, szNonceCount, pszCNonce, pszQop,
        pszMethod, pszURI, HA2, Response);
    printf("Response = %s\n", Response);
};

```

Language Tags

Language tags are short, standardized strings that name spoken languages—for example, “fr” (French) and “en-GB” (Great Britain English). Each tag has one or more parts, separated by hyphens, called *subtags*. Language tags were described in detail in the section “Language Tags and HTTP” in Chapter 16.

This appendix summarizes the rules, standardized tags, and registration information for language tags. It contains the following reference material:

- Rules for the first (primary) subtag are summarized in “First Subtag Rules.”
- Rules for the second subtag are summarized in “Second Subtag Rules.”
- IANA-registered language tags are shown in Table G-1.
- ISO 639 language codes are shown in Table G-2.
- ISO 3166 country codes are shown in Table G-3.

First Subtag Rules

If the first subtag is:

- Two characters long, it’s a language code from the ISO 639* and 639-1 standards
- Three characters long, it’s a language code listed in the ISO 639-2† standard
- The letter “i,” the language tag is explicitly IANA-registered
- The letter “x,” the language tag is a private, nonstandard, extension subtag

The ISO 639 and 639-2 names are summarized in Table G-2.

* See ISO standard 639, “Codes for the representation of names of languages.”

† See ISO 639-2, “Codes for the representation of names of languages—Part 2: Alpha-3 code.”

Second Subtag Rules

If the second subtag is:

- Two characters long, it's a country/region defined by ISO 3166*
- Three to eight characters long, it may be registered with the IANA
- One character long, it is illegal

The ISO 3166 country codes are summarized in Table G-3.

IANA-Registered Language Tags

Table G-1. Language tags

IANA language tag	Description
i-bnn	Bunun
i-default	Default language context
i-hak	Hakka
i-klingon	Klingon
i-lux	Luxembourgish
i-mingo	Mingo
i-navajo	Navajo
i-pwn	Paiwan
i-tao	Tao
i-tay	Tayal
i-tsu	Tsou
no-bok	Norwegian "Book language"
no-nyn	Norwegian "New Norwegian"
zh-gan	Kan or Gan
zh-guoyu	Mandarin or Standard Chinese
zh-hakka	Hakka
zh-min	Min, Fuzhou, Hokkien, Amoy, or Taiwanese
zh-wuu	Shanghaiese or Wu
zh-xiang	Xiang or Hunanese
zh-yue	Cantonese

* The country codes AA, QM–QZ, XA–XZ and ZZ are reserved by ISO 3166 as user-assigned codes. These must not be used to form language tags.

ISO 639 Language Codes

Table G-2. ISO 639 and 639-2 language codes

Language	ISO 639	ISO 639-2
Abkhazian	ab	abk
Achinese		ace
Acoli		ach
Adangme		ada
Afar	aa	aar
Afrihili		afh
Afrikaans	af	afr
Afro-Asiatic (Other)		afa
Akan		aka
Akkadian		akk
Albanian	sq	alb/sqi
Aleut		ale
Algonquian languages		alg
Altaic (Other)		tut
Amharic	am	amh
Apache languages		apa
Arabic	ar	ara
Aramaic		arc
Arapaho		arp
Araucanian		arn
Arawak		arw
Armenian	hy	arm/hye
Artificial (Other)		art
Assamese	as	asm
Athapaskan languages		ath
Austronesian (Other)		map
Avaric		ava
Avestan		ave
Awadhi		awa
Aymara	ay	aym
Azerbaijani	az	aze
Aztec		nah
Balinese		ban
Baltic (Other)		bat

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Baluchi		bal
Bambara		bam
Bamileke languages		bai
Banda		bad
Bantu (Other)		bnt
Basa		bas
Bashkir	ba	bak
Basque	eu	baq/eus
Beja		bej
Bemba		bem
Bengali	bn	ben
Berber (Other)		ber
Bhojpuri		bho
Bihari	bh	bih
Bikol		bik
Bini		bin
Bislama	bi	bis
Braj		bra
Breton	be	bre
Buginese		bug
Bulgarian	bg	bul
Buriat		bua
Burmese	my	bur/mya
Byelorussian	be	bel
Caddo		cad
Carib		car
Catalan	ca	cat
Caucasian (Other)		cau
Cebuano		ceb
Celtic (Other)		cel
Central American Indian (Other)		cai
Chagatai		chg
Chamorro		cha
Chechen		che
Cherokee		chr
Cheyenne		chy

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Chibcha		chb
Chinese	zh	chi/zho
Chinook jargon		chn
Choctaw		cho
Church Slavic		chu
Chuvash		chv
Coptic		cop
Cornish		cor
Corsican	co	cos
Cree		cre
Creek		mus
Creoles and Pidgins (Other)		crp
Creoles and Pidgins, English-based (Other)		cpe
Creoles and Pidgins, French-based (Other)		cpf
Creoles and Pidgins, Portuguese-based (Other)		cpp
Cushitic (Other)		cus
Croatian	hr	
Czech	cs	ces/cze
Dakota		dak
Danish	da	dan
Delaware		del
Dinka		din
Divehi		div
Dogri		doi
Dravidian (Other)		dra
Duala		dua
Dutch	nl	dut/nla
Dutch, Middle (ca. 1050-1350)		dum
Dyula		dyu
Dzongkha	dz	dzo
Efik		efi
Egyptian (Ancient)		egy
Ekajuk		eka
Elamite		elx
English	en	eng
English, Middle (ca. 1100-1500)		enm

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
English, Old (ca. 450-1100)		ang
Eskimo (Other)		esk
Esperanto	eo	epo
Estonian	et	est
Ewe		ewe
Ewondo		ewo
Fang		fan
Fanti		fat
Faroese	fo	fao
Fijian	fj	fij
Finnish	fi	fin
Finno-Ugrian (Other)		fiu
Fon		fon
French	fr	fra/fre
French, Middle (ca. 1400-1600)		frm
French, Old (842- ca. 1400)		fro
Frisian	fy	fry
Fulah		ful
Ga		gaa
Gaelic (Scots)		gae/gdh
Gallegan	gl	glg
Ganda		lug
Gayo		gay
Geez		gez
Georgian	ka	geo/kat
German	de	deu/ger
German, Middle High (ca. 1050-1500)		gmh
German, Old High (ca. 750-1050)		goh
Germanic (Other)		gem
Gilbertese		gil
Gondi		gon
Gothic		got
Grebo		grb
Greek, Ancient (to 1453)		grc
Greek, Modern (1453-)	el	ell/gre
Greenlandic	kl	kal

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Guarani	gn	grn
Gujarati	gu	guj
Haida		hai
Hausa	ha	hau
Hawaiian		haw
Hebrew	he	heb
Herero		her
Hiligaynon		hil
Himachali		him
Hindi	hi	hin
Hiri Motu		hmo
Hungarian	hu	hun
Hupa		hup
Iban		iba
Icelandic	is	ice/isl
Igbo		ibo
Ijo		ijo
Iloko		ilo
Indic (Other)		inc
Indo-European (Other)		ine
Indonesian	id	ind
Interlingua (IALA)	ia	ina
Interlingue	ie	ine
Inuktitut	iu	iku
Inupiak	ik	ipk
Iranian (Other)		ira
Irish	ga	gai/iri
Irish, Old (to 900)		sga
Irish, Middle (900 - 1200)		mga
Iroquoian languages		iro
Italian	it	ita
Japanese	ja	jpn
Javanese	jav/jw	jav/jaw
Judeo-Arabic		jrb
Judeo-Persian		jpr
Kabyle		kab

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Kachin		kac
Kamba		kam
Kannada	kn	kan
Kanuri		kau
Kara-Kalpak		kaa
Karen		kar
Kashmiri	ks	kas
Kawi		kaw
Kazakh	kk	kaz
Khasi		kha
Khmer	km	khm
Khoisan (Other)		khi
Khotanese		kho
Kikuyu		kik
Kinyarwanda	rw	kin
Kirghiz	ky	kir
Komi		kom
Kongo		kon
Konkani		kok
Korean	ko	kor
Kpelle		kpe
Kru		kro
Kuanyama		kua
Kumyk		kum
Kurdish	ku	kur
Kurukh		kru
Kusaie		kus
Kutenai		kut
Ladino		lad
Lahnda		lah
Lamba		lam
Langue d'Oc (post-1500)	oc	oci
Lao	lo	lao
Latin	la	lat
Latvian	lv	lav
Letzeburgesch		ltz

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Lezghian		lez
Lingala	ln	lin
Lithuanian	lt	lit
Lozi		loz
Luba-Katanga		lub
Luiseno		lui
Lunda		lun
Luo (Kenya and Tanzania)		luo
Macedonian	mk	mac/mak
Madurese		mad
Magahi		mag
Maithili		mai
Makasar		mak
Malagasy	mg	mlg
Malay	ms	may/msa
Malayalam		mal
Maltese	ml	mlt
Mandingo		man
Manipuri		mni
Manobo languages		mno
Manx		max
Maori	mi	mao/mri
Marathi	mr	mar
Mari		chm
Marshall		mah
Marwari		mwr
Masai		mas
Mayan languages		myn
Mende		men
Micmac		mic
Minangkabau		min
Miscellaneous (Other)		mis
Mohawk		moh
Moldavian	mo	mol
Mon-Kmer (Other)		mkh
Mongo		lol

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Mongolian	mn	mon
Mossi		mos
Multiple languages		mul
Munda languages		mun
Nauru	na	nau
Navajo		nav
Ndebele, North		nde
Ndebele, South		nbl
Ndongo		ndo
Nepali	ne	nep
Newari		new
Niger-Kordofanian (Other)		nic
Nilo-Saharan (Other)		ssa
Niuean		niu
Norse, Old		non
North American Indian (Other)		nai
Norwegian	no	nor
Norwegian (Nynorsk)		nno
Nubian languages		nub
Nyamwezi		nym
Nyanja		nya
Nyankole		nyn
Nyoro		nyo
Nzima		nzi
Ojibwa		oji
Oriya	or	ori
Oromo	om	orm
Osage		osa
Ossetic		oss
Otomian languages		oto
Pahlavi		pal
Palauan		pau
Pali		pli
Pampanga		pam
Pangasinan		pag
Panjabi	pa	pan

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Papiamentu		pap
Papuan-Australian (Other)		paa
Persian	fa	fas/per
Persian, Old (ca 600 - 400 B.C.)		peo
Phoenician		phn
Polish	pl	pol
Ponape		pon
Portuguese	pt	por
Prakrit languages		pra
Provençal, Old (to 1500)		pro
Pushto	ps	pus
Quechua	qu	que
Rhaeto-Romance	rm	roh
Rajasthani		raj
Rarotongan		rar
Romance (Other)		roa
Romanian	ro	ron/rum
Romany		rom
Rundi	rn	run
Russian	ru	rus
Salishan languages		sal
Samaritan Aramaic		sam
Sami languages		smi
Samoan	sm	smo
Sandawe		sad
Sango	sg	sag
Sanskrit	sa	san
Sardinian		srd
Scots		sco
Selkup		sel
Semitic (Other)		sem
Serbian	sr	
Serbo-Croatian	sh	scr
Serer		srr
Shan		shn
Shona	sn	sna

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Sidamo		sid
Siksika		bla
Sindhi	sd	snd
Singhalese	si	sin
Sino-Tibetan (Other)		sit
Siouan languages		sio
Slavic (Other)		sla
Siswant	ss	ssw
Slovak	sk	slk/slo
Slovenian	sl	slv
Sogdian		sog
Somali	so	som
Songhai		son
Sorbian languages		wen
Sotho, Northern		nso
Sotho, Southern	st	sot
South American Indian (Other)		sai
Spanish	es	esl/spa
Sukuma		suk
Sumerian		sux
Sudanese	su	sun
Susu		sus
Swahili	sw	swa
Swazi		ssw
Swedish	sv	sve/swe
Syriac		syr
Tagalog	tl	tgl
Tahitian		tah
Tajik	tg	tgk
Tamashek		tmh
Tamil	ta	tam
Tatar	tt	tat
Telugu	te	tel
Tereno		ter
Thai	th	tha
Tibetan	bo	bod/tib

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Tigre		tig
Tigrinya	ti	tir
Timne		tem
Tivi		tiv
Tlingit		tli
Tonga (Nyasa)	to	tog
Tonga (Tonga Islands)		ton
Truk		tru
Tsimshian		tsi
Tsonga	ts	tso
Tswana	tn	tsn
Tumbuka		tum
Turkish	tr	tur
Turkish, Ottoman (1500–1928)		ota
Turkmen	tk	tuk
Tuvinian		tyv
Twi	tw	twi
Ugaritic		uga
Uighur	ug	uig
Ukrainian	uk	ukr
Umbundu		umb
Undetermined		und
Urdu	ur	urd
Uzbek	uz	uzb
Vai		vai
Venda		ven
Vietnamese	vi	vie
Volapük	vo	vol
Votic		vot
Wakashan languages		wak
Walamo		wal
Waray		war
Washo		was
Welsh	cy	cym/wel
Wolof	wo	wol
Xhosa	xh	xho

Table G-2. ISO 639 and 639-2 language codes (continued)

Language	ISO 639	ISO 639-2
Yakut		sah
Yao		yao
Yap		yap
Yiddish	yi	yid
Yoruba	yo	yor
Zapotec		zap
Zenaga		zen
Zhuang	za	zha
Zulu	zu	zul
Zuni		zun

ISO 3166 Country Codes

Table G-3. ISO 3166 country codes

Country	Code
Afghanistan	AF
Albania	AL
Algeria	DZ
American Samoa	AS
Andorra	AD
Angola	AO
Anguilla	AI
Antarctica	AQ
Antigua and Barbuda	AG
Argentina	AR
Armenia	AM
Aruba	AW
Australia	AU
Austria	AT
Azerbaijan	AZ
Bahamas	BS
Bahrain	BH
Bangladesh	BD
Barbados	BB
Belarus	BY
Belgium	BE

Table G-3. ISO 3166 country codes (continued)

Country	Code
Belize	BZ
Benin	BJ
Bermuda	BM
Bhutan	BT
Bolivia	BO
Bosnia and Herzegovina	BA
Botswana	BW
Bouvet Island	BV
Brazil	BR
British Indian Ocean Territory	IO
Brunei Darussalam	BN
Bulgaria	BG
Burkina Faso	BF
Burundi	BI
Cambodia	KH
Cameroon	CM
Canada	CA
Cape Verde	CV
Cayman Islands	KY
Central African Republic	CF
Chad	TD
Chile	CL
China	CN
Christmas Island	CX
Cocos (Keeling) Islands	CC
Colombia	CO
Comoros	KM
Congo	CG
Congo (Democratic Republic of the)	CD
Cook Islands	CK
Costa Rica	CR
Cote D'Ivoire	CI
Croatia	HR
Cuba	CU
Cyprus	CY
Czech Republic	CZ

Table G-3. ISO 3166 country codes (continued)

Country	Code
Denmark	DK
Djibouti	DJ
Dominica	DM
Dominican Republic	DO
East Timor	TP
Ecuador	EC
Egypt	EG
El Salvador	SV
Equatorial Guinea	GQ
Eritrea	ER
Estonia	EE
Ethiopia	ET
Falkland Islands (Malvinas)	FK
Faroe Islands	FO
Fiji	FJ
Finland	FI
France	FR
French Guiana	GF
French Polynesia	PF
French Southern Territories	TF
Gabon	GA
Gambia	GM
Georgia	GE
Germany	DE
Ghana	GH
Gibraltar	GI
Greece	GR
Greenland	GL
Grenada	GD
Guadeloupe	GP
Guam	GU
Guatemala	GT
Guinea	GN
Guinea-Bissau	GW
Guyana	GY
Haiti	HT

Table G-3. ISO 3166 country codes (continued)

Country	Code
Heard Island and McDonald Islands	HM
Holy See (Vatican City State)	VA
Honduras	HN
Hong Kong	HK
Hungary	HU
Iceland	IS
India	IN
Indonesia	ID
Iran (Islamic Republic of)	IR
Iraq	IQ
Ireland	IE
Israel	IL
Italy	IT
Jamaica	JM
Japan	JP
Jordan	JO
Kazakhstan	KZ
Kenya	KE
Kiribati	KI
Korea (Democratic People's Republic of)	KP
Korea (Republic of)	KR
Kuwait	KW
Kyrgyzstan	KG
Lao People's Democratic Republic	LA
Latvia	LV
Lebanon	LB
Lesotho	LS
Liberia	LR
Libyan Arab Jamahiriya	LY
Liechtenstein	LI
Lithuania	LT
Luxembourg	LU
Macao	MO
Macedonia (The Former Yugoslav Republic of)	MK
Madagascar	MG
Malawi	MW

Table G-3. ISO 3166 country codes (continued)

Country	Code
Malaysia	MY
Maldives	MV
Mali	ML
Malta	MT
Marshall Islands	MH
Martinique	MQ
Mauritania	MR
Mauritius	MU
Mayotte	YT
Mexico	MX
Micronesia (Federated States of)	FM
Moldova (Republic of)	MD
Monaco	MC
Mongolia	MN
Montserrat	MS
Morocco	MA
Mozambique	MZ
Myanmar	MM
Namibia	NA
Nauru	NR
Nepal	NP
Netherlands	NL
Netherlands Antilles	AN
New Caledonia	NC
New Zealand	NZ
Nicaragua	NI
Niger	NE
Nigeria	NG
Niue	NU
Norfolk Island	NF
Northern Mariana Islands	MP
Norway	NO
Oman	OM
Pakistan	PK
Palau	PW
Palestinian Territory (Occupied)	PS
Panama	PA

Table G-3. ISO 3166 country codes (continued)

Country	Code
Papua New Guinea	PG
Paraguay	PY
Peru	PE
Philippines	PH
Pitcairn	PN
Poland	PL
Portugal	PT
Puerto Rico	PR
Qatar	QA
Reunion	RE
Romania	RO
Russian Federation	RU
Rwanda	RW
Saint Helena	SH
Saint Kitts and Nevis	KN
Saint Lucia	LC
Saint Pierre and Miquelon	PM
Saint Vincent and the Grenadines	VC
Samoa	WS
San Marino	SM
Sao Tome and Principe	ST
Saudi Arabia	SA
Senegal	SN
Seychelles	SC
Sierra Leone	SL
Singapore	SG
Slovakia	SK
Slovenia	SI
Solomon Islands	SB
Somalia	SO
South Africa	ZA
South Georgia and the South Sandwich Islands	GS
Spain	ES
Sri Lanka	LK
Sudan	SD
Suriname	SR
Svalbard and Jan Mayen	SJ

Table G-3. ISO 3166 country codes (continued)

Country	Code
Swaziland	SZ
Sweden	SE
Switzerland	CH
Syrian Arab Republic	SY
Taiwan, Province of China	TW
Tajikistan	TJ
Tanzania (United Republic of)	TZ
Thailand	TH
Togo	TG
Tokelau	TK
Tonga	TO
Trinidad and Tobago	TT
Tunisia	TN
Turkey	TR
Turkmenistan	TM
Turks and Caicos Islands	TC
Tuvalu	TV
Uganda	UG
Ukraine	UA
United Arab Emirates	AE
United Kingdom	GB
United States	US
United States Minor Outlying Islands	UM
Uruguay	UY
Uzbekistan	UZ
Vanuatu	VU
Venezuela	VE
Viet NAM	VN
Virgin Islands (British)	VG
Virgin ISLANDS (U.S.)	VI
Wallis and Futuna	WF
Western Sahara	EH
Yemen	YE
Yugoslavia	YU
Zambia	ZM

Language Administrative Organizations

ISO 639 defines a maintenance agency for additions to and changes in the list of languages in ISO 639. This agency is:

International Information Centre for Terminology (Infoterm)
P.O. Box 130
A-1021 Wien
Austria

Phone: +43 1 26 75 35 Ext. 312
Fax: +43 1 216 32 72

ISO 639-2 defines a maintenance agency for additions to and changes in the list of languages in ISO 639-2. This agency is:

Library of Congress
Network Development and MARC Standards Office
Washington, D.C. 20540
USA

Phone: +1 202 707 6237
Fax: +1 202 707 0115
URL: <http://www.loc.gov/standards/iso639/>

The maintenance agency for ISO 3166 (country codes) is:

ISO 3166 Maintenance Agency Secretariat
c/o DIN Deutsches Institut fuer Normung
Burggrafenstrasse 6
Postfach 1107
D-10787 Berlin
Germany

Phone: +49 30 26 01 320
Fax: +49 30 26 01 231
URL: <http://www.din.de/gremien/nas/nabd/iso3166ma/>

MIME Charset Registry

This appendix describes the MIME charset registry maintained by the Internet Assigned Numbers Authority (IANA). A formatted table of charsets from the registry is provided in Table H-1.

MIME Charset Registry

MIME charset tags are registered with the IANA (<http://www.iana.org/numbers.htm>). The charset registry is a flat-file text database of records. Each record contains a charset name, reference citations, a unique MIB number, a source description, and a list of aliases. A name or alias may be flagged “preferred MIME name.”

Here is the record for US-ASCII:

```
Name: ANSI_X3.4-1968                                [RFC1345, KXS2]^
MIBenum: 3
Source: ECMA registry
Alias: iso-ir-6
Alias: ANSI_X3.4-1986
Alias: ISO_646.irv:1991
Alias: ASCII
Alias: ISO646-US
Alias: US-ASCII (preferred MIME name)
Alias: us
Alias: IBM367
Alias: cp367
Alias: csASCII
```

The procedure for registering a charset with the IANA is documented in RFC 2978 (<http://www.ietf.org/rfc/rfc2978.txt>).

Preferred MIME Names

Of the 235 charsets registered at the time of this writing, only 20 include “preferred MIME names”—common charsets used by email and web applications. These are:

Big5	EUC-JP	EUC-KR
GB2312	ISO-2022-JP	ISO-2022-JP-2
ISO-2022-KR	ISO-8859-1	ISO-8859-2
ISO-8859-3	ISO-8859-4	ISO-8859-5
ISO-8859-6	ISO-8859-7	ISO-8859-8
ISO-8859-9	ISO-8859-10	KOI8-R
Shift-JIS	US-ASCII	

Registered Charsets

Table H-1 lists the contents of the charset registry as of March 2001. Refer directly to <http://www.iana.org> for more information about the contents of this table.

Table H-1. IANA MIME charset tags

Charset tag	Aliases	Description	References
US-ASCII	ANSI_X3.4-1968, iso-ir-6, ANSI_X3.4-1986, ISO_646.irv:1991, ASCII, ISO646-US, us, IBM367, cp367, csASCII	ECMA registry	RFC1345, KXS2
ISO-10646-UTF-1	csISO10646UTF1	Universal Transfer Format (1)—this is the multibyte encoding that subsets ASCII-7; it does not have byte-ordering issues	
ISO_646.basic:1983	ref, csISO646basic1983	ECMA registry	RFC1345, KXS2
INVARIANT	csINVARIANT		RFC1345, KXS2
ISO_646.irv:1983	iso-ir-2, irv, csISO2IntlRefVersion	ECMA registry	RFC1345, KXS2
BS_4730	iso-ir-4, ISO646-GB, gb, uk, csISO4UnitedKingdom	ECMA registry	RFC1345, KXS2
NATS-SEFI	iso-ir-8-1, csNATSSEFI	ECMA registry	RFC1345, KXS2
NATS-SEFI-ADD	iso-ir-8-2, csNATSSEFIADD	ECMA registry	RFC1345, KXS2
NATS-DANO	iso-ir-9-1, csNATSDANO	ECMA registry	RFC1345, KXS2
NATS-DANO-ADD	iso-ir-9-2, csNATSDANOADD	ECMA registry	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
SEN_850200_B	iso-ir-10, FI, ISO646-FI, ISO646-SE, se, csISO10Swedish	ECMA registry	RFC1345, KXS2
SEN_850200_C	iso-ir-11, ISO646-SE2, se2, csISO11SwedishForNames	ECMA registry	RFC1345, KXS2
KS_C_5601-1987	iso-ir-149, KS_C_5601-1989, KSC_5601, korean, csKSC56011987	ECMA registry	RFC1345, KXS2
ISO-2022-KR	csISO2022KR	RFC 1557 (see also KS_C_5601-1987)	RFC1557, Choi
EUC-KR	csEUCKR	RFC 1557 (see also KS_C_5861-1992)	RFC1557, Choi
ISO-2022-JP	csISO2022JP	RFC 1468 (see also RFC 2237)	RFC1468, Murai
ISO-2022-JP-2	csISO2022JP2	RFC 1554	RFC1554, Ohta
ISO-2022-CN		RFC 1922	RFC1922
ISO-2022-CN-EXT		RFC 1922	RFC1922
JIS_C6220-1969-jp	JIS_C6220-1969, iso-ir-13, katakana, x0201-7, csISO13JISC6220jp	ECMA registry	RFC1345, KXS2
JIS_C6220-1969-ro	iso-ir-14, jp, ISO646-JP, csISO14JISC6220ro	ECMA registry	RFC1345, KXS2
IT	iso-ir-15, ISO646-IT, csISO15Italian	ECMA registry	RFC1345, KXS2
PT	iso-ir-16, ISO646-PT, csISO16Portuguese	ECMA registry	RFC1345, KXS2
ES	iso-ir-17, ISO646-ES, csISO17Spanish	ECMA registry	RFC1345, KXS2
greek7-old	iso-ir-18, csISO18Greek7Old	ECMA registry	RFC1345, KXS2
latin-greek	iso-ir-19, csISO19LatinGreek	ECMA registry	RFC1345, KXS2
DIN_66003	iso-ir-21, de, ISO646-DE, csISO21German	ECMA registry	RFC1345, KXS2
NF_Z_62-010_(1973)	iso-ir-25, ISO646-FR1, csISO25French	ECMA registry	RFC1345, KXS2
Latin-greek-1	iso-ir-27, csISO27LatinGreek1	ECMA registry	RFC1345, KXS2
ISO_5427	iso-ir-37, csISO5427Cyrillic	ECMA registry	RFC1345, KXS2
JIS_C6226-1978	iso-ir-42, csISO42JISC62261978	ECMA registry	RFC1345, KXS2
BS_viewdata	iso-ir-47, csISO47BSViewdata	ECMA registry	RFC1345, KXS2
INIS	iso-ir-49, csISO49INIS	ECMA registry	RFC1345, KXS2
INIS-8	iso-ir-50, csISO50INIS8	ECMA registry	RFC1345, KXS2
INIS-cyrillic	iso-ir-51, csISO51INISCyrillic	ECMA registry	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
ISO_5427:1981	iso-ir-54, ISO5427Cyrillic1981	ECMA registry	RFC1345, KXS2
ISO_5428:1980	iso-ir-55, csISO5428Greek	ECMA registry	RFC1345, KXS2
GB_1988-80	iso-ir-57, cn, ISO646-CN, csISO57GB1988	ECMA registry	RFC1345, K5, KXS2
GB_2312-80	iso-ir-58, chinese, csISO58GB231280	ECMA registry	RFC1345, KXS2
NS_4551-1	iso-ir-60, ISO646-NO, no, csISO60DanishNorwegian, csISO60Norwegian1	ECMA registry	RFC1345, KXS2
NS_4551-2	ISO646-NO2, iso-ir-61, no2, csISO61Norwegian2	ECMA registry	RFC1345, KXS2
NF_Z_62-010	iso-ir-69, ISO646-FR, fr, csISO69French	ECMA registry	RFC1345, KXS2
videotex-suppl	iso-ir-70, csISO70VideotexSupp1	ECMA registry	RFC1345, KXS2
PT2	iso-ir-84, ISO646-PT2, csISO84Portuguese2	ECMA registry	RFC1345, KXS2
ES2	iso-ir-85, ISO646-ES2, csISO85Spanish2	ECMA registry	RFC1345, KXS2
MSZ_7795.3	iso-ir-86, ISO646-HU, hu, csISO86Hungarian	ECMA registry	RFC1345, KXS2
JIS_C6226-1983	iso-ir-87, x0208, JIS_X0208-1983, csISO87JISX0208	ECMA registry	RFC1345, KXS2
greek7	iso-ir-88, csISO88Greek7	ECMA registry	RFC1345, KXS2
ASMO_449	ISO_9036, arabic7, iso-ir-89, csISO89ASMO449	ECMA registry	RFC1345, KXS2
iso-ir-90	csISO90	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-a	iso-ir-91, jp-ocr-a, csISO91JISC62291984a	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-b	iso-ir-92, ISO646-JP-OCR-B, jp-ocr-b, csISO92JISC62991984b	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-b-add	iso-ir-93, jp-ocr-b-add, csISO93JIS62291984badd	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-hand	iso-ir-94, jp-ocr-hand, csISO94JIS62291984hand	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-hand-add	iso-ir-95, jp-ocr-hand-add, csISO95JIS62291984handadd	ECMA registry	RFC1345, KXS2
JIS_C6229-1984-kana	iso-ir-96, csISO96JISC62291984kana	ECMA registry	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
ISO_2033-1983	iso-ir-98, e13b, csISO2033	ECMA registry	RFC1345, KXS2
ANSI_X3.110-1983	iso-ir-99, CSA_T500-1983, NAPLPS, csISO99NAPLPS	ECMA registry	RFC1345, KXS2
ISO-8859-1	ISO_8859-1:1987, iso-ir-100, ISO_8859-1, latin1, l1, IBM819, CP819, csISOLatin1	ECMA registry	RFC1345, KXS2
ISO-8859-2	ISO_8859-2:1987, iso-ir-101, ISO_8859-2, latin2, l2, csISOLatin2	ECMA registry	RFC1345, KXS2
T.61-7bit	iso-ir-102, csISO102T617bit	ECMA registry	RFC1345, KXS2
T.61-8bit	T.61, iso-ir-103, csISO103T618bit	ECMA registry	RFC1345, KXS2
ISO-8859-3	ISO_8859-3:1988, iso-ir-109, ISO_8859-3, latin3, l3, csISOLatin3	ECMA registry	RFC1345, KXS2
ISO-8859-4	ISO_8859-4:1988, iso-ir-110, ISO_8859-4, latin4, l4, csISOLatin4	ECMA registry	RFC1345, KXS2
ECMA-cyrillic	iso-ir-111, csISO111ECMACyrillic	ECMA registry	RFC1345, KXS2
CSA_Z243.4-1985-1	iso-ir-121, ISO646-CA, csa7-1, ca, csISO121Canadian1	ECMA registry	RFC1345, KXS2
CSA_Z243.4-1985-2	iso-ir-122, ISO646-CA2, csa7-2, csISO122Canadian2	ECMA registry	RFC1345, KXS2
CSA_Z243.4-1985-gr	iso-ir-123, csISO123CSAZ24341985gr	ECMA registry	RFC1345, KXS2
ISO-8859-6	ISO_8859-6:1987, iso-ir-127, ISO_8859-6, ECMA-114, ASMO-708, arabic, csISOLatinArabic	ECMA registry	RFC1345, KXS2
ISO_8859-6-E	csISO88596E	RFC 1556	RFC1556, IANA
ISO_8859-6-I	csISO88596I	RFC 1556	RFC1556, IANA
ISO-8859-7	ISO_8859-7:1987, iso-ir-126, ISO_8859-7, ELOT_928, ECMA-118, greek, greek8, csISOLatinGreek	ECMA registry	RFC1947, RFC1345, KXS2
T.101-G2	iso-ir-128, csISO128T101G2	ECMA registry	RFC1345, KXS2
ISO-8859-8	ISO_8859-8:1988, iso-ir-138, ISO_8859-8, hebrew, csISOLatinHebrew	ECMA registry	RFC1345, KXS2
ISO_8859-8-E	csISO88598E	RFC 1556	RFC1556, Nussbacher

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
ISO_8859-8-I	csISO88598I	RFC 1556	RFC1556, Nussbacher
CSN_369103	iso-ir-139, csISO139CSN369103	ECMA registry	RFC1345, KXS2
JUS_I.B1.002	iso-ir-141, ISO646-YU, js, yu, csISO141JUSIB1002	ECMA registry	RFC1345, KXS2
ISO_6937-2-add	iso-ir-142, csISOTextComm	ECMA registry and ISO 6937-2:1983	RFC1345, KXS2
IEC_P27-1	iso-ir-143, csISO143IECP271	ECMA registry	RFC1345, KXS2
ISO-8859-5	ISO_8859-5:1988, iso-ir-144, ISO_8859-5, cyrillic, csISOLatinCyrillic	ECMA registry	RFC1345, KXS2
JUS_I.B1.003-serb	iso-ir-146, serbian, csISO146Serbian	ECMA registry	RFC1345, KXS2
JUS_I.B1.003-mac	macedonian, iso-ir-147, csISO147Macedonian	ECMA registry	RFC1345, KXS2
ISO-8859-9	ISO_8859-9:1989, iso-ir-148, ISO_8859-9, latin5, IS, csISOLatin5	ECMA registry	RFC1345, KXS2
greek-ccitt	iso-ir-150, csISO150, csISO150GreekCCITT	ECMA registry	RFC1345, KXS2
NC_NC00-10:81	cuba, iso-ir-151, ISO646-CU, csISO151Cuba	ECMA registry	RFC1345, KXS2
ISO_6937-2-25	iso-ir-152, csISO6937Add	ECMA registry	RFC1345, KXS2
GOST_19768-74	ST_SEV_358-88, iso-ir-153, csISO153GOST1976874	ECMA registry	RFC1345, KXS2
ISO_8859-supp	iso-ir-154, latin1-2-5, csISO8859Supp	ECMA registry	RFC1345, KXS2
ISO_10367-box	iso-ir-155, csISO10367Box	ECMA registry	RFC1345, KXS2
ISO-8859-10	iso-ir-157, I6, ISO_8859-10:1992, csISOLatin6, latin6	ECMA registry	RFC1345, KXS2
latin-lap	lap, iso-ir-158, csISO158Lap	ECMA registry	RFC1345, KXS2
JIS_X0212-1990	x0212, iso-ir-159, csISO159JISX02121990	ECMA registry	RFC1345, KXS2
DS_2089	DS2089, ISO646-DK, dk, csISO646Danish	Danish Standard, DS 2089, February 1974	RFC1345, KXS2
us-dk	csUSDK		RFC1345, KXS2
dk-us	csDKUS		RFC1345, KXS2
JIS_X0201	X0201, csHalfWidthKatakana	JIS X 0201-1976—1 byte only; this is equivalent to JIS/Roman (similar to ASCII) plus 8-bit half-width katakana	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
KSC5636	ISO646-KR, csKSC5636		RFC1345, KXS2
ISO-10646-UCS-2	csUnicode	The 2-octet Basic Multilingual Plane, a.k.a. Unicode—this needs to specify network byte order; the standard does not specify it (it is a 16-bit integer space)	
ISO-10646-UCS-4	csUCS4	The full code space (same comment about byte order; these are 31-bit numbers)	
DEC-MCS	dec, csDECMCS	VAX/VMS User's Manual, Order Number: AI-Y517A-TE, April 1986	RFC1345, KXS2
hp-roman8	roman8, r8, csHPRoman8	LaserJet IIP Printer User's Manual, HP part no 33471-90901, Hewlett-Packard, June 1989	HP-PCL5, RFC1345, KXS2
macintosh	mac, csMacintosh	The Unicode Standard v1.0, ISBN 0201567881, Oct 1991	RFC1345, KXS2
IBM037	cp037, ebcdic-cp-us, ebcdic-cp-ca, ebcdic-cp-wt, ebcdic-cp-nl, csIBM037	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM038	EBCDIC-INT, cp038, csIBM038	IBM 3174 Character Set Ref, GA27-3831-02, March 1990	RFC1345, KXS2
IBM273	CP273, csIBM273	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM274	EBCDIC-BE, CP274, csIBM274	IBM 3174 Character Set Ref, GA27-3831-02, March 1990	RFC1345, KXS2
IBM275	EBCDIC-BR, cp275, csIBM275	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM277	EBCDIC-CP-DK, EBCDIC-CP-NO, csIBM277	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM278	CP278, ebcdic-cp-fi, ebcdic-cp-se, csIBM278	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM280	CP280, ebcdic-cp-it, csIBM280	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM281	EBCDIC-JP-E, cp281, csIBM281	IBM 3174 Character Set Ref, GA27-3831-02, March 1990	RFC1345, KXS2
IBM284	CP284, ebcdic-cp-es, csIBM284	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM285	CP285, ebcdic-cp-gb, csIBM285	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM290	cp290, EBCDIC-JP-kana, csIBM290	IBM 3174 Character Set Ref, GA27-3831-02, March 1990	RFC1345, KXS2
IBM297	cp297, ebcdic-cp-fr, csIBM297	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
IBM420	cp420, ebcdic-cp-ar1, csIBM420	IBM NLS RM Vol2 SE09-8002-01, March 1990, IBM NLS RM p 11-11	RFC1345, KXS2
IBM423	cp423, ebcdic-cp-gr, csIBM423	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM424	cp424, ebcdic-cp-he, csIBM424	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM437	cp437, 437, csPC8CodePage437	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM500	CP500, ebcdic-cp-be, ebcdic-cp-ch, csIBM500	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM775	cp775, csPC775Baltic	HP PCL 5 Comparison Guide (P/N 5021-0329) pp B-13, 1996	HP-PCL5
IBM850	cp850, 850, csPC850Multilingual	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM851	cp851, 851, csIBM851	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM852	cp852, 852, csPCp852	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM855	cp855, 855, csIBM855	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM857	cp857, 857, csIBM857	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM860	cp860, 860, csIBM860	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM861	cp861, 861, cp-is, csIBM861	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM862	cp862, 862, csPC862LatinHebrew	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM863	cp863, 863, csIBM863	IBM keyboard layouts and code pages, PN 07G4586, June 1991	RFC1345, KXS2
IBM864	cp864, csIBM864	IBM keyboard layouts and code pages, PN 07G4586, June 1991	RFC1345, KXS2
IBM865	cp865, 865, csIBM865	IBM DOS 3.3 Ref (Abridged), 94X9575, Feb 1987	RFC1345, KXS2
IBM866	cp866, 866, csIBM866	IBM NLDG Vol2 SE09-8002-03, August 1994	Pond
IBM868	CP868, cp-ar, csIBM868	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM869	cp869, 869, cp-gr, csIBM869	IBM keyboard layouts and code pages, PN 07G4586, June 1991	RFC1345, KXS2
IBM870	CP870, ebcdic-cp-roece, ebcdic-cp-yu, csIBM870	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
IBM871	CP871, ebcdic-cp-is, csIBM871	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM880	cp880, EBCDIC-Cyrillic, csIBM880	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM891	cp891, csIBM891	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM903	cp903, csIBM903	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM904	cp904, 904, csIBBM904	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM905	CP905, ebcdic-cp-tr, csIBM905	IBM 3174 Character Set Ref, GA27-3831-02, March 1990	RFC1345, KXS2
IBM918	CP918, ebcdic-cp-ar2, csIBM918	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
IBM1026	CP1026, csIBM1026	IBM NLS RM Vol2 SE09-8002-01, March 1990	RFC1345, KXS2
EBCDIC-AT-DE	csIBMBCDICATDE	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-AT-DE-A	csEBCDICATDEA	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-CA-FR	csEBCDICCFAFR	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-DK-NO	csEBCDICDKNO	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-DK-NO-A	csEBCDICDKNOA	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-FI-SE	csEBCDICFISE	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-FI-SE-A	csEBCDICFISEA	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-FR	csEBCDICFR	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-IT	csEBCDICIT	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-PT		IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-ES	csEBCDICES	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-ES-A	csEBCDICESA	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
EBCDIC-ES-S	csEBCDICES	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-UK	csEBCDICUK	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
EBCDIC-US	csEBCDICUS	IBM 3270 Char Set Ref Ch 10, GA27-2837-9, April 1987	RFC1345, KXS2
UNKNOWN-8BIT	csUnknown8BiT		RFC1428
MNEMONIC	csMnemonic	RFC 1345, also known as "mnemonic+ascii+38"	RFC1345, KXS2
MNEM	csMnem	RFC 1345, also known as "mnemonic+ascii+8200"	RFC1345, KXS2
VISCII	csVISCII	RFC 1456	RFC1456
VIQR	csVIQR	RFC 1456	RFC1456
KOI8-R	csKOI8R	RFC 1489, based on GOST-19768-74, ISO-6937/8, INIS-Cyrillic, ISO-5427	RFC1489
KOI8-U		RFC 2319	RFC2319
IBM00858	CCSID00858, CP00858, PC-Multilingual-850+euro	IBM (see .../assignments/character-set-info/IBM00858) [Mahdi]	
IBM00924	CCSID00924, CP00924, ebcdic-Latin9--euro	IBM (see .../assignments/character-set-info/IBM00924) [Mahdi]	
IBM01140	CCSID01140, CP01140, ebcdic-us-37+euro	IBM (see .../assignments/character-set-info/IBM01140) [Mahdi]	
IBM01141	CCSID01141, CP01141, ebcdic-de-273+euro	IBM (see .../assignments/character-set-info/IBM01141) [Mahdi]	
IBM01142	CCSID01142, CP01142, ebcdic-dk-277+euro, ebcdic-no-277+euro	IBM (see .../assignments/character-set-info/IBM01142) [Mahdi]	
IBM01143	CCSID01143, CP01143, ebcdic-fi-278+euro, ebcdic-se-278+euro	IBM (see .../assignments/character-set-info/IBM01143) [Mahdi]	
IBM01144	CCSID01144, CP01144, ebcdic-it-280+euro	IBM (see .../assignments/character-set-info/IBM01144) [Mahdi]	
IBM01145	CCSID01145, CP01145, ebcdic-es-284+euro	IBM (see .../assignments/character-set-info/IBM01145) [Mahdi]	
IBM01146	CCSID01146, CP01146, ebcdic-gb-285+euro	IBM (see .../assignments/character-set-info/IBM01146) [Mahdi]	
IBM01147	CCSID01147, CP01147, ebcdic-fr-297+euro	IBM (see .../assignments/character-set-info/IBM01147) [Mahdi]	
IBM01148	CCSID01148, CP01148, ebcdic-international-500+euro	IBM (see .../assignments/character-set-info/IBM01148) [Mahdi]	

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
IBM01149	CCSID01149, CP01149, ebcdic-is-871+euro	IBM (see .../assignments/character-set-info/IBM01149) [Mahdi]	
Big5-HKSCS	None	See (.../assignments/character-set-info/Big5-HKSCS) [Yick]	
UNICODE-1-1	csUnicode11	RFC 1641	RFC1641
SCSU	None	SCSU (see .../assignments/character-set-info/SCSU) [Scherer]	
UTF-7	None	RFC 2152	RFC2152
UTF-16BE	None	RFC 2781	RFC2781
UTF-16LE	None	RFC 2781	RFC2781
UTF-16	None	RFC 2781	RFC2781
UNICODE-1-1-UTF-7	csUnicode11UTF7	RFC 1642	RFC1642
UTF-8		RFC 2279	RFC2279
iso-8859-13		ISO (see ...assignments/character-set-info/iso-8859-13)[Tumasonis]	
iso-8859-14	iso-ir-199, ISO_8859-14:1998, ISO_8859-14, latin8, iso-celtic, l8	ISO (see ...assignments/character-set-info/iso-8859-14) [Simonsen]	
ISO-8859-15	ISO_8859-15	ISO	
JIS_Encoding	csJISEncoding	JIS X 0202-1991; uses ISO 2022 escape sequences to shift code sets, as documented in JIS X 0202-1991	
Shift_JIS	MS_Kanji, csShiftJIS	This charset is an extension of csHalfWidthKatakana—it adds graphic characters in JIS X 0208. The CCSs are JIS X0201:1997 and JIS X0208:1997. The complete definition is shown in Appendix 1 of JISX0208:1997. This charset can be used for the top-level media type “text”.	
EUC-JP	Extended_UNIX_Code_Packed_Format_for_Japanese, csEUCPkdFmtJapanese	Standardized by OSF, UNIX International, and UNIX Systems Laboratories Pacific. Uses ISO 2022 rules to select code set. code set 0: US-ASCII (a single 7-bit byte set); code set 1: JIS X0208-1990 (a double 8-bit byte set) restricted to A0–FF in both bytes; code set 2: half-width katakana (a single 7-bit byte set) requiring SS2 as the character prefix; code set 3: JIS X0212-1990 (a double 7-bit byte set) restricted to A0–FF in both bytes requiring SS3 as the character prefix.	

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
Extended_UNIX_Code_Fixed_Width_for_Japanese	csEUCFixWidJapanese	Used in Japan. Each character is 2 octets. code set 0: US-ASCII (a single 7-bit byte set), 1st byte = 00, 2nd byte = 20–7E; code set 1: JIS X0208-1990 (a double 7-bit byte set) restricted to A0–FF in both bytes; code set 2: half-width katakana (a single 7-bit byte set), 1st byte = 00, 2nd byte = A0–FF; code set 3: JIS X0212-1990 (a double 7-bit byte set) restricted to A0–FF in the first byte and 21–7E in the second byte.	
ISO-10646-UCS-Basic	csUnicodeASCII	ASCII subset of Unicode. Basic Latin = collection 1. See ISO 10646, Appendix A.	
ISO-10646-Unicode-Latin1	csUnicodeLatin1, ISO-10646	ISO Latin-1 subset of Unicode. Basic Latin and Latin-1. Supplement = collections 1 and 2. See ISO 10646, Appendix A, and RFC 1815.	
ISO-10646-J-1		ISO 10646 Japanese. See RFC 1815.	
ISO-Unicode-IBM-1261	csUnicodeIBM1261	IBM Latin-2, -3, -5, Extended Presentation Set, GCSGID: 1261	
ISO-Unicode-IBM-1268	csUnicodeIBM1268	IBM Latin-4 Extended Presentation Set, GCSGID: 1268	
ISO-Unicode-IBM-1276	csUnicodeIBM1276	IBM Cyrillic Greek Extended Presentation Set, GCSGID: 1276	
ISO-Unicode-IBM-1264	csUnicodeIBM1264	IBM Arabic Presentation Set, GCSGID: 1264	
ISO-Unicode-IBM-1265	csUnicodeIBM1265	IBM Hebrew Presentation Set, GCSGID: 1265	
ISO-8859-1-Windows-3.0-Latin-1	csWindows30Latin1	Extended ISO 8859-1 Latin-1 for Windows 3.0. PCL Symbol Set ID: 9U.	HP-PCL5
ISO-8859-1-Windows-3.1-Latin-1	csWindows31Latin1	Extended ISO 8859-1 Latin-1 for Windows 3.1. PCL Symbol Set ID: 19U.	HP-PCL5
ISO-8859-2-Windows-Latin-2	csWindows31Latin2	Extended ISO 8859-2. Latin-2 for Windows 3.1. PCL Symbol Set ID: 9E.	HP-PCL5
ISO-8859-9-Windows-Latin-5	csWindows31Latin5	Extended ISO 8859-9. Latin-5 for Windows 3.1. PCL Symbol Set ID: 5T.	HP-PCL5
Adobe-Standard-Encoding	csAdobeStandardEncoding	PostScript Language Reference Manual. PCL Symbol Set ID: 10J.	Adobe
Ventura-US	csVenturaUS	Ventura US-ASCII plus characters typically used in publishing, such as pilcrow, copyright, registered, trademark, section, dagger, and double dagger in the range A0 (hex) to FF (hex). PCL Symbol Set ID: 14J.	HP-PCL5

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
Ventura-International	csVenturaInternational	Ventura International. ASCII plus coded characters similar to Roman8. PCL Symbol Set ID: 13J.	HP-PCL5
PC8-Danish-Norwegian	csPC8DanishNorwegian	PC Danish Norwegian 8-bit PC set for Danish Norwegian. PCL Symbol Set ID: 11U.	HP-PCL5
PC8-Turkish	csPC8Turkish	PC Latin Turkish. PCL Symbol Set ID: 9T.	HP-PCL5
IBM-Symbols	csIBMSymbols	Presentation Set, CPGID: 259	IBM-CIDT
IBM-Thai	csIBMThai	Presentation Set, CPGID: 838	IBM-CIDT
HP-Legal	csHPLegal	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 1U.	HP-PCL5
HP-Pi-font	csHPPiFont	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 15U.	HP-PCL5
HP-Math8	csHPMath8	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 8M.	HP-PCL5
Adobe-Symbol-Encoding	csHPPSMath	PostScript Language Reference Manual. PCL Symbol Set ID: 5M.	Adobe
HP-DeskTop	csHPDesktop	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 7J.	HP-PCL5
Ventura-Math	csVenturaMath	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 6M.	HP-PCL5
Microsoft-Publishing	csMicrosoftPublishing	PCL 5 Comparison Guide, Hewlett-Packard, HP part number 5961-0510, October 1992. PCL Symbol Set ID: 6J.	HP-PCL5
Windows-31J	csWindows31J	Windows Japanese. A further extension of Shift_JIS to include NEC special characters (Row 13), NEC selection of IBM extensions (Rows 89 to 92), and IBM extensions (Rows 115 to 119). The CCSs are JIS X0201:1997, JIS X0208:1997, and these extensions. This charset can be used for the top-level media type "text", but it is of limited or specialized use (see RFC 2278). PCL Symbol Set ID: 19K.	
GB2312	csGB2312	Chinese for People's Republic of China (PRC) mixed 1-byte, 2-byte set: 20–7E = 1-byte ASCII; A1–FE = 2-byte PRC Kanji. See GB 2312-80. PCL Symbol Set ID: 18C.	

Table H-1. IANA MIME charset tags (continued)

Charset tag	Aliases	Description	References
Big5	csBig5	Chinese for Taiwan Multibyte set. PCL Symbol Set id: 18T.	[Tantsetthi]
windows-1250		Microsoft (see ../character-set-info/windows-1250) [Lazhintseva]	
windows-1251		Microsoft (see ../character-set-info/windows-1251) [Lazhintseva]	
windows-1252		Microsoft (see ../character-set-info/windows-1252) [Wendt]	
windows-1253		Microsoft (see ../character-set-info/windows-1253) [Lazhintseva]	
windows-1254		Microsoft (see ../character-set-info/windows-1254) [Lazhintseva]	
windows-1255		Microsoft (see ../character-set-info/windows-1255) [Lazhintseva]	
windows-1256		Microsoft (see ../character-set-info/windows-1256) [Lazhintseva]	
windows-1257		Microsoft (see ../character-set-info/windows-1257) [Lazhintseva]	
windows-1258		Microsoft (see ../character-set-info/windows-1258) [Lazhintseva]	
TIS-620		Thai Industrial Standards Institute (TISI)	
HZ-GB-2312		RFC 1842, RFC 1843 [RFC1842, RFC1843]	

Symbols

: (colon), use in headers, 47
= (equals sign), base-64 encoding, 572
/~ (slash-tilde), 122

Numbers

8-bit identity encoding, 382
100 Continue status code, 59, 60
100-199 status codes, 59–60, 505
200-299 status codes, 61, 505
300-399 status codes, 61–64, 506
400-499 status codes, 65–66, 506
500-599 status codes, 66, 507
2MSL (maximum segment lifetime), 85

A

absolute URLs, 30
Accept headers, 69, 508
 robots and, 225
Accept-Charset headers, 371, 375, 509
 MIME charset encoding tags and, 374
Accept-Encoding headers, 509
Accept-Instance-Manipulation headers, 367
Accept-Language headers, 371, 385, 510
 content negotiation and, 398
Accept-Ranges headers, 510
access controls, 124
 proxy authentication, 156
access proxies, 137
advertising, hit counts and caches, 194–196
age and freshness lifetime, 188
Age headers, 510
agents, 19

algorithms

aging and freshness, 187–194
document age calculation, 189–194
instance-manipulation algorithms, 367
LM-Factor, 184
message digest algorithms, 291–294
 symmetric authentication, 298
Nagle's algorithm, 84
redirection, enhanced DNS-based, 457
resource-discovery algorithm
 (WPAD), 143, 465
RSA, 317
aliases (URLs), 219
Allow headers, 159, 511
<allprop> element, 437
anonymizers, 136
anycast addressing, 457
Apache web servers, 110
 content negotiation, 399
 MultiViews directive, 400
 type-map files, 399
DirectoryIndex configuration
 directive, 123
document root, setting, 121
HostnameLookups configuration
 directive, 115
HTTP headers, control of, 186
IdentityCheck configuration
 directive, 116
magic typing, 126
APIs (application programming
 interfaces), 203
 server extensions, 205
 web services and, 205

We'd like to hear your suggestions for improving our indexes. Send email to index@oreilly.com.

- application/* MIME types, 540–557
- application programming interfaces (see APIs)
- application servers, 123, 203
- ASCII character set, 379
- asymmetric cryptography, 315
- attacks, 303–306
 - batched brute-force attacks, 305
 - chosen plaintext attacks, 305
 - dictionary attacks, 304
 - enumeration, 313
 - evidence of, 301
 - header tampering, 303
 - hostile proxies, 304
 - man-in-the-middle attacks, 304
 - replay attacks, 284, 303
 - preventing, 289
- audio/* MIME types, 557–559
- authentication, 277–280
 - basic (see basic authentication)
 - challenge/response framework, 278
 - digest (see digest authentication)
 - headers, 278
 - HTCP, 480
 - multiple authentication schemes, risks of, 303
 - protocols, 278
 - proxy servers, 156
 - server, using digital certificates, 321 (see also HTTPS)
- Authentication-Info directives, 576
- Authorization headers, 281, 511
 - directives, 575
 - preemptive generation, 295
- automatic expansion of URLs, 30

B

- bandwidth
 - bottlenecks, 161
 - transfer times and, 162
- base URLs, 32
- base-64 encoding, 570–572
 - alphabet, 571
 - equals sign (=), 572
 - HTTP, compatibility with, 570
 - padding, 572
 - Perl implementation, 572
 - purpose, 570
 - username/password, 282
- bases, 31

- basic authentication, 281–284
 - example, 281
 - headers, 281, 300
 - insecurity of, 283, 286
 - protection space, 302
 - by proxy servers, 283
 - username/password encoding, 282
 - web server vs. proxy, 283 (see also authentication)
- batched brute-force attacks, 305
- Binary Wire Protocol, 250, 252
- blind relays, 94
- browsers
 - Host headers and older versions of, 419
 - HTTP, use of, 13
 - parallel connections, maximum, 90
 - URLs, automatic expansion of, 34
- byte hit rate, 167

C

- Cache Array Routing Protocol (see CARP)
- Cache-Control headers, 175–177, 182–186, 189, 511
 - directives, 361
- caches, 18, 133, 161–196
 - advertising and, 194–196
 - aging and freshness algorithms, 187–194
 - byte hit rate, 167
 - cache busting, 492
 - cache hits and misses, 165, 168
 - cache meshes, 170
 - cache validators, 181
 - Cache-Control headers, 175–177, 182–186, 189, 511
 - directives, 361
 - cookies and, 273
 - digest authentication and, 302
 - distance delays and, 163
 - DNS caching and load balancing, 456
 - document expiration, 175
 - exclusion of documents from, 182
 - expiration time, setting, 182
 - Expires headers, 175, 183
 - flash crowds and, 163
 - freshness check, 173
 - heuristic expiration, 184
 - hit logs, 195
 - hit rate, 167
 - HTTP-EQUIV tag, 187

- logging, 174
- lookup, 173
- maintaining currency, 175
- message truncation and, 344
- network bottlenecks and, 161
- parent caches, 169
- parsing, 172
- peering, 171
- processing steps, 171–175
 - flowchart, 175
- receiving, 172
- response creation, 174
- revalidate hits, 165
- revalidations, 165–166
- sending, 174
- server revalidation, 175, 177
- setting controls, 186
- sibling caches, 171
- surrogate caches, 421
- topologies, 168
- uses, 161
- caching headers, 68
- caching proxies (see caches)
- caching proxy servers, 169
- canonicalizing of URLs, 37, 220
- CARP (Cache Array Routing Protocol), 475–478
 - disadvantages, 476
 - ICP vs., 475
 - redirection method, 476
- CAs (certificate authorities), 327
- case sensitivity, language tags, 386
- CDNs (content distribution networks), 421
- certificate authorities (CAs), 327
- certificates (see digital certificates)
- CGI (Common Gateway Interface), 203, 204
- challenge/response authentication
 - model, 278
 - challenge headers, 72
 - multiple challenges, 301
- character encoding schemes, 377, 381
- character repertoire, 377
- character sets, 35, 371–376
 - encoding, 370
 - mechanisms, 36
 - restricted characters, 36
- characters, 378
 - URLs, legal in, 35
- charset tags, 371
 - IANA MIME character set registry and, 371
- chemical/* MIME types, 559–561
- child filters, 131
- chosen plaintext attacks, 305
- chunked encodings, 345
- ciphers (see under cryptography)
- ciphertext, 310
- cleartext, 310
- client error status codes (400–499), 65–66, 505
- client hostname identification, 115
- client proxy configuration, 141–144
 - manual, 142
 - PAC (Proxy Auto-configuration) protocol, 142, 463
 - WPAD (Web Proxy Autodiscovery Protocol), 143
- client-driven negotiation, 396
 - disadvantages, 396
- Client-ip headers, 258, 260, 512
- clients
 - 100 Continue status code and, 59
 - freshness constraints, 185
 - Host header requirements, 418
 - identification, 257–276
 - fat URLs, using for, 262
 - IP addresses, using for, 259
 - user logins, using for, 260 (see also cookies)
 - supported character sets, 375
- client-side gateways, 199
 - security accelerator gateways, 202
- client-side state, 265
- code width, 377
- coded character sets, 377, 379
- coded characters, 376
- coding space, 376
- collections, 439
- collisions, one-way digests, 288
- colon (:), use in headers, 47
- Combined Log Format, 485
- Common Gateway Interface (CGI), 203, 204
- Common Log Format, 484
- composite MIME types, 534
- conditional requests, 362
 - headers, 70, 178–181
- configuration URLs (CURLs), 465
- CONNECT method, 206–208, 336
- connection handshake delays, 82
- Connection headers, 86, 512
- connections (see HTTP connections)
- content distribution networks (CDNs), 421

- content encodings, 345, 351–354
- content injection, 405
- content negotiation, 395–403
 - on Apache web servers, 399
 - client-driven negotiation, 396
 - headers, 397
 - other protocols and, 405
 - performance limitations, 405
 - quality values, 398
 - server-driven negotiation, 397–400
 - techniques, 395
 - transparent negotiation, 400–403
- content routers, 134, 170
- Content-Base headers, 513
- Content-Encoding headers, 513
- content-injection, 405
- Content-Language headers, 371, 384, 513
- Content-Length headers, 344–347, 514
 - content encoding and, 345
 - persistent connections and, 345
- Content-Location headers, 514
- Content-MD5 headers, 347, 514
- Content-Range headers, 515
- Content-Type headers, 348–351, 515
 - character encodings, 349
 - charset parameter, 371
 - META tags and, 375
 - MIME charset encoding tags and, 374
 - multipart form submissions, 349
- continuation lines, in headers, 51
- Continue status code (100), 59–60
- Cookie headers, 516
- Cookie2 headers, 516
- cookies, 263–276
 - browsers, storage on, 264
 - caching, 273
 - domain attributes, 267
 - functioning, 264
 - information contained in, 264
 - Path attributes, 268
 - privacy and, 275
 - security and, 275
 - session tracking, 272
 - Set-Cookie2 headers, 271
 - specifications, 268
 - third-party vendors, use by, 267
 - types, 264
 - Version 0, 269
 - Version 1, 270–272
 - headers, 272
 - version negotiation, 272
 - web site specificity of, 266
- COPY method, 442
- country codes, 388
- country tokens, 388
- crawlers, 215–224
 - aliasing, 219
 - canonicalizing of URLs, 220
 - checkpoints, 219
 - cycles, avoiding, 217–218, 222–224
 - dups, 218
 - filesystem link cycles, 220
 - hash tables, 218
 - loops, 217
 - lossy presence bit maps, 218
 - partitioning, 219
 - root set, 216
 - search trees, 218
 - tracking of visited sites, 218
 - traps, 220–224
- CRLF, 44
 - in entities, 343
- cryptographic checksums, 289
- cryptography, 309–317
 - asymmetrically keyed ciphers, 315
 - cipher machines, 311
 - ciphers, 310–315
 - digital, 311
 - ciphertext, 310
 - cleartext, 310
 - enumeration attacks, 313
 - hybrid cryptosystems, 317
 - keyed ciphers, 311
 - keys, 311, 312
 - key length, 313
 - sharing, logistical aspects of, 315
 - public-key cryptography, 315–317
 - computation speed, 317
 - digital signing with, 318
 - RSA algorithm, 317
 - symmetric-key ciphers, 313
- CURLs (configuration URLs), 465
- cycles, avoiding (web robots), 217–218, 222–224
 - filesystem link cycles, 220

D

- data formats, conversion, 135
- date formats, 392
- Date headers, 516
- DAV headers, 431
 - compliance classes, 445
- decomposing of URLs, 33

- dedicated web hosting, 412
- delayed acknowledgements, 83
- DELETE method, 58, 441
- delta encodings, 359, 365–367
 - server disk space and, 368
- delta generators and appliers, 368
- <depth> element, 434
- Depth headers, 431
- Destination headers, 431
- dictionary attacks, 304
- digcalc.c file, 578
- digcalc.h file, 577
- digest authentication, 286–306, 574–580
 - algorithms, 291–295
 - input data, 291
 - authentication process, 287
 - Authentication-Info directives, 576
 - Authorization directives, 575
 - basic authentication, compared to, 286
 - caching and, 302
 - digest calculations, 291
 - error handling, 301
 - H(A1) and H(A2) reference code, 577
 - handshakes, 290
 - headers, 300
 - MD5 and MD5-sess, 291
 - message-related data (A2), 293, 298
 - nonces, 289
 - password files, vulnerabilities of, 305
 - preemptive authorization, 296
 - protection space, 302
 - request and response digest reference code, 577
 - revalidating a session, 295
 - rewriting URIs, 302
 - security, 286
 - security-related data (A1), 293
 - session, 295
 - symmetric authentication, 298
 - WWW-Authenticate directives, 574–575
 - (see also authentication)
- digests, 288
 - algorithm input data, 291
 - collisions, 288
- digital certificates, 319–322
 - public-key cryptography and, 320
 - server authentication, use for, 321
 - universal standard, lack of, 320
 - virtual hosting and, 328
 - X.509v3 certificates, 320
- digital cryptography (see cryptography)
- digital signatures, 317–319
 - example, 318
- digtest.c file, 580
- directory listings, 122
 - disabling, 123
- discrete MIME types, 534
- distance delays, 163
- distributed objects (HTTP:NG), 249
- DNS
 - caching, 456
 - DNS A record lookup, 467
 - redirection, 453–457
 - enhanced algorithms for, 457
 - multiple addresses and round-robin address rotation, 455
 - resolvers, 453
 - round robin, 454, 455
 - load balancing with, 456
- docroots (document roots), 120
 - private, 122
 - user home directory, 122
 - virtually hosted, 121
- document access control, 132
- document expiration, 175
 - setting, 182
- document hit rate, 167
- document roots (see docroots)
- documents
 - age and freshness lifetime, 188
 - age-calculation algorithms, 189–194
 - caching, preventing, 182
 - freshness and aging algorithms, 187–194
 - heuristic expiration, 184
- Domain Name Service (see DNS)
- domain names, internationalization of, 392
- downstream message flow, 44
- dups (web robots), 218
- dynamic content resource mapping, 123

E

- egress proxies, 137
- embedded web servers, 111
- encodings, 372
 - chunked encodings, 345
 - content encodings, 345, 351–354
 - delta encodings, 359, 365–367
 - impact on server disk space, 368
 - fixed-width, 381
 - transfer encodings, 354–359
- end-of-line sequence, 44

- entities, 342
 - body length, determining, 346
 - Content-Length headers, 344–347
 - CRLF line, 343
- entity bodies, 44, 47, 52, 343
 - MIME types, 348
- entity digests, 347
- entity headers (see under headers)
- entity tags (see ETags)
- enumeration attacks, 313
- equals sign (=), base-64 encoding, 572
- escape sequences, 36
- ETags (entity tags), 180, 298
 - headers, 517
 - using, 181
- euc-jp encoding, 383
- Expect headers, 517
- experimental MIME types, 569
- expiration of documents, setting, 182
- Expires headers, 175, 183, 517
- explicit MIME typing, 126
- Extensible Markup Language (see XML)
- extension APIs, 205
- extension headers, 51, 68
- extension methods, 58

F

- Fast CGI, 205
- fat URLs, 262
 - limitations of, 263
- file scheme, 39
- filesystem link cycles, 220
- FindProxyForURL() method, 143
- fingerprint functions, 289
- first subtag, 387
- fixed-width encodings, 381
- flash crowds, 163
- format conversion, 404
- FPAdminScriptUrl, 426
- FPAuthorScriptUrl, 426
- FPShtmlScriptUrl, 426
- Fpsrvadm, 428
- frag or fragment component, URLs, 30
- freshness and aging algorithms, 187–194
- freshness lifetime, 188
- From headers, 258, 517
 - robots and, 225
- FrontPage, 424–429
 - client and server extension communication, 426

- FrontPage Server Extensions (FPSE), 424
- HTTP POST requests and, 425
- listExploreDocs element, POST request
 - body, 427
- listHiddenDocs element, POST request
 - body, 427
- root web, 425
- RPC protocol, 426
- security, 428
- server administrator utility (Fpsrvadm), 428
- service_name element, POST request
 - body, 427
- subweb, 426
- virtual servers, 425
- ftp scheme, 39
- full NAT, 461
- full-text indexes, 243

G

- gateways, 19, 197–205
 - client- and server-side, 199
 - client-side security accelerator gateways, 202
 - examples, 198
 - protocol gateways, 200
 - proxies, contrasted with, 130
 - resource gateways, 203
 - server-side security gateways, 202
 - server-side web gateways, 200
 - Via headers and, 153
- general headers (see under headers)
- general-purpose software web servers, 110
- Generic Router Encapsulation (GRE), 472
- GET command, virtual hosting issues, 414
- GET messages, processing steps, 171–175
- GET method, 53
- getpeername function, 259
- glyphs, 378
- GRE (Generic Router Encapsulation), 472

H

- H(A1) and H(A2) reference code, 577
- half NAT, 461
- handshake delays, 82
- handshakes, digest authentication, 290
- hash tables, 218
- H(d), one-way hash, 291
- HEAD method, 54
- HEAD response, 346

- headers, 47, 51, 67–73, 508–532
 - Accept, 69, 508
 - robots and, 225
 - Accept-Charset, 371, 375, 509
 - MIME charset encoding tags and, 374
 - Accept-Encoding, 509
 - Accept-Language, 371, 385, 510
 - content negotiation and, 398
 - Accept-Ranges, 510
 - Age, 510
 - Allow, 159, 511
 - authentication, 278
 - basic, 281, 300
 - digest, 300
 - Authentication-Info directives, 576
 - Authorization, 281, 511
 - directives, 575
 - preemptive generation, 295
 - Cache-Control, 175–177, 182–186, 189, 511
 - directives, 361
 - character set requirements, 392
 - classification, 51
 - client identification using, 258
 - Client-ip, 258, 260, 512
 - Connection, 86, 512
 - content negotiation, 397
 - Content-Base, 513
 - Content-Encoding, 513
 - Content-Language, 371, 384, 513
 - Content-Length, 344–347, 514
 - Content-Location, 514
 - Content-MD5, 347, 514
 - Content-Range, 515
 - Content-Type, 348–351, 515
 - charset parameter, 371
 - continuation lines, 51
 - Cookie, 516
 - Cookie2, 516
 - Date, 516
 - DAV, 431
 - compliance classes, 445
 - Depth, 431
 - Destination, 431
 - entity headers, 51, 67, 72
 - content headers, 72
 - entity caching headers, 73
 - HTTP/1.1, 342
 - ETag, 517
 - examples, 51
 - Expect, 517
 - Expires, 175, 183, 517
 - extension headers, 68
 - From, 517
 - robots and, 225
 - general headers, 51, 67, 68
 - caching headers, 68
 - Heuristic Expiration Warning, 184
 - Host, 417, 418, 419, 518
 - robots and, 225
 - HTCP cache headers, 480
 - If, 431
 - If-Match, 519
 - If-Modified-Since, 166, 178, 518
 - If-None-Match, 180, 519
 - If-Range, 519
 - If-Unmodified-Since, 520
 - Last-Modified, 520
 - Location, 520
 - Lock-Token, 431
 - max-age, 183
 - Max-Forwards, 155, 521
 - for media types, 348
 - Meter, 196, 493
 - MIME-Version, 521
 - must-revalidate, 183
 - no-cache, 182
 - no-store, 182
 - Overwrite, 432, 442
 - Pragma, 68, 182, 521
 - Proxy-Authenticate, 522
 - Proxy-Authorization, 522
 - Proxy-Connection, 96, 523
 - Public, 523
 - Range, 524
 - Referer, 259, 524
 - robots and, 225
 - request headers, 51, 67, 69–71
 - accept headers, 69
 - client identification using, 258
 - conditional request headers, 70
 - proxy request headers, 70
 - request security headers, 70
 - response headers, 51, 67, 71–72
 - negotiation headers, 71
 - response security headers, 72
 - Retry-After, 525
 - Server, 525
 - Set-Cookie, 264, 525
 - caching and, 273
 - domain attributes, 267
 - Set-Cookie2, 271, 526

- headers (*continued*)
 - syntax, 51
 - tampering attack, 303
 - TE, 526
 - Timeout, 432, 435
 - Title, 527
 - Trailer, 526
 - Transfer-Encoding, 527
 - UA-, 527
 - for uncacheable documents, 182
 - unsupported headers, handling, 158
 - Upgrade, 528
 - User-Agent, 225, 259, 528
 - Vary, 402, 529
 - Via, 151–154, 529
 - Want-Digest, 348
 - Warning, 530
 - Warning 13, 184
 - for WebDAV, 431
 - WWW-Authenticate, 281, 531, 574–575
 - X-Cache, 531
 - X-Forwarded-For, 260, 531
 - X-Pad, 531
 - X-Serial-Number, 532
- heartbeat messages, 473
- heuristic expiration of documents, 184
- Heuristic Expiration Warning headers, 184
- history expansion, 34
- hit logs, 195
- hit metering, 492–494
 - Meter headers, 493
- hit rate, 167
- host component, URLs, 27
- Host headers, 417, 418, 518
 - clients, requirements for, 418
 - missing host headers, 419
 - proxies and, 418, 419
 - robots and, 225
 - web servers, interpretation by, 419
- hostile proxies, 304
- hosting services, 411
 - dedicated web hosting, 412
- hostname expansion, 34
- hostnames, 13
- <href> element, 438
- .htaccess, 428
- HTCP (Hyper Text Caching Protocol), 478–481
 - authentication, 480
 - cache headers, 480
 - caching policies, setting, 480
 - data components, 479
 - message structure, 478
 - opcodes, 480
- HTML (Hypertext Markup Language)
 - displaying resources using HTTP, 13
 - documents, relative URLs in, 31
 - fragments, referencing, 30
 - robot-control META tags, 237
- HTTP (Hypertext Transfer Protocol), xiii, 3–11, 247
 - authentication, challenge/response framework, 278
 - (see also authentication)
 - authentication schemes, security risks, 303
 - base-64 encoding, compatibility with, 570
 - caching (see caches)
 - character sets (see character sets)
 - clients and servers, 4
 - commands, 8
 - CONNECT method, 206–208
 - connections (see HTTP connections)
 - entities (see entities)
 - headers (see headers)
 - hit metering extension, 492
 - HTTP-NG (see HTTP-NG)
 - informational resources, 21
 - instance manipulations, 359
 - international content support, 370
 - limitations, 248
 - messages (see HTTP messages)
 - methods, 8
 - performance considerations (see under TCP)
 - proxy servers (see HTTP proxy servers)
 - redirection, 452–453
 - relays, 212
 - reliability of, 3
 - revalidations, 165–166
 - robots, standards for, 225
 - secure HTTP (see HTTPS)
 - status codes, 9, 505–507
 - TCP, dependency on, 80
 - textual basis of, 10
 - transactions, 8
 - delays, causes of, 80
 - truncation detection, 344
 - versions, 16
- HTTP connections, 74, 75, 86–104
 - closing, 101–104
 - Connection headers, 86, 512
 - establishing, 13, 15–16

- keep-alive connections
 - (HTTP/1.0+), 91–96
 - parallel (see parallel connections)
 - persistent (see persistent connections)
 - pipelined connections, 99
 - Proxy-Connection headers, 96, 523
 - serial loading, 87
 - (see also TCP)
 - HTTP messages, 8, 10, 43–73
 - entity bodies, 47, 52
 - example, 11
 - flow, 43
 - GET messages, processing steps, 171–175
 - headers (see headers)
 - methods (see methods)
 - reason phrases, 47, 50
 - redirection of, 450
 - request lines, 48
 - robots, setting conditions for, 226
 - request URLs, 46
 - response lines, 48
 - robots, handling by, 227
 - start lines, 47–51
 - status codes (see status codes)
 - structure, 44
 - syntax, 45
 - tracing across proxies, 150–157
 - Via headers, 151–154
 - Version 0.9 messages, 52
 - versions, 46, 50
 - HTTP proxy servers, 129–160
 - authentication, 156
 - client proxy configuration, 141–144
 - client traffic acquisition, 140
 - deploying, 137
 - interoperation, 157–160
 - messages, tracing, 150–157
 - proxy and server requests, handling, 146
 - proxy hierarchies, 138
 - public and private proxies, 130
 - TRACE method and network
 - diagnosis, 155
 - unsupported headers and methods,
 - handling, 158
 - URIs, in-flight modification of, 147
 - uses, 131–136
 - http scheme, 38
 - HTTP State Management Mechanism, 265
 - HTTP: The Next Generation (see HTTP-NG)
 - HTTP/0.9, 16
 - HTTP/1.0, 16
 - server requests to virtual hosts, problems
 - with, 413
 - HTTP/1.0+, 17
 - HTTP/1.1, 17
 - enhanced methods, 444
 - entity header fields, 342
 - Host headers (see Host headers)
 - limitations, 248
 - request pipelining, 99
 - TRACE method, 155
 - HTTP/2.0, 17
 - HTTP-EQUIV tag, 187
 - HTTP-NG (HTTP: The Next Generation), 17, 248–253
 - current status, 252
 - message transport layer, 249, 250
 - modularization, 248
 - object types, 252
 - remote invocation layer, 249, 250
 - web application layer, 249, 251
 - HTTPS, 76, 308–309, 322–336
 - authentication, 326
 - clients, 328–335
 - OpenSSL example, 329–335
 - CONNECT method, HTTP, 206–208, 336
 - connecting, 324
 - default port, 323
 - OpenSSL, 328–335
 - schemes, 308, 323
 - site certificate validation, 327
 - SSL handshake, 324–326
 - tunnels (see tunnels)
 - https scheme, 38
 - Hyper Text Caching Protocol (see HTCP)
 - Hypertext Markup Language (see HTML)
 - Hypertext Transfer Protocol (see HTTP)
- I**
- IANA (Internet Assigned Numbers Authority)
 - instance manipulations, registered
 - types, 367
 - MIME charset registry, 602–615
 - MIME type registration, 537–539
 - registered language tags, 386, 582
 - ICP (Internet Cache Protocol), 473
 - vs. CARP, 475
 - ident protocol, 115

- If headers, 431
- If-Match headers, 519
- If-Modified-Since headers, 166, 178, 518
- If-None-Match headers, 180, 519
- If-Range headers, 519
- If-Unmodified-Since headers, 520
- image/* MIME types, 561–562
- inbound messages, 43
- inbound proxies, 138
- indexes, full-text, 243
- informational status codes (100-199), 59–60, 505
- ingress proxies, 137
- instance manipulations, 359, 367–369
 - delta encodings, 365
 - IANA registered types, 367
 - range requests, 364
- integrity protection, 299
- intercepting proxies, 140, 146
 - URI resolution with, 149
- internationalization
 - date formats, 392
 - domain names, 392
 - headers, character set for, 392
 - ISO 3166 country codes, 594–600
 - ISO 639 language codes, 583–594
 - languages, administrative
 - organizations, 601
 - URL variants, 395
- Internet Assigned Numbers Authority (see IANA)
- Internet Cache Protocol (ICP), 473
- Internet search engines (see search engines)
- IP address forwarding, 460
- IP (Internet protocol) addresses, 13
 - clients, identification using, 259
 - virtual hosting and, 416
- IP MAC forwarding, 459
- IP packets, 76
- iPlanet web servers, 110
- ISO 3166 country codes, 594–600
- ISO 639 language codes, 583–594
- iso-2022-jp encoding, 382
- iso-8859 character set, 380

J

- Japanese encodings, 382, 383
- JIS X 0201, 0208, and 0212 character sets, 380
- Joe's Hardware Store web site, xiv

K

- KD(s,d) digest, 291
- keep-alive connections (HTTP/1.0+), 91–96
 - (see also persistent connections)
- keyed ciphers, 311
- keys, 311, 312
 - key length, 313

L

- language preferences, configuring, 389
- language tags, 370, 384, 581–600
 - case sensitivity, 386
 - first subtag rules, 581
 - IANA-registered tags, 386, 582
 - reference tables, 389
 - second subtag rules, 582
 - subtags, 386
 - syntax, 385
- Last-Modified dates, using, 181
- Last-Modified headers, 520
- layering of protocols, 12
- layout delay, preventing, 88
- ligatures, 378
- listExploreDocs element, POST request
 - body, 427
- listHiddenDocs element, POST request
 - body, 427
- LM-Factor algorithm, 184
- load balancing, 449
 - DNS round robin, 454–457
 - single clients and, 456
- loading, serial, 87
- Location headers, 520
- LOCK method, 433
 - status codes, 436
- lock refreshes, 435
- <lockdiscovery> element, 435
- <lockinfo> element, 434
- locking, 433
- <locktoken> element, 434
- Lock-Token headers, 431
- logging, 483–492
 - commonly logged fields, 483
 - interpretation, 484
 - log formats, 484–492
 - Combined Log Format, 485
 - Common Log Format, 484
 - Netscape Extended 2 Log Format, 487–489

- Netscape Extended Log Format, 486
- Squid Proxy Log Format, 489–492
- privacy concerns, 495
- loops (web robots), 217

M

- MAC (Media Access Control) addresses, 459
- magic typing, 126
- mailto scheme, 38
- man-in-the-middle attacks, 304
- manual client proxy configuration, 142
- master origin server, 420
- max-age response headers, 183
- Max-Forwards headers, 155, 521
- MD5, 288, 291, 293, 347
- MD5-sess, 291, 293
- Media Access Control (MAC) addresses, 459
- media types, 348
 - multipart, 349
- message body, 44
- message digest algorithms, 291–294
 - symmetric authentication, 298
- message integrity protection, 299
- message/* MIME types, 563
- message transport layer (HTTP-NG), 249, 250
- message truncation, 344
- messages (see HTTP messages)
- <META HTTP-EQUIV> tag, 187
- META tag directives, 239
- meta-information, 43
- Meter headers, 196, 493
- methods, 46, 48, 53–59
 - CONNECT, 206–208, 336
 - DELETE, 58, 441
 - extension methods, 58
 - GET, 53
 - HEAD, 54
 - OPTIONS, 57, 159, 445
 - POST, 55
 - PUT, 54, 444
 - redirection status codes and, 64
 - TRACE, 55
 - unsupported, handling, 158
- Microsoft FrontPage (see FrontPage)
- Microsoft Internet Explorer
 - cookie storage, 266
 - language preference configuration, 389
- Microsoft web servers, 110

- MIME (Multipurpose Internet Mail Extensions)
 - charset encoding tags, 374
 - charset registry, 602–615
 - preferred MIME names, 603
 - “multipart” email messages, 349 (see also MIME types)
- MIME types, 5, 533–569
 - application/* types, 540–557
 - audio/* types, 557–559
 - chemical/* types, 559–561
 - composite types, 534
 - discrete types, 534
 - documentation, 534
 - experimental types, 569
 - IANA registration, 537–539
 - media type registry, 539
 - process, 537
 - registration trees, 537
 - rules, 538
 - template, 538
 - image/* types, 561–562
 - message/* types, 563
 - model/* types, 563
 - multipart/* types, 535, 564
 - primary types, 536
 - structure, 534
 - syntax, 536
 - tables, 539–569
 - text/* types, 565–568
 - video/* types, 568
- MIME typing, 125
- MIME::Base64 Perl module, 572
- MIME-Version headers, 521
- mirrored server farms, 420
- MKCOL method, 440
- mod_cern_meta module, Apache web server, 186
- model/* MIME types, 563
- mod_expires module, Apache web server, 186
- mod_headers module, Apache web server, 186
- MOVE method, 442
- multi-homed servers, 425
- multipart form-data encodings, 349
- multipart/* MIME types, 535, 564
- multiplexed architectures, 119
 - I/O web servers, 119
 - multithreaded web servers, 119

- multiprocess, multithreaded web servers, 118
- Multipurpose Internet Mail Extensions (see MIME; MIME types)
- <multistatus> element, 438
- MultiViews directive, 400
- must-revalidate response headers, 183

N

- Nagle's algorithm, 84
- namespace management, 439–444
 - methods used for, 440
 - status codes, 443
- namespaces, 388
 - language subtags, 387–389
 - XML, 430
- NAT (Network Address Translation), 460
- NECP (Network Element Control Protocol), 461
- negotiation headers, 71
- Netscape Extended 2 Log Format, 487–489
- Netscape Extended Log Format, 486
- Netscape Navigator
 - cookies
 - storage, 265
 - Version 0, 269
 - language preference configuration, 389
- Network Address Translation (NAT), 460
- network bottlenecks, 161
- Network Element Control Protocol (NECP), 461
- network exchange proxies, 137
- news scheme, 39
- no-cache response headers, 182
- nonces, 289–298
 - next nonce pregeneration, 297
 - reuse, 297
 - selection, 298
 - time-synchronized generation, 297
- no-store response headers, 182
- .nsconfig, 428

O

- object types, HTTP-NG, 252
- one-way digests, 288
- one-way hashes, 291
 - functions, 289
- opaquelocktoken scheme, 433, 434
- OpenSSL, 328–335
 - example client, 329–335

- OPTIONS method, 57, 445
 - requests, 159
 - response headers to, 445
- origin servers, 420
- outbound messages, 43
- outbound proxies, 138
- Overwrite headers, 432, 442

P

- PAC files, 142
 - autodiscovery, 465
- PAC (Proxy Auto-Configuration) protocol, 463
- parallel connections, 88–90
 - impression of speed, 90
 - loading speed, 88
 - open connection limits, 90
 - persistent connections vs., 91
- parameters component, URLs, 28
- parent and child relationships, 138
- parent caches, 169
- password component, URLs, 27
- passwords
 - digest authentication password file, risks, 305
 - digest authentication, security, 287
- path component, URLs, 28
- Perl code for interaction with robots.txt files, 235
- Perl web server, 111
- persistent connections, 90–99
 - Content-Length headers and, 345
 - keep-alive connections (HTTP/1.0+), 91–96
 - parallel connections vs., 91
 - restrictions and rules, 98
- persistent uniform resource locators (PURLs), 40
- pipelined connections, 99
- plaintext, security and, 310
- port component, URLs, 27
- port exhaustion, 85
- port numbers, 13, 77
 - default values, 13
 - virtual hosting and, 415
- POST method, 55
- POST requests, FrontPage and, 425
- Pragma headers, 68, 182, 521
 - Pragma: no-cache headers, 182
- precompiled dictionary attacks, 305

- preemptive authorization, 295
- presence bit arrays (web robots), 218
- presentation forms, 378
- primary subtags, 386
- privacy, 495
 - cookies and, 275
 - robots and, 229
- private caches, 168
- private docroots, 122
- private proxies, 130
- <prop> element, 437
- <propertyupdate> element, 438
- PROPFIND method, 437
 - server response elements, 438
 - XML elements, used with, 437
- <propname> element, 437
- PROPPATCH method, 438–439
 - XML elements, used with, 438
- <propstat> element, 438
- “protecting the header”, 87
- protection spaces, 295, 301
- protocol gateways, 200
- protocol stack, 76
- protocols, layering of, 12
- proxies
 - 100 Continue status code and, 60
 - authentication, 156
 - deploying, 137
 - egress proxies, 137
 - gateways, contrasted with, 130
 - hostile, 304
 - HTTP proxies (see HTTP proxy servers)
 - ingress proxies, 137
 - intercepting proxies, 140, 146
 - URI resolution with, 149
 - interoperation, 157–160
 - messages, tracing, 150–157
 - “missing scheme/host/port”
 - problem, 146
 - proxy and server requests, handling, 146
 - proxy hierarchies, 138
 - content routing, 139
 - dynamic parent selection, 140
 - public and private, 130
 - redirection and, 449
 - surrogates, 146
 - transparent negotiation, 400
 - transparent proxies, 140
 - tunnels and, 335
 - URIs, in-flight modification of, 147
 - uses, 131–136
 - (see also HTTP proxy servers)
- Proxy-Authenticate headers, 522
- proxy authentication, 283
- Proxy-Authorization headers, 522
- Proxy Auto-configuration (PAC)
 - protocol, 142, 463
 - (see also PAC files)
- proxy cache hierarchies, 169
- proxy caches, 169
- Proxy-Connection headers, 96, 523
- proxy redirection
 - CARP (Cache Array Routing Protocol), 475–478
 - HTCP (Hyper Text Caching Protocol), 478–481
 - ICP (Internet Cache Protocol), 473
- proxy redirection methods, 462–469
 - explicit browser configuration, 463
 - disadvantages, 463
 - PAC (Proxy Auto-configuration)
 - protocol, 463
 - WPAD (see WPAD)
- proxy request headers, 70
- proxy servers, 18
 - networks, use in securing, 335
 - (see also HTTP proxy servers)
- proxy URIs vs. server URIs, 144
- public caches, 169
- Public headers, 523
- public proxies, 130
- public-key cryptography (see under cryptography)
- publishing systems
 - FrontPage (see FrontPage)
 - WebDAV (see WebDAV)
 - (see also web publishing)
- PURLs (persistent uniform resource locators), 40
- PUT method, 54, 444

Q

- qop (quality of protection), 293
 - enhancements, 299
- quality factors, 371
- quality of protection (see qop)
- quality values, 398
- query component, URLs, 29

R

- Range headers, 524
- range requests, 363
- realm directive, 280
- realms (protection spaces), 301
- reason phrases, 9, 47, 50, 505–507
- redirection, 126, 448–481
 - anycast addressing, 457
 - enhanced DNS-based algorithms, 457
 - IP address forwarding, 460
 - IP MAC forwarding, 459
 - load balancing and, 127, 449
 - methods, 450
 - DNS redirection, 453–457
 - HTTP redirection, 452–453
 - proxy methods, 462–469
 - proxy techniques, 451
 - NECP (Network Element Control Protocol), 461
 - protocols, 450
 - proxies, role in, 449
 - purpose, 449
 - techniques, 448
 - temporary redirect, 126
 - transparent redirection, 469
 - URL augmentation, 126
 - WCCP (Web Cache Coordination Protocol), 470–473
- redirection status codes (300–399), 61–64, 505
- Referer headers, 259, 524
 - robots and, 225
- relative URLs, 30–34
 - bases, 31
 - resolving, 33
- relays, 212
 - keep-alive connections and, 212
- relevancy ranking, 245
- reliable bit pipe, 75
- remote invocation layer (HTTP-NG), 249, 250
- remote procedure call (RPC) protocol, FrontPage, 426
- <remove> element, 439
- replay attacks, 284, 289, 303
 - preventing, 289
- replica origin servers, 420
- request digest reference code, 577
- request digests, 294
- request headers (see under headers)
- request lines, 48
- request messages, 10, 45, 47
 - methods, 46
 - request URLs, 46
- request method (HTTP), 294
- request pipelining, 99
- request security headers, 70
- reserved characters (see restricted characters)
- resource gateways, 203
- resource locator servers, 40
- resource paths, 24
- resources, mapping and accessing of, 120
- response digest reference code, 577
- response entities, 125
- response headers (see under headers)
- response lines, 48
- response messages, 10, 45, 125
- restricted characters, 36
- Retry-After headers, 525
- revalidate hits, 165
- revalidate misses, 166
- revalidations, 165–166
- reverse proxies, 134, 137
- RobotRules object, 235
- robots
 - conditional HTTP requests, 226
 - entities and, 227
 - etiquette, 239–241
 - excluding from web sites, 229–239
 - HTML robot control META tags, 237
 - HTTP and, 225
 - request headers, identifying, 225
 - META directives, 237
 - Meta HTML tags and, 227
 - privacy and, 229
 - problems caused by, 228
 - response handling, 227
 - search engines, 242–246
 - status codes, handling of, 227 (see also robots.txt files)
- Robots Exclusion Standard, 230
- robots.txt files, 229
 - caching and expiration, 234
 - comments, 234
 - disallow and allow lines, 233
 - example, 236
 - fetching, 231
 - records, 232
 - specification, changes in, 234
 - status codes for retrievals, 231
 - syntax, 232
 - User-Agent line, 232
 - web sites and, 231

- root set, 216
- root web, 425
- round-robin load balancing, 453
 - DNS round robin, 454–457
- routers and anycast addressing, 457
- RPC protocol, FrontPage, 426
- RSA algorithm, 317
- rtsp, rtspu schemes, 39

S

- schemes, 7, 24, 27
 - common formats, 38
 - URLs for, 499–504
- search engines, 242–246
 - architecture, 242
 - full-text indexes, 243
 - queries, 244
 - relevancy ranking, 245
 - results, sorting and formatting, 244
 - spoofing, 245
- search trees, 218
- second subtag, 388
- Secure Sockets Layer (see SSL)
- security
 - basic authentication and, 283
 - cookies and, 275
 - digest authentication and, 286
 - firewalls, 132
 - FrontPage, security model, 428
 - HTTP authentication schemes, associated risks, 303
 - key length and, 314
 - multiple authentication schemes, risks, 303
 - security realms, 280
 - WPAD security hole, 468
- segments, 76
- “sender silly window syndrome”, 84
- serial loading, 87
- serial transaction delays, 87
- server error status codes (500-599), 66, 505
- server farms, 413
 - mirrored servers, 420
- Server headers, 525
- Server response header field, 154
- server URIs vs. proxy URIs, 144
- server-driven negotiation, 397–400
 - server-side extensions, 400
- servers
 - 100 Continue status codes and, 60
 - accelerators, 134
 - certificates, 326
 - delta encodings, impact on, 368
 - error status codes (500-599), 66
 - extension APIs, 205
 - FrontPage Server Extensions (FPSE), 424
 - Host headers, interpreting, 419
 - multi-homed servers, 425
 - revalidation, 177
 - server farms, 413
 - master origin servers, 420
 - replica origin servers, 420
 - supported functionality, identifying, 159
 - validation, 175
- server-side extensions, 400
- server-side gateways, 199
 - security gateways, 202
 - web gateways, 200
- server-side includes (SSIs), 124
- service groups, 472
- service_name element, POST request body, 427
- sessions, cookies and, 264
 - tracking with, 272
- <set> element, 438
- Set-Cookie headers, 264, 525
 - caching and, 273
 - domain attributes, 267
- Set-Cookie2 headers, 271, 526
- shared hosting, 413–419
- shared keys, 315
- shared proxies, 130
- sibling caches, 171
- Simple Object Access Protocol (SOAP), 206
- single-threaded web servers, 118
- site certificate validation, 327
- slow hits, 165
- s-maxage response headers, 183
- SOAP (Simple Object Access Protocol), 206
- sockets API, 78
 - calls, 78
- software web servers, 110
- spiders, 215
- spoofing, 245
- Squid Proxy Log Format, 489–492
- SSIs (server-side includes), 124
- SSL (Secure Sockets Layer), 308
 - authentication, 326
 - handshakes, 324–326
 - HTTPS, integration in, 323
 - OpenSSL, 328–335
 - site certificate validation, 327
 - tunnels, 209
 - vs. HTTP/HTTPS gateways, 210

- SSLay, 329
 - (see also OpenSSL)
- start lines, 47–51
- status codes, 9, 49, 59–67
 - classes, 49
 - client error codes (400-499), 65–66, 505
 - HTTP codes, 505–507
 - informational status codes (100-199), 59–60, 505
 - LOCK method, 436
 - namespace management methods, 443
 - redirection status codes (300-399), 61–64, 505
 - robots, handling by, 227
 - server error codes (500-599), 66, 505
 - success status codes (200-299), 61, 505
 - UNLOCK method, 436
- <status> element, 438
- strong validators, 181, 363
- subtags, 386, 389
 - first subtag, 387
 - second subtag, 388
- subweb, 426
- success status codes (200-299), 61, 505
- surrogate caches, 421
- surrogate proxies, 137
- surrogates, 134, 146
- symbolic links and cycles, 220
- symmetric authentication, 298
- symmetric-key ciphers, 313
- syntax, headers, 51

T

- TCP slow start (see under TCP)
- TCP (Transmission Control Protocol), 74–86
 - connections, 75
 - distinguishing values, 77
 - establishing, 13, 79
 - web server handling of, 115 (see also HTTP connections)
 - network delays, causes, 80
 - performance considerations, 80–86
 - connection handshake delays, 82
 - delayed acknowledgements, 83
 - delays, most common causes, 81–86
 - Nagle's algorithm, 84
 - port exhaustion, 85
 - TCP slow start, 83
 - TCP_NODELAY, 84
 - TIME_WAIT accumulation, 85
 - port numbers and, 77
 - reliability, 74
 - serial loading, 87
 - sockets API, 78
 - TCP slow start, 83
- TCP/IP (Transmission Control Protocol/Internet Protocol), 11
- TCP_NODELAY parameter, 84
- TE headers, 526
- Telnet example, 15–16
- telnet scheme, 40
- text/* MIME types, 565–568
- <timeout> element, 434
- Timeout headers, 432, 435
- TIME_WAIT accumulation, 85
- Title headers, 527
- TLS (Transport Layer Security), 308
 - (see also SSL)
- TRACE method, 55, 155–157
 - Max-Forwards headers and, 155
- Trailer headers, 526
- transactional direction, messages, 43
- transactions, 8
- transcoders, 135
- transcodings, 395, 403–406
 - content injection, 405
 - format conversion, 404
 - information synthesis, 404
 - types, 404
 - vs. static pregenerated content, 405
- transfer encodings, 354–359
- Transfer-Encoding headers, 527
- Transmission Control Protocol/Internet Protocol (TCP/IP), 11
- transparent negotiation, 400–403
 - caching, 401
 - Vary headers, 402
- transparent proxies, 140
- transparent redirection, 469
- Transport Layer Security (see TLS)
- trees, 218
- truncation detection, 344
- tunnels, 19, 206–212
 - authentication, 211
 - HTTPS SSL, 335–336
 - security, 211
 - SSL tunnels vs. HTTP/HTTPS gateways, 210
- type negotiation, 126
- type-map files, 399
- type-o-serve web server, 112

U

- UA- headers, 527
- UCS (Universal Character Set), 381
- uncachable documents, 182
- uniform resource identifiers (see URIs)
- uniform resource locators (see URLs)
- uniform resource names (see URNs)
- Universal Character Set (UCS), 381
- UNLOCK method, 435
 - status codes, 436
- Upgrade headers, 528
- upstream message flow, 44
- uri-directive-value, 294
- URIs (uniform resource identifiers), 6, 24
 - client autoexpansion and hostname resolution, 147
 - intercepting proxies, resolution with, 149
 - internationalization, 389–391
 - resolution, 144–150
 - proxy vs. server, 144
 - with a proxy, explicit, 149
 - without a proxy, 148
 - rewriting, 302
 - schemes, 499–504
- URLs (uniform resource locators), 6, 23–42
 - advantages of, 25
 - aliases, 219
 - augmentation, 126
 - automatic expansion, 34
 - canonical form, 37
 - canonicalizing, 220
 - character sets, 35, 35–38
 - CURLs, 465
 - examples, 13
 - fat URLs, client identification using, 262
 - informational resources, 41
 - portability, 35
 - PURLs, 40
 - relative URLs, 30–34
 - restricted characters, 36
 - schemes, 7
 - schemes (see schemes)
 - shortcuts, 30
 - structure, 24
 - syntax, 26–30
 - URIs, as a subset of, 24
 - variants, 395
 - virtual hosting, paths, 415
- URNs (uniform resource names), 7, 40
 - standardization, 41
- US-ASCII character set, 379

- user agents, 19
- user component, URLs, 27
- user home directory docroots, 122
- User-Agent headers, 225, 259, 528
- user-tracking systems, content injection
 - and, 405
- UTF-8 encoding, 382

V

- validators, 362
 - Last-Modified dates, using, 181
 - strong and weak, 181, 363
- variable-length codes, 372
- variable-width modal encodings, 381
- variable-width nonmodal encodings, 381
- Vary headers, 402, 529
- Vermeer Technologies, Inc., 424
- version numbers, HTTP messages, 50
- Via headers, 151–154, 529
 - gateways and, 153
 - privacy and security, 154
 - request and response paths, 153
 - Server response header fields and, 154
 - syntax, 152
- video/* MIME types, 568
- virtual hosting, 225, 413–419
 - digital certificates and, 328
 - docroots, 121
 - GET command, problems with, 414
 - Host headers, by, 417
 - (see also Host headers)
 - IP addresses, by, 416
 - problems for hosters, 417
 - port numbers, by, 415
 - server requests, absence of host information in HTTP/1.0, 413
 - fixes, 415
 - URL paths, by, 415
- virtual servers, 425

W

- Want-Digest headers, 348
- Warning headers, 530
 - Heuristic Expiration Warning, 184
- WCCP (Web Cache Coordination Protocol), 470–473
 - GRE packet encapsulation, 472
 - heartbeat messages, 473
 - load balancing, 473
 - operation, 470

- WCCP (*continued*)
 - service groups, 472
 - versions, 470
 - WCCP2 messages, 470–472
 - header and components, 471
- weak validators, 181, 363
- web application layer (HTTP-NG), 249, 251
- web architecture, 17–20
- Web Cache Coordination Protocol (see WCCP)
- web caches, 18, 161
 - (see also caches)
- web clients and servers, 4
- Web Distributed Authoring and Versioning (see WebDAV)
- web hosting, 411–422
 - hosting services, 411
 - shared or virtual hosting, 413–419
- web pages, 9
- Web Proxy Autodiscovery Protocol (see WPAD)
- web proxy servers, 129
- web publishing
 - collaborative authoring, 429
 - publishing systems, 424
- web resources, 4
- web robots, 215–246
 - crawlers (see crawlers)
 - examples, 215
 - spiders, 215
- web servers, 109
 - access controls, 124
 - appliances, 111
 - client hostname identification, 115
 - client identification, 257–276
 - cookies, using, 263–276
 - fat URLs, using, 262
 - headers, using, 258
 - IP address, using, 259
 - user login, using, 260
 - connection input/output processing
 - architectures, 117
 - connections, handling new, 115
 - directory listings, 122
 - docroots, 120
 - dynamic content resource mapping, 123
 - embedded web servers, 111
 - explicit typing, 126
 - HTTP proxy servers (see HTTP proxy servers)
 - ident protocol, 115
 - implementations, 109
 - logging, 127
 - MIME typing, 125
 - multiplexed I/O servers, 119
 - multiplexed multithreaded servers, 119
 - multiprocess, multithreaded servers, 118
 - Perl example, 111
 - redirection responses, 126
 - request handling, 449
 - request messages
 - receiving, 116
 - structure of, 117
 - resources, mapping and accessing of, 120
 - response entities, 125
 - response messages, 125
 - responses, sending, 127
 - single-threaded servers, 118
 - software web servers, 110
 - SSIs (server-side includes), 124
 - tasks of, 113–114
 - type negotiation, 126
 - type-o-serve, 112
 - user authentication (see authentication)
- web services, 206
- web sites
 - personalizing of user experience, 257
 - reliability, improving, 419–422
 - mirrored server farms, 420
 - robots, exclusion, 229–239
 - robots.txt files, 231
 - speed, improving, 422
- web tunnels (see tunnels)
- WebDAV (Web Distributed Authoring and Versioning), 429–446
 - collaborative authoring and, 429
 - collections, 439–444
 - DAV header, 431
 - Depth header, 431
 - Destination header, 431
 - enhanced HTTP/1.1 methods, 444
 - headers, 431
 - If header, 431
 - LOCK method, 433
 - locking, 432
 - Lock-Token headers, 431
 - META data, embedding, 436–439
 - methods, 429
 - namespace management, 439–444
 - OPTIONS method, 445
 - Overwrite headers, 432
 - overwrites, preventing, 432

- PROPATCH method, 438–439
- properties, 436–439
- PROPFIND method, 437
- PUT method, 444
- Timeout headers, 432
- UNLOCK method, 435
- version management, 446
- XML and, 430
- WebMUX protocol, 250, 251
- WPAD (Web Proxy Autodiscovery Protocol), 143, 464–469
 - administration, 469
 - CURLs, 465
 - DHCP discovery, 467
 - DNS A record lookup, 467
 - PAC file retrieval, 467
 - PAC file autodiscovery, 465
 - resource-discovery algorithm, 143, 465
 - spoofing, 468

- timeouts, 468
- timing, 468
- WWW-Authenticate headers, 281, 531
 - directives, 574–575
- WWW::RobotRules object, 235

X

- X.509v3 certificates, 320
- X-Cache headers, 531
- X-Forwarded-For headers, 260, 531
- XML (Extensible Markup Language), 206, 430
 - elements used in locking, 434
 - namespace, 430
 - schema definition, XML documents, 430
 - WebDAV and, 430
- X-Pad headers, 531
- X-Serial-Number headers, 532

About the Authors

David Gourley is the Chief Technology Officer of Endeca, where he leads the research and development of Endeca's products. Endeca develops Internet and intranet information-access solutions that provide new ways to navigate and explore enterprise data. Prior to working at Endeca, David was a member of the founding engineering team at Inktomi, where he helped develop Inktomi's Internet search database and was a key developer of Inktomi's web caching products.

David earned a B.A. in Computer Science from the University of California at Berkeley, and he holds several patents in web technologies.

Brian Totty was most recently the Vice President of R&D at Inktomi Corporation (a company he helped found in 1996), where he led research and development of web caching, streaming media, and Internet search technologies. Formerly, he was a scientist at Silicon Graphics, where he designed and optimized software for high-performance networking and supercomputing systems. Before that, he held an engineering position at Apple Computer's Advanced Technology Group.

Brian holds a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign and a B.S. degree in Computer Science and Electrical Engineering from MIT, where he received the Organick award for computer systems research. He also has developed and taught award-winning courses on Internet technology for the University of California Extension system.

Marjorie Sayer writes about network caching software at Inktomi Corporation. After earning M.A. and Ph.C. degrees in Mathematics at the University of California at Berkeley, she worked on mathematics curriculum reform. Since 1990 she has written about energy resource management, parallel systems software, telephony, and networking.

Sailu Reddy currently leads the development of embedded performance-enhancing HTTP proxies at Inktomi Corporation. Sailu has been developing complex software systems for 12 years and has been deeply involved in web infrastructure research and development since 1995. He was a core engineer of Netscape's first web server and web proxy products and of several following generations. His technical experience includes HTTP applications, data compression techniques, database engines, and collaboration management. Sailu earned an M.S. in Information Systems from the University of Arizona and holds several patents in web technologies.

Anshu Aggarwal is a Director of Engineering at Inktomi Corporation. He leads the protocol-processing engineering teams for Inktomi's web caching products, and he has been involved in the design of web technologies at Inktomi since 1997. Anshu holds M.S. and Ph.D. degrees in Computer Science from the University of Colorado at Boulder, specializing in memory-consistent techniques for distributed multiprocessor machines. He also holds M.S. and B.S. degrees in Electrical Engineering. Anshu is the author of several technical papers and holds two patents.

Colophon

Our look is the result of reader comments, our own experimentation, and feedback from distribution channels. Distinctive covers complement our distinctive approach to technical topics, breathing personality and life into potentially dry subjects.

The animal on the cover of *HTTP: The Definitive Guide* is a thirteen-lined ground squirrel (*Spermophilus tridecemlineatus*), common to central North America. True to its name, the thirteen-lined ground squirrel has thirteen stripes with rows of light spots that run the length of its back. Its color pattern blends into its surroundings, protecting it from predators. Thirteen-lined ground squirrels are members of the squirrel family, which includes chipmunks, ground squirrels, tree squirrels, prairie dogs, and woodchucks. They are similar in size to the eastern chipmunk but smaller than the common gray squirrel, averaging about 11 inches in length (including a 5–6 inch tail).

Thirteen-lined ground squirrels go into hibernation in October and emerge in late March or early April. Each female usually produces one litter of 7–10 young each May. The young leave the burrows at four to five weeks of age and are fully grown at six weeks. Ground squirrels prefer open areas with short grass and well-drained sandy or loamy soils for burrows, and they avoid wooded areas—mowed lawns, golf courses, and parks are common habitats.

Ground squirrels can cause problems when they create burrows, dig up newly planted seeds, and damage vegetable gardens. However, they are important prey to several predators, including badgers, coyotes, hawks, weasels, and various snakes, and they benefit humans directly by feeding on many harmful weeds, weed seeds, and insects.

Rachel Wheeler was the production editor and copyeditor for *HTTP: The Definitive Guide*. Leanne Soylemez, Sarah Sherman, and Mary Anne Weeks Mayo provided quality control, and Derek Di Matteo and Brian Sawyer provided production assistance. John Bickelhaupt wrote the index.

Ellie Volckhausen designed the cover of this book, based on a series design by Edie Freedman. The cover image is an original illustration created by Lorrie LeJeune. Emma Colby produced the cover layout with QuarkXPress 4.1 using Adobe's ITC Garamond font.

David Futato and Melanie Wang designed the interior layout, based on a series design by David Futato. Joe Wizda prepared the files for production in FrameMaker 5.5.6. The text font is Linotype Birka; the heading font is Adobe Myriad Condensed; and the code font is LucasFont's TheSans Mono Condensed. The illustrations that appear in the book were produced by Robert Romano and Jessamyn Read using Macromedia FreeHand 9 and Adobe Photoshop 6. This colophon was written by Rachel Wheeler.