HTTP State Management Mechanism

Abstract

   This document defines the HTTP Cookie and Set-Cookie header fields.
   These header fields can be used by HTTP servers to store state
   (called cookies) at HTTP user agents, letting the servers maintain a
   stateful session over the mostly stateless HTTP protocol. Although
   cookies have many historical infelicities that degrade their security
   and privacy, the Cookie and Set-Cookie header fields are widely used
   on the Internet. This document obsoletes RFC 2965.

Status of This Memo

   This is an Internet Standards Track document.

   This document is a product of the Internet Engineering Task Force
   (IETF). It represents the consensus of the IETF community. It has
   received public review and has been approved for publication by the
   Internet Engineering Steering Group (IESG). Further information on
   Internet Standards is available in Section 2 of RFC 5741.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   http://www.rfc-editor.org/info/rfc6265.

   This document may contain material from IETF Documents or IETF
   Contributions published or made publicly available before November
   10, 2008. The person(s) controlling the copyright in some of this
   material may not have granted the IETF Trust the right to allow
   modifications of such material outside the IETF Standards Process.
   Without obtaining an adequate license from the person(s) controlling
   the copyright in such materials, this document may not be modified
   outside the IETF Standards Process, and derivative works of it may
   not be created outside the IETF Standards Process, except to format
   it for publication as an RFC or to translate it into languages other
   than English.

Table of Contents

1. Introduction

     This document defines the HTTP Cookie and Set-Cookie header fields.
     Using the Set-Cookie header field, an HTTP server can pass name/value
     pairs and associated metadata (called cookies) to a user agent. When
     the user agent makes subsequent requests to the server, the user
     agent uses the metadata and other information to determine whether to
     return the name/value pairs in the Cookie header.

     Although simple on their surface, cookies have a number of
     complexities. For example, the server indicates a scope for each
     cookie when sending it to the user agent. The scope indicates the
     maximum amount of time in which the user agent should return the
     cookie, the servers to which the user agent should return the cookie,
     and the URI schemes for which the cookie is applicable.

     For historical reasons, cookies contain a number of security and
     privacy infelicities. For example, a server can indicate that a
     given cookie is intended for "secure" connections, but the Secure
     attribute does not provide integrity in the presence of an active
     network attacker. Similarly, cookies for a given host are shared
     across all the ports on that host, even though the usual "same-origin
     policy" used by web browsers isolates content retrieved via different
     ports.

     There are two audiences for this specification: developers of cookie-
     generating servers and developers of cookie-consuming user agents.

To maximize interoperability with user agents, servers SHOULD limit
themselves to the well-behaved profile defined in Section 4 when
generating cookies.

User agents MUST implement the more liberal processing rules defined
in Section 5, in order to maximize interoperability with existing
servers that do not conform to the well-behaved profile defined in
Section 4.

This document specifies the syntax and semantics of these headers as
they are actually used on the Internet. In particular, this document
does not create new syntax or semantics beyond those in use today.
The recommendations for cookie generation provided in Section 4
represent a preferred subset of current server behavior, and even the
more liberal cookie processing algorithm provided in Section 5 does
not recommend all of the syntactic and semantic variations in use
today. Where some existing software differs from the recommended
protocol in significant ways, the document contains a note explaining
the difference.

Prior to this document, there were at least three descriptions of
cookies: the so-called "Netscape cookie specification" [Netscape],
RFC 2109 [RFC2109], and RFC 2965 [RFC2965]. However, none of these
documents describe how the Cookie and Set-Cookie headers are actually
used on the Internet (see [Kri2001] for historical context). In
relation to previous IETF specifications of HTTP state management
mechanisms, this document requests the following actions:

1. Change the status of [RFC2109] to Historic (it has already been
    obsoleted by [RFC2965]).

2. Change the status of [RFC2965] to Historic.

3. Indicate that [RFC2965] has been obsoleted by this document.

In particular, in moving RFC 2965 to Historic and obsoleting it, this
document deprecates the use of the Cookie2 and Set-Cookie2 header
fields.

2. Conventions

2.1. Conformance Criteria

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in [RFC2119].

Requirements phrased in the imperative as part of algorithms (such as "strip any leading space characters" or "return false and abort these steps") are to be interpreted with the meaning of the key word ("MUST", "SHOULD", "MAY", etc.) used in introducing the algorithm.

Conformance requirements phrased as algorithms or specific steps can be implemented in any manner, so long as the end result is equivalent. In particular, the algorithms defined in this specification are intended to be easy to understand and are not intended to be performant.

2.2. Syntax Notation

This specification uses the Augmented Backus-Naur Form (ABNF) notation of [RFC5234].

The following core rules are included by reference, as defined in [RFC5234], Appendix B.1: ALPHA (letters), CR (carriage return), CRLF (CR LF), CTLs (controls), DIGIT (decimal 0-9), DQUOTE (double quote), HEXDIG (hexadecimal 0-9/A-F/a-f), LF (line feed), NUL (null octet), OCTET (any 8-bit sequence of data except NUL), SP (space), HTAB (horizontal tab), CHAR (any [USASCII] character), VCHAR (any visible [USASCII] character), and WSP (whitespace).

The OWS (optional whitespace) rule is used where zero or more linear whitespace characters MAY appear:

```
OWS            = *( [ obs-fold ] WSP )
                 ; "optional" whitespace
obs-fold       = CRLF
```

OWS SHOULD either not be produced or be produced as a single SP character.

2.3. Terminology

The terms user agent, client, server, proxy, and origin server have the same meaning as in the HTTP/1.1 specification ([RFC2616], Section 1.3).

The request-host is the name of the host, as known by the user agent, to which the user agent is sending an HTTP request or from which it is receiving an HTTP response (i.e., the name of the host to which it sent the corresponding HTTP request).

The term request-uri is defined in Section 5.1.2 of [RFC2616].

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.

fastcase
Smarter legal research.