

HTTPbis Working Group
Internet-Draft
Obsoletes: [2616](#) (if approved)
Updates: [2617](#) (if approved)
Intended status: Standards Track
Expires: May 3, 2012

R. Fielding, Ed.
Adobe
J. Gettys
Alcatel-Lucent
J. Mogul
HP
H. Frystyk
Microsoft
L. Masinter
Adobe
P. Leach
Microsoft
T. Berners-Lee
W3C/MIT
Y. Lafon, Ed.
W3C
J. Reschke, Ed.
greenbytes
October 31, 2011

HTTP/1.1, part 7: Authentication
draft-ietf-httpbis-p7-auth-17

Abstract

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. HTTP has been in use by the World Wide Web global information initiative since 1990. This document is Part 7 of the seven-part specification that defines the protocol referred to as "HTTP/1.1" and, taken together, obsoletes [RFC 2616](#).

Part 7 defines the HTTP Authentication framework.

Editorial Note (To be removed by RFC Editor)

Discussion of this draft should take place on the HTTPBIS working group mailing list (ietf-http-wg@w3.org), which is archived at <http://lists.w3.org/Archives/Public/ietf-http-wg/>.

The current issues list is at <http://tools.ietf.org/wg/httpbis/trac/report/3> and related documents (including fancy diffs) can be found at <http://tools.ietf.org/wg/httpbis/>.

The changes in this draft are summarized in [Appendix C.18](#).

Internet-Draft

HTTP/1.1, Part 7

October 2011

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 3, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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	4
1.1.	Conformance and Error Handling	4

Fielding, et al.

Expires May 3, 2012

[Page 2]

Internet-Draft

HTTP/1.1, Part 7

October 2011

1.2.	Syntax Notation	4
1.2.1.	Core Rules	5
2.	Access Authentication Framework	5
2.1.	Challenge and Response	5
2.2.	Protection Space (Realm)	7
2.3.	Authentication Scheme Registry	7
2.3.1.	Considerations for New Authentication Schemes	8
3.	Status Code Definitions	9
3.1.	401 Unauthorized	9
3.2.	407 Proxy Authentication Required	9
4.	Header Field Definitions	9
4.1.	Authorization	9
4.2.	Proxy-Authenticate	10
4.3.	Proxy-Authorization	11
4.4.	WWW-Authenticate	11
5.	IANA Considerations	12
5.1.	Authenticaton Scheme Registry	12
5.2.	Status Code Registration	12
5.3.	Header Field Registration	12
6.	Security Considerations	12
6.1.	Authentication Credentials and Idle Clients	13
7.	Acknowledgments	13
8.	References	13
8.1.	Normative References	13
8.2.	Informative References	14
Appendix A.	Changes from RFCs 2616 and 2617	14
Appendix B.	Collected ABNF	15
Appendix C.	Change Log (to be removed by RFC Editor before publication)	16
C.1.	Since RFC 2616	16
C.2.	Since draft-ietf-httpbis-p7-auth-00	16
C.3.	Since draft-ietf-httpbis-p7-auth-01	16
C.4.	Since draft-ietf-httpbis-p7-auth-02	16
C.5.	Since draft-ietf-httpbis-p7-auth-03	16
C.6.	Since draft-ietf-httpbis-p7-auth-04	16
C.7.	Since draft-ietf-httpbis-p7-auth-05	17

C.8.	Since draft-ietf-httpbis-p7-auth-06	17
C.9.	Since draft-ietf-httpbis-p7-auth-07	17
C.10.	Since draft-ietf-httpbis-p7-auth-08	17
C.11.	Since draft-ietf-httpbis-p7-auth-09	17
C.12.	Since draft-ietf-httpbis-p7-auth-10	17
C.13.	Since draft-ietf-httpbis-p7-auth-11	17
C.14.	Since draft-ietf-httpbis-p7-auth-12	18
C.15.	Since draft-ietf-httpbis-p7-auth-13	18
C.16.	Since draft-ietf-httpbis-p7-auth-14	18
C.17.	Since draft-ietf-httpbis-p7-auth-15	18
C.18.	Since draft-ietf-httpbis-p7-auth-16	19
Index	19

[1.](#) Introduction

This document defines HTTP/1.1 access control and authentication. It includes the relevant parts of [RFC 2616](#) with only minor changes, plus the general framework for HTTP authentication, as previously defined in "HTTP Authentication: Basic and Digest Access Authentication" ([\[RFC2617\]](#)).

HTTP provides several OPTIONAL challenge-response authentication mechanisms which can be used by a server to challenge a client request and by a client to provide authentication information. The "basic" and "digest" authentication schemes continue to be specified in [RFC 2617](#).

[1.1.](#) Conformance and Error Handling

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 [\[RFC2119\]](#).

This document defines conformance criteria for several roles in HTTP communication, including Senders, Recipients, Clients, Servers, User-Agents, Origin Servers, Intermediaries, Proxies and Gateways. See Section 2 of [\[Part1\]](#) for definitions of these terms.

An implementation is considered conformant if it complies with all of the requirements associated with its role(s). Note that SHOULD-level requirements are relevant here, unless one of the documented exceptions is applicable.

This document also uses ABNF to define valid protocol elements ([Section 1.2](#)). In addition to the prose requirements placed upon them, Senders MUST NOT generate protocol elements that are invalid.

Unless noted otherwise, Recipients MAY take steps to recover a usable protocol element from an invalid construct. However, HTTP does not define specific error handling mechanisms, except in cases where it has direct impact on security. This is because different uses of the protocol require different error handling strategies; for example, a Web browser may wish to transparently recover from a response where the Location header field doesn't parse according to the ABNF, whereby in a systems control protocol using HTTP, this type of error recovery could lead to dangerous consequences.

[1.2](#). Syntax Notation

This specification uses the ABNF syntax defined in Section 1.2 of [[Part1](#)] (which extends the syntax defined in [[RFC5234](#)] with a list

rule). [Appendix B](#) shows the collected ABNF, with the list rule expanded.

The following core rules are included by reference, as defined in [[RFC5234](#)], [Appendix B.1](#): ALPHA (letters), CR (carriage return), CRLF (CR LF), CTL (controls), DIGIT (decimal 0-9), DQUOTE (double quote), HEXDIG (hexadecimal 0-9/A-F/a-f), LF (line feed), OCTET (any 8-bit sequence of data), SP (space), and VCHAR (any visible US-ASCII character).

[1.2.1](#). Core Rules

The core rules below are defined in [[Part1](#)]:

BWS	= <BWS, defined in [Part1], Section 1.2.2>
OWS	= <OWS, defined in [Part1], Section 1.2.2>
quoted-string	= <quoted-string, defined in [Part1], Section 3.2.3>
token	= <token, defined in [Part1], Section 3.2.3>

[2](#). Access Authentication Framework

[2.1](#). Challenge and Response

HTTP provides a simple challenge-response authentication mechanism that can be used by a server to challenge a client request and by a client to provide authentication information. It uses an extensible, case-insensitive token to identify the authentication scheme, followed by additional information necessary for achieving authentication via that scheme. The latter can either be a comma-separated list of attribute-value pairs or a single sequence of characters capable of holding base64-encoded information.

auth-scheme = token

auth-param = token BWS "=" BWS (token / quoted-string)

b64token = 1*(ALPHA / DIGIT /
"-" / "." / "_" / "~" / "+" / "/") * "="

The "b64token" syntax allows the 66 unreserved URI characters ([[RFC3986](#)]), plus a few others, so that it can hold a base64, base64url (URL and filename safe alphabet), base32, or base16 (hex) encoding, with or without padding, but excluding whitespace ([[RFC4648](#)]).

The 401 (Unauthorized) response message is used by an origin server to challenge the authorization of a user agent. This response MUST include a WWW-Authenticate header field containing at least one

challenge applicable to the requested resource.

The 407 (Proxy Authentication Required) response message is used by a proxy to challenge the authorization of a client and MUST include a Proxy-Authenticate header field containing at least one challenge applicable to the proxy for the requested resource.

challenge = auth-scheme [1*SP (b64token / #auth-param)]

Note: User agents will need to take special care in parsing the WWW-Authenticate and Proxy-Authenticate header field values because they can contain more than one challenge, or if more than one of each is provided, since the contents of a challenge can itself contain a comma-separated list of authentication parameters.

Note: Many browsers fail to parse challenges containing unknown schemes. A workaround for this problem is to list well-supported schemes (such as "basic") first.

A user agent that wishes to authenticate itself with an origin server -- usually, but not necessarily, after receiving a 401 (Unauthorized) -- MAY do so by including an Authorization header field with the request.

A client that wishes to authenticate itself with a proxy -- usually, but not necessarily, after receiving a 407 (Proxy Authentication Required) -- MAY do so by including a Proxy-Authorization header field with the request.

Both the Authorization field value and the Proxy-Authorization field value consist of credentials containing the authentication information of the client for the realm of the resource being requested. The user agent MUST choose to use one of the challenges with the strongest auth-scheme it understands and request credentials from the user based upon that challenge.

credentials = auth-scheme [1*SP (b64token / #auth-param)]

If the origin server does not wish to accept the credentials sent with a request, it SHOULD return a 401 (Unauthorized) response. The response MUST include a WWW-Authenticate header field containing at least one (possibly new) challenge applicable to the requested resource.

If a proxy does not accept the credentials sent with a request, it SHOULD return a 407 (Proxy Authentication Required). The response MUST include a Proxy-Authenticate header field containing a (possibly

new) challenge applicable to the proxy for the requested resource.

The HTTP protocol does not restrict applications to this simple challenge-response mechanism for access authentication. Additional mechanisms MAY be used, such as encryption at the transport level or via message encapsulation, and with additional header fields specifying authentication information. However, such additional mechanisms are not defined by this specification.

Proxies MUST forward the WWW-Authenticate and Authorization headers unmodified and follow the rules found in [Section 4.1](#).

[2.2](#). Protection Space (Realm)

The authentication parameter realm is reserved for use by authentication schemes that wish to indicate the scope of protection:

```
realm          = "realm" BWS "=" BWS realm-value
realm-value    = quoted-string
```

A protection space is defined by the canonical root URI (the scheme and authority components of the effective request URI; see [Section 4.3](#) of [\[Part1\]](#)) of the server being accessed, in combination with the realm value if present. These realms allow the protected resources on a server to be partitioned into a set of protection spaces, each with its own authentication scheme and/or authorization database. The realm value is a string, generally assigned by the origin server, which can have additional semantics specific to the authentication scheme. Note that there can be multiple challenges with the same auth-scheme but different realms.

The protection space determines the domain over which credentials can be automatically applied. If a prior request has been authorized, the same credentials MAY be reused for all other requests within that protection space for a period of time determined by the authentication scheme, parameters, and/or user preference. Unless otherwise defined by the authentication scheme, a single protection space cannot extend outside the scope of its server.

[2.3](#). Authentication Scheme Registry

The HTTP Authentication Scheme Registry defines the name space for the authentication schemes in challenges and credentials.

Registrations MUST include the following fields:

- o Authentication Scheme Name

- o Pointer to specification text

- o Notes (optional)

Values to be added to this name space are subject to IETF review ([\[RFC5226\]](#), [Section 4.1](#)).

The registry itself is maintained at
<<http://www.iana.org/assignments/http-authschemes>>.

[2.3.1](#). Considerations for New Authentication Schemes

There are certain aspects of the HTTP Authentication Framework that put constraints on how new authentication schemes can work:

- o Authentication schemes need to be compatible with the inherent constraints of HTTP; for instance, that messages need to keep their semantics when inspected in isolation, thus an authentication scheme can not bind information to the TCP session over which the message was received (see Section 2.2 of [\[Part1\]](#)).
- o The authentication parameter "realm" is reserved for defining Protection Spaces as defined in [Section 2.2](#). New schemes MUST NOT use it in a way incompatible with that definition.
- o The "b64token" notation was introduced for compatibility with existing authentication schemes and can only be used once per challenge/credentials. New schemes thus ought to use the "auth-param" syntax instead, because otherwise future extensions will be impossible.
- o The parsing of challenges and credentials is defined by this specification, and cannot be modified by new authentication schemes. When the auth-param syntax is used, all parameters ought to support both token and quoted-string syntax, and syntactical constraints ought to be defined on the field value after parsing (i.e., quoted-string processing). This is necessary so that recipients can use a generic parser that applies to all authentication schemes.

Note: the fact that the value syntax for the "realm" parameter is restricted to quoted-string was a bad design choice not to be repeated for new parameters.

- o Authentication schemes need to document whether they are usable in origin-server authentication (i.e., using WWW-Authenticate), and/or proxy authentication (i.e., using Proxy-Authenticate).

- o The credentials carried in an Authorization header field are specific to the User Agent, and therefore have the same effect on HTTP caches as the "private" Cache-Control response directive, within the scope of the request they appear in.

Therefore, new authentication schemes which choose not to carry credentials in the Authorization header (e.g., using a newly defined header) will need to explicitly disallow caching, by mandating the use of either Cache-Control request directives (e.g., "no-store") or response directives (e.g., "private").

[3.](#) Status Code Definitions

[3.1.](#) 401 Unauthorized

The request requires user authentication. The response MUST include a WWW-Authenticate header field ([Section 4.4](#)) containing a challenge applicable to the target resource. The client MAY repeat the request with a suitable Authorization header field ([Section 4.1](#)). If the request already included Authorization credentials, then the 401 response indicates that authorization has been refused for those credentials. If the 401 response contains the same challenge as the prior response, and the user agent has already attempted authentication at least once, then the user SHOULD be presented the representation that was given in the response, since that representation might include relevant diagnostic information.

[3.2.](#) 407 Proxy Authentication Required

This code is similar to 401 (Unauthorized), but indicates that the client ought to first authenticate itself with the proxy. The proxy MUST return a Proxy-Authenticate header field ([Section 4.2](#)) containing a challenge applicable to the proxy for the target resource. The client MAY repeat the request with a suitable Proxy-Authorization header field ([Section 4.3](#)).

[4.](#) Header Field Definitions

This section defines the syntax and semantics of HTTP/1.1 header fields related to authentication.

[4.1.](#) Authorization

The "Authorization" header field allows a user agent to authenticate itself with a server -- usually, but not necessarily, after receiving a 401 (Unauthorized) response. Its value consists of credentials

containing information of the user agent for the realm of the resource being requested.

Authorization = credentials

If a request is authenticated and a realm specified, the same credentials SHOULD be valid for all other requests within this realm (assuming that the authentication scheme itself does not require otherwise, such as credentials that vary according to a challenge value or using synchronized clocks).

When a shared cache (see Section 1.2 of [\[Part6\]](#)) receives a request containing an Authorization field, it MUST NOT return the corresponding response as a reply to any other request, unless one of the following specific exceptions holds:

1. If the response includes the "s-maxage" cache-control directive, the cache MAY use that response in replying to a subsequent request. But (if the specified maximum age has passed) a proxy cache MUST first revalidate it with the origin server, using the header fields from the new request to allow the origin server to authenticate the new request. (This is the defined behavior for s-maxage.) If the response includes "s-maxage=0", the proxy MUST always revalidate it before re-using it.
2. If the response includes the "must-revalidate" cache-control directive, the cache MAY use that response in replying to a subsequent request. But if the response is stale, all caches MUST first revalidate it with the origin server, using the header fields from the new request to allow the origin server to authenticate the new request.
3. If the response includes the "public" cache-control directive, it MAY be returned in reply to any subsequent request.

[4.2.](#) Proxy-Authenticate

The "Proxy-Authenticate" header field consists of a challenge that indicates the authentication scheme and parameters applicable to the proxy for this effective request URI (Section 4.3 of [\[Part1\]](#)). It MUST be included as part of a 407 (Proxy Authentication Required) response.

Proxy-Authenticate = 1#challenge

Unlike WWW-Authenticate, the Proxy-Authenticate header field applies only to the current connection and SHOULD NOT be passed on to downstream clients. However, an intermediate proxy might need to obtain its own credentials by requesting them from the downstream client, which in some circumstances will appear as if the proxy is forwarding the Proxy-Authenticate header field.

Fielding, et al.

Expires May 3, 2012

[Page 10]

Internet-Draft

HTTP/1.1, Part 7

October 2011

[4.3.](#) Proxy-Authorization

The "Proxy-Authorization" header field allows the client to identify itself (or its user) to a proxy which requires authentication. Its value consists of credentials containing the authentication information of the user agent for the proxy and/or realm of the resource being requested.

Proxy-Authorization = credentials

Unlike Authorization, the Proxy-Authorization header field applies only to the next outbound proxy that demanded authentication using the Proxy-Authenticate field. When multiple proxies are used in a chain, the Proxy-Authorization header field is consumed by the first outbound proxy that was expecting to receive credentials. A proxy MAY relay the credentials from the client request to the next proxy if that is the mechanism by which the proxies cooperatively authenticate a given request.

[4.4.](#) WWW-Authenticate

The "WWW-Authenticate" header field consists of at least one challenge that indicates the authentication scheme(s) and parameters applicable to the effective request URI (Section 4.3 of [[Part1](#)]).

It MUST be included in 401 (Unauthorized) response messages and MAY be included in other response messages to indicate that supplying credentials (or different credentials) might affect the response.

WWW-Authenticate = 1#challenge

User agents are advised to take special care in parsing the WWW-

Authenticate field value as it might contain more than one challenge, or if more than one WWW-Authenticate header field is provided, the contents of a challenge itself can contain a comma-separated list of authentication parameters.

For instance:

```
WWW-Authenticate: Newauth realm="apps", type=1,
                  title="Login to \"apps\"", Basic realm="simple"
```

This header field contains two challenges; one for the "Newauth" scheme with a realm value of "apps", and two additional parameters "type" and "title", and another one for the "Basic" scheme with a realm value of "simple".

[5.](#) IANA Considerations

[5.1.](#) Authenticon Scheme Registry

The registration procedure for HTTP Authentication Schemes is defined by [Section 2.3](#) of this document.

The HTTP Method Authentication Scheme shall be created at <http://www.iana.org/assignments/http-authschemes>.

[5.2.](#) Status Code Registration

The HTTP Status Code Registry located at <http://www.iana.org/assignments/http-status-codes> shall be updated with the registrations below:

Value	Description	Reference
401	Unauthorized	Section 3.1
407	Proxy Authentication Required	Section 3.2

[5.3.](#) Header Field Registration

The Message Header Field Registry located at <<http://www.iana.org/assignments/message-headers/message-header-index.html>> shall be updated with the permanent registrations below (see [RFC3864]):

Header Field Name	Protocol	Status	Reference
Authorization	http	standard	Section 4.1
Proxy-Authenticate	http	standard	Section 4.2
Proxy-Authorization	http	standard	Section 4.3
WWW-Authenticate	http	standard	Section 4.4

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

6. Security Considerations

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

6.1. Authentication Credentials and Idle Clients

Existing HTTP clients and user agents typically retain authentication information indefinitely. HTTP/1.1 does not provide a method for a server to direct clients to discard these cached credentials. This is a significant defect that requires further extensions to HTTP. Circumstances under which credential caching can interfere with the application's security model include but are not limited to:

- o Clients which have been idle for an extended period following which the server might wish to cause the client to reprompt the user for credentials.
- o Applications which include a session termination indication (such as a "logout" or "commit" button on a page) after which the server side of the application "knows" that there is no further reason for the client to retain the credentials.

This is currently under separate study. There are a number of work-arounds to parts of this problem, and we encourage the use of password protection in screen savers, idle time-outs, and other methods which mitigate the security problems inherent in this problem. In particular, user agents which cache credentials are encouraged to provide a readily accessible mechanism for discarding cached credentials under user control.

[7.](#) Acknowledgments

This specification takes over the definition of the HTTP Authentication Framework, previously defined in [RFC 2617](#). We thank John Franks, Phillip M. Hallam-Baker, Jeffery L. Hostetler, Scott D. Lawrence, Paul J. Leach, Ari Luotonen, and Lawrence C. Stewart for their work on that specification. See [Section 6 of \[RFC2617\]](#) for further acknowledgements.

See Section 11 of [\[Part1\]](#) for the Acknowledgments related to this document revision.

[8.](#) References

[8.1.](#) Normative References

- [Part1] Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., and J. Reschke, Ed., "HTTP/1.1, part 1: URIs, Connections, and Message Parsing", [draft-ietf-httpbis-p1-messaging-17](#) (work in progress), October 2011.

Fielding, et al.	Expires May 3, 2012	[Page 13]
------------------	---------------------	-----------

Internet-Draft	HTTP/1.1, Part 7	October 2011
----------------	------------------	--------------

- [Part6] Fielding, R., Ed., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T., Lafon, Y., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1, part 6: Caching", [draft-ietf-httpbis-p6-cache-17](#) (work in progress), October 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.

[8.2.](#) Informative References

- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", [RFC 2617](#), June 1999.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), September 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), October 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

[Appendix A.](#) Changes from RFCs 2616 and 2617

The "realm" parameter isn't required anymore in general; consequently, the ABNF allows challenges without any auth parameters. ([Section 2](#))

The "b64token" alternative to auth-param lists has been added for consistency with legacy authentication schemes such as "Basic". ([Section 2](#))

Fielding, et al.

Expires May 3, 2012

[Page 14]

Internet-Draft

HTTP/1.1, Part 7

October 2011

Change ABNF productions for header fields to only define the field value. ([Section 4](#))

[Appendix B.](#) Collected ABNF

Authorization = credentials

BWS = <BWS, defined in [\[Part1\]](#), Section 1.2.2>

OWS = <OWS, defined in [\[Part1\]](#), Section 1.2.2>

Proxy-Authenticate = *("," OWS) challenge *(OWS "," [OWS
challenge])

Proxy-Authorization = credentials

WWW-Authenticate = *("," OWS) challenge *(OWS "," [OWS challenge
])

auth-param = token BWS "=" BWS (token / quoted-string)
auth-scheme = token

b64token = 1*(ALPHA / DIGIT / "-" / "." / "_" / "~" / "+" / "/")
* "="

challenge = auth-scheme [1*SP (b64token / [("," / auth-param) *(
OWS "," [OWS auth-param])])]

credentials = auth-scheme [1*SP (b64token / [("," / auth-param)
*(OWS "," [OWS auth-param])])]

quoted-string = <quoted-string, defined in [\[Part1\]](#), Section 3.2.3>

realm = "realm" BWS "=" BWS realm-value
realm-value = quoted-string

token = <token, defined in [\[Part1\]](#), Section 3.2.3>

ABNF diagnostics:

; Authorization defined but not used
; Proxy-Authenticate defined but not used
; Proxy-Authorization defined but not used
; WWW-Authenticate defined but not used
; realm defined but not used

[Appendix C](#). Change Log (to be removed by RFC Editor before publication)

[C.1](#). Since [RFC 2616](#)

Extracted relevant partitions from [[RFC2616](#)].

[C.2](#). Since [draft-ietf-httpbis-p7-auth-00](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/35>>: "Normative and Informative references"

[C.3](#). Since [draft-ietf-httpbis-p7-auth-01](#)

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

- o Explicitly import BNF rules for "challenge" and "credentials" from [RFC2617](#).
- o Add explicit references to BNF syntax and rules imported from other parts of the specification.

[C.4](#). Since [draft-ietf-httpbis-p7-auth-02](#)

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

- o Reference [RFC 3984](#), and update header field registrations for header fields defined in this document.

[C.5](#). Since [draft-ietf-httpbis-p7-auth-03](#)

None.

[C.6](#). Since [draft-ietf-httpbis-p7-auth-04](#)

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

- o Use "/" instead of "|" for alternatives.
- o Introduce new ABNF rules for "bad" whitespace ("BWS"), optional whitespace ("OWS") and required whitespace ("RWS").
- o Rewrite ABNFs to spell out whitespace rules, factor out header

field value format definitions.

C.7. Since [draft-ietf-httpbis-p7-auth-05](#)

Final work on ABNF conversion

(<<http://tools.ietf.org/wg/httpbis/trac/ticket/36>>):

- o Add appendix containing collected and expanded ABNF, reorganize ABNF introduction.

C.8. Since [draft-ietf-httpbis-p7-auth-06](#)

None.

C.9. Since [draft-ietf-httpbis-p7-auth-07](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/198>>: "move IANA registrations for optional status codes"

C.10. Since [draft-ietf-httpbis-p7-auth-08](#)

No significant changes.

C.11. Since [draft-ietf-httpbis-p7-auth-09](#)

Partly resolved issues:

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

C.12. Since [draft-ietf-httpbis-p7-auth-10](#)

None.

C.13. Since [draft-ietf-httpbis-p7-auth-11](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/130>>: "introduction to part 7 is work-in-progress"

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/195>>: "auth-param syntax"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/224>>: "Header Classification"

Fielding, et al.

Expires May 3, 2012

[Page 17]

Internet-Draft

HTTP/1.1, Part 7

October 2011

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/237>>: "absorbing the auth framework from 2617"

Partly resolved issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/141>>: "should we have an auth scheme registry"

C.14. Since [draft-ietf-httpbis-p7-auth-12](#)

None.

C.15. Since [draft-ietf-httpbis-p7-auth-13](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/276>>: "untangle ABNFs for header fields"

C.16. Since [draft-ietf-httpbis-p7-auth-14](#)

None.

C.17. Since [draft-ietf-httpbis-p7-auth-15](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/78>>: "Relationship between 401, Authorization and WWW-Authenticate"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/177>>: "Realm required on challenges"

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/195>>: "auth-param syntax"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/257>>: "Considerations for new authentications schemes"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/287>>: "LWS in auth-param ABNF"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/309>>: "credentials ABNF missing SP (still using implied LWS?)"

Fielding, et al.

Expires May 3, 2012

[Page 18]

Internet-Draft

HTTP/1.1, Part 7

October 2011

[C.18.](#) Since [draft-ietf-httpbis-p7-auth-16](#)

Closed issues:

- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/186>>: "Document HTTP's error-handling philosophy"
- o <<http://tools.ietf.org/wg/httpbis/trac/ticket/320>>: "add advice on defining auth scheme parameters"

Index

4

401 Unauthorized (status code) 9

407 Proxy Authentication Required (status code) 9

A

auth-param 5

auth-scheme 5

Authorization header field 9

B

b64token 5

C

challenge 6

credentials 6

G

Grammar

auth-param 5
auth-scheme 5
Authorization 10
b64token 5
challenge 6
credentials 6
Proxy-Authenticate 10
Proxy-Authorization 11
realm 7
WWW-Authenticate 11

H

Header Fields

Authorization 9
Proxy-Authenticate 10
Proxy-Authorization 11
WWW-Authenticate 11

Fielding, et al.

Expires May 3, 2012

[Page 19]

Internet-Draft

HTTP/1.1, Part 7

October 2011

P

Protection Space 7
Proxy-Authenticate header field 10
Proxy-Authorization header field 11

R

Realm 7
realm 7
realm-value 7

S

Status Codes

401 Unauthorized 9
407 Proxy Authentication Required 9

W

WWW-Authenticate header field 11

Authors' Addresses

Roy T. Fielding (editor)
Adobe Systems Incorporated
345 Park Ave
San Jose, CA 95110
USA

EEmail: fielding@gbiv.com
URI: <http://roy.gbiv.com/>

Jim Gettys
Alcatel-Lucent Bell Labs
21 Oak Knoll Road
Carlisle, MA 01741
USA

EEmail: jg@freedesktop.org
URI: <http://gettys.wordpress.com/>

Fielding, et al.

Expires May 3, 2012

[Page 20]

Internet-Draft

HTTP/1.1, Part 7

October 2011

Jeffrey C. Mogul
Hewlett-Packard Company
HP Labs, Large Scale Systems Group
1501 Page Mill Road, MS 1177
Palo Alto, CA 94304
USA

EEmail: JeffMogul@acm.org

Henrik Frystyk Nielsen
Microsoft Corporation

1 Microsoft Way
Redmond, WA 98052
USA

E-Mail: henrikn@microsoft.com

Larry Masinter
Adobe Systems Incorporated
345 Park Ave
San Jose, CA 95110
USA

E-Mail: LMM@acm.org
URI: <http://larry.masinter.net/>

Paul J. Leach
Microsoft Corporation
1 Microsoft Way
Redmond, WA 98052

E-Mail: paulle@microsoft.com

Tim Berners-Lee
World Wide Web Consortium
MIT Computer Science and Artificial Intelligence Laboratory
The Stata Center, Building 32
32 Vassar Street
Cambridge, MA 02139
USA

E-Mail: timbl@w3.org
URI: <http://www.w3.org/People/Berners-Lee/>

France

EMail: ylafon@w3.org

URI: <http://www.raubacapeu.net/people/yves/>

Julian F. Reschke (editor)

greenbytes GmbH

Hafenweg 16

Muenster, NW 48155

Germany

Phone: +49 251 2807760

Fax: +49 251 2807761

EMail: julian.reschke@greenbytes.de

URI: <http://greenbytes.de/tech/webdav/>