

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: August 23, 2010

Y. Oiwa  
H. Watanabe  
H. Takagi  
RCIS, AIST  
H. Suzuki  
Yahoo! Japan  
February 19, 2010

# Mutual Authentication Protocol for HTTP

## draft-oiwa-http-mutualauth-06

### Abstract

This document specifies the "Mutual authentication protocol for Hyper-Text Transport Protocol". This protocol provides true mutual authentication between HTTP clients and servers using simple password-based authentication. Unlike The Basic and Digest HTTP access authentication protocol, this protocol ensures that the server knows the user's entity (encrypted password) upon successful authentication. This prevents common phishing attacks: phishing attackers cannot convince users that the user has been authenticated to the genuine website. Furthermore, even when a user has been authenticated against an illegitimate server, the server cannot gain any bit of information about user's passwords. The protocol is designed as an extension to the HTTP protocol, and the protocol design intends to replace existing authentication mechanism such as Basic/Digest access authentications and form-based authentications.

### Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 23, 2010.

## Copyright Notice

Copyright (c) 2010 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 BSD License.

## Table of Contents

1. Introduction
  - 1.1. Requirements Language
2. Protocol Overview
3. Message Syntax
  - 3.1. Tokens and Extensive-tokens
  - 3.2. Numbers
  - 3.3. Strings
4. Messages
  - 4.1. 401-B0
  - 4.2. 401-B0-stale
  - 4.3. req-A1
  - 4.4. 401-B1
  - 4.5. req-A3
  - 4.6. 200-B4
  - 4.7. 200-Optional-B0
5. Decision procedure for the client
6. Decision procedure for the server
7. Authentication-Control header
  - 7.1. Location-when-unauthenticated field
  - 7.2. Location-when-logout field
  - 7.3. Logout-timeout
8. Authentication Algorithms
  - 8.1. Common functions
  - 8.2. Functions for discrete-logarithm settings
  - 8.3. Functions for elliptic-curve settings
9. Authentication Realms
  - 9.1. Resolving ambiguities
10. Validation Methods
11. Session Management
12. Methods to extend this protocol
13. IANA Considerations
14. Security Considerations
  - 14.1. General Assumptions
  - 14.2. Implementation Considerations
  - 14.3. Usage Considerations
15. Notice on intellectual properties
16. Acknowledgement

## 17. References

### 17.1. Normative References

### 17.2. Informative References

Appendix A. Group parameters for discrete-logarithm based algorithms

Appendix B. Derived numerical values

Appendix C. Draft Remarks from the Authors

Appendix D. Draft Change Log

D.1. Changes in revision 06

D.2. Changes in revision 05

D.3. Changes in revision 04

D.4. Changes in revision 03

D.5. Changes in revision 02

§ Authors' Addresses

## **1. Introduction**

This document specifies the "Mutual authentication protocol for Hyper-Text Transport Protocol". This protocol provides true mutual authentication between HTTP clients and servers using simple password-based authentication. Unlike Basic and Digest HTTP access authentication protocol [RFC2617], the protocol ensures that server knows the user's entity (encrypted password) upon successful authentication. This prevents common phishing attacks: phishing attackers cannot convince users that the user has been authenticated to the genuine website. Furthermore, even when a user has been authenticated against an illegitimate server, the server cannot gain any bit of information about user's passwords.

Recently, phishing attacks are getting more and more sophisticated. Phishers not only steal user's password directly, but imitate successful authentication to steal the user's sensitive information, check the password validity by forwarding the password to the legitimate server, or employ a man-in-the-middle attack to hijack user's login session. Existing countermeasures such as one-time passwords cannot completely solve these problems.

The protocol prevents such attacks by providing users a way to discriminate between true and fake web servers using their own passwords. Even when a user inputs his/her password to a fake website, using this authentication method, any information about the password does not leak to the phisher, and the user certainly notices that the mutual authentication has failed. Phishers cannot make such authentication attempt succeed, even if they forward received data from a user to the legitimate server or vice versa. Users can input sensitive data to the web forms after confirming that the mutual authentication has succeeded, without fear of phishing attacks.

To achieve this goal, this protocol uses a mechanism in ISO/IEC 11770-4 [ISO.11770-4.2006], a kind of PAKE (Password-Authenticated Key Exchange) authentication algorithms as a basis. The use of PAKE mechanism allows users to use familiar ID/password based accesses, without fear of leaking any password information to the communication peer. The protocol, as a whole, is designed as a natural extension to the HTTP protocol [RFC2616]. The protocol is designed to be used both with HTTP and with HTTPS, the latter of which will prevent transport-level eavesdropping.

The design also considers to replace current form-based Web authentication, which is very vulnerable against phishing attacks. To this purpose, several extensions to current HTTP authentication mechanism [RFC2617] are introduced. These include:

- multi-host single authentication within an Internet domain ([Section 9](#)),
- non-mandatory, optional authentication on HTTP ([Section 4.7](#)),
- log out from both server and client side ([Section 7](#)), and
- finer control for redirection depending on authentication status ([Section 7](#)).

## 1.1. Requirements Language

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 [\[RFC2119\]](#).

## 2. Protocol Overview

The following sequence is a typical sequence for the first access to the resource.

- If the server (S) has received a request for mutual-authentication protected resources from the Client (C) (which is not a req-A1 nor a req-A3 message), it sends a 401-B0 message to C.  
When C has received a 401-B0 message, C SHOULD check validity of the message. If succeed, C processes the body of the message, and enables the password entry field.
- If the user has input the username and password as a response to the 401-B0 message, C creates a value s\_A, calculates the value w\_A, and then constructs and sends a req-A1 message.
- If S has received a req-A1 message, S should check validity of w\_A, record the received w\_A value, and then look up the username from the user table. if the user is found, S prepares a new session id (sid), records it into a session table, and then constructs s\_B, calculates w\_B, and sends a 401-B1 message.  
If there is no matching user found, the server SHOULD construct a fake w\_B value and send an 401-B1 message.
- When C has received a 401-B1 message as a response for a req-A1 message, C should check validity of w\_B, and compute z and o\_A, and send a req-A3 message.  
If C receives any messages other than 401-B1, C MUST NOT process the message body and treat it as a fatal communication error condition. This case includes the reception of HTTP OK (200-status) message.
- If S has received a req-A3 message, S should look up the received sid from the session table. If there is no matching sid, or if S has not received the corresponding req-A1 message beforehand, S SHOULD send a 401-B0-stale message.  
Otherwise, S should compute o\_A and check its value. If the validation has failed, it means that the authentication has failed. The server SHOULD send a 401-B0 message.  
If the validation has succeeded, the server SHOULD calculate o\_B, and send a 200-B4 message.
- In a response to a req-B1 message, when C has received a 401-B0 message, it means that the authentication has failed, possibly due to the wrong password being given. C MAY ignore the body of the 401-B0 message in this case.  
When C has received a 200-B4 message, C MUST first compute the value of o\_B and validate the value o\_B sent from the server. If it has not verified successfully, C MUST ignore the body of the message, and treat the situation as a fatal communication error condition. If the verification has succeeded, C will process the body of the message.  
If C receives any messages other than 401-B0 or valid 200-B4, C MUST NOT process the message body and other headers and treat it as a fatal communication error condition. This case includes the reception of usual HTTP OK (200-status) messages.

For the second or later request to the server, if the client knows that the resource is likely to require the authentication, the client MAY omit first unauthenticated request and send req-A1 message immediately. In this case, the first (and only the first) response from the server MAY be a normal, unauthenticated message, and client MAY accept such messages.

Furthermore, if client owns a valid session ID (sid), the client MAY send a req-A3 message using existing sid. In such cases, the server MAY have been thrown out the corresponding sessions from the session table. In this case, the server SHOULD send a 401-B0-stale message as a response to req-A3 message, and C SHOULD retry from constructing a req-A1 message.

For more detail, see [Section 5](#) and [Section 6](#).

### 3. Message Syntax

The Mutual authentication protocol uses five headers: WWW-Authenticate (in responses with status code 401), Optional-WWW-Authenticate (in responses with positive status codes), Authentication-Control (in responses), Authorization (in requests), and Authentication-info (in positive responses). These five headers share the common syntax described in [Figure 1](#). The syntax is denoted in the augmented BNF syntax defined in [\[RFC5234\]](#). The syntax is a subset of the one described in [\[RFC2617\]](#).

---

```

header           = header-name ":" [spaces] auth-scheme spaces fields
header-name      = "WWW-Authenticate" / "Optional-WWW-Authenticate"
                  / "Authorization" / "Authentication-info"
                  / "Authentication-Control"
spaces           = 1*( " " / %x09 / %x0D.0A ( " " / %x09) ) ; LWSP
auth-scheme      = "Mutual" ; see HTTP for other values
fields          = field *([spaces] "," spaces field)
field           = key "=" value
key             = extensive-token
extensive-token = token / extension-token
extension-token = token "@" token
token           = 1*(%x30-39 / %x41-5A / %x61-7A / "." / "-" / "_")
value          = extensive-token / integer / hex-integer
                / hex-fixed-number
                / base64-fixed-number / string
integer         = "0" / (%x31-39 *%x30-39) ; no leading zeros
hex-integer     = "0"
                / ((%x31-39 / %x41-46 / %x61-66) ; no leading zeros
                  *(%x30-39 / %x41-46 / %x61-66))
hex-fixed-number = 1*(%x30-39 / %x41-46 / %x61-66)
base64-fixed-number = string
string         = %x22 *(%x20-21 / %x23-5B / %x5D-FF
                  / %x5C.22 / "\\\" / "\,") %x22

```

Figure 1: the BNF syntax for the headers used in the protocol

---

#### 3.1. Tokens and Extensive-tokens

The tokens MUST be interpreted case-insensitively, and SHOULD be sent in the same case as shown in the specification. When these are used as (partial) inputs to any hash or other mathematical functions, it MUST be used in lower-case. All hex-fixed-number or hex-integer numbers are also case-insensitive, and SHOULD be sent in lower-case.

Extensive-tokens are used where the set of acceptable tokens are extensible. Any non-standard extensions of this protocol **MUST** use the extension-tokens of format "<token>@<domain-name>", where domain-name is the valid registered (sub-)domain name on the Internet owned by the party who defines extensions.

## 3.2. Numbers

The syntax definitions of integer and hex-integer only allow representations which do not contain extra leading 0s.

The numbers represented as a hex-fixed-number **MUST** have even characters (i.e. multiple of eight bits). When these are generated from cryptographic values, those **SHOULD** have the natural length: if these are generated from a hash function, these lengths **SHOULD** correspond to the hash size; if these are representing elements of a mathematical group, its lengths **SHOULD** be the shortest which can represent all elements in the group. See [Appendix B](#) for information about the length of the fields used in this specification. Other values such as session-id are represented in any (even) length determined by the side who generates it first, and the same length **SHALL** be used throughout the whole communications by both peers.

The numbers represented as a base64-fixed-number **SHALL** be generated as follows: first, the number is converted to a big-endian octet-string representation. The length of the representation is determined in the same way as above. Then, the string is encoded by [the Base 64 encoding](#) [RFC4648], and then enclosed by two double-quotations.

## 3.3. Strings

All strings outside ASCII or equivalent character sets **SHOULD** be encoded using [UTF-8 encoding](#) [RFC3629] of the [ISO 10646-1 character set](#) [ISO.10646-1.1993]. Both peers **SHOULD** reject any invalid UTF-8 sequences which cause decoding ambiguities (e.g. containing <"> in the second or later byte of the UTF-8 encoded characters). To encode character strings, these will first be encoded according to UTF-8 without leading BOM, then all occurrences of characters <"> and "\" will be escaped by prepending "\", and two <">s will be put around the string. If the contents of the strings are comma-separated values, the commas in the values are also quoted by "\".

If strings are representing a domain name or URI which contains non-ASCII characters, the host parts **SHOULD** be encoded using puny-code defined in [\[RFC3492\]](#) instead of UTF-8, and **SHOULD** use lower-case ASCII characters.

For Base64-fixed-numbers, which use the string syntax, see the previous section.

## 4. Messages

In this section we define seven kinds of messages used in the authentication protocol, Along with formats and requirements of the headers for each message. Messages are distinguished by the fields contained in a header corresponding to the authentication. The procedure below will determine the kind of a message which each HTTP request/response belongs to.

- If the message is a response with a "401" status:
  - If it does not contain any WWW-Authenticate header, it is an error.
  - If the WWW-Authenticate header specifies a scheme other than "Mutual", it is a normal response in this draft's scheme.

- Otherwise, the response contains a "WWW-Authenticate: Mutual" header. If the header contains both sid and stale fields, it is an error.
- If the header contains a stale field with a value of 0, it is a 401-B0 message.
- If the header contains a stale field with a value of 1, it is a 401-B0-stale message.
- If the header contains a sid field, it is a 401-B1 message.
- If the message is a response with a positive status (such as 200, 206, 302 or others):
  - If it contains both Authentication-Info and Optional-WWW-Authenticate headers, it is an error.
  - If it contains a Authentication-Info header with a scheme "Mutual", it is a 200-B4 message.
  - If it contains a Optional-WWW-Authenticate header with "Mutual" scheme, it is a 200-Optional-B0 message.
  - If it contains a Optional-WWW-Authenticate header with a scheme other than "Mutual", it is either an error or a normal response, and the behavior is not defined in this specification.
  - Otherwise, it is a normal response.
- If the message is a request:
  - If it do not contain a Authorization header, or it contains a Authorization header with a scheme other than Mutual, it is a normal request.
  - Otherwise, the request contains a "Authorization: Mutual" header. If the header contains a sid field, it is a req-A3 message.
  - If the header do not contain a sid field, it is a req-A1 message.

Implementations MAY perform checks stricter than the procedure above, according to the definitions in the rest of this section. If there is an error, the implementation MAY either treat it as a fatal communication error, or ignore the corresponding header and treat it as a normal message. Please be noted, however, when Mutual authentication is performed, there are several cases where client SHOULD NOT accept normal responses, as specified later in [\(Section 5\)](#).

To determine which message are expected to be sent, see [Section 5](#) and [Section 6](#).

In the descriptions below, the allowed type of values for each header field is shown in parenthesis after the key names. The type "algorithm-determined" means that the acceptable value type for the field is one of the types defined in [Section 3](#), and is determined by the value of the "algorithm" field.

Note: The term "optional" here means that omitting the field is allowed and has specific meanings in communications (i.e. it is not generally "OPTIONAL" defined in [\[RFC2119\]](#)).

#### 4.1. 401-B0

Every 401-B0 message SHALL be a valid HTTP 401 (Authentication Required) message containing one (and only one: hereafter not explicitly noticed) "WWW-Authenticate" header of the following format.

```
WWW-Authenticate: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", stale=0,
version=-draft05
```

The header SHALL contain the fields with the following keys:

version:

(extensive-token) should be the token "-draft05" in this specification. The behavior when other values are specified is undefined.

algorithm:

(extensive-token) specifies the authentication algorithm to be used. The value MUST be one of the tokens described in [Section 8](#), or the tokens specified in other supplemental specification documentations.

validation:

(extensive-token) specifies the method of host validation. The value MUST be one of the tokens described in [Section 10](#), or the tokens specified in other supplemental specification documentations.

auth-domain:

(optional, string) specifies authentication domain, the set of hosts on which authentication credentials are valid. It MUST be one of the strings described in [Section 9](#). If the value is omitted, it is assumed to be the host part of the requested URI.

realm:

(string) is a UTF-8 encoded string representing the name of the authentication realm inside the authentication domain.

pwd-hash:

(optional, extensive-token) specifies the hash algorithm (referred to by ph) used for additionally hashing the password. The valid tokens are

- none:  $ph(p) = p$
- md5:  $ph(p) = MD5(p)$
- digest-md5:  $ph(p) = MD5(\text{username} | ":" | \text{realm} | ":" | p)$ , the same value as MD5(A1) for "MD5" algorithm in [\[RFC2617\]](#).
- sha1:  $ph(p) = SHA1(p)$

If omitted, the value "none" is assumed. The use of "none" is recommended.

stale:

(token) MUST be "0".

Any additional fields SHOULD NOT be contained in the header, except those explicitly specified in supplement specifications of the "authentication algorithm".

The algorithm will determine the types and the values for w\_A, w\_B, o\_A and o\_B.

## 4.2. 401-B0-stale

A 401-B0-stale message is a variant of 401-B0 message, which means that the client has sent a request message which is not for any active session.

WWW-Authenticate: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", stale=1, version=-draft05

The header MUST contain the same fields as in 401-B0, except that stale field holds the integer 1.

## 4.3. req-A1

Every req-A1 message SHALL be a valid HTTP request message containing a "Authorization" header of the following format.

Authorization: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", user="xxxx", wa=xxxx, version=-draft05



The header SHALL contain the fields with the following keys:

version:

(extensive-token) should be the token "-draft05" in this specification. The behavior when other values are specified is undefined.

algorithm, validation, auth-domain, realm:

MUST be the same value as it is received from the server.

user:

(string) is the UTF-8 encoded name of the user.

wa:

(algorithm-determined) is the value of w\_A specified by the used algorithm (see [Section 8](#)).

#### 4.4. 401-B1

Every 401-B1 message SHALL be a valid HTTP 401 (Authentication Required) message containing a "WWW-Authenticate" header of the following format.

WWW-Authenticate: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", sid=xxxx, wb=xxxx, nc-max=x, nc-window=x, time=x, path="xxxx", version=-draft05

The header SHALL contain the fields with the following keys:

version:

(extensive-token) should be the token "-draft05" in this specification. The behavior when other values are specified is undefined.

algorithm, validation, auth-domain, realm:

MUST be the same value as it is received from the client.

sid:

(hex-fixed-number) MUST be a session id, which is a random integer. The sid SHOULD have uniqueness of at least 80 bits or the square of the maximal estimated transactions concurrently available in the session table, whichever is larger. Sids are local to each authentication realm concerned: the same sids for different authentication realms SHOULD be treated as independent ones.

wb:

(algorithm-determined) is the value of w\_B specified by the algorithm (see [Section 8](#)).

nc-max:

(hex-integer) is the maximal value of nonce counts which S accepts.

nc-window:

(hex-integer) the number of available nonce slots which the server will accept. The value of nc-window is RECOMMENDED to be thirty-two ("20" in hex-integer) or more.

time:

(integer) represents the suggested time (in seconds) which the client can reuse the session represented by sid. It is RECOMMENDED to be at least 60. The value of this field is not directly linked to the duration that the server keeps track of the session represented by sid.

path:

(optional, string) specifies for which path in the URI space the same authentication is expected to apply. The value is in the same format as it is specified in [\[RFC2617\]](#) for the Digest authentications, and clients are RECOMMENDED to recognize it. The all path elements contained in the field MUST be inside the specified auth-domain: if not, client SHOULD ignore such elements.

## 4.5. req-A3

Every req-A3 message SHALL be a valid HTTP request message containing a "Authorization" header of the following format.

Authorization: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", sid=xxxx, nc=x, oa=xxxx, version=-draft05

The fields contained in the header are as follows:

version:

(extensive-token) should be the token "-draft05" in this specification. The behavior when other values are specified is undefined.

algorithm, validation, auth-domain, realm:

MUST be the same value as it is received from the server for the session.

sid:

(hex-fixed-number) MUST be one of the sid values which has been received from the server for the same authentication realm.

nc:

(hex-integer) is a nonce value which is unique among the requests sharing the same sid. The value of nc SHOULD satisfy the following properties:

- It is not larger than the nc-max value which has been sent from the server in the session represented by the sid.
- It has not yet sent previously in the same session.
- It is not smaller than (largest-nc - nc-window), where largest-nc is the maximal value of nc which has previously been sent in the session, and nc-window is the value of the nc-window field which has been sent from the server in the session.

oa:

(algorithm-determined) is the value of o\_A specified by the algorithm (see [Section 8](#)).

## 4.6. 200-B4

Every 200-B1 message SHALL be a valid HTTP message which is not 401 (Authentication Required) type, containing an "Authentication-Info" header of the following format.

Authentication-Info: Mutual sid=xxxx, ob=xxxx, version=-draft05

The fields contained in the header are as follows:

version:

(extensive-token) should be the token "-draft05" in this specification. The behavior when other values are specified is undefined.

sid:

(hex-fixed-number) MUST be the value received from the client.

ob:

(algorithm-determined) is the value of o\_B specified by the algorithm (see [Section 8](#)).

logout-timeout:

(optional, integer) is a number of seconds after which the client should re-validate the user's password for the current authentication realm. As a special case, the value 0 means that the client SHOULD automatically forget the user-inputted password to the current authentication realm and revert to the unauthenticated state (i.e. server-initiated logout). This

does not, however, mean that the long-term memories for the passwords (such as password reminders and auto fill-ins) should be removed. If a new value of timeout is received for the same authentication realm, it overrides the previous timeout.

## 4.7. 200-Optional-B0

The 200-Optional-B0 messages will allow a non-mandatory authentication which is not possible under current HTTP authentication mechanism. In several Web applications, users can access the same contents both as a guest user and as a authenticated users. In usual Web applications, it is implemented using Cookies and custom form-based authentications. The new method of authentication using this message will provide a replacement for those authentication systems. The support for this message is RECOMMENDED, unless an authentication is mandatory for some specific applications.

Servers MAY send HTTP successful responses (response code 200, 206 and others) containing the Optional-WWW-Authenticate header, when it is allowed to send 401-B0 responses (with one exception described below). Such responses are hereafter called 200-Optional-B0 responses.

HTTP/1.1 200 OK

Optional-WWW-Authenticate: Mutual algorithm=xxxx, validation=xxxx, realm="xxxx", stale=0

The fields contained in the Optional-WWW-Authenticate header is the same as the 401-B0 message described in [Section 4.1](#). For authentication-related matters, a 200-Optional-B0 message will have the same meaning as a 401-B0 message with a corresponding WWW-Authenticate header. (The behavior for other matters, such as caching, MAY be different between 200-Optional-B0 and 401-B0 messages.)

The 200-Optional-B0 message is a only place where an Optional-WWW-Authenticate header is allowed. If a server is to send a 401-B1 or a 401-B0-stale responses, it SHALL NOT replace it with 200-Optional-B0 or similar responses. Furthermore, if a server is going to send a 401-B0 message as a responses to req-A3 message with a correct realm, the server MUST send a 401-B0 message, not a 200-Optional-B0 message.

Servers requesting optional mutual authentication SHOULD send the path field in 401-B1 messages with an appropriate value. Client software supporting optional mutual authentication MUST recognize the field, and MUST send either req-A1 or req-A3 request for the URI space inside the specified paths, instead of unauthenticated requests.

## 5. Decision procedure for the client

To securely implement the protocol, the user client must be careful for accepting authenticated responses from the server. This also holds upon reception of "normal responses" (responses which do not contain Mutual-related headers) from HTTP servers.

Clients SHOULD implement the decision procedure equivalent to the one shown below. (Unless implementers understand what is required for the security, they should not alter this.) The labels on the steps are for informational purpose only.

Step 1 (step\_new\_request):

If the client software needs to get a new Web resource, check whether the resource is expected to be inside some authentication realm for which the user has already authenticated by the Mutual authentication scheme. If yes, go to Step 2. Otherwise, go to Step 5.

Step 2:

Check whether there is an available sid for the authentication realm you expect. If there is one, go to Step 3. Otherwise, go to Step 4.

Step 3 (step\_send\_a3\_1):

Send a req-A3 request.

- If you receive a 401-B0 message with a different authentication realm than expected, go to Step 6.
- If you receive a 200-Optional-B0 message with a different authentication realm than expected, go to Step 6.
- If you receive a 401-B0-stale message, go to Step 9.
- If you receive a 401-B0 message, go to Step 13.
- If you receive a 200-B4 message, go to Step 14.
- If you receive a normal response, go to Step 11.

Step 4 (step\_send\_a1\_1):

Send a req-A1 request.

- If you receive a 401-B0 message with a different authentication realm than expected, go to Step 6.
- If you receive a 200-Optional-B0 message with a different authentication realm than expected, go to Step 6.
- If you receive a 401-B1 message, go to Step 10.
- If you receive a normal response, go to Step 11.

Step 5 (step\_send\_normal\_1):

Send a request without any Mutual authentication headers.

- If you receive a 401-B0 message, go to Step 6.
- If you receive a 200-Optional-B0 message, go to Step 6.
- If you receive a normal response, go to Step 11.

Step 6 (step\_rcvd\_b0):

Check whether you know the user's password for the requested authentication realm. If yes, go to Step 7. Otherwise, go to Step 12.

Step 7:

Check whether there is an available sid for the authentication realm you expect. If there is one, go to Step 8. Otherwise, go to Step 9.

Step 8 (step\_send\_a3):

Send a req-A3 request.

- If you receive a 401-B0-stale message, go to Step 9.
- If you receive a 401-B0 message, go to Step 13.
- If you receive a 200-B4 message, go to Step 14.

Step 9 (step\_send\_a1):

Send a req-A1 request.

- If you receive a 401-B1 message, go to Step 10.

Step 10 (step\_rcvd\_b1):

Send a req-A3 request.

- If you receive a 401-B0 message, go to Step 13.
- If you receive a 200-B4 message, go to Step 14.

Step 11 (step\_rcvd\_normal):

This case means that the resource requested is out of the authenticated area. The client will be in the "UNAUTHENTICATED" status. If the response contains a request for authentications other than Mutual, it MAY be handled normally.

Step 12 (step\_rcvd\_b0\_unknown):

This case means that the resource requested requires Mutual authentication, and the user is not authenticated yet. The client will be in the "AUTH\_REQUESTED" status, is RECOMMENDED to process the content sent from the server and ask user a username and a password. If the user has inputted those, go to Step 9.

Step 13 (step\_rcvd\_b0\_failed):

This case means that in some reason the authentication failed: possibly the password or the username is invalid for the authenticated resource. Forget the password for the authentication realm and go to Step 12.

Step 14 (step\_rcvd\_b4):

Check the validity of the received o\_b value. If it is equal to the expected value, it means that the mutual authentication has been succeeded. The client will be in the "AUTH\_SUCCEEDED" status.

If the value is unexpected, it is a fatal communication error.

If a user requests to log out explicitly (via user interfaces), the client MUST forget user's password, go to step 5 and reload the current resource without authentication credential.

Any kind of response (including a normal response) other than those shown in the above procedure SHOULD be interpreted as fatal communication error, and in such cases user clients MUST NOT process any data (response body and other content-related headers) sent from the server.

The client software SHOULD display the three client status to the end-user. For an interactive client, however, if a request is a sub-request for a resource included to another page (e.g. embedded images, style sheets, frames etc.), its status MAY be omitted from being shown, and any "AUTH\_REQUESTED" statuses MAY be treated in the same way as an "UNAUTHENTICATED" status.

Figure 2 shows the full client-side state diagram.

---



- If the received user name is invalid, send a fake 401-B1 response.
- Otherwise, send a 401-B1 response.
- When the server receives a req-A3 request:
  - If the requested resource is not protected by the Mutual Authentication, send a normal response.
  - If the authentication realm specified in the req-A3 request is not the expected one, send either a 401-B0 or a 200-Optional-B0 response.
  - If the received sid correspond to a fake session generated for an unknown user, send a 401-B0 response.
  - If the received sid is either inactive or unknown, send a 401-B0-stale response.
  - If the received oa is incorrect, send a 401-B0 response.
  - If the received oa is correct, send a 200-B4 response.

## 7. Authentication-Control header

The Authentication-Control header gives more precise control for the client behavior for Web applications using Mutual Access Control Protocol. This headers may usually be generated in an application layer, as opposed to WWW-Authenticate headers which will be generated by Web servers.

Support of this header is RECOMMENDED for interactive clients and not required for non-interactive clients. Web applications SHOULD consider security impacts of behavior of clients which do not support this header.

The "auth-scheme" of this header and other authentication-related headers within the same message MUST be equal. This document does not define any behavior associated with this header, when the "auth-scheme" of this header is not "Mutual".

### 7.1. Location-when-unauthenticated field

```
Authentication-Control: Mutual
location-when-unauthenticated="http://www.example.com/login.html"
```

The field "location-when-unauthenticated" specifies a location which any unauthenticated users of clients should be redirected to. This header may be used, for example, when there is a central login page for the whole Web application. The value of this field MUST be a string that contains an absolute URL location. If a given URL is not absolute, clients MAY consider it as a relative URL from the current location.

This field MAY be used with 401-B0 and 200-Optional-B0 messages; however, use of this with 200-optional-B0 messages is not recommended. If there is a 200-B4, 401-B0-stale or 401-B1 message with this field, clients MUST ignore this field.

When a client receives a message with this field, if and only if the client's state after the processing the response is either Step 12 or Step 13 (i.e., a state in which the client will process response body and ask user's password), the client will treat the whole response as if it were a 303 "See Other" response with a Location header with the value of this field (i.e., client will be redirected to the specified location with a GET request). Unlike a normal 303 response, if the client can proceed authentication without user's interaction (like steps 3, 4, 8, 9 and 10), this field is ignored.

The specified location SHOULD be included in a set of locations specified in the "auth-domain" field of the corresponding 401-B0 message. If this is not satisfied, clients MAY ignore this field.

## 7.2. Location-when-logout field

Authentication-Control: Mutual location-when-logout="http://www.example.com/byebye.html"

The field "location-when-logout" specifies a location where the client is to be redirected when users request logout explicitly. The value of this field MUST be a string that contains an absolute URL location. If a given URL is not absolute, clients MAY consider it as a relative URL from the current location.

This field MAY be used with 200-B4 messages. If there is a 401-B0, 401-B1, 401-B0-stale, 200-optional-B0 or normal 200 message with this field, clients MUST ignore this field.

When users of a client request to terminate an authentication session, and if the client currently displays a page supplied by a response with this field, the client will be redirected to the specified location by a new GET request (like received a 303 response), instead of reloading the page without authentication credentials. It is recommendable for Web applications to send this field with an appropriate value for any responses (except those with redirection (3XX) statuses) for non-GET requests.

## 7.3. Logout-timeout

Authentication-Control: Mutual logout-timeout=300

The field "logout-timeout" has the same meaning as the field of the same name in "Authentication-info" headers. This field will be used with 200-B4 messages. If both are specified, clients are recommended to use the one with the smaller value.

## 8. Authentication Algorithms

This document specifies only one family of the authentication algorithm. The family consists of four authentication algorithms, which only differ in underlying mathematical groups and security parameters. The algorithms do not add any additional fields. The tokens for algorithms are

- "iso11770-4-ec-p256" for the 256-bit prime-field elliptic-curve setting.
- "iso11770-4-ec-p521" for the 521-bit prime-field elliptic-curve setting.
- "iso11770-4-dl-2048" for the 2048-bit discrete-logarithm setting.
- "iso11770-4-dl-4096" for the 4096-bit discrete-logarithm setting.

For the elliptic-curve settings, the underlying fields and the curves used for elliptic-curve cryptography are the prime field and the Curve P-256 and P-521, respectively, specified in the appendix of [FIPS PUB 186-2](#) [FIPS.186-2.2000] specification. The hash functions H are SHA-256 for P-256 curve and SHA-512 for P-521 curve, respectively, defined in [FIPS PUB 180-2](#) [FIPS.180-2.2002]. The representation of fields wa, wb, oa, and ob is hex-fixed-number.

For discrete-logarithm settings, the underlying groups are 2048-bit and 4096-bit MODP groups defined in [\[RFC3526\]](#) respectively. See [Appendix A](#) for the exact specification of the group and associated parameters. The hash functions H are SHA-256 for the 2048-bit field and SHA-512 for the 4096-bit field, respectively. The representation of fields wa, wb, oa, and ob is base64-fixed-number.



The clients SHOULD support at least "iso11770-4-dl-2048" algorithm, and are advised to support all of the above four algorithms whenever possible. The server software implementations SHOULD support at least "iso11770-4-dl-2048" algorithm, unless it is known that users will not use it.

This algorithm uses Key Agreement Mechanism 3 (KAM3) defined in Section 6.3 of [ISO/IEC-11770-4](#) [ISO.11770-4.2006] as a basis.

## 8.1. Common functions

The password-based string  $\pi$  used by this authentication is derived in the following manner:

$$\pi = H(\text{VS}(\text{algorithm}) \mid \text{VS}(\text{auth-domain}) \mid \text{VS}(\text{realm}) \mid \text{VS}(\text{username}) \mid \text{VS}(\text{ph}(\text{password}))).$$

The values of algorithm, realm and auth-domain are taken from the values contained in the 401-B0 (or 200-optional-B0, hereafter implied) message. When  $\pi$  is used in the context of an octet string, it SHALL have the natural length derived from the size of the output of function H (e.g. 32 octets for SHA-256). The function ph is defined by the value of the pwd-hash field given in a 401-B0 message.

The function VI encodes natural numbers into octet strings in the following manner: integers are represented in big-endian radix-128 string, where each digit is represented by a octet 0x80–0xff except the last digit represented by 0x00–0x7f. The first octet MUST NOT be 0x80. For example,  $\text{VI}(i) = \text{octet}(i)$  for  $i < 128$ , and  $\text{VI}(i) = \text{octet}(0x80 \mid (i \gg 7)) \mid \text{octet}(i \& 127)$  for  $128 \leq i < 16384$ . This encoding is the same as the one used for subcomponents of object identifiers in [the ASN.1 encoding](#) [ITU.X690.1994].

The function VS encodes variable-length octet string into decodable octet string, as in the following manner:

$$\text{VS}(s) = \text{VI}(\text{length}(s)) \mid s$$

where  $\text{length}(s)$  is a number of octets (not characters) in  $s$ .

The function OCTETS converts an integer to corresponding radix-256 big-endian octet string having its natural length: See [Section 3.2](#) for the definition of the "natural length". Note that this is different from the function GE2OS\_x in [\[ISO.11770-4.2006\]](#), which takes the shortest representation.

The equations for J, w\_A, T, z, and w\_B are specified differently for the discrete-logarithm setting and the elliptic-curve setting based on [\[ISO.11770-4.2006\]](#). These equations are defined later in this section.

The values o\_A and o\_B are derived by the following equation. Note that these equations are different from ones specified in [\[ISO.11770-4.2006\]](#).

$$\begin{aligned} o\_A &= H(\text{octet}(04) \mid \text{OCTETS}(w\_A) \mid \text{OCTETS}(w\_B) \mid \text{OCTETS}(z) \mid \text{VI}(\text{nc}) \mid \text{VS}(v)) \\ o\_B &= H(\text{octet}(03) \mid \text{OCTETS}(w\_A) \mid \text{OCTETS}(w\_B) \mid \text{OCTETS}(z) \mid \text{VI}(\text{nc}) \mid \text{VS}(v)) \end{aligned}$$

## 8.2. Functions for discrete-logarithm settings

In this section, the equation  $(x / y \text{ mod } z)$  denotes a natural number  $w$  less than  $z$  which satisfies  $(w * y) \text{ mod } z = x \text{ mod } z$ .

For the discrete-logarithm, we refer some of the domain parameters by the following symbols:

- q: for "the prime" of the group.
- g: for "the generator" associated with the group.
- r: for the order of the subgroup generated by g.

The function J is defined as

$$J(\pi) = g^{\pi} \bmod q.$$

The value of w\_A is derived as

$$w_A = g^{s_A} \bmod q,$$

where s\_A is a random integer within range [1, r-1] and r is the size of the subgroup generated by g. In addition, s\_A MUST be larger than  $\log(q)/\log(g)$  (so that  $g^{s_A} > q$ ).

The value of w\_A SHALL satisfy  $1 < w_A < q-1$ . The server MUST check this condition upon reception.

The value of w\_B is derived from J(pi) and w\_A as:

$$w_B = (J(\pi) * w_A^{(H(\text{octet}(1) | \text{OCTETS}(w_A)))))^{s_B} \bmod q,$$

where s\_B is a random number within range [1, r-1]. The value of w\_B MUST satisfy  $1 < w_B < q-1$ . If this condition is not hold, the server MUST retry with another value of s\_B. The client MUST check this condition upon reception.

The value z in the client side is derived by the following equation:

$$z = w_B^{(s_A + H(\text{octet}(2) | \text{OCTETS}(w_A) | \text{OCTETS}(w_B))) / (s_A * H(\text{octet}(1) | w_A) + \pi) \bmod r) \bmod q.$$

The value z in the server side is derived by the following equation:

$$z = (w_A * g^{(H(\text{octet}(2) | \text{OCTETS}(w_A) | \text{OCTETS}(w_B)))))^{s_B} \bmod q.$$

### 8.3. Functions for elliptic-curve settings

For the elliptic-curve setting, we refer some of the domain parameters by the following symbols:

- q: for the prime used to define the field,
- G: for the defined point called the generator,
- r: for the order of the subfield generated by G.

The function P(p) converts a curve point p to an integer representing the point p, by computing  $x * 2 + (y \bmod 2)$ , where (x, y) are the coordinates of the point p. P'(z) is the inverse of function P, that is, it converts an integer z to a point p which satisfies  $P(p) = z$ . If such p is exist, it is uniquely defined. Otherwise, z does not represent a valid curve point. The operation  $[x] * p$  denotes an integer-multiplication of point p: it calculates  $p + p + \dots$  (x times)  $\dots + p$ . See literatures on elliptic-curve cryptography for the exact algorithms for those. 0\_E represents the infinity point. The equation  $(x / y \bmod z)$  denotes a natural number w less than z which satisfies  $(w * y) \bmod z = x \bmod z$ .

the function J is defined as

$$J(\text{pi}) = [\text{pi}] * G.$$

The value of  $w\_A$  is derived as

$$w\_A = P(W\_A), \text{ where } W\_A = [s\_A] \times G.$$

where  $s\_A$  is a random number within range [1, r-1]. The value of  $w\_A$  MUST represent a valid curve point, and  $W\_A$  SHALL NOT be  $0\_E$ . The server MUST check this condition upon reception.

The value of  $w\_B$  is derived from  $J(\text{pi})$  and  $W\_A = P'(w\_A)$  as:

$$w\_B = P(W\_B), \text{ where } W\_B = [s\_B] * (J(\text{pi}) + [H(\text{octet}(1) | \text{OCTETS}(w\_A))] * W\_A).$$

where  $s\_B$  is a random number within range [1, r-1]. The value of  $w\_B$  MUST represent a valid curve point and satisfy  $[4] * P'(w\_B) \neq 0\_E$ . If this condition does not hold, the server MUST retry with another value of  $s\_B$ . The client MUST check this condition upon reception.

The value  $z$  in the client side is derived by the following equation:

$$z = P([(s\_A + H(\text{octet}(2) | \text{OCTETS}(w\_A) | \text{OCTETS}(w\_B))) / (s\_A * H(\text{octet}(1) | \text{OCTETS}(w\_A)) + \text{pi}) \bmod r] * W\_B), \text{ where } W\_B = P'(w\_B).$$

The value  $z$  in the server side is derived by the following equation:

$$z = P([s\_B] * (W\_A + [H(\text{octet}(2) | \text{OCTETS}(w\_A) | \text{OCTETS}(w\_B))] * G)), \text{ where } W\_A = P'(w\_A).$$

## 9. Authentication Realms

In this protocol, an "authentication realm" is defined as a set of resources (URIs) for which the same set of user names and passwords is valid for. If the server requests authentication for the authentication realm which the client is already authenticated, the client will automatically perform authentication using the already-known secrets. On the contrary, for the different authentication realms, clients SHOULD NOT automatically reuse the usernames and passwords for another realm.

Just like Basic and Digest access authentication protocol, Mutual authentication protocol supports multiple, separate authentication realms to be set up inside each host. Furthermore, the protocol supports that a single authentication realm spans over several hosts in the same Internet domain.

Each authentication realm is defined and distinguished by the triple of an "authentication algorithm", an "authentication domain", a "realm" parameter. Server operators are NOT RECOMMENDED to use the same pair of an authentication domain and a realm for different authentication algorithms, however.

Authentication algorithms are defined in [Section 4](#) and [Section 8](#). Realm parameters are just a string, as defined in [Section 4](#). Authentication domains are described in the rest of this section.

An authentication domain specifies the range of hosts which the authentication realm spans over. In the protocol, it MUST currently be one of the following strings.

- the string in format "<scheme>://<host>:<port>", where scheme, host and port are the URI parts of the requested URI. Even if the request-URI does not have a port part, the string will include the one (i.e. 80 for http and 443 for https). Use this when authentication is only valid for specific protocol (such as https).
- The "host" part of the requested URI. This is the default value. Authentication realms in this kind of authentication domain will span over several protocols (i.e. http and https) and ports, but not over different hosts.
- String in format "\*.<domain-postfix>", where "domain-postfix" is either the host part of the requested URI, or any domain in which the requested host is included (this means that the specification "\*.example.com" is valid for all of hosts "www.example.com", "web.example.com" and "example.com"). The domain-postfix must be equal to or included in a valid Internet domain assigned to specific organization: if the clients can know by some way (such as blacklists for HTTP cookies) that the specified domain is not to be assigned to any specific organization (e.g. "\*.com" or "\*.jp"), the client is RECOMMENDED to reject the authentication request.

In the above specifications, every "scheme", "host" and "domain" MUST be in lower-case, and IDNs MUST be represented in puny-code [RFC3492]. All "port"s MUST be in the shortest, unsigned, decimal number notation. Not obeying these requirements will cause failure of valid authentication attempts.

## 9.1. Resolving ambiguities

In the above definition of authentication domains, several domains will overlap each other. Depending on the "path" parameters given in the "401-B1" message (see Section 4), There may be several candidates when the client is to send a request with authentication credentials included (at the Steps 3 and 4 of the decision procedure shown in Section 5).

If such choices are required, the following procedure SHOULD be followed.

- If the client has previously sent a request to the same URI, and it remembers the authentication realm requested by 401-B0 messages at that time, use that realm.
- In other cases, use one of authentication realms which represents most-specific authentication domains. In the list of possible domain specifications shown above, one described earlier has priority over ones described after that.  
If there are several choices with different domain-postfix specifications, the one which has longer domain possible has priority over ones with shorter domain-postfix.
- If there are realms with the same specifications of authentication domain, there is no defined priority: client MAY choose any one of possible choices.

If possible, server operators are recommended to avoid such ambiguities by setting "path" parameters properly.

## 10. Validation Methods

The "validation method" specifies a method to "relate" the mutual authentication processed by this protocol with other authentications already performed in the underlying layers and to prevent man-in-the-middle attacks. It decides the value of v which is an input to authentication protocols.

The valid tokens for the validation field and corresponding values of v are as follows:

host:

hostname validation: v will be the ASCII string in the following format: "scheme://host:port", where scheme, host and port are the URI parts correspond to the currently accessing resource. The scheme and host are lower-case, and the port is in a shortest decimal representation. Even if the request-URI does not have a port part, v will include the one.

tls-cert:

TLS certificate validation: v will be the octet string of the hash value of the public key certificate used in underlying TLS [RFC5246] (or SSL) connection. The hash value is defined as the value of the whole signed certificate (specified as "Certificate" in [RFC5280]), hashed by the hash algorithm specified by the authentication algorithm used.

tls-key:

TLS shared-key validation: v will be the octet string of the shared master secret negotiated in underlying TLS (or SSL) connection.

If the HTTP protocol is used on unencrypted channel, the validation type MUST be "host". If HTTP/TLS [RFC2818] (https) protocol is used with server certificates, the validation type MUST be either "tls-cert" or "tls-key". If HTTP/TLS protocol is used with anonymous Diffie-Hellman key exchange, the validation type MUST be "tls-key" (but see the note below).

Clients MUST validate this field upon reception of 401-B0 messages.

However, when the protocol is used on web browsers with any scripting capabilities, the anonymous Diffie-Hellman family of TLS (or SSL) cipher-suite MUST NOT be used even if "tls-key" validated Mutual authentication has been employed, and the certificate shown in TLS (or SSL) negotiation MUST be verified using PKI. For other systems, if the "tls-key" validation is used on TLS (or SSL) protocol without certificate verification using PKI, those systems MUST ensure that all transactions with authenticated peer servers MUST use and be validated by the Mutual authentication protocol, regardless of the existence of the 401-B0 responses.

The protocol defines two variants for validation on TLS connections. The method "tls-key" method is more secure. However, there are some situations where tls-cert is more preferable.

- When TLS accelerating proxies are used. In this case, it is difficult for the authenticating server to acquire the TLS key information which is used between the client and the proxy. It is not the case for client-side "tunneling" proxies using CONNECT method extension of HTTP.
- When a black-box implementation of the TLS protocol is used on either peer.

Implementations supporting Mutual authentication over https protocol SHOULD support "tls-cert" validation unless it is not applicable. Support for "tls-key" validation is OPTIONAL for both servers and clients.

## 11. Session Management

In the Mutual authentication protocol, a session represented by a sid is generated By the first 4 messages (first request, 401-B0, req-A1 and 401-B1). This session can be used for one or more requests for resources protected by the same realm in the same server. Note that the session management is only an inside detail of the protocol and usually not visible to normal users. If a session expires, the client and server will automatically reestablish another session without telling it to the users.

The server SHOULD accept at least one req-A3 request for each session, given that the request reaches the server in a time window specified by the timeout field in the 401-B1 message, and that there are no emergent reasons (such as flooding attacks) to forget the sessions. After that, the server MAY discard any session at any time and MAY send 401-B0-stale messages for any req-A3 requests.

The client MAY send two or more requests using a single session specified by the sid. However, for all such requests, the values of the nonce-counter (nc field) MUST be different from each other. The server MUST check for duplication of the received nonces, and if any duplication is detected, the server MUST discard the session and respond by a 401-B0-stale message.

In addition, for each session, if the client has already sent a request with nonce value  $x$ , it SHOULD NOT send requests with a nonce value not larger than  $(x - \text{nc-window})$ . The server MAY reject any requests with nonces violating this rule with 401-B0-stale responses. This restriction enables servers to implement duplicated nonce detection in a constant memory.

Values of nonces and nonce-related values MUST always be treated as natural numbers within infinite range. Implementations using fixed-width integers or fixed-precision floating numbers MUST handle integer overflow correctly and carefully. Such implementations are RECOMMENDED to accept any larger values which cannot be represented in the fixed-width integer representations, as long as other limits such as internal header-length restrictions are not involved. The protocol is designed carefully so that both clients and servers can implement the protocol only with fixed-width integers, by rounding any overflowed values to the maximum possible value.

## 12. Methods to extend this protocol

If a non-standard extension to the this protocol is implemented, it MUST use the extension-tokens defined in [Section 3](#) to avoid conflicts with this protocol and other extensions.

Authentication algorithms other than those defined in this document MAY use other representations for keys "wa", "wb", "oa" and "ob", replace those keys, and/or add fields to the messages containing those fields by supplemental specifications. If those specifications use keys other than shown above, it is RECOMMENDED to use extension-tokens to avoid any key-name conflict with the future extension of this protocol.

Extension-tokens MAY be freely used for any non-standard, private and/or experimental uses for those fields provided that the domain part in the token is appropriately used.

## 13. IANA Considerations

The tokens used for authentication-algorithm, pwd-hash, and validation fields MUST be allocated by IANA. To acquire registered tokens, a specification for the use of such tokens MUST be available as an RFC, as outlined in [\[RFC5226\]](#).

Note: More formal declarations will be added in future drafts to meet RFC 5226 requirements.

## 14. Security Considerations

## 14.1. General Assumptions

- The protocol is secure against passive eavesdropping and replay attacks. However, the protocol relies on transport security including DNS security for active attacks. HTTP/TLS SHOULD be used where transport security is not assured and data secrecy is important.
- Used with HTTP/TLS, if TLS server certificates are reliably verified, the protocol gives true protection against active man-in-the-middle attacks.
- Even if the server certificate is not used or is unreliable, the protocol gives protection against active man-in-the-middle attacks for each HTTP request/response pair. However, in such cases, JavaScript or similar scripting facilities can be used to affect Mutually-authenticated contents from other contents not protected by this authentication mechanism. This is the reason why this protocol requires that valid TLS server certificates MUST be presented ([Section 10](#)).

## 14.2. Implementation Considerations

- To securely implement the protocol, the Authentication-Info headers in the 200-B4 messages MUST always be validated by the client. If the validation fails, the client MUST NOT process any content sent with the message, including the body part. Non-compliance to this will enable phishing attacks.
- The authentication status on the client-side SHOULD be visible to the users of the client. In addition, the method for asking user's name and passwords SHOULD be carefully designed so that (1) the user can easily distinguish request of this authentication methods from other existing authentication methods such as Basic and Digest methods, and (2) the Web contents cannot imitate the user-interfaces of this protocol.  
An informational memo regarding user-interface considerations and recommendations for implementing this protocol will be separately published.
- For HTTP/TLS communications, when a web form is submitted from Mutually-authenticated pages with the validation methods of "tls-cert" to a URI which is protected by the same realm (so indicated by the path field), if server certificate has been changed since the pages has been received, the peer is RECOMMENDED to be revalidated using a req-A1 message with an "Expect: 100-continue" header. The same applies when the page is received with the validation methods of "tls-key", and when the TLS session has been expired.
- Server-side storage of user passwords are advised to have the values encrypted by one-way function  $J(\pi)$ , instead of the real passwords, those hashed by  $ph$ , or  $\pi$ .

## 14.3. Usage Considerations

- The user-names inputted by user may be sent automatically to any servers sharing the same auth-domain. This means that when host-type auth-domain is used for authentication in HTTPS site, and when an HTTP server on the same host requests Mutual authentication with the same realm, the client will send the user-name in a clear text. If user-names have to be kept secret against eavesdropping, the server must use full-scheme-type auth-domain parameter. On the contrary, passwords are not exposed to eavesdroppers even on HTTP requests.
- "Pwd\_hash" field is only provided for backward compatibility for password databases, and using "none" function is the mostly secure choice and RECOMMENDED. If values other than "none" is used, you must ensure that the hash values of the passwords were not exposed to the public. Note that hashed password databases for plain-text authentications are usually not considered secret.
- If the server provides several ways of storing server-side password database, it is advised to store the values encrypted by one-way function  $J(\pi)$ , instead of the real passwords, those hashed by  $ph$ ,

or pi.

## 15. Notice on intellectual properties

The National Institute of Advanced Industrial Science and Technology (AIST) and Yahoo! Japan, Inc. has jointly submitted a patent application about the protocol proposed in this documentation to the Patent Office of Japan. The patent is intended to be open to any implementors of this protocol and its variants under non-exclusive royalty-free manner. For the detail of the patent application and its status, please contact the author of this document.

The elliptic-curve based authentication algorithms might involve several existing patents of third-parties. The authors of the document take no position regarding the validity or scope of such patents, and other patents as well.

## 16. Acknowledgement

We gratefully acknowledge Lepidum, Co. Ltd. for support on design and trial implementation of this protocol.

## 17. References

### 17.1. Normative References

- [FIPS.180-2.2002] National Institute of Standards and Technology, "Secure Hash Standard," FIPS PUB 180-2, August 2002.
- [FIPS.186-2.2000] National Institute of Standards and Technology, "Digital Signature Standard (DSS)," FIPS PUB 186-2, January 2000.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TEXT, HTML, XML).
- [RFC2818] Rescorla, E., "HTTP Over TLS," RFC 2818, May 2000 (TEXT).
- [RFC3526] Kivinen, T. and M. Kojo, "More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)," RFC 3526, May 2003 (TEXT).
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646," STD 63, RFC 3629, November 2003 (TEXT).
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings," RFC 4648, October 2006 (TEXT).
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," STD 68, RFC 5234, January 2008 (TEXT).
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246, August 2008 (TEXT).



## 17.2. Informative References

- [I-D.altman-tls-channel-bindings] Altman, J., Williams, N., and L. Zhu, "Channel Bindings for TLS," draft-altman-tls-channel-bindings-07 (work in progress), October 2009 (TXT).
- [ISO.10646-1.1993] International Organization for Standardization, "Information Technology - Universal Multiple-octet coded Character Set (UCS) - Part 1: Architecture and Basic Multilingual Plane," ISO Standard 10646-1, May 1993.
- [ISO.11770-4.2006] International Organization for Standardization, "Information technology – Security techniques – Key management – Part 4: Mechanisms based on weak secrets," ISO Standard 11770-4, May 2006.
- [ITU.X690.1994] International Telecommunications Union, "Information Technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)," ITU-T Recommendation X.690, 1994.
- [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 (TXT, PS, PDF, HTML, XML).
- [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 (TXT, HTML, XML).
- [RFC3492] Costello, A., "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)," RFC 3492, March 2003 (TXT).
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," BCP 26, RFC 5226, May 2008 (TXT).
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280, May 2008 (TXT).

## Appendix A. Group parameters for discrete-logarithm based algorithms

The MODP group used for the iso11770-4-dl-2048 algorithm is defined by the following parameters.

The prime is:

```
q = 0xFFFFFFFF FFFFFFFF C90FDAA2 2168C234 C4C6628B 80DC1CD1
    29024E08 8A67CC74 020BBEA6 3B139B22 514A0879 8E3404DD
    EF9519B3 CD3A431B 302B0A6D F25F1437 4FE1356D 6D51C245
    E485B576 625E7EC6 F44C42E9 A637ED6B 0BFF5CB6 F406B7ED
    EE386BFB 5A899FA5 AE9F2411 7C4B1FE6 49286651 ECE45B3D
    C2007CB8 A163BF05 98DA4836 1C55D39A 69163FA8 FD24CF5F
    83655D23 DCA3AD96 1C62F356 208552BB 9ED52907 7096966D
    670C354E 4ABC9804 F1746C08 CA18217C 32905E46 2E36CE3B
    E39E772C 180E8603 9B2783A2 EC07A28F B5C55DF0 6F4C52C9
    DE2BCBF6 95581718 3995497C EA956AE5 15D22618 98FA0510
    15728E5A 8AACAA68 FFFFFFFF FFFFFFFF.
```

The generator is:

$g = 2$ .

The size of the subgroup generated by  $g$  is:

```
r = (q - 1) / 2 =
0x7FFFFFFF FFFFFFFF E487ED51 10B4611A 62633145 C06E0E68
    94812704 4533E63A 0105DF53 1D89CD91 28A5043C C71A026E
    F7CA8CD9 E69D218D 98158536 F92F8A1B A7F09AB6 B6A8E122
    F242DABB 312F3F63 7A262174 D31BF6B5 85FFAE5B 7A035BF6
    F71C35FD AD44CFD2 D74F9208 BE258FF3 24943328 F6722D9E
    E1003E5C 50B1DF82 CC6D241B 0E2AE9CD 348B1FD4 7E9267AF
    C1B2AE91 EE51D6CB 0E3179AB 1042A95D CF6A9483 B84B4B36
    B3861AA7 255E4C02 78BA3604 650C10BE 19482F23 171B671D
    F1CF3B96 0C074301 CD93C1D1 7603D147 DAE2AEF8 37A62964
    EF15E5FB 4AAC0B8C 1CCAA4BE 754AB572 8AE9130C 4C7D0288
    0AB9472D 45565534 7FFFFFFF FFFFFFFF.
```

The MODP group used for the iso11770-4-dl-4096 algorithm is defined by the following parameters.

The prime is:

```
q = 0xFFFFFFFF FFFFFFFF C90FDAA2 2168C234 C4C6628B 80DC1CD1
    29024E08 8A67CC74 020BBEA6 3B139B22 514A0879 8E3404DD
    EF9519B3 CD3A431B 302B0A6D F25F1437 4FE1356D 6D51C245
    E485B576 625E7EC6 F44C42E9 A637ED6B 0BFF5CB6 F406B7ED
    EE386BFB 5A899FA5 AE9F2411 7C4B1FE6 49286651 ECE45B3D
    C2007CB8 A163BF05 98DA4836 1C55D39A 69163FA8 FD24CF5F
    83655D23 DCA3AD96 1C62F356 208552BB 9ED52907 7096966D
    670C354E 4ABC9804 F1746C08 CA18217C 32905E46 2E36CE3B
    E39E772C 180E8603 9B2783A2 EC07A28F B5C55DF0 6F4C52C9
    DE2BCBF6 95581718 3995497C EA956AE5 15D22618 98FA0510
    15728E5A 8AAAC42D AD33170D 04507A33 A85521AB DF1CBA64
    ECFB8504 58DBEF0A 8AEA7157 5D060C7D B3970F85 A6E1E4C7
    ABF5AE8C DB0933D7 1E8C94E0 4A25619D CEE3D226 1AD2EE6B
    F12FFA06 D98A0864 D8760273 3EC86A64 521F2B18 177B200C
    BBE11757 7A615D6C 770988C0 BAD946E2 08E24FA0 74E5AB31
    43DB5BFC E0FD108E 4B82D120 A9210801 1A723C12 A787E6D7
    88719A10 BDBA5B26 99C32718 6AF4E23C 1A946834 B6150BDA
    2583E9CA 2AD44CE8 DBBBC2DB 04DE8EF9 2E8EFC14 1FBECAA6
    287C5947 4E6BC05D 99B2964F A090C3A2 233BA186 515BE7ED
    1F612970 CEE2D7AF B81BDD76 2170481C D0069127 D5B05AA9
    93B4EA98 8D8FDDC1 86FFB7DC 90A6C08F 4DF435C9 34063199
    FFFFFFFF FFFFFFFF.
```

The generator is:

$$g = 2.$$

The size of the subgroup generated by g is:

```

r = (q - 1) / 2 =
0x7FFFFFFF FFFFFFFF E487ED51 10B4611A 62633145 C06E0E68
 94812704 4533E63A 0105DF53 1D89CD91 28A5043C C71A026E
 F7CA8CD9 E69D218D 98158536 F92F8A1B A7F09AB6 B6A8E122
 F242DABB 312F3F63 7A262174 D31BF6B5 85FFAE5B 7A035BF6
 F71C35FD AD44CFD2 D74F9208 BE258FF3 24943328 F6722D9E
 E1003E5C 50B1DF82 CC6D241B 0E2AE9CD 348B1FD4 7E9267AF
 C1B2AE91 EE51D6CB 0E3179AB 1042A95D CF6A9483 B84B4B36
 B3861AA7 255E4C02 78BA3604 650C10BE 19482F23 171B671D
 F1CF3B96 0C074301 CD93C1D1 7603D147 DAE2AEF8 37A62964
 EF15E5FB 4AAC0B8C 1CCAA4BE 754AB572 8AE9130C 4C7D0288
 0AB9472D 45556216 D6998B86 82283D19 D42A90D5 EF8E5D32
 767DC282 2C6DF785 457538AB AE83063E D9CB87C2 D370F263
 D5FAD746 6D8499EB 8F464A70 2512B0CE E771E913 0D697735
 F897FD03 6CC50432 6C3B0139 9F643532 290F958C 0BBD9006
 5DF08BAB BD30AEB6 3B84C460 5D6CA371 047127D0 3A72D598
 A1EDADFE 707E8847 25C16890 54908400 8D391E09 53C3F36B
 C438CD08 5EDD2D93 4CE1938C 357A711E 0D4A341A 5B0A85ED
 12C1F4E5 156A2674 6DDDE16D 826F477C 97477E0A 0FDF6553
 143E2CA3 A735E02E CCD94B27 D04861D1 119DD0C3 28ADF3F6
 8FB094B8 67716BD7 DC0DEEBB 10B8240E 68034893 EAD82D54
 C9DA754C 46C7EEEE C37FDBEE 48536047 A6FA1AE4 9A0318CC
 FFFFFFFF FFFFFFFF .

```

## Appendix B. Derived numerical values

This section gives several numerical values for implementing this protocol, derived from the above specifications. The values shown in this section are for informative purpose only.

	<b>dl-2048</b>	<b>dl-4096</b>	<b>ec-p256</b>	<b>ec-p521</b>	
Size of w_A etc.	2048	4096	257	522	(bits)
Size of H(...)	256	512	256	512	(bits)
length of OCTETS(w_A) etc.	256	512	33	66	(octets)
length of wa, wb field values.	346 *	686 *	66	132	(octets)
length of oa, ob field values.	46 *	90 *	64	128	(octets)
minimum allowed s_A	2048	4096	1	1	

(The numbers marked with \* include enclosing quotation marks.)

## Appendix C. Draft Remarks from the Authors

The following items are currently under consideration for future revisions by the authors.

- Restructuring of the draft, possibly separating it to several parts, e.g. introduction, general HTTP extensions and Mutual authentication.
- Format of the "Authentication-Control" header and other header fields extending the general

HTTP authentication scheme, and harmonization of those with other draft proposals.

- Whether to keep TLS-key validation or not.
- When keeping tls-key validation, whether to use "TLS channel binding" [I-D.altman-tls-channel-bindings] for "tls-key" verification (Section 10). Note that existing implementations of TLS should be considered to determine this.
- Adding test vectors for ensuring implementation correctness.
- Possibly adding a method for servers to detect availability of Mutual authentication on client-side.
- Applying the protocol for proxy authentication/authorization.

## Appendix D. Draft Change Log

### D.1. Changes in revision 06

- Integrated Optional Mutual Authentication to the main part.
- Clarified the decision procedure for message recognitions.
- Clarified that a new authentication request for any sub-requests in interactive clients may be silently discarded.
- Typos and confusing phrases are fixed.
- Several "future considerations" are added.

The field "version" is NOT changed from the previous draft, as the semantics has not been changed.

### D.2. Changes in revision 05

- A new field "version" is added for supporting future incompatible changes with a single implementation. In the (first) final specification its value will be changed to 1.
- A new header "Authentication-Control" added for precise control of application-level authentication behavior.

### D.3. Changes in revision 04

- Changed text of patent licenses: the phrase "once the protocol is accepted as an Internet standard" is removed so that the sentence also covers the draft versions of this protocol.
- The "tls-key" verification is now OPTIONAL.
- Several description fixes and clarifications.

### D.4. Changes in revision 03

- Wildcard domain specifications (e.g. "\*.example.com") is allowed for auth-domain parameters (Section 4.1).
- Specification of the "tls-host" verification is updated (incompatible change).
- State transitions fixed.
- Requirements for servers about w\_a values clarified.
- RFC references are updated.

## D.5. Changes in revision 02

- Auth-realm is extended to allow full-scheme type.
- A decision diagram for clients and decision procedures for servers are added.
- 401-B1 and req-A3 messages is changed to have authentication realm information.
- Bugs on equations for o\_A and o\_B is fixed.
- Detailed equations for the whole algorithm is included.
- Elliptic-curve algorithms are updated.
- Several clarifications and other minor updates.

## Authors' Addresses

Yutaka Oiwa  
National Institute of Advanced Industrial Science and Technology  
Research Center for Information Security  
Akihabara Daibiru #1003  
1-18-13 Sotokanda  
Chiyoda-ku, Tokyo  
JP

Phone: +81 3-5298-4722

Email: [mutual-auth-contact@m.aist.go.jp](mailto:mutual-auth-contact@m.aist.go.jp)

Hajime Watanabe  
National Institute of Advanced Industrial Science and Technology

Hiromitsu Takagi  
National Institute of Advanced Industrial Science and Technology

Hirofumi Suzuki  
Yahoo! Japan, Inc.  
Midtown Tower  
9-7-1 Akasaka  
Minato-ku, Tokyo  
JP

Phone: +81 3-6440-6290