

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: January 16, 2014

S. Farrell  
Trinity College Dublin  
P. Hoffman  
VPN Consortium  
M. Thomas  
Phresheez  
July 15, 2013

**HTTP Origin-Bound Authentication (HOBA)**  
**draft-ietf-httpauth-hoba-01**

**Abstract**

HTTP Origin-Bound Authentication (HOBA) is a design for an HTTP authentication method with credentials that are not vulnerable to phishing attacks, and that does not require a server-side password database. The design can also be used in Javascript-based authentication embedded in HTML. HOBA is an alternative to HTTP authentication schemes that require passwords with all the negative attributes that come with password-based systems.

**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 January 16, 2014.

**Copyright Notice**

Copyright (c) 2013 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.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">4</a>
<a href="#">1.1.</a>	<a href="#">Interfacing to Applications (Cookies) . . . . .</a>	<a href="#">5</a>
<a href="#">1.2.</a>	<a href="#">Terminology . . . . .</a>	<a href="#">6</a>
<a href="#">2.</a>	<a href="#">The HOBA Authentication Scheme . . . . .</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">HOBA HTTP Authentication Mechanism . . . . .</a>	<a href="#">8</a>
<a href="#">4.</a>	<a href="#">Using HOBA-http . . . . .</a>	<a href="#">9</a>
<a href="#">4.1.</a>	<a href="#">CPK Preparation Phase . . . . .</a>	<a href="#">10</a>
<a href="#">4.2.</a>	<a href="#">Signing Phase . . . . .</a>	<a href="#">10</a>
<a href="#">4.3.</a>	<a href="#">Authentication Phase . . . . .</a>	<a href="#">10</a>
<a href="#">4.4.</a>	<a href="#">Logging in on a New User Agent . . . . .</a>	<a href="#">11</a>
<a href="#">5.</a>	<a href="#">Using HOBA-js . . . . .</a>	<a href="#">11</a>
<a href="#">5.1.</a>	<a href="#">Key Storage . . . . .</a>	<a href="#">12</a>
<a href="#">5.2.</a>	<a href="#">User Join . . . . .</a>	<a href="#">12</a>
<a href="#">5.3.</a>	<a href="#">User Login . . . . .</a>	<a href="#">12</a>
<a href="#">5.4.</a>	<a href="#">Enrolling a New User Agent . . . . .</a>	<a href="#">12</a>
<a href="#">5.5.</a>	<a href="#">Replay Protection . . . . .</a>	<a href="#">13</a>
<a href="#">5.6.</a>	<a href="#">Signature Parameters . . . . .</a>	<a href="#">13</a>
<a href="#">5.7.</a>	<a href="#">Session Management . . . . .</a>	<a href="#">15</a>
<a href="#">5.8.</a>	<a href="#">Multiple Accounts on One User Agent . . . . .</a>	<a href="#">15</a>
<a href="#">5.9.</a>	<a href="#">Oddities . . . . .</a>	<a href="#">16</a>
<a href="#">6.</a>	<a href="#">Additional Services . . . . .</a>	<a href="#">16</a>
<a href="#">6.1.</a>	<a href="#">Registration . . . . .</a>	<a href="#">16</a>
<a href="#">6.2.</a>	<a href="#">Associating Additional Keys to an Existing Account . . . . .</a>	<a href="#">17</a>
<a href="#">6.3.</a>	<a href="#">Logging Out . . . . .</a>	<a href="#">18</a>
<a href="#">6.4.</a>	<a href="#">Getting a Fresh Challenge . . . . .</a>	<a href="#">18</a>
<a href="#">7.</a>	<a href="#">Mandatory-to-Implement Algorithms . . . . .</a>	<a href="#">18</a>
<a href="#">8.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">18</a>
<a href="#">8.1.</a>	<a href="#">localStorage Security for Javascript . . . . .</a>	<a href="#">19</a>
<a href="#">9.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">19</a>
<a href="#">9.1.</a>	<a href="#">HOBA Authentication Scheme . . . . .</a>	<a href="#">20</a>
<a href="#">9.2.</a>	<a href="#">.well-known URLs . . . . .</a>	<a href="#">20</a>
<a href="#">9.3.</a>	<a href="#">Algorithm Names . . . . .</a>	<a href="#">20</a>
<a href="#">9.4.</a>	<a href="#">Key Identifier Types . . . . .</a>	<a href="#">20</a>
<a href="#">9.5.</a>	<a href="#">Device Identifier Types . . . . .</a>	<a href="#">20</a>
<a href="#">10.</a>	<a href="#">Implementation Status . . . . .</a>	<a href="#">20</a>
<a href="#">11.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">21</a>
<a href="#">12.</a>	<a href="#">References . . . . .</a>	<a href="#">21</a>
<a href="#">12.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">21</a>
<a href="#">12.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">22</a>
<a href="#">Appendix A.</a>	<a href="#">Problems with Passwords . . . . .</a>	<a href="#">23</a>
<a href="#">Appendix B.</a>	<a href="#">Examples . . . . .</a>	<a href="#">23</a>
<a href="#">Appendix C.</a>	<a href="#">Changes . . . . .</a>	<a href="#">25</a>
<a href="#">C.1.</a>	<a href="#">WG-01 . . . . .</a>	<a href="#">25</a>
<a href="#">C.2.</a>	<a href="#">WG-00 . . . . .</a>	<a href="#">25</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">25</a>



## 1. Introduction

[[ Commentary is in double-square brackets, like this. As you'll see there are a bunch of details still to be figured out. Feedback on those is very welcome. Also note that the authors fully expect that the description of HOBA-http and HOBA-js to be mostly merged in the draft; they're both here now so readers can see some alternatives and maybe support particular proposals. ]]

HTTP Origin-Bound Authentication (HOBA) is an authentication design that can be used as an HTTP authentication scheme and for Javascript-based authentication embedded in HTML. The main goal of HOBA is to offer an easy-to-implement authentication scheme that is not based on passwords, but that can easily replace HTTP or HTML forms-based password authentication. If deployment of HOBA reduces the number of password entries in databases by any appreciable amount, then it would be worthwhile. As an HTTP authentication scheme, it would work in the current HTTP 1.0 and HTTP 1.1 authentication framework, and will very likely work with whatever changes are made to the HTTP authentication scheme in HTTP 2.0. As a JavaScript design, HOBA demonstrates a way for clients and servers to interact using the same credentials that are used by the HTTP authentication scheme.

The HTTP specification defines basic and digest authentication methods for HTTP that have been in use for many years, but which, being based on passwords, are susceptible to theft of server-side databases. (See [[RFC2617](#)] for the original specification, and [[I-D.ietf-httpbis-p7-auth](#)] for clarifications and updates to the authentication mechanism.) Even though few large web sites use basic and digest authentication, they still use username/password authentication and thus have large susceptible server-side databases of passwords.

Instead of passwords, HOBA uses digital signatures as an authentication mechanism. HOBA also adds useful features such as credential management and session logout. In HOBA, the client creates a new public-private key pair for each host ("web-origin") to which it authenticates; web-origins are defined in [[RFC6454](#)]. These keys are used in HOBA for HTTP clients to authenticate themselves to servers in the HTTP protocol or in a Javascript authentication program.

HOBA keys need not be stored in public key certificates. Because these are generally "bare keys", there is none of the semantic overhead of PKIX certificates, particularly with respect to naming and trust anchors. Thus, client public keys ("CPKs") do not have any publicly-visible identifier for the user who possesses the corresponding private key, nor the web-origin with which the client



is using the CPK. HOBA keys are stored in `subjectPublicKeyInfo` structures from PKIX [[RFC5280](#)]

HOBA also defines some services that are required for modern HTTP authentication:

- o Servers can bind a CPK with an identifier, such as an account name. HOBA allows servers to define their own policies for binding CPKs with accounts during account registration.
- o Users are likely to use more than one device or user agent (UA) for the same HTTP based service, so HOBA gives a way to associate more than one CPK to the same account, but without having to register for each separately.
- o Users are also likely to lose a private key, or the client's memory of which key pair is associated with which origin. For example if a user loses the computer or mobile device in which state is stored. HOBA allows for clients to tell servers to delete the association between an existing CPK and an account.
- o Logout features can be useful for UAs, so HOBA defines a way to close a current HTTP "session", and also a way to close all current sessions, even if more than one session is currently active from different UAs for the same account.
- o Since there are always devices and applications in which state of the art digital signature mechanism runtimes are significant, and since HTTP authentication in theory requires that every HTTP request to a given realm have a signature in an "Authorization" header field, and since HOBA is a challenge response scheme, we also define a way in which HTTP servers can indicate the duration for which they will consider a given challenge value to be valid. As a consequence we also define a way for UAs to fetch a fresh challenge.

### **1.1. Interfacing to Applications (Cookies)**

HOBA can be used as a drop-in replacement for password-based user authentication schemes used in common web applications. The simplest way in which this can be done is to (re-)direct the UA to a HOBA "Login" URL and for the response to a successful HTTP request containing a HOBA signature to set a session cookie [[RFC6265](#)]. Further interactions with the web application will then be secured via the session cookie, as is commonly done today.

While cookies are bearer tokens, and thus weaker than HOBA signatures, they are currently ubiquitously used. If non-bearer





token session continuation schemes are developed in future in the IETF or elsewhere, then those can interface to HOBA as easily as with any password based authentication scheme.

## **1.2. Terminology**

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 [RFC 2119](#) [[RFC2119](#)].

A client public key ("CPK") is the public key and associated cryptographic parameters needed for a server to validate a signature.

The term "account" is (loosely) used to refer to whatever data structure(s) the server maintains that are associated with an identity. That will contain of at least one CPK and a web-origin; it will also optionally include an HTTP "realm" as defined in the HTTP authentication specification. It might also involve many other non-standard pieces of data that the server accumulates as part of account creation processes. An account may have many CPKs that are considered equivalent in terms of being usable for authentication, but the meaning of "equivalent" is really up to the server and is not defined here.

When describing something that is specific to HOBA as an HTTP authentication mechanism or HOBA as a JavaScript implementation, this document uses the terms "HOBA-http" and "HOBA-js", respectively.

Web client: the content and javascript code that run within the context of a single UA instance (such as a tab in a web browser).

User agent (UA): typically, but not always, a web browser.

User: a person who is running a UA. In this document, "user" does not mean "user name" or "account name".

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

## **2. The HOBA Authentication Scheme**

A UA that implements HOBA maintains a list of web-origins and realms. The UA also maintains one or more client credentials for each web-origin/realm combination for which it has created a CPK.

On receipt of a challenge (and optional realm) from a server, the client marshals an HOBA to-be-signed (TBS) blob that includes a



client generated nonce, the web-origin, the realm, an identifier for the CPK and the challenge string; and signs that hashed blob with the private key corresponding to the CPK for that web-origin. The formatting chosen for this TBS blob is chosen so as to make server-side signature verification as simple as possible for a wide-range of current server tooling.

Figure 1 specifies the ABNF for the signature input.

```
HOBA-TBS = nonce alg origin realm kid challenge
nonce = unreserved
alg = 1*2DIGIT
origin = scheme authority port
realm = unreserved
kid = unreserved
challenge = unreserved
```

Figure 1: To-be-signed data for HOBA

The fields above contain the following:

- o nonce: is a random value chosen by the UA and MUST be base64url encoded before being included in the HOBA-TBS value. UAs MUST be able to use at least 32 bits of randomness in generating a nonce. UAs SHOULD be able to use up to 64 bits of randomness for nonces.
- o alg: specifies the signature algorithm being used encoded as an ASCII character as defined in [Section 9.3](#). RSA-SHA256 MUST be supported, RSA-SHA1 MAY be supported. The IANA registered algorithm values are encoded as ASCII numbers; for example, the encoding of RSA-SHA256 is 0x30.
- o origin: is the web origin expressed as the concatenation of the scheme, authority and port are from [\[RFC3986\]](#). These are not base64 encoded as they will be most readily available to the server in plain text. For example, if accessing the URL "https://www.example.com:8080/foo" then the bytes input to the signature process will be "httpswww.example.com8080"
- o realm: is similarly just a string with the syntactic restrictions defined in [\[I-D.ietf-httpbis-p7-auth\]](#). If no realm is specified for this authentication then this is absent. (A missing field here is no problem since both sides know when it needs to be there.)



- o kid: is a key identifier - this MUST be a base64url encoded value that is presented to the server in the HOBA client result (see below).
- o challenge: MUST be a base64url encoded challenge value that the server chose to send to the client

The HOBA-TBS string is the input to the client's signing process, but is not itself sent over the network since some fields are already inherent in the HTTP exchange. The challenge however is sent over the network so as to make it simpler for a server to be stateless. (One form of stateless challenge might be a ciphertext that the server decrypts and checks, but that is an implementation detail.) The value that is sent over the network is the HOBA "client result" which we now define.

The HOBA "client result" is a dot-separated string that includes the signature and is sent in the HTTP Authorized header field value using the value syntax defined in Figure 2. The "sig" value is the base64url encoded version of the binary output of the signing process. The kid, challenge and nonce are as defined above and are also base64url encoded. [[Expect more changes here. This is very like JOSE's compact form and maybe ought be an instance of that.]]

```
HOBA-RES = kid "." challenge "." nonce "." sig
sig = unreserved
```

Figure 2: HOBA Client Result value

The HOBA scheme is far from new, for example, the basic idea is pretty much identical to the first two messages from "Mechanism R" on page 6 of [[MI93](#)] which predates HOBA by 20 years.

### **3. HOBA HTTP Authentication Mechanism**

An HTTP server that supports HOBA authentication includes the "HOBA" auth-scheme value in a WWW-Authenticate header field when it wants the client to authenticate with HOBA.

- o If the "HOBA" scheme is listed, it MUST be followed by two or more auth-param values. The auth-param attributes defined by this specification are below. Other auth-param attributes MAY be used as well. Unknown auth-param attributes MUST be ignored by clients, if present.



- o The "challenge" attribute MUST be included. The challenge is the string made up of the base64url encoded octets that the server wants the client to sign in its response. The challenge SHOULD be unique for every HTTP 401 response in order to prevent replay attacks from passive observers.
- o An "expires" attribute MUST be included that specifies the number of seconds from the time the HTTP response is emitted for which responses to this challenge can be accepted.
- o A "realm" attribute MAY be included to indicate the scope of protection in the manner described in HTTP/1.1, Part 7 [[I-D.ietf-httpbis-p7-auth](#)]. The "realm" attribute MUST NOT appear more than once.

When the "client response" is created, the UA encodes the HOBA client-result (a string matching the HOBA-RES production in Figure 2 as an auth-param with the name "result" and returns that in the Authorization header.

Note that a HOBA signature is good for however long the expires attribute allows. This means that replay is potentially possible within the time window specified by the "expires" value chosen by the server. Servers SHOULD attempt to detect any such replay and MAY react to such replays by responding with a second (or subsequent) 401-status HTTP response containing a new challenge.

UAs MAY optimise their use of challenges by pre-fetching a challenge value, for example after "expires"/2 seconds have elapsed using the ".well-known/hoba/getChal" scheme described below. This also allows for pre-calculation of HOBA signatures, if that is required in order to produce a responsive user interface.

#### **4. Using HOBA-http**

[[A lot of this is similar to the HOBA-js discussion below. At some point some nuclear fusion might be nice, but for now it might be best to keep them separate until we understand better what can be merged, and what is different.]]

The interaction between an HTTP client and HTTP server using HOBA happens in three phases: the CPK preparation phase, the signing phase, and the authentication phase. The first and second phase are done in a standard fashion; the third is done using site-specific methods. In addition, we provide a mechanism for pre-fetching challenges.





#### **4.1. CPK Preparation Phase**

In the CPK preparation phase, the client determines if it already has a CPK for the web-origin it is going to. If the has a CPK, the client will use it; if the client does not have a CPK, it generates one in anticipation of the server asking for one.

#### **4.2. Signing Phase**

In the signing phase, the client connects to the server, the server asks for HOBA-based authentication, and the client authenticates by signing a blob of information as described in the previous sections.

The UA tries to access a protected resource on the server. The server sends the HOBA WWW-Authenticate challenge. The UA receives the challenge and signs the challenge using the CPK it either already had or just generated. The server validates the signature. If validation fails, or if the server chooses to do so for other reasons, then server aborts the transaction via a 403 Unauthorized HTTP response.

#### **4.3. Authentication Phase**

In the authentication phase, the server extracts the CPK from the signing phase and decides if it recognizes the CPK. If the server recognizes the CPK, the server may finish the client authentication process. If this stage of the process involves additional information for authentication, such as asking the user which account she wants to use (in the case where a UA is used for multiple accounts on a site), the server can prompt the user for account identifying information or the user could choose based on HTML offered by the server before the 401 is triggered. None of this is standardized: it all follows the server's security policy and session flow. At the end of this, the server probably assigns or updates a session cookie for the client.

If the server does not recognize the CPK the server might send the client through a either a join or login-new-UA (see below) process. This process is completely up to the server, and probably entails using HTML and JavaScript to ask the user some questions in order to assess whether or not the server wants to give the client an account. Completion of the joining process might require confirmation by email, SMS, Captcha, and so on.

Note that there is no necessity for the server to initiate a joining or login process upon completion of the signing phase. Indeed, the server may desire to challenge the UA even for unprotected resources and set a session cookie for later use in a join or login process as



it becomes necessary. For example, a server might only want to offer an account to someone who had been to a few pages on the web site; in such a case, the server could use the CPK from an associated session cookie as a way of building reputation for the user until the server wants the user to join.

After the UA is authenticated (if the user had to join, this could be the last step of joining), the server gives the UA access to the protected resource that was originally requested at the beginning of the signing phase. It is quite likely that the server would also update the UA's session cookie for the web site.

#### **4.4. Logging in on a New User Agent**

When a user wants to use a new UA for an existing account, the flows are similar to logging in with an already-joined UA or joining for the first time. In fact, the CPK preparation phase (with the UA knowing that it needs to create a new CPK) and the signing phase are identical.

During the authentication phase, the server could use HTML and JavaScript to ask the user if they are really a new user or want to associate this new CPK with an already-joined CPK. The server can then use some out-of-band method (such as a confirmation email round trip, SMS, or an UA that is already enrolled) to verify that the "new" user is the same as the already-enrolled one.

### **5. Using HOBA-js**

[[ A description of how to use the same HOBA semantics, but doing everything in Javascript in a web page. This is more of a demonstration that you could get the similar semantics via JS rather than a normative section.]]

Web sites using javascript can also perform origin-bound authentication without needing to involve the http layer, and by inference not needing HOBA-specific support in browsers. One element is required: localStorage (see <http://www.w3.org/TR/webstorage/>), and one when it is available will be highly desirable: WebCrypto (see <http://www.w3.org/TR/WebCryptoAPI>). In lieu of WebCrypto, javascript crypto libraries can be employed with the known deficiencies of PRNG, and the general immaturity of those libraries. The following section outlines a mechanism for Javascript HOBA clients to initially enroll, subsequent enrollment on new clients, login, and how HOBA-js relates to web based session management. As with HOBA-http, a pure Javascript implementation retains the property that only CPKs are stored on the server, so that server compromise doesn't suffer the



multiplier affect that the various recent password exposure debacles have vividly demonstrated.

### **5.1. Key Storage**

We use HTML 5's `localStorage` feature for key storage. Conceptually an implementation stores the dictionary account identifier, public key, private key tuples in the origin's `localStorage` for subsequent authentication requests. How this is actually stored in `localStorage` is an implementation detail. We rely on the security properties of the same-origin policy that `localStorage` enforces. See the security considerations for discussion about attacks on `localStorage`.

### **5.2. User Join**

To join a web site, the HOBA-js client generates a public/private key pair and takes as input the account identifier to which the key pair should be bound. The key pair and account identifier are stored in `localStorage` for later use. The UA then signs the join information (see below) using the private key, and forms a message with the public key (CPK) and the signed data. The server receives the message and verifies the signed data using the supplied key. The server creates the account and adds the public key to a list of public keys associated with this account.

### **5.3. User Login**

Each time the user needs to log in to the server, it creates a login message (see below) and signs the message using the relevant private key stored in `localStorage`. The signed login message along with the associated CPK identifier is sent to the server. The server receives the message and verifies the signed data. If the supplied public key is amongst the set of valid public keys for the supplied account, then the login proceeds. See below for a discussion about replay.

### **5.4. Enrolling a New User Agent**

When a user wants to start using a different UA, the website has two choices: use a currently enrolled UA to permit the enrollment or use a trusted out of band mechanism (eg email, sms, etc). To enroll a new UA using an existing UA, the web site can display a one-time password on the currently enrolled UA. This password is a one-time password and expires in a fixed amount of time (say, 30 minutes). It doesn't need to be an overly fussy password since it's one-time and times out quickly. The user then inputs the one-time password and the new UA generates a new asymmetric key pair and includes the one-time password in the login message to the server (see below).



Alternatively if an enrolled UA is not available, and the site has an out of band communication mechanism (eg, sms, email, etc) a user can request that a one-time password be sent to the user. The server generates and stores the one-time password as above. The user receives the one-time password, inputs as above on the new UA, and the HOBA-js client forms the login message as above.

In both cases, when the server receives a login message with a one-time password, it checks to see if the password supplied is in a list of unexpired one-time passwords associated with that account. If the password matches, the server verifies the signature, expires or deletes the one-time password and adds the supplied public key to the list of public keys associated with the user assuming the signature verified correctly. Subsequent logins proceed as above in User Login.

### **5.5. Replay Protection**

To guard against replay of a legitimate login/join message, we use Kerberos-like timestamps in the expectation of synchronization between the browser's and server's clocks is sufficiently reliable. This saves an HTTP round trip which is desirable, though a challenge-response mechanism as in HOBA-http could also be used. The client puts the current system time into the URL, and the server side vets it against its system time. Like Kerberos, a replay cache covering a signature timeout window is required on the server. This can be done using a database table that is keyed (in the database sense of the term) using the signature bits. If the signature is in the replay table, it ought be rejected. If the timestamp in the signature is outside the current replay cache window then it also gets rejected.

[[ An addition of the ability for the server to reject a client with potential time skew and give it a nonce (as with HOBA-http) would allow the size of the replay cache to be set to just a few minutes rather than a much longer period. Or the HOBA server could always use a nonce method. This is worthy of more discussion. ]].

### **5.6. Signature Parameters**

Since we only require agreement between the server and the client where the client is under the control of the server, the actual url parameter names here are only advisory. For each signed url, the client forms a url with the necessary login/join information. For example, suppose example.com has login and join scripts with various parameters:

- o `http://example.com/site/login.php?username=Mike`





- o `http://example.com/site/  
join.php?username=Mike&email=mike@example.com&sms=555.1212`

The client then appends a signature parameter block to the url:

- o `curtime`: the time in milliseconds since unix epoch (ie, `new Date().getTime()`).
- o `pubkey`: the url encoded public key. See DKIM for the format of the base64 encoded PEM formatted key.
- o `temppass`: an optional url encoded one-time password for subsequent enrollment.
- o `keyalg`: currently RSA. 2048 bit keys should be use if WebCrypto is available
- o `digestalg`: currently SHA1. SHA256 should be used if WebCrypto is available.
- o `signature`: empty for signing canonicalization purposes

[[ Signing the full url is problematic with PHP; we should take a clue from what OAUTH does here; we almost certainly need to add some host identifying information...]] To create the signature, the canonical text includes the path portion, the site-specific url parameters and appends a signature block onto the end of the url. The signature block consists of the parameters listed above with an empty signature parameter (ie, `signature=`), eg:

- o Login: `/site/  
login.php?username=Mike&curtime=1234567890.1234&keyalg=RSA&  
digestalg=SHA1&signature=`
- o Join: `/site/  
join.php?username=Mike&email=mike@  
example.com&curtime=1234567890.1234&keyalg=RSA&digestalg=SHA1&  
signature=`
- o Login New User Agent: `/site/  
login.php?username=Mike&curtime=1234567890.1234&temppass=1239678&  
keyalg=RSA&digestalg=SHA1&signature=`

The canonical signature text is then signed with the private key associated with the account. The signature is then base64 encoded and appended to the full url, and sent to the server using XMLHttpRequest as usual. On receipt of the login request, the server first extracts the timestamp (`curtime`) and determines whether the



timestamp is fresh (see above) rejecting the request if stale. The server then removes the scheme and domain:port portion of the incoming url, and removes the signature value only to create the canonical signature text. The server then extracts the public key along with the account and verifies the signature. If the signature verifies, the server then determines whether this is an enrolled public key for the user. If it is, login/join succeeds. If the key is not enrolled, the server then checks to see if a one-time password was supplied. If not, login/join fails. If a one-time password was supplied, the server checks to see if a one-time password is valid and fails if not. If valid, the server disables the one-time password (eg, deletes it from its database) and adds the new public key to the list of enrolled public keys for this user.

Once verified, the server may start up normal cookie-based session management (see below). The server should send back status to the HOBA-js client to determine whether the login/join was successful. The details are left as an implementation detail.

Note: the client SHOULD use an HTTP POST for the XMLHttpRequest as both the public key and signature blocks may exhaust the maximum size for a GET request (typically around 2KB).

### **5.7. Session Management**

Session Management is identical to username/password session management. That is, the session management tool (such as PHP, Python CGI, and so on) inserts a session cookie into the output to the browser, and logging out simply removes the session cookie. HOBA-js does nothing to help or hurt session cookie hijacking -- TLS is still our friend.

### **5.8. Multiple Accounts on One User Agent**

A shared UA with multiple accounts is possible if the account identifier is stored along with the asymmetric key pair binding them to one another. Multiple entries can be kept, one for each account, and selected by the current user. This, of course, is fraught with the possibility for abuse, since you're enrolling the device potentially long-term. A couple of things can possibly be done to combat that. First, the user can request that the credential be erased from keystore. Similarly, in the enrollment phase, a user could request that the key pair only be kept for a certain amount of time, or that it not be stored at all. Last, it's probably best to just not use shared devices at all since that's never especially safe.



### **5.9. Oddities**

With the same-origin policy, subdomains do not have access to the same `localStorage` as parent domains do. For larger/more complex sites this could be an issue that requires enrollment into subdomains with the requisite hassle for users. One way to get around this is to use session cookies as they can be used across subdomains. That is, login using a single well-known domain, and then use session cookies to navigate around a site.

## **6. Additional Services**

HOBA uses a well-known URL [[RFC5785](#)] "hoba" as a base URI for performing many tasks: "https://www.example.com/.well-known/hoba". These URLs are based on the name of the host that the HTTP client is accessing. There are many use cases for these URLs to redirect to other URLs: a site that does registration through a federated site, a site that only does registration under HTTPS, and so on. Like any HTTP client, HOBA clients MUST be able to handle redirection of these URLs. [[There are a bunch of security issues to consider related to cases where a re-direct brings you off-origin.]]

All additional services MUST be done in TLS-protected sessions ([[RFC5246](#)]).

### **6.1. Registration**

Normally, a registration is expected to happen after a UA receives a WWW-Authenticate for a web-origin and realm for which it has no associated CPK. The (protocol part of the) process of registration for a HOBA account on a server is relatively light-weight. The UA generates a new key pair, and associates it with the web-origin/realm in question. The UA sets up a TLS-protected session, goes to the registration URL ".well-known/hoba/register", and submits the CPK using a POST message as described below. It is up to the server to decide what kind of user interaction is required before the account is finally set up.

If the UA has a CPK associated with the web-origin, but not for the realm concerned, then a new registration is REQUIRED. If the server did not wish for that outcome, then it ought not use a different realm.

The POST message sent to the registration URL contains an HTML form (x-www-form-encoded) with one mandatory field (pub) and some optional fields that allow the UA to specify the type and value of key and device identifiers that the UA wishes to use. [[The device stuff is



just a thought.]]

- o pub: is a mandatory field containing the PEM formatted public key of the client. See [Appendix C of \[RFC6376\]](#) for an example of how to generate this key format.
- o kidtype: contains the type of key identifier, this is a numeric value intended to contain one of the values from [Section 9.4](#). If this is not present then the mandatory to implement "DANE-hash" option MUST be used.
- o kid: contains the key identifier as a base64url encoded string that is of the type indicated in the kidtype. If the kid is a hash of a public key then the correct (base64url encoded) hash value MUST be provided and the server SHOULD check that and refuse the registration if an incorrect value was supplied.
- o didtype: specifies a kind of device identifier intended to contain one of the values from [Section 9.5](#), if absent then the "string" form of device identifier MUST be used.
- o did: A UTF8 string that specifies the device identifier. This can be used to help a user be confident that authentication has worked, e.g., following authentication some web content might say "You last logged in from device 'did' at time T."

## **[6.2.](#) Associating Additional Keys to an Existing Account**

It is common for a user to have multiple UAs, and to want all those UAs to be able to authenticate to a single account. One method to allow a user who has an existing account to be able to authenticate on a second device is to securely transport the private and public keys and the origin information from the first device to the second. Previous history with such key transport has been spotty at best. As an alternative, HOBA allows associating a CPK from the second device to the account created on the first device.

Instead of registering on the new device, the UA generates a new key pair, associates it with the web-origin/realm in question, goes to the URL for starting an association, ".well-known/hoba/associate-start" in a TLS-protected session, and submits the new CPK using a POST message. [[ More description is clearly needed here. ]] The server's response to this request is a nonce with at least 128 bits of entropy. That nonce SHOULD be easy for the user to copy and type, such as using Base32 encoding (see [\[RFC4648\]](#)). The user then uses the first UA to log into the origin, goes to the URL for finishing an association, ".well-known/hoba/associate-finish", and submits the nonce using a POST message. [[ More description is





clearly needed here. ]]. The server then knows that the authenticated user is associated with the second CPK. The server can choose to associate the two CPKs with one account. Whether to do so is entirely at the server's discretion however, but the server SHOULD make the outcome clear to the user.

### **6.3. Logging Out**

When the user wishes to logout, the UA simply goes to ".well-known/hoba/logout". The UA MAY also delete session cookies associated with the session. [[Is that right?, maybe a SHOULD- or MUST-delete would be better]]

The server-side MUST NOT allow TLS session resumption for any logged out session and SHOULD also revoke or delete any cookies associated with the session.

### **6.4. Getting a Fresh Challenge**

If the UA would like a "fresh" challenge then it can send a POST or GET message to ".well-known/hoba/getchal". A successful (200 status) response MUST include a fresh (base64url encoded) HOBA challenge for this origin in the body of the response.

## **7. Mandatory-to-Implement Algorithms**

RSA-SHA256 MUST be supported. RSA-SHA1 MAY be used. RSA modulus lengths of at least 2048 bits SHOULD be used.

[[Maybe we should add ECDSA with P256 for shorter signatures.]]

## **8. Security Considerations**

If key binding was server-selected then a bad actor could bind different accounts belonging to the user from the network with possible bad consequences, especially if one of the private keys was compromised somehow.

Binding my CPK with someone else's account would be fun and profitable so SHOULD be appropriately hard. In particular the string generated by the server MUST be hard to guess, for whatever level of difficulty is chosen by the server. The server SHOULD NOT allow a random guess to reveal whether or not an account exists.

[[The potential impact on privacy of HOBA needs to be addressed. If a site can use a 401 and a CPK to track users without permission that



would be not-so-nice so some guidance on how a UA could indicate to a user that HOBA stuff is going on might be needed.]]

[[lots more TBD, be nice to your private keys etc. etc.]]

### **8.1. localStorage Security for Javascript**

Our use of localStorage will undoubtedly be a cause for concern. localStorage uses the same-origin model which says that the scheme, domain and port define a localStorage instance. Beyond that, any code executing will have access to private keying material. Of particular concern are XSS attacks which could conceivably take the keying material and use it to create UAs under the control of an attacker. But XSS attacks are in reality across the board devastating since they can and do steal credit card information, passwords, perform illicit acts, etc, etc. It's not clear that we introduce unique threats from which clear text passwords don't already suffer.

Another source of concern is local access to the keys. That is, if an attacker has access to the UA itself, they could snoop on the key through a javascript console, or find the file(s) that implement localStorage on the host computer. Again it's not clear that we are worse in this regard because the same attacker could get at browser password files, etc too. One possible mitigation is to encrypt the keystore with a password/pin the user supplies. This may sound counter intuitive, but the object here is to keep passwords off of servers to mitigate the multiplier effect of a large scale compromise ala LinkedIn because of shared passwords across sites.

It's worth noting that HOBA uses asymmetric keys and not passwords when evaluating threats. As various password database leaks have shown, the real threat of a password breach is not just to the site that was breached, it's all of the sites a user used the same password on too. That is, the collateral damage is severe because password reuse is common. Storing a password in localStorage would also have a similar multiplier effect for an attacker, though perhaps on a smaller scale than a server-side compromise: one successful crack gains the attacker potential access to hundreds if not thousands of sites the user visits. HOBA does not suffer from that attack multiplier since each asymmetric key pair is unique per site/UA/user.

## **9. IANA Considerations**



### **9.1. HOBA Authentication Scheme**

Authentication Scheme Name: HOBA

Pointer to specification text: [[ this document ]]

Notes (optional): The HOBA scheme can be used with either HTTP servers or proxies. [[But we need to figure out the proxy angle;-)]]

### **9.2. .well-known URLs**

We probably want a new registry for the labels beneath .well-known/hoba so that other folks can add additional features in a controlled way, e.g. for CPK/account revocation or whatever.

### **9.3. Algorithm Names**

TBD, hopefully re-use and existing registry

"0" means RSA-SHA256

"1" means RSA-SHA1

### **9.4. Key Identifier Types**

"0" means a hashed public key, as done in DANE. [[RFC6698](#)]

"1" means a URI, such as a mailto: or acct: URI, but anything conforming to [[RFC3986](#)] is ok.i

"2" means an unformatted string, at the user's/UA's whim

### **9.5. Device Identifier Types**

"0" means an unformatted nickname, at the user's/UA's whim

## **10. Implementation Status**

[[Note to RFC editor - please delete this section before publication.]]

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [[I-D.sheffer-running-code](#)]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any



individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC Editor: replace by a reference to this document], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, by considering the running code as evidence of valuable experimentation and feedback that has made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

At the time of writing there are two known implementations. One done by Stephen Farrell of HOBA-HTTP and a HOBA-JS variant and another by Michael Thomas of an HOBA-JS variant. More details will be provided in future drafts as those implementations mature and e.g. open-source licensing terms and release schedules are figured out.

## **11. Acknowledgements**

Thanks to the following for good comments received during the preparation of this specification: Julian Reschke [[and many more to be added]. All errors and stupidities are of course the editors' fault.

## **12. References**

### **12.1. Normative References**

- [I-D.ietf-httpbis-p7-auth]  
Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Authentication", [draft-ietf-httpbis-p7-auth-23](#) (work in progress), July 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.





- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", [RFC 5785](#), April 2010.
- [RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), December 2011.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", [RFC 6698](#), August 2012.

## **12.2. Informative References**

- [I-D.sheffer-running-code] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: the Implementation Status Section", [draft-sheffer-running-code-06](#) (work in progress), June 2013.
- [MI93] Mitchell and Thomas, "Standardising Authentication Protocols Based on Public-Key Techniques.", Journal of Computer Security 2 (1993): 23-36. , 1993.
- [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.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), October 2006.
- [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.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), April 2011.
- [RFC6376] Crocker, D., Hansen, T., and M. Kucherawy, "DomainKeys Identified Mail (DKIM) Signatures", STD 76, [RFC 6376](#), July 2013.



## **Appendix A. Problems with Passwords**

By far the most common mechanism for web authentication is passwords that can be remembered by the user, called "memorizable passwords". There is plenty of good research on how users typically use memorizable passwords ([[ handful of citations goes here ]]), but some of the highlights are that users typically try hard to reuse passwords on as many web sites as possible, and that web sites often use either email addresses or users' names as the identifier that goes with these passwords.

If an attacker gets access to the database of memorizable passwords, that attacker can impersonate any of the users. Even if the breach is discovered, the attacker can still impersonate users until every password is changed. Even if all the passwords are changed or at least made unusable, the attacker now possesses a list of likely username/password pairs that might exist on other sites.

Using memorizable passwords on unencrypted channels also poses risks to the users. If a web site uses either the HTTP Plain authentication method, or an HTML form that does no cryptographic protection of the password in transit, a passive attacker can see the password and immediately impersonate the user. If a hash-based authentication scheme such as HTTP Digest authentication is used, a passive attacker still has a high chance of being able to determine the password using a dictionary of known passwords.

[[ Say a bit about non-memorizable passwords. Still subject to database attack, although that doesn't give the attacker knowledge for other systems. Safe if digest authentication is used, but that's rare. ]]

## **Appendix B. Examples**

[[Will add more later and probably use example.org.]]

The following values show an example of HOBA-HTTP authentication to the origin <https://hoba-local.ie>. Carriage-returns have been added and need to be removed to validate the example.



-----BEGIN PRIVATE KEY-----

```
MIIEVwIBADANBgkqhkiG9w0BAQEFAASCBBKkwggSlAgEAAoIBAQDNF6tZUbsM7ZrO
5Lyzvn15lJAfOz7j7xdc3hmeS0fh/DCiJWwE5qqffr0v0vXYN+qUTlsXPeBYdrz/
my0YYC02u2QFhDbFRvpM/EuMzZUWTQzkKyU7nSjtPq1LZJJ/Rh6PnjTgImoIMn92
JZgJZgl/vzd29K5Z94JzdJ4Z5bNmQ/gCpjlv0Wvi+GpJ3lDo1csDEyxATxyTUx1J
K73+/RPOngUF9nrXN6kqeH7RkERRz+PFhFGM6r3tU5povJx0P0bzVl6R4kptWLn2
GqDb5LS9iwszsk6YHWQcLOIAMVZ9ITPN2PSuQX0ym81B9qB1V6fTfo2r0Yo1Djn6C
YRX62p4dAgMBAECggEBAJiyWlcFrPhxJ20G1iAVYaJVxAAcwjQ+X0ydyAEbUtnk
Q+1VZ1k2zc43zVxXz5aN+y80L4ncXd4/eXPteu09J6yqVEvvJkA3GkCbTzykC64w
67otjWkXF90bZbXmQrtRXokzRzbhKIS15ER4tPu6ZrQgEBGXFWCQ0SVY3CV36dvm
xU2Y90wf98BGQF5V047S9h8N/VtBn1ttI/6BhQHGsM+TbRR0b6oT9Mp/kqh5jZ2P
dH/17Q6aj6zeFY1jvI6HUpjc0mVOSKPA00qYLZNYC0sqXxw9vCtEODziFZwr02tL
jm14/H8+AZvtR6ZvbzP8h3G95PNdG4Z1eGZv5p0oJoECgYEA8omsyssjVgLIon60
Q6h3Ej0i8I+mfVXYXcAZhv/1MnojFZgMLmuZ13Lgi3dIHfYFNYGFRi8g9iYU/Ljtf
G3B9o6PQsvL0yym85TZzF/yn+kbInHR2yQUiA1Ck1rnrFx8dla/BjAp1Pxumyh/z
HyNFplVb0eRQS4K5F72wScuB9DECgYEA2Hnq10SLetl4lrGAEHDoiIKw++ACQ/Tn
w50+xAx0Xr9i0mczwpHls+zFSQHTy4za8C/tBLxg6HCHKZS04lh+p5SR21UNgwfr
X6xBZ4xoH9zHTFdD2yrHdxAoItD5oMTAUNHxDS8kcSY1+YcjRs2T4kc0NIxeBlp
0++dxK9I6a0CgYEAp52WGSCCbzLFTeeA1RdcEuW0s2PTgPK0cVwNSEskPZpDH01T
ndEnJmpzfG8XG6z8uJsJLD1a9eu4Wk8Vz3TSn4Da/pEBtFZIAXC74dvuivzqJ44l
eY9ejkPxZ6RdYEFUXNoOPKYCiralchLahq5tuCJNRZkQFN9m441og9dtHEECgYEA
qnvRqlpXUqfEZYFi5w/UwfwTJrozbouInylT0piqe8njtTUjuV8ndPzKHoYrXXwP
zMshsfIdq9E7UU7S/IVPMfE6sW6ZVpE9GDrTw5X7RuSb/I5ZPVjCgA00XsQQKmEd
7YesFGSoAXDAIn/yClrc+eR0WneHSBtTGkXKjWSyWn0CgYAlHdG3wj/sOL+JI7p8
fgSoRpkrbjfrzcqSCDVfbT84sUz3dX0bZecrI3v0XUZySakIDha5ZbxXNST/bfea
kReKsoeFY9J7os0hbfgWxiCQwD3wN9tS2yJbJYl1mCFc45YfylkXBzB0S2V/i6rv
vCWw1nijRb+WhpUF8HyKW64bxA==
```

-----END PRIVATE KEY-----

Key Identifier: Zhh5vD5ovE0NqTD0ufSUFRu5dzZMe-K0rMX2cN3vqWw

Challenge: zzrYL7Ba0Qt1z0sl4fMY+EYcG4eT2h+JXi+jEGzozQ0=

Nonce: xXSFdZ-7ahM

Tbsorigin: httpshoba-local.ie443

The resulting signature is:

```
r1ZXAWPxpkd9iyI9TvwNYb0LT6Nth4WRYL4ciLZD6Wvvsni8AYLduUEPdo5ezfo
K_W_Hi4nyHmtRzPpAW9YSGhsyY0d7GSZH7Kd6ncCPVBQuHQdHI5n60Js1itD7hK
t4bCtP3zGxkg_w71KGU2RXcQDfcTNmFcs2ice8RrrvNh1lzRViHO-scV0VNBk19J
LTUyaisiNKq-sNK14_RIG6AeivAGxLLXnN_RzttNe5d0XrXJ1nRUSFmeN6ZfVHE7
qf6l0RqaMeyqsDoJe1MIyISn6sCGzMZmplizNw_eg2QJIJX3Ttat9mTfT5UZyYUq
8meaqRXMhoQWHLweFTNYw
```

The final HTTP header field sent with a request is then:



Authorization: result="Zhh5vD5ovE0NqTD0ufSUFRu5dzZMe-KOrMX2cN3vq  
Ww.zzrYL7BaOQtIzOs14fMY+EYcG4eT2h+JXi+jEGzozQ0=.xXSfZ-7ahM.r1ZX  
AWPXpzkD9iyI9TvwNYb0LT6Nth4WRYL4ciLZD6Wvvsni8AYLduUEPdo5ezfoK\_\_W  
\_Hi4nyHmtrZpPAW9YSGhsyY0d7GSZH7Kd6ncCPVBQuHQdHI5n60JsIitD7hKt4bC  
tP3zGxkg\_W71KGU2RXcQDfcTNmFcs2ice8RrrvNh1lzRViH0-scV0VNBk19JLTUy  
aisiNKq-sNK14\_RIG6AeivAGxLLXnN\_RzttNe5d0XrXJ1nRUSFmeN6ZfVHE7qf6l  
ORqaMeyqsDoJe1MIyISn6sCGzMZmplizNw\_eg2QJIJX3Txat9mTfT5UZyYUq8mea  
qRXMhoQWHGLweFTNYw"

## **Appendix C. Changes**

[[Note to RFC editor - please delete this section before publication.]]

### **C.1. WG-01**

- o A few clarifications/fixes.
- o Added getchal interface and expires attribute
- o Added Implementation Status section
- o Added examples

### **C.2. WG-00**

- o First WG draft, replacing [draft-farrell-httpbis-hoba-02](#)
- o Fleshed out HTTP scheme some more.
- o Fleshed out registration form more.

## Authors' Addresses

Stephen Farrell  
Trinity College Dublin  
Dublin, 2  
Ireland

Phone: +353-1-896-2354  
Email: [stephen.farrell@cs.tcd.ie](mailto:stephen.farrell@cs.tcd.ie)





Paul Hoffman  
VPN Consortium

Email: paul.hoffman@vpnc.org

Michael Thomas  
Phresheez

Email: mike@phresheez.com