

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 19, 2017

B. Campbell
J. Bradley
Ping Identity
H. Tschofenig
ARM
November 15, 2016

Resource Indicators for OAuth 2.0
draft-campbell-oauth-resource-indicators-02

Abstract

This straw-man specification defines an extension to The OAuth 2.0 Authorization Framework that enables the client and authorization server to more explicitly to communicate about the protected resource(s) to be accessed.

Status of This Memo

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

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

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

This Internet-Draft will expire on May 19, 2017.

Copyright Notice

Copyright (c) 2016 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

1. Introduction	2
1.1. Requirements Notation and Conventions	3
1.2. Terminology	3
2. Resource Parameter	3
3. IANA Considerations	5
3.1. OAuth Parameters Registration	5
3.1.1. Registry Contents	5
3.2. OAuth Extensions Error Registration	6
3.2.1. Registry Contents	6
4. Security Considerations	6
5. References	7
5.1. Normative References	7
5.2. Informative References	7
Appendix A. Acknowledgements	8
Appendix B. Document History	8
Authors' Addresses	8

1. Introduction

Several years of deployment and implementation experience with OAuth 2.0 [RFC6749] has uncovered a need, in some circumstances, for the client to explicitly signal to the authorization sever where it intends to use the access token it is requesting.

Knowing which resource server will process the access token enables the authorization server to construct the token as necessary for that entity. Properly encrypting the token (or content within the token) to a particular resource server, for example, requires knowing which resource server will receive and decrypt the token. Furthermore, various resource servers oftentimes have different requirements with respect to the data contained in, or referenced by, the token and knowing the resource server where the client intends to use the s token allows the the authorization server to mint the token accordingly.

Specific knowledge of the intended recipient(s) of the access token also helps facilitate improved security characteristics of the token itself. Bearer tokens, currently the only defined type of OAuth access token, allow any party in possession of a token to get access to the associated resources. To prevent misuse, two important security assumptions must hold: bearer tokens must be protected from disclosure in storage and in transit and the access token must only be valid for use at a specific resource server and for a specific

scope. When the authorization server is informed of the resource server that will process the access token, it can restrict the intended audience of that token such that it cannot be used at other resource servers. Section 5.2 of OAuth 2.0 Authorization Framework: Bearer Token Usage [RFC6750] prescribes including the token's intended recipients within the token to prevent token redirect.

Scope, from Section 3.3 of OAuth 2.0 [RFC6749], sometimes is overloaded to convey the location or identity of the resource server, however, doing so isn't always feasible or desirable. Scope is typically about what access is being requested rather than where that access will be redeemed (e.g. "email", "user:follow", "user_photos", and "channels:read" are a small sample of scope values in use).

A means for the client to signal to the authorization sever where it intends to use the access token it's requesting is important and useful. A number of implementations and deployments of OAuth 2.0 have already employed proprietary parameters toward that end. This specification aims to provide a standardized and interoperable alternative to the proprietary approaches going forward.

1.1. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Terminology

This specification uses the terms "access token", "refresh token", "authorization server", "resource server", "authorization endpoint", "authorization request", "authorization response", "token endpoint", "grant type", "access token request", "access token response", and "client" defined by The OAuth 2.0 Authorization Framework [RFC6749].

2. Resource Parameter

The client may indicate the resource server(s) for which it is requesting an access token by including the following parameter in the request.

resource

OPTIONAL. The value of the "resource" parameter indicates a resource server where the requested access token will be used. It MUST be an absolute URI, as specified by Section 4.3 of [RFC3986], and MUST NOT include a query or fragment component. If the authorization server fails to parse the provided value or does not

consider the resource server acceptable, it MUST reject the request and provide an error response with the error code "invalid_resource". Multiple "resource" parameters may be used to indicate that the issued token is intended to be used at multiple resource servers.

When an access token will be returned from the authorization endpoint, the "resource" parameter is used in the authorization request to the authorization endpoint as defined in Section 4.2.1 of OAuth 2.0 [RFC6749]. An example of an authorization request where the client tells the authorization server that it wants a token for use at "https://rs.example.com/" is shown in Figure 1 below.

```
GET /as/authorization.oauth2?response_type=token
    &client_id=s6BhdRkqt3&state=laeb
    &redirect_uri=https%3A%2F%2Fclient%2Eexample%2Ecom%2Fcb
    &resource=https%3A%2F%2Frs.example.com%2F HTTP/1.1
Host: authorization-server.example.com
```

Figure 1: Protected Resource Request

When the access token is returned from the token endpoint, the request parameter is included in the token request to the token endpoint. Sections 4.1.1, 4.3.1, 4.4.2, 4.5 and 6 of OAuth 2.0 [RFC6749] define requests to the token endpoint with different grant types. An example of a token request, using a refresh token, where the client tells the authorization server that it wants a token for use at "https://rs.example.com/" is shown in Figure 2 below.

```
POST /as/token.oauth2 HTTP/1.1
Host: authorization-server.example.com
Authorization: Basic czZCaGRSa3F0Mzpsc3FFelFsVW9lQUU5cHg0RlNyNHlJ
Content-Type: application/x-www-form-urlencoded

grant_type=refresh_token
&refresh_token=4LTC8lb0acc6Oy4esc1Nk9BWC0imAwH
&resource=https%3A%2F%2Frs.example.com%2F
```

Figure 2: Protected Resource Request

The "resource" parameter indicates the physical location of resource server, typically as an https URL, where the client intends to use the requested access token. This enables the authorization server to apply policy as appropriate for the resource, such as determining the type and content of the token to be issued, if and how the token is to be encrypted, and applying appropriate audience restrictions to the token.

The client SHOULD provide the most specific URI that it can for the set of resources or API it intends to access. In practice a client will know a base URI for the resource server application that it interacts with, which is appropriate to use as the value of the "resource" parameter. The client SHOULD use the base URI for the API unless specific knowledge of resource server dictates the client use a shorter path. For example, the value "https://rs.example.com/" would be used for a resource server that is the exclusive application on that host, however, if the resource server is one of many applications on that host, something like "https://rs.example.com/application/" would be used. Another example, for an API like SCIM [RFC7644] that has multiple endpoints such as "https://rs.example.com/scim/Users", "https://rs.example.com/scim/Groups", and "https://rs.example.com/scim/Schemas" The client should use "https://rs.example.com/scim/" as the resource so that the issued access token is valid for all the endpoints of the SCIM API.

The authorization server SHOULD audience restrict the access token to the resource server(s) indicated by the "resource" parameter. Audience restrictions can be communicated in JSON Web Tokens [RFC7519] with the "aud" claim and the top-level member of the same name provides the audience restriction information in a Token Introspection [RFC7662] response. The authorization server may use the exact "resource" value as the audience or it may map from that value to a more general URI or abstract identifier for the resource server.

The requested resource pertains to the access token that is the expected result of the request and not to the underlying access granted by the resource owner.

3. IANA Considerations

3.1. OAuth Parameters Registration

This specification registers the following value in the IANA "OAuth Parameters" registry [IANA.OAuth.Parameters] established by [RFC6749].

3.1.1. Registry Contents

- o Parameter name: resource
- o Parameter usage location: authorization request, token request
- o Change controller: IESG
- o Specification document(s): Section 2 of [[this specification]]

3.2. OAuth Extensions Error Registration

This specification registers the following error in the IANA "OAuth Extensions Error Registry" [IANA.OAuth.Parameters] established by [RFC6749].

3.2.1. Registry Contents

- o Error name: `invalid_resource`
- o Error usage location: `implicit grant error response, token error response`
- o Related protocol extension: `resource parameter`
- o Change controller: IESG
- o Specification document(s): Section 2 of [[this specification]]

4. Security Considerations

An access token that is audience restricted to a resource server, which obtains the token legitimately, cannot be used to access resources on behalf of the resource owner at other resource servers. The "resource" parameter enables a client to indicate the resource server where the requested access token will be used, which in turn enables the authorization server to apply the appropriate audience restrictions to the token.

Some Resource servers may host user content or be multi-tenant. In order to avoid attacks that might confuse a client into sending a AT to a user controlled resource it is important to use the a specific resource URI including path and not use just a host with no path. This will cause any AT issued for accessing the user controlled resource to have a invalid audience if replayed against the legitimate resource API.

Although multiple occurrences of the "resource" parameter may be included in a request, using only a single "resource" parameter is encouraged. A bearer token that has multiple intended recipients (audiences) can be used by any one of those recipients at any other. Thus, a high degree of trust between the involved parties is needed when using access tokens with multiple audiences. Furthermore an authorization server may be unwilling or unable to fulfill a token request with multiple resources.

[[TODO: I continue to question the value of allowing multiple resources vs the functional and security complexity that comes with doing so. Writing the preceding paragraph just underscores that concern. So just noting it here.]]

5. References

5.1. Normative References

- [IANA.OAuth.Parameters]
IANA, "OAuth Parameters",
<<http://www.iana.org/assignments/oauth-parameters>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997,
<<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005,
<<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012,
<<http://www.rfc-editor.org/info/rfc6749>>.

5.2. Informative References

- [I-D.draft-tschofenig-oauth-audience]
Tschofenig, H., "OAuth 2.0: Audience Information", draft-tschofenig-oauth-audience (work in progress), February 2013.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012,
<<http://www.rfc-editor.org/info/rfc6750>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015,
<<http://www.rfc-editor.org/info/rfc7519>>.
- [RFC7644] Hunt, P., Ed., Grizzle, K., Ansari, M., Wahlstroem, E., and C. Mortimore, "System for Cross-domain Identity Management: Protocol", RFC 7644, DOI 10.17487/RFC7644, September 2015, <<http://www.rfc-editor.org/info/rfc7644>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015,
<<http://www.rfc-editor.org/info/rfc7662>>.

Appendix A. Acknowledgements

The following individuals contributed to discussions relating to and giving rise to this draft specification:

George Fletcher, Hans Zandbelt, Justin Richer, Michael Jones, Nat Sakimura, Phil Hunt, Sergey Beryozkin, and Anthony "no go" Nadalin.

Appendix B. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-01

- o Move Hannes Tschofenig, who wrote [I-D.draft-tschofenig-oauth-audience] in '13, from Acknowledgements to Authors.
- o Added IANA Considerations to register the "resource" parameter and "invalid_resource" error code.

-00

- o Initial draft to define a resource parameter for OAuth 2.0.

Authors' Addresses

Brian Campbell
Ping Identity

Email: brian.d.campbell@gmail.com

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com

Hannes Tschofenig
ARM

Email: hannes.tschofenig@gmx.net

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2016

P. Hunt
Oracle
A. Nadalin
Microsoft
March 13, 2016

OAuth 2.0 Bound Configuration Lookup
draft-hunt-oauth-bound-config-00

Abstract

This specification defines a mechanism for the client of an OAuth 2.0 protected resource service to obtain the configuration details of an OAuth 2.0 authorization server that is capable of authorizing access to a specific resource service. The information includes the OAuth 2.0 component endpoint location URIs and as well as authorization server capabilities.

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 September 14, 2016.

Copyright Notice

Copyright (c) 2016 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

1.	Introduction	2
1.1.	Requirements Notation and Conventions	3
1.2.	Terminology	4
2.	Authorization Server WebFinger Discovery	4
3.	Authorization Server Metadata	8
4.	String Operations	11
5.	Compatibility Notes	12
6.	Security Considerations	12
6.1.	TLS Requirements	12
6.2.	Impersonation Attacks	12
7.	IANA Considerations	13
7.1.	Registration of OAuth2 Relation Type	13
7.2.	OAuth Discovery Metadata Registry	13
7.2.1.	Metadata Registry Procedures	13
7.2.2.	Registration Template	14
7.2.3.	Initial Registry Contents	15
7.3.	Updated Registration Instructions	18
8.	References	18
8.1.	Normative References	19
8.2.	Informative References	21
Appendix A.	Acknowledgements	22
Appendix B.	Document History	22
Authors' Addresses	22

1. Introduction

This specification generalizes the configuration metadata format defined by "OpenID Connect Discovery 1.0" [OpenID.Discovery] to support a wider set of OAuth 2.0 use cases.

This is intentionally parallel to the way that the "OAuth 2.0 Dynamic Client Registration Protocol" [RFC7591] specification generalized the dynamic client registration mechanisms defined by "OpenID Connect Dynamic Client Registration 1.0" [OpenID.Registration] in a way that was compatible with it.

[[Editors note: This draft is intended as a replacement alternative for consideration to draft-ietf-oauth-discovery-01 and as such it contains the same metadata attribute definitions and registry requirements rather than references the original draft.]]

This specification enables clients, having previously discovered an OAuth 2.0 protected resource and an OAuth 2.0 Authorization service

(through other means such as an administrative process, a configuration file, use of ".well-known" [RFC5785]), to discover the OAuth 2.0 configuration metadata and to confirm the relationship between the previously discovered authorization server and the protected resource endpoints through the use of query mechanism. This specification uses WebFinger [RFC7033] to query the authorization server configuration discovery to locate the assigned authorization server (if any) for a specified resource endpoint URI and its configuration properties as an authoritative set. This process is described in Section 2.

The use WebFinger serves the following purposes:

- o To allow service providers to have some flexibility in the relationship and number of authorization servers able to issue access tokens for use with OAuth 2.0 protected resources.
- o Is the resource that the client is requesting authorization for supported by the authorization server?
- o To inform clients about improperly discovered endpoints for valid resources. For example, that would enable an attacker to insert a proxy for a "man-in-the-middle" attack.
- o To bind together the resource service endpoint as part of a valid set along with the endpoints described in the OAuth configuration metadata.

Once the authorization server has been identified, this specification provides the configuration information that is retrieved from a well-known location as a JSON [RFC7159] document, including its OAuth 2.0 endpoint locations and authorization server capabilities. This process is described in Section 2.

1.1. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

All uses of JSON Web Signature (JWS) [JWS] and JSON Web Encryption (JWE) [JWE] data structures in this specification utilize the JWS Compact Serialization or the JWE Compact Serialization; the JWS JSON Serialization and the JWE JSON Serialization are not used.

1.2. Terminology

This specification uses the terms "Access Token", "Authorization Code", "Authorization Endpoint", "Authorization Grant", "Authorization Server", "Client", "Client Authentication", "Client Identifier", "Client Secret", "Grant Type", "Protected Resource", "Redirection URI", "Refresh Token", "Resource Owner", "Resource Server", "Response Type", and "Token Endpoint" defined by OAuth 2.0 [RFC6749], the terms "Claim Name", "Claim Value", and "JSON Web Token (JWT)" defined by JSON Web Token (JWT) [JWT], and the term "Response Mode" defined by OAuth 2.0 Multiple Response Type Encoding Practices [OAuth.Responses].

This specification also defines the following terms:

Resource

Entity that is the target of a request in WebFinger.

Host

Server where a WebFinger service is hosted.

2. Authorization Server WebFinger Discovery

Authorization server WebFinger discovery is a means of determining the location of the authorization server's configuration information.

WebFinger discovery is RECOMMENDED; if a client knows the authorization server's configuration information location through an out-of-band mechanism, it can skip this step and proceed to Section 2.

WebFinger discovery uses the following information to make a discovery request:

resource

A REQUIRED URI that represents the actual resource URL the client intends to use to access the protected resource. This URL may be for a specific sub-set of data (e.g. associated with a resource owner), or it MAY be the base URL for a set of protected resources accessible within the scope of the client's anticipated access authorization. The URI also be the base URL for a service when the end-user is a new user and no specific user-owned resource is defined.

host

The server where the WebFinger service is hosted.

rel

The URI identifying the type of service whose location is being requested.

OAuth 2.0 configuration discovery uses the following "rel" value in WebFinger [RFC7033]:

Rel Type	URI
OAuth 2.0 Configuration Information	oauth2

To start discovery of OAuth 2.0 configuration information, the client first discover's the resource endpoint and the associated OAuth URI through a mechanism defined by the resource service provider. The parameter "resource", whose value is the protected resource's endpoint URL MUST be included.

The client then makes an HTTP "GET" request to the OAuth host's WebFinger [RFC7033] endpoint using the URI as the "resource" parameter value and the "rel" value "oauth2" to obtain the authorization server's configuration information location.

Upon receiving a request, the WebFinger service MAY match the specified "resource" value through specific value matches or the service MAY match a subset of the URI such as the host name and an initial path. For example, if the client is using a URI to a specific subset of resources, the subset can be ignored.

In the response, one or more matches MAY be returned in the "links" array element. Each array element MUST contain a "rel" attribute, an "href" attribute, and a `_properties_` attribute.

The value of the "href" shall contain a value equivalent to the OAuth "issuer" (see Section 3). The value of "rel" is "oauth2".

The "properties" attribute contains the authorization server configuration metadata (see Section 3) that corresponds to the OAuth 2.0 server identified in the "href" attribute. This information includes a set of claims about the authorization server's configuration, including all necessary endpoints and public key location information that contains a set of claims as its members that are a subset of the metadata values defined in Section 3. Other claims MAY also be returned.

Claims that return multiple values are represented as JSON arrays. Claims with zero elements MUST be omitted from the response.

An example WebFinger discovery request follows. To find the authorization server's configuration information for the resource "https://finance.example.com/abcpath", the WebFinger parameters are as follows:

WebFinger Parameter	Value
resource	https://finance.example.com/abcpath
host	example.com
rel	oauth2

The client would make the following WebFinger request to discover the authorization server's configuration information location (with line wraps within lines for display purposes only):

```
GET /.well-known/webfinger
  ?resource=https%3A%2F%2Ffinance.example.com%2Fabcpath
  &rel=oauth2
HTTP/1.1
Host: example.com
```

```
HTTP/1.1 200 OK
Content-Type: application/jrd+json
```

```
{
  "subject": "https://finance.example.com",
  "links":
  [
    {
      "rel": "oauth2",
      "href": "https://server.example.com",
      "properties":
      {
        "issuer":
          "https://server.example.com",
        "authorization_endpoint":
          "https://server.example.com/oauth/authorize",
        "token_endpoint":
          "https://server.example.com/oauth/token",
        "token_endpoint_auth_methods_supported":
          ["client_secret_basic", "private_key_jwt"],
        "token_endpoint_auth_signing_alg_values_supported":
          ["RS256", "ES256"],
        "userinfo_endpoint":
          "https://server.example.com/oauth/userinfo",
        "jwks_uri":
          "https://server.example.com/jwks.json",
        "registration_endpoint":
          "https://server.example.com/oauth/register",
        "scopes_supported":
          ["openid", "profile", "email", "address",
           "phone", "offline_access"],
        "response_types_supported":
          ["code", "code token"],
        "service_documentation":
          "http://server.example.com/oauth/service_documentation.html",
        "ui_locales_supported":
          ["en-US", "en-GB", "en-CA", "fr-FR", "fr-CA"]
      }
    }
  ]
}
```

The discovered authorization server is identified as "https://server.example.com".

In the event that the authorization server does not support the requested resource, the server MUST indicate that it was unable to match the request as per Section 4.2 [RFC7033].

If any of the validation procedures defined in this specification fail, any operations requiring the information that failed to correctly validate MUST be aborted and the information that failed to validate MUST NOT be used.

The "issuer" value returned MUST be identical to the configuration information location URL that was directly used to retrieve the configuration information.

3. Authorization Server Metadata

Authorization servers can have metadata describing their configuration. These authorization server metadata values are used by this specification:

issuer

REQUIRED. The authorization server's issuer identifier, which is a URL that uses the "https" scheme and has no query or fragment components. This is the location where ".well-known" RFC 5785 [RFC5785] resources containing information about the authorization server are published. Using these well-known resources is described in Section 2. The issuer identifier is used to prevent authorization server mix-up attacks, as described in "OAuth 2.0 Mix-Up Mitigation" [I-D.jones-oauth-mix-up-mitigation].

authorization_endpoint

REQUIRED. URL of the authorization server's authorization endpoint [RFC6749].

token_endpoint

URL of the authorization server's token endpoint [RFC6749]. This is REQUIRED unless only the implicit grant type is used.

jwtks_uri

RECOMMENDED. URL of the authorization server's JWK Set [JWK] document. This contains the signing key(s) the client uses to validate signatures from the authorization server. The JWK Set MAY also contain the Server's encryption key(s), which are used by RPs to encrypt requests to the Server. When both signing and encryption keys are made available, a "use" (public key use) parameter value is REQUIRED for all keys in the referenced JWK Set

to indicate each key's intended usage. Although some algorithms allow the same key to be used for both signatures and encryption, doing so is NOT RECOMMENDED, as it is less secure. The JWK "x5c" parameter MAY be used to provide X.509 representations of keys provided. When used, the bare key values MUST still be present and MUST match those in the certificate.

`registration_endpoint`

RECOMMENDED. URL of the authorization server's OAuth 2.0 Dynamic Client Registration endpoint [RFC7591].

`scopes_supported`

RECOMMENDED. JSON array containing a list of the OAuth 2.0 [RFC6749] "scope" values that this authorization server supports. Servers MAY choose not to advertise some supported scope values even when this parameter is used.

`response_types_supported`

REQUIRED. JSON array containing a list of the OAuth 2.0 "response_type" values that this authorization server supports.

`response_modes_supported`

OPTIONAL. JSON array containing a list of the OAuth 2.0 "response_mode" values that this authorization server supports, as specified in OAuth 2.0 Multiple Response Type Encoding Practices [OAuth.Responses]. If omitted, the default is ["query", "fragment"]. The response mode value "form_post" is also defined in OAuth 2.0 Form Post Response Mode [OAuth.Post].

`grant_types_supported`

OPTIONAL. JSON array containing a list of the OAuth 2.0 grant type values that this authorization server supports. If omitted, the default value is ["authorization_code", "implicit"].

`token_endpoint_auth_methods_supported`

OPTIONAL. JSON array containing a list of client authentication methods supported by this token endpoint. Client authentication method values are used in the "token_endpoint_auth_method" parameter defined in Section 2 of [RFC7591]. If omitted, the default is "client_secret_basic" -- the HTTP Basic Authentication Scheme specified in Section 2.3.1 of OAuth 2.0 [RFC6749].

`token_endpoint_auth_signing_alg_values_supported`

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the token endpoint for the signature on the JWT [JWT] used to authenticate the client at the token endpoint for the "private_key_jwt" and "client_secret_jwt"

authentication methods. Servers SHOULD support "RS256". The value "none" MUST NOT be used.

`service_documentation`

OPTIONAL. URL of a page containing human-readable information that developers might want or need to know when using the authorization server. In particular, if the authorization server does not support Dynamic Client Registration, then information on how to register clients needs to be provided in this documentation.

`ui_locales_supported`

OPTIONAL. Languages and scripts supported for the user interface, represented as a JSON array of BCP47 [RFC5646] language tag values.

`op_policy_uri`

OPTIONAL. URL that the authorization server provides to the person registering the client to read about the authorization server's requirements on how the client can use the data provided by the authorization server. The registration process SHOULD display this URL to the person registering the client if it is given. As described in Section 5, despite the identifier "op_policy_uri", appearing to be OpenID-specific, its usage in this specification is actually referring to a general OAuth 2.0 feature that is not specific to OpenID Connect.

`op_tos_uri`

OPTIONAL. URL that the authorization server provides to the person registering the client to read about authorization server's terms of service. The registration process SHOULD display this URL to the person registering the client if it is given. As described in Section 5, despite the identifier "op_tos_uri", appearing to be OpenID-specific, its usage in this specification is actually referring to a general OAuth 2.0 feature that is not specific to OpenID Connect.

`revocation_endpoint`

OPTIONAL. URL of the authorization server's OAuth 2.0 revocation endpoint [RFC7009].

`revocation_endpoint_auth_methods_supported`

OPTIONAL. JSON array containing a list of client authentication methods supported by this revocation endpoint. The valid client authentication method values are those registered in the IANA "OAuth Token Endpoint Authentication Methods" registry [IANA.OAuth.Parameters].

`revocation_endpoint_auth_signing_alg_values_supported`

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the revocation endpoint for the signature on the JWT [JWT] used to authenticate the client at the revocation endpoint for the "private_key_jwt" and "client_secret_jwt" authentication methods. The value "none" MUST NOT be used.

`introspection_endpoint`

OPTIONAL. URL of the authorization server's OAuth 2.0 introspection endpoint [RFC7662].

`introspection_endpoint_auth_methods_supported`

OPTIONAL. JSON array containing a list of client authentication methods supported by this introspection endpoint. The valid client authentication method values are those registered in the IANA "OAuth Token Endpoint Authentication Methods" registry [IANA.OAuth.Parameters] or those registered in the IANA "OAuth Access Token Types" registry [IANA.OAuth.Parameters]. (These values are and will remain distinct, due to Section 7.3.)

`introspection_endpoint_auth_signing_alg_values_supported`

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the introspection endpoint for the signature on the JWT [JWT] used to authenticate the client at the introspection endpoint for the "private_key_jwt" and "client_secret_jwt" authentication methods. The value "none" MUST NOT be used.

`code_challenge_methods_supported`

OPTIONAL. JSON array containing a list of PKCE [RFC7636] code challenge methods supported by this authorization server. Code challenge method values are used in the "code_challenge_method" parameter defined in Section 4.3 of [RFC7636]. The valid code challenge method values are those registered in the IANA "PKCE Code Challenge Methods" registry [IANA.OAuth.Parameters].

Additional authorization server metadata parameters MAY also be used. Some are defined by other specifications, such as OpenID Connect Discovery 1.0 [OpenID.Discovery].

4. String Operations

Processing some OAuth 2.0 messages requires comparing values in the messages to known values. For example, the member names in the configuration information response might be compared to specific member names such as "issuer". Comparing Unicode [UNICODE] strings, however, has significant security implications.

Therefore, comparisons between JSON strings and other Unicode strings MUST be performed as specified below:

1. Remove any JSON applied escaping to produce an array of Unicode code points.
2. Unicode Normalization [USA15] MUST NOT be applied at any point to either the JSON string or to the string it is to be compared against.
3. Comparisons between the two strings MUST be performed as a Unicode code point to code point equality comparison.

5. Compatibility Notes

The identifiers "op_policy_uri", and "op_tos_uri" contain strings referring to the OpenID Connect [OpenID.Core] family of specifications that were originally defined by "OpenID Connect Discovery 1.0" [OpenID.Discovery]. Despite the reuse of these identifiers that appear to be OpenID-specific, their usage in this specification is actually referring to general OAuth 2.0 features that are not specific to OpenID Connect.

6. Security Considerations

6.1. TLS Requirements

Implementations MUST support TLS. Which version(s) ought to be implemented will vary over time, and depend on the widespread deployment and known security vulnerabilities at the time of implementation. The authorization server MUST support TLS version 1.2 [RFC5246] and MAY support additional transport-layer security mechanisms meeting its security requirements. When using TLS, the client MUST perform a TLS/SSL server certificate check, per RFC 6125 [RFC6125]. Implementation security considerations can be found in Recommendations for Secure Use of TLS and DTLS [BCP195].

To protect against information disclosure and tampering, confidentiality protection MUST be applied using TLS with a ciphersuite that provides confidentiality and integrity protection.

6.2. Impersonation Attacks

TLS certificate checking MUST be performed by the client, as described in Section 6.1, when making an authorization server configuration information request. Checking that the server certificate is valid for the configuration information location URL prevents man-in-middle and DNS-based attacks. These attacks could

cause a client to be tricked into using an attacker's keys and endpoints, which would enable impersonation of the legitimate authorization server. If an attacker can accomplish this, they can access the resources that the affected client has access to using the authorization server that they are impersonating.

An attacker may also attempt to impersonate an authorization server by publishing a discovery document that contains an "issuer" claim using the configuration information location URL of the authorization server being impersonated, but with its own endpoints and signing keys. This would enable it to impersonate that authorization server, if accepted by the client. To prevent this, RPs MUST ensure that the configuration information location URL they are using for the configuration information request exactly matches the value of the "issuer" metadata value in the authorization server configuration information document received by the client.

7. IANA Considerations

7.1. Registration of OAuth2 Relation Type

This section registers the WebFinger link relation type "oauth2" as per Section 6.2.1 of [RFC5988].

Relation Name: oauth2

Description: Refers to discovery of OAuth2 configuration metadata.

Reference: This specification.

7.2. OAuth Discovery Metadata Registry

This specification establishes the IANA "OAuth Discovery Metadata" registry for OAuth 2.0 authorization server metadata names. The registry records the authorization server metadata member and a reference to the specification that defines it.

7.2.1. Metadata Registry Procedures

The following registration procedure is used for the registry established by this specification.

Values are registered on a Specification Required [RFC5226] basis after a two-week review period on the `oauth-ext-review@ietf.org` mailing list, on the advice of one or more Designated Experts. However, to allow for the allocation of values prior to publication, the Designated Experts may approve registration once they are satisfied that such a specification will be published.

Registration requests sent to the mailing list for review should use an appropriate subject (e.g., "Request to register OAuth Discovery Metadata: example").

Within the review period, the Designated Experts will either approve or deny the registration request, communicating this decision to the review list and IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the request successful. Registration requests that are undetermined for a period longer than 21 days can be brought to the IESG's attention (using the iesg@ietf.org mailing list) for resolution.

Criteria that should be applied by the Designated Experts includes determining whether the proposed registration duplicates existing functionality, determining whether it is likely to be of general applicability or whether it is useful only for a single application, and whether the registration makes sense.

IANA must only accept registry updates from the Designated Experts and should direct all requests for registration to the review mailing list.

It is suggested that multiple Designated Experts be appointed who are able to represent the perspectives of different applications using this specification, in order to enable broadly-informed review of registration decisions. In cases where a registration decision could be perceived as creating a conflict of interest for a particular Expert, that Expert should defer to the judgment of the other Experts.

7.2.2. Registration Template

Discovery Metadata Name:

The name requested (e.g., "issuer"). This name is case-sensitive. Names may not match other registered names in a case-insensitive manner unless the Designated Experts state that there is a compelling reason to allow an exception.

Discovery Metadata Description:

Brief description of the discovery metadata (e.g., "Issuer URL").

Change Controller:

For Standards Track RFCs, list the "IESG". For others, give the name of the responsible party. Other details (e.g., postal address, email address, home page URI) may also be included.

Specification Document(s):

Reference to the document or documents that specify the parameter, preferably including URIs that can be used to retrieve copies of the documents. An indication of the relevant sections may also be included but is not required.

7.2.3. Initial Registry Contents

- o Discovery Metadata Name: "issuer"
- o Discovery Metadata Description: URL of the authorization server's configuration information location
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "authorization_endpoint"
- o Discovery Metadata Description: URL of the authorization server's authorization endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "token_endpoint"
- o Discovery Metadata Description: URL of the authorization server's token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "jwks_uri"
- o Discovery Metadata Description: URL of the authorization server's JWK Set document
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "registration_endpoint"
- o Discovery Metadata Description: URL of the authorization server's OAuth 2.0 Dynamic Client Registration Endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "scopes_supported"
- o Discovery Metadata Description: JSON array containing a list of the OAuth 2.0 "scope" values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "response_types_supported"
- o Discovery Metadata Description: JSON array containing a list of the OAuth 2.0 "response_type" values that this authorization server supports

- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "response_modes_supported"
- o Discovery Metadata Description: JSON array containing a list of the OAuth 2.0 "response_mode" values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "grant_types_supported"
- o Discovery Metadata Description: JSON array containing a list of the OAuth 2.0 grant type values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "token_endpoint_auth_methods_supported"
- o Discovery Metadata Description: JSON array containing a list of client authentication methods supported by this token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "token_endpoint_auth_signing_alg_values_supported"
- o Discovery Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the token endpoint for the signature on the JWT used to authenticate the client at the token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "service_documentation"
- o Discovery Metadata Description: URL of a page containing human-readable information that developers might want or need to know when using the authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "ui_locales_supported"
- o Discovery Metadata Description: Languages and scripts supported for the user interface, represented as a JSON array of BCP47 language tag values
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "op_policy_uri"

- o Discovery Metadata Description: URL that the authorization server provides to the person registering the client to read about the authorization server's requirements on how the client can use the data provided by the authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "op_tos_uri"
- o Discovery Metadata Description: URL that the authorization server provides to the person registering the client to read about authorization server's terms of service
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "revocation_endpoint"
- o Discovery Metadata Description: URL of the authorization server's OAuth 2.0 revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "revocation_endpoint_auth_methods_supported"
- o Discovery Metadata Description: JSON array containing a list of client authentication methods supported by this revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "revocation_endpoint_auth_signing_alg_values_supported"
- o Discovery Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the revocation endpoint for the signature on the JWT used to authenticate the client at the revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "introspection_endpoint"
- o Discovery Metadata Description: URL of the authorization server's OAuth 2.0 introspection endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "introspection_endpoint_auth_methods_supported"
- o Discovery Metadata Description: JSON array containing a list of client authentication methods supported by this introspection endpoint

- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "introspection_endpoint_auth_signing_alg_values_supported"
- o Discovery Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the introspection endpoint for the signature on the JWT used to authenticate the client at the introspection endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

- o Discovery Metadata Name: "code_challenge_methods_supported"
- o Discovery Metadata Description: PKCE code challenge methods supported by this authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 3 of [[this specification]]

7.3. Updated Registration Instructions

This specification adds to the instructions for the Designated Experts of the following IANA registries, both of which are in the "OAuth Parameters" registry [IANA.OAuth.Parameters]:

- o OAuth Access Token Types
- o OAuth Token Endpoint Authentication Methods

IANA has added a link to this specification in the Reference sections of these registries. [[RFC Editor: The above sentence is written in the past tense as it would appear in the final specification, even though these links won't actually be created until after the IESG has requested publication of the specification. Please delete this note after the links are in place.]]

For these registries, the designated experts must reject registration requests in one registry for values already occurring in the other registry. This is necessary because the "introspection_endpoint_auth_methods_supported" parameter allows for the use of values from either registry. That way, because the values in the two registries will continue to be mutually exclusive, no ambiguities will arise.

8. References

8.1. Normative References

- [BCP195] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/bcp195>>.
- [IANA.OAuth.Parameters] IANA, "OAuth Parameters", <<http://www.iana.org/assignments/oauth-parameters>>.
- [JWA] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <<http://tools.ietf.org/html/rfc7518>>.
- [JWE] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<http://tools.ietf.org/html/rfc7516>>.
- [JWK] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<http://tools.ietf.org/html/rfc7517>>.
- [JWS] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<http://tools.ietf.org/html/rfc7515>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<http://tools.ietf.org/html/rfc7519>>.
- [OAuth.Post] Jones, M. and B. Campbell, "OAuth 2.0 Form Post Response Mode", April 2015, <http://openid.net/specs/oauth-v2-form-post-response-mode-1_0.html>.
- [OAuth.Responses] de Medeiros, B., Ed., Scurtescu, M., Tarjan, P., and M. Jones, "OAuth 2.0 Multiple Response Type Encoding Practices", February 2014, <http://openid.net/specs/oauth-v2-multiple-response-types-1_0.html>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, DOI 10.17487/RFC2246, January 1999, <<http://www.rfc-editor.org/info/rfc2246>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<http://www.rfc-editor.org/info/rfc5646>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<http://www.rfc-editor.org/info/rfc5785>>.
- [RFC5988] Nottingham, M., "Web Linking", RFC 5988, DOI 10.17487/RFC5988, October 2010, <<http://www.rfc-editor.org/info/rfc5988>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<http://www.rfc-editor.org/info/rfc6125>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.
- [RFC7009] Lodderstedt, T., Ed., Dronia, S., and M. Scurtescu, "OAuth 2.0 Token Revocation", RFC 7009, DOI 10.17487/RFC7009, August 2013, <<http://www.rfc-editor.org/info/rfc7009>>.

- [RFC7033] Jones, P., Salgueiro, G., Jones, M., and J. Smarr, "WebFinger", RFC 7033, DOI 10.17487/RFC7033, September 2013, <<http://www.rfc-editor.org/info/rfc7033>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [RFC7565] Saint-Andre, P., "The 'acct' URI Scheme", RFC 7565, DOI 10.17487/RFC7565, May 2015, <<http://www.rfc-editor.org/info/rfc7565>>.
- [RFC7591] Richer, J., Ed., Jones, M., Bradley, J., Machulak, M., and P. Hunt, "OAuth 2.0 Dynamic Client Registration Protocol", RFC 7591, DOI 10.17487/RFC7591, July 2015, <<http://www.rfc-editor.org/info/rfc7591>>.
- [RFC7636] Sakimura, N., Ed., Bradley, J., and N. Agarwal, "Proof Key for Code Exchange by OAuth Public Clients", RFC 7636, DOI 10.17487/RFC7636, September 2015, <<http://www.rfc-editor.org/info/rfc7636>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<http://www.rfc-editor.org/info/rfc7662>>.
- [UNICODE] The Unicode Consortium, "The Unicode Standard", <<http://www.unicode.org/versions/latest/>>.
- [USA15] Davis, M. and K. Whistler, "Unicode Normalization Forms", Unicode Standard Annex 15, June 2015, <<http://www.unicode.org/reports/tr15/>>.

8.2. Informative References

- [I-D.jones-oauth-mix-up-mitigation]
Jones, M. and J. Bradley, "OAuth 2.0 Mix-Up Mitigation", draft-jones-oauth-mix-up-mitigation-01 (work in progress), January 2016.
- [OpenID.Core]
Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", November 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.

[OpenID.Discovery]

Sakimura, N., Bradley, J., Jones, M., and E. Jay, "OpenID Connect Discovery 1.0", November 2014, <http://openid.net/specs/openid-connect-discovery-1_0.html>.

[OpenID.Registration]

Sakimura, N., Bradley, J., and M. Jones, "OpenID Connect Dynamic Client Registration 1.0", November 2014, <http://openid.net/specs/openid-connect-registration-1_0.html>.

Appendix A. Acknowledgements

This specification is an alternate proposal to the OAuth 2.0 Discovery working group draft (draft-ietf-oauth-discovery-01) and contains substantial text from this draft. The author would like to thank the original draft authors Mike Jones, Nat Sakimura, and John Bradley as well as the OAuth Working Group.

This specification is based in part on the OpenID Connect Discovery 1.0 specification, which was produced by the OpenID Connect working group of the OpenID Foundation.

Appendix B. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-00

- o Created the initial working group version based on draft-jones-oauth-discovery-01, with no normative changes.

Authors' Addresses

Phil Hunt
Oracle

Email: phil.hunt@yahoo.com
URI: <http://www.independentid.com>

Tony Nadalin
Microsoft

Email: tonynad@microsoft.com

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

M. Jones
Microsoft
P. Hunt
Oracle
A. Nadalin
Microsoft
March 13, 2017

Authentication Method Reference Values
draft-ietf-oauth-amr-values-08

Abstract

The "amr" (Authentication Methods References) claim is defined and registered in the IANA "JSON Web Token Claims" registry but no standard Authentication Method Reference values are currently defined. This specification establishes a registry for Authentication Method Reference values and defines an initial set of Authentication Method Reference values.

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 September 14, 2017.

Copyright Notice

Copyright (c) 2017 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

1. Introduction	2
1.1. Requirements Notation and Conventions	4
1.2. Terminology	4
2. Authentication Method Reference Values	4
3. Relationship to "acr" (Authentication Context Class Reference)	6
4. Privacy Considerations	6
5. Security Considerations	7
6. IANA Considerations	7
6.1. Authentication Method Reference Values Registry	7
6.1.1. Registration Template	8
6.1.2. Initial Registry Contents	9
7. References	11
7.1. Normative References	11
7.2. Informative References	12
Appendix A. Examples	13
Appendix B. Acknowledgements	14
Appendix C. Document History	14
Authors' Addresses	15

1. Introduction

The "amr" (Authentication Methods References) claim is defined and registered in the IANA "JSON Web Token Claims" registry [IANA.JWT.Claims] but no standard Authentication Method Reference values are currently defined. This specification establishes a registry for Authentication Method Reference values and defines an initial set of Authentication Method Reference values.

For context, the "amr" (Authentication Methods References) claim is defined by Section 2 of the OpenID Connect Core 1.0 specification [OpenID.Core] as follows:

amr

OPTIONAL. Authentication Methods References. JSON array of strings that are identifiers for authentication methods used in the authentication. For instance, values might indicate that both password and OTP authentication methods were used. The definition of particular values to be used in the "amr" Claim is beyond the scope of this specification. Parties using this claim will need to agree upon the meanings of the values used, which may be

context-specific. The "amr" value is an array of case sensitive strings.

Each "amr" value typically provides an identifier for a family of closely-related authentication methods. For example, the "otp" identifier intentionally covers both time-based and HMAC-based OTPs. Many relying parties will be content to know that an OTP has been used in addition to a password; the distinction between which kind of OTP was used is not useful to them. Thus, there's a single identifier that can be satisfied in two or more nearly equivalent ways.

Similarly, there's a whole range of nuances between different fingerprint matching algorithms. They differ in false positive and false negative rates over different population samples and also differ based on the kind and model of fingerprint sensor used. Like the OTP case, many relying parties will be content to know that a fingerprint match was made, without delving into and differentiating based on every aspect of the implementation of fingerprint capture and match. The "fpt" identifier accomplishes this.

Ultimately, the relying party is depending upon the identity provider to do reasonable things. If it does not trust the identity provider to do so, it has no business using it. The "amr" value lets the identity provider signal to the relying party additional information about what it did, for the cases in which that information is useful to the relying party.

The "amr" values defined by this specification are not intended to be an exhaustive set covering all use cases. Additional values can and will be added to the registry by other specifications. Rather, the values defined herein are an intentionally small set that are already actually being used in practice.

The values defined by this specification only make distinctions that are known to be useful to relying parties. Slicing things more finely than would be used in practice would actually hurt interop, rather than helping it, because it would force relying parties to recognize that several or many different values actually mean the same thing to them.

For context, while the claim values registered pertain to authentication, note that OAuth 2.0 [RFC6749] is designed for resource authorization and cannot be used for authentication without employing appropriate extensions, such as those defined by OpenID Connect Core 1.0 [OpenID.Core]. The existence of the "amr" claim and values for it should not be taken as encouragement to try to use

OAuth 2.0 for authentication without employing extensions enabling secure authentication to be performed.

When used with OpenID Connect, if the identity provider supplies an "amr" claim in the ID Token resulting from a successful authentication, the relying party can inspect the values returned and thereby learn details about how the authentication was performed. For instance, the relying party might learn that only a password was used or it might learn that iris recognition was used in combination with a hardware-secured key. Whether "amr" values are provided and which values are understood by what parties are both beyond the scope of this specification. The OpenID Connect MODRNA Authentication Profile 1.0 [OpenID.MODRNA] is one example of an application context that uses "amr" values defined by this specification.

1.1. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Terminology

This specification uses the terms defined by JSON Web Token (JWT) [JWT] and OpenID Connect Core 1.0 [OpenID.Core].

2. Authentication Method Reference Values

The following is a list of Authentication Method Reference values defined by this specification:

face

Biometric authentication [RFC4949] using facial recognition

fpt

Biometric authentication [RFC4949] using a fingerprint

geo

Use of geolocation information for authentication, such as that provided by [W3C.REC-geolocation-API-20161108]

hwk

Proof-of-possession (PoP) of a hardware-secured key. See Appendix C of [RFC4211] for a discussion on PoP.

iris

Biometric authentication [RFC4949] using an iris scan

- kba**
Knowledge-based authentication [NIST.800-63-2] [ISO29115]
- mca**
Multiple-channel authentication [MCA]. The authentication involves communication over more than one distinct communication channel. For instance, a multiple-channel authentication might involve both entering information into a workstation's browser and providing information on a telephone call to a pre-registered number.
- mfa**
Multiple-factor authentication [NIST.800-63-2] [ISO29115]. When this is present, specific authentication methods used may also be included.
- otp**
One-time password [RFC4949]. One-time password specifications that this authentication method applies to include [RFC4226] and [RFC6238].
- pin**
Personal Identification Number (PIN) [RFC4949] or pattern (not restricted to containing only numbers) that a user enters to unlock a key on the device. This mechanism should have a way to deter an attacker from obtaining the PIN by trying repeated guesses.
- pwd**
Password-based authentication [RFC4949]
- rba**
Risk-based authentication [JECM]
- retina**
Biometric authentication [RFC4949] using a retina scan
- sc**
Smart card [RFC4949]
- sms**
Confirmation using SMS [SMS] text message to the user at a registered number
- swk**
Proof-of-possession (PoP) of a software-secured key. See Appendix C of [RFC4211] for a discussion on PoP.

tel

Confirmation by telephone call to the user at a registered number. This authentication technique is sometimes also referred to as "call back" [RFC4949].

user

User presence test. Evidence that the end-user is present and interacting with the device. This is sometimes also referred to as "test of user presence" [W3C.WD-webauthn-20170216].

vbm

Biometric authentication [RFC4949] using a voiceprint

wia

Windows integrated authentication [MSDN]

3. Relationship to "acr" (Authentication Context Class Reference)

The "acr" (Authentication Context Class Reference) claim and "acr_values" request parameter are related to the "amr" (Authentication Methods References) claim, but with important differences. An Authentication Context Class specifies a set of business rules that authentications are being requested to satisfy. These rules can often be satisfied by using a number of different specific authentication methods, either singly or in combination. Interactions using "acr_values" request that the specified Authentication Context Classes be used and that the result should contain an "acr" claim saying which Authentication Context Class was satisfied. The "acr" claim in the reply states that the business rules for the class were satisfied -- not how they were satisfied.

In contrast, interactions using the "amr" claim make statements about the particular authentication methods that were used. This tends to be more brittle than using "acr", since the authentication methods that may be appropriate for a given authentication will vary over time, both because of the evolution of attacks on existing methods and the deployment of new authentication methods.

4. Privacy Considerations

The list of "amr" claim values returned in an ID Token reveals information about the way that the end-user authenticated to the identity provider. In some cases, this information may have privacy implications.

While this specification defines identifiers for particular kinds of credentials, it does not define how these credentials are stored or protected. For instance, ensuring the security and privacy of

biometric credentials that are referenced by some of the defined Authentication Method Reference values is beyond the scope of this specification.

5. Security Considerations

The security considerations in OpenID Connect Core 1.0 [OpenID.Core] and OAuth 2.0 [RFC6749] and the OAuth 2.0 Threat Model [RFC6819] apply to applications using this specification.

As described in Section 3, taking a dependence upon particular authentication methods may result in brittle systems since the authentication methods that may be appropriate for a given authentication will vary over time.

6. IANA Considerations

6.1. Authentication Method Reference Values Registry

This specification establishes the IANA "Authentication Method Reference Values" registry for "amr" claim array element values. The registry records the Authentication Method Reference value and a reference to the specification that defines it. This specification registers the Authentication Method Reference values defined in Section 2.

Values are registered on an Expert Review [RFC5226] basis after a three-week review period on the `jwt-reg-review@ietf.org` mailing list, on the advice of one or more Designated Experts. To increase potential interoperability, the experts are requested to encourage registrants to provide the location of a publicly-accessible specification defining the values being registered, so that their intended usage can be more easily understood.

Registration requests sent to the mailing list for review should use an appropriate subject (e.g., "Request to register Authentication Method Reference value: otp").

Within the review period, the Designated Experts will either approve or deny the registration request, communicating this decision to the review list and IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the request successful. Registration requests that are undetermined for a period longer than 21 days can be brought to the IESG's attention (using the `iesg@ietf.org` mailing list) for resolution.

IANA must only accept registry updates from the Designated Experts and should direct all requests for registration to the review mailing list.

It is suggested that the same Designated Experts evaluate these registration requests as those who evaluate registration requests for the IANA "JSON Web Token Claims" registry [IANA.JWT.Claims].

Criteria that should be applied by the Designated Experts includes determining whether the proposed registration duplicates existing functionality, whether it is likely to be of general applicability or whether it is useful only for a single application, whether the value is actually being used, and whether the registration description is clear.

6.1.1.1. Registration Template

Authentication Method Reference Name:

The name requested (e.g., "otp") for the authentication method or family of closely-related authentication methods. Because a core goal of this specification is for the resulting representations to be compact, it is RECOMMENDED that the name be short -- that is, not to exceed 8 characters without a compelling reason to do so. To facilitate interoperability, the name must use only printable ASCII characters excluding double quote ('"') and backslash ('\') (the Unicode characters with code points U+0021, U+0023 through U+005B, and U+005D through U+007E). This name is case sensitive. Names may not match other registered names in a case-insensitive manner unless the Designated Experts state that there is a compelling reason to allow an exception.

Authentication Method Reference Description:

Brief description of the Authentication Method Reference (e.g., "One-time password").

Change Controller:

For Standards Track RFCs, state "IESG". For others, give the name of the responsible party. Other details (e.g., postal address, email address, home page URI) may also be included.

Specification Document(s):

Reference to the document or documents that specify the parameter, preferably including URIs that can be used to retrieve copies of the documents. An indication of the relevant sections may also be included but is not required.

6.1.2. Initial Registry Contents

- o Authentication Method Reference Name: "face"
- o Authentication Method Reference Description: Facial recognition
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "fpt"
- o Authentication Method Reference Description: Fingerprint biometric
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "geo"
- o Authentication Method Reference Description: Geolocation
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "hwk"
- o Authentication Method Reference Description: Proof-of-possession of a hardware-secured key
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "iris"
- o Authentication Method Reference Description: Iris scan biometric
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "kba"
- o Authentication Method Reference Description: Knowledge-based authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "mca"
- o Authentication Method Reference Description: Multiple-channel authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "mfa"
- o Authentication Method Reference Description: Multiple-factor authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "otp"
- o Authentication Method Reference Description: One-time password

- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "pin"
- o Authentication Method Reference Description: Personal Identification Number or pattern
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "pwd"
- o Authentication Method Reference Description: Password-based authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "rba"
- o Authentication Method Reference Description: Risk-based authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "retina"
- o Authentication Method Reference Description: Retina scan biometric
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "sc"
- o Authentication Method Reference Description: Smart card
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "sms"
- o Authentication Method Reference Description: Confirmation using SMS
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "swk"
- o Authentication Method Reference Description: Proof-of-possession of a software-secured key
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "tel"
- o Authentication Method Reference Description: Confirmation by telephone call
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "user"
- o Authentication Method Reference Description: User presence test
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "vbm"
- o Authentication Method Reference Description: Voice biometric
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Authentication Method Reference Name: "wia"
- o Authentication Method Reference Description: Windows integrated authentication
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

7. References

7.1. Normative References

- [IANA.JWT.Claims] IANA, "JSON Web Token Claims", <<http://www.iana.org/assignments/jwt>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, May 2015, <<http://www.rfc-editor.org/info/rfc7519>>.
- [OpenID.Core] Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", November 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.

7.2. Informative References

- [ISO29115] International Organization for Standardization, "ISO/IEC 29115:2013 -- Information technology - Security techniques - Entity authentication assurance framework", ISO/IEC 29115:2013, April 2013, <http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=45138>.
- [JECM] Williamson, G., "Enhanced Authentication In Online Banking", Journal of Economic Crime Management 4.2: 18-19, 2006, <<http://utica.edu/academic/institutes/ecii/publications/articles/51D6D996-90F2-F468-AC09C4E8071575AE.pdf>>.
- [MCA] ldapwiki.com, "Multiple-channel Authentication", August 2016, <<https://www.ldapwiki.com/wiki/Multiple-channel%20Authentication>>.
- [MSDN] Microsoft, "Integrated Windows Authentication with Negotiate", September 2011, <<http://blogs.msdn.com/b/benjaminperkins/archive/2011/09/14/iis-integrated-windows-authentication-with-negotiate.aspx>>.
- [NIST.800-63-2] National Institute of Standards and Technology (NIST), "Electronic Authentication Guideline", NIST Special Publication 800-63-2, August 2013, <<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-2.pdf>>.
- [OpenID.MODRMA] Connotte, J. and J. Bradley, "OpenID Connect MODRMA Authentication Profile 1.0", March 2017, <http://openid.net/specs/openid-connect-modrna-authentication-1_0.html>.
- [RFC4211] Schaad, J., "Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)", RFC 4211, DOI 10.17487/RFC4211, September 2005, <<http://www.rfc-editor.org/info/rfc4211>>.
- [RFC4226] M'Raihi, D., Bellare, M., Hoornaert, F., Naccache, D., and O. Ranen, "HOTP: An HMAC-Based One-Time Password Algorithm", RFC 4226, DOI 10.17487/RFC4226, December 2005, <<http://www.rfc-editor.org/info/rfc4226>>.

- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.
- [RFC6238] M'Raihi, D., Machani, S., Pei, M., and J. Rydell, "TOTP: Time-Based One-Time Password Algorithm", RFC 6238, DOI 10.17487/RFC6238, May 2011, <<http://www.rfc-editor.org/info/rfc6238>>.
- [RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<http://www.rfc-editor.org/info/rfc6819>>.
- [SMS] 3rd Generation Partnership Project, "Technical realization of the Short Message Service (SMS)", 3GPP Technical Specification (TS) 03.40 V7.5.0 (2001-12), January 2002, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=141>>.
- [W3C.REC-geolocation-API-20161108] Popescu, A., "Geolocation API Specification 2nd Edition", World Wide Web Consortium Recommendation REC-geolocation-API-20161108, November 2016, <<https://www.w3.org/TR/2016/REC-geolocation-API-20161108>>.
- [W3C.WD-webauthn-20170216] Bharadwaj, V., Le Van Gong, H., Balfanz, D., Czeskis, A., Birgisson, A., Hodges, J., Jones, M., Lindemann, R., and J. Jones, "Web Authentication: An API for accessing Scoped Credentials", World Wide Web Consortium Working Draft WD-webauthn-20170216, February 2017, <<http://www.w3.org/TR/2017/WD-webauthn-20170216/>>.

Appendix A. Examples

In some cases, the "amr" claim value returned may contain a single Authentication Method Reference value. For example, the following "amr" claim value indicates that the authentication performed used an iris scan biometric:

```
"amr": ["iris"]
```

In other cases, the "amr" claim value returned may contain multiple Authentication Method Reference values. For example, the following "amr" claim value indicates that the authentication performed used a password and knowledge-based authentication:

```
"amr": ["pwd", "kba"]
```

Appendix B. Acknowledgements

Caleb Baker participated in specifying the original set of "amr" values. Jari Arkko, John Bradley, Ben Campbell, Brian Campbell, William Denniss, Linda Dunbar, Stephen Farrell, Paul Kyzivat, Elaine Newton, James Manger, Catherine Meadows, Alexey Melnikov, Kathleen Moriarty, Nat Sakimura, and Mike Schwartz provided reviews of the specification.

Appendix C. Document History

```
[[ to be removed by the RFC editor before publication as an RFC ]]
```

-08

- o Added text in the IANA Registration Template saying that names can be for families of closely-related authentication methods, as suggested by Stephen Farrell.

-07

- o Clarified that the values are intended to provide identifiers for families of closely-related authentication methods.
- o Updated the MODRNA Authentication Profile reference.

-06

- o Addressed IESG comments. Identifiers are now restricted to using only printable JSON-friendly ASCII characters. Additional references to documentation relevant to specific AMR values were added.

-05

- o Specified characters allowed in "amr" values, reusing the IANA Considerations language on this topic from RFC 7638.

-04

- o Added examples with single and multiple values.
- o Clarified that the actual credentials referenced are not part of this specification to avoid additional privacy concerns for biometric data.
- o Clarified that the OAuth 2.0 Threat Model [RFC6819] applies to applications using this specification.

-03

- o Addressed shepherd comments.

-02

- o Addressed working group last call comments.

-01

- o Distinguished between retina and iris biometrics.
- o Expanded the introduction to provide additional context to readers.
- o Referenced the OpenID Connect MODRMA Authentication Profile 1.0 specification, which uses "amr" values defined by this specification.

-00

- o Created the initial working group draft from draft-jones-oauth-amr-values-05 with no normative changes.

Authors' Addresses

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

Phil Hunt
Oracle

Email: phil.hunt@yahoo.com

Anthony Nadalin
Microsoft

Email: tonynad@microsoft.com

OAuth Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: August 7, 2016

J. Bradley, Ed.
Ping Identity
A. Sanso
Adobe Systems
H. Tschofenig
February 04, 2016

OAuth 2.0 Security: Closing Open Redirectors in OAuth
draft-ietf-oauth-closing-redirectors-00.txt

Abstract

This document gives additional security considerations for OAuth, beyond those in the OAuth 2.0 specification and in the OAuth 2.0 Threat Model and Security Considerations.

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 August 7, 2016.

Copyright Notice

Copyright (c) 2016 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

1.	Introduction	2
1.1.	Notational Conventions	2
1.2.	Terminology	2
2.	Authorization Server Error Response	3
2.1.	Abuse: The Authorization Server As Open Redirector	3
2.2.	Security Compromise: The Authorization Server As Open Redirector	4
2.3.	Mitigation	5
3.	Acknowledgements	6
4.	Normative References	6
	Appendix A. Document History	6
	Authors' Addresses	7

1. Introduction

This document gives additional security considerations for OAuth, beyond those in the OAuth 2.0 specification [RFC6749] and in the OAuth 2.0 Threat Model and Security Considerations [RFC6819]. In particular focuses its attention on the risk of abuse the Authorization Server (AS) (Section 1.2) as an open redirector.

It contains the following content:

- o Describes the Authorization Server Error Response as defined in [RFC6749].
- o Describes the risk of abuse the Authorization Server as an open redirector.
- o Gives some mitigation details on how to hinder the risk of open redirector in the ?AS?.

1.1. Notational Conventions

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].

Unless otherwise noted, all the protocol parameter names and values are case sensitive.

1.2. Terminology

Authorization Server (AS)

The server issuing access tokens to the client after successfully authenticating the resource owner and obtaining authorization.

Redirection endpoint

Used by the authorization server to return responses containing authorization credentials to the client via the resource owner user-agent.

2. Authorization Server Error Response

The OAuth 2.0 specification [RFC6749] defines the Error Response associated with the Authorization Code Grant flow and the Implicit Grant flow. Both flows use a redirection endpoint where the resource owner's user agent is directed after the resource owner has completed interacting with the authorization server. The redirection endpoint is also used in the error response scenario. As per RFC6749 Section 4.1.2.1 and 4.2.2.1 [RFC6749] if the resource owner denies the access request or if the request fails for reasons other than a missing or invalid redirection URI, the ?AS? redirects the user-agent by sending the following HTTP response:

```
HTTP/1.1 302 Found Location: https://client.example.com/  
cb?error=access_denied
```

2.1. Abuse: The Authorization Server As Open Redirector

As described in [RFC6819] an attacker could utilize a user's trust in an ?AS? to launch a phishing attack. The attack described here though is not mitigated using the countermeasures listed in [RFC6819]. In this scenario the attacker:

- o Performs a client registration as per the core specification [RFC6749]. The provided redirection URI is a malicious one e.g. `https://attacker.com` (namely the one where the victim's user agent will land without any validation)
- o Prepare a forged URI using the assumption that the ?AS? complies with the OAuth 2.0 specification [RFC6749]. In particular with the ?AS? Error Response described in the previous section (Section 2). As an example he can use a wrong or not existing scope e.g.

```
https://AUTHORIZATION_SERVER/authorize?response_type=code&client_i  
d=s6BhdRkqt3&state=xyz&redirect_uri=https%3A%2F%2Fattacker%2Ecom&s  
cope=INVALID_SCOPE
```

- o Attempt the phishing attack trying to have the victim clicking the forged URI prepared on the previous step. Should the attack succeeds the victim's user agent is redirected to `https://attacker.com` (all with any user interaction) The HTTP

referer header will be set to the AS domain perhaps allowing manipulation of the user.

2.2. Security Compromise: The Authorization Server As Open Redirector

The attacker can use a redirect error redirection to intercept redirect based protocol messages via the Referer header and URI fragment. In this scenario the attacker:

- o Performs a registration of a malicious client as per the core specification [RFC6749]. The provided redirection URI is a malicious one e.g. `https://attacker.com` (This URI will capture the fragment and referer header sent as part of the error)
- o Creates a invalid Authentication request URI for the malicious client. As an example he can use a wrong or not existing scope e.g.

```
https://AUTHORIZATION_SERVER/authorize?response_type=code&client_id=malicious_client&redirect_uri=https%3A%2F%2Fattacker%2Ecom&scope=INVALID_SCOPE
```

- o If the AS supports sticky grants (not re-prompting for consent based on a previous grant) a valid authentication request for the user may also be used to trigger a 30x redirect.
- o Performs a OAuth Authorization request using the invalid Authorization request as the `redirect_uri`. This works if the AS is pattern matching `redirect_uri` and has a public client that shares the same domain as the AS.

(line breaks for display only)

```
https://AUTHORIZATION_SERVER/authorize?response_type=token
&client_id=good-client&scope=VALID_SCOPE
&redirect_uri=https%3A%2F%2FAUTHORIZATION_SERVER%FAuthorize
%3Fresponse_type%3Dcode
%26client_id%3Dattacker-client-id
%26scope%3DINVALID_SCOPE
%26redirect_uri%3Dhttps%253A%252F%252Fattacker.com
```

Figure 1

- o Receive the response redirected to `https://attacker.Com`

The legitimate OAuth Authorization response will include an access token in the URI fragment.

Most web browsers will append the fragment to the URI sent in the location header of a 302 response if no fragment is included in the location URI.

If the Authorization request is code instead of token, the same technique is used, but the code is leaked by the browser in the referer header rather than the fragment.

This causes the access token from a successful authorization to be leaked across the redirect to the malicious client. This is due to browser behaviour and not because the AS has included any information in the redirect URI other than the error code.

Protocols other than OAuth may be particularly vulnerable to this if they are only verifying the domain of the redirect. Performing exact redirect URI matching in OAuth will protect the AS, but not other protocols.

It should be noted that a legitimate OAuth client registered with a AS might be compromised and used as a redirect target by an attacker, perhaps without the knowledge of the client site. This increases a the attack surface for a ?AS?.

2.3. Mitigation

In order to defend against the attacks described in Section 2.1 and Section 2.2 the ?AS? can either:

- o Respond with an HTTP 400 (Bad Request) status code.
- o Perform a redirect to an intermediate URI under the control of the AS to clear referer information in the browser that may contain security token information. This page SHOULD provide notice to the resource owner that an error occurred, and request permission to redirect them to an external site.

If redirected, a fragment "#" MUST be appended to the error redirect URI. This prevents the browser from reattaching the fragment from a previous URI to the new location URI.

Some

When redirecting via 30x a Content Security Policy header SHOULD be added:

Content-Security-Policy: referrer origin;

Figure 2

When redirecting via a form post the following tag SHOULD be included:

```
<meta name="referrer" content="origin"/>
```

Figure 3

Only newer browsers support these headers, so users with older browsers will be vulnerable to leaking referer information unless a intermediate redirect is used.s

3. Acknowledgements

We would like to thank all the people that participated to the discussion, namely Bill Burke, Hans Zandbelt, Justin P. Richer, Phil Hunt, Takahiko Kawasaki, Torsten Lodderstedt, Sergey Beryozkin.

4. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.
- [RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<http://www.rfc-editor.org/info/rfc6819>>.

Appendix A. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-01

- o Added information on HTTP headers to include to set referrer to origin

-00

- o Wrote the first draft.
- o Changed Document name to conform to WG naming convention
- o Added Section on redirect leaking security information
- o Added Terminology section
- o fixed file name
- o cleaned up mitigations a bit

Authors' Addresses

John Bradley (editor)
Ping Identity

Phone: +1 202-630-5272
Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Antonio Sanso
Adobe Systems

Email: asanso@adobe.com

Hannes Tschofenig

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

OAuth
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

W. Denniss
Google
J. Bradley
Ping Identity
M. Jones
Microsoft
H. Tschofenig
ARM Limited
March 11, 2019

OAuth 2.0 Device Authorization Grant
draft-ietf-oauth-device-flow-15

Abstract

The OAuth 2.0 Device Authorization Grant is designed for internet-connected devices that either lack a browser to perform a user-agent based authorization, or are input-constrained to the extent that requiring the user to input text in order to authenticate during the authorization flow is impractical. It enables OAuth clients on such devices (like smart TVs, media consoles, digital picture frames, and printers) to obtain user authorization to access protected resources without using an on-device user-agent.

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 <https://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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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

1. Introduction	3
2. Terminology	5
3. Protocol	5
3.1. Device Authorization Request	5
3.2. Device Authorization Response	7
3.3. User Interaction	8
3.3.1. Non-textual Verification URI Optimization	9
3.4. Device Access Token Request	10
3.5. Device Access Token Response	11
4. Discovery Metadata	12
5. Security Considerations	12
5.1. User Code Brute Forcing	13
5.2. Device Code Brute Forcing	13
5.3. Device Trustworthiness	14
5.4. Remote Phishing	14
5.5. Session Spying	15
5.6. Non-confidential Clients	15
5.7. Non-Visual Code Transmission	15
6. Usability Considerations	15
6.1. User Code Recommendations	16
6.2. Non-Browser User Interaction	17
7. IANA Considerations	17
7.1. OAuth Parameters Registration	17
7.1.1. Registry Contents	17
7.2. OAuth URI Registration	17
7.2.1. Registry Contents	17
7.3. OAuth Extensions Error Registration	17
7.3.1. Registry Contents	18
7.4. OAuth 2.0 Authorization Server Metadata	18
7.4.1. Registry Contents	18
8. Normative References	18
Appendix A. Acknowledgements	19
Appendix B. Document History	20
Authors' Addresses	22

1. Introduction

This OAuth 2.0 [RFC6749] protocol extension, sometimes referred to as "device flow", enables OAuth clients to request user authorization from applications on devices that have limited input capabilities or lack a suitable browser. Such devices include those smart TVs, media console, picture frames and printers which lack an easy input method or suitable browser required for traditional OAuth interactions. The authorization flow defined by this specification instructs the user to review the authorization request on a secondary device, such as a smartphone which does have the requisite input and browser capabilities to complete the user interaction.

The Device Authorization Grant is not intended to replace browser-based OAuth in native apps on capable devices like smartphones. Those apps should follow the practices specified in OAuth 2.0 for Native Apps [RFC8252].

The operating requirements to be able to use this authorization grant type are:

- (1) The device is already connected to the Internet.
- (2) The device is able to make outbound HTTPS requests.
- (3) The device is able to display or otherwise communicate a URI and code sequence to the user.
- (4) The user has a secondary device (e.g., personal computer or smartphone) from which they can process the request.

As the device authorization grant does not require two-way communication between the OAuth client and the user-agent (unlike other OAuth 2 grant types such as the Authorization Code and Implicit grant types), it supports several use cases that cannot be served by those other approaches.

Instead of interacting with the end user's user agent, the client instructs the end user to use another computer or device and connect to the authorization server to approve the access request. Since the protocol supports clients that can't receive incoming requests, clients poll the authorization server repeatedly until the end user completes the approval process.

The device typically chooses the set of authorization servers to support (i.e., its own authorization server, or those by providers it has relationships with). It is not uncommon for the device application to support only a single authorization server, such as

with a TV application for a specific media provider that supports only that media provider's authorization server. The user may not have an established relationship yet with that authorization provider, though one can potentially be set up during the authorization flow.

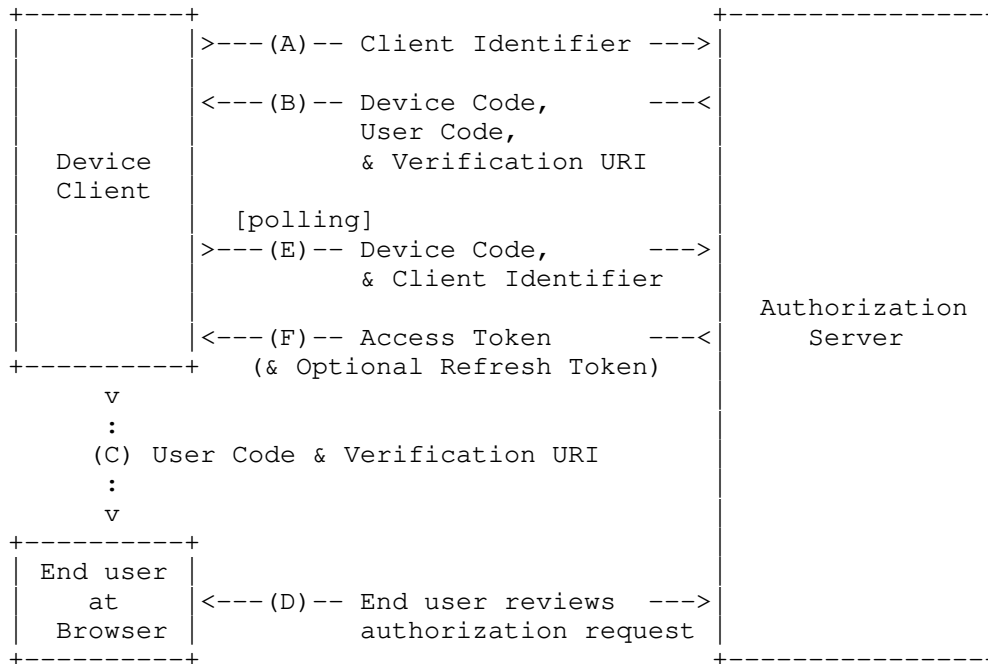


Figure 1: Device Authorization Flow

The device authorization flow illustrated in Figure 1 includes the following steps:

- (A) The client requests access from the authorization server and includes its client identifier in the request.
- (B) The authorization server issues a device code, an end-user code, and provides the end-user verification URI.
- (C) The client instructs the end user to use its user agent (on another device) and visit the provided end-user verification URI. The client provides the user with the end-user code to enter in order to review the authorization request.
- (D) The authorization server authenticates the end user (via the user agent) and prompts the user to grant the client's access

request. If the user agrees to the client's access request, the user enters the user code provided by the client. The authorization server validates the user code provided by the user.

(E) While the end user reviews the client's request (step D), the client repeatedly polls the authorization server to find out if the user completed the user authorization step. The client includes the verification code and its client identifier.

(F) The authorization server validates the verification code provided by the client and responds back with the access token if the user granted access, an error if they denied access, or indicates that the client should continue to poll.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Device Authorization Endpoint:

The authorization server's endpoint capable of issuing device verification codes, user codes, and verification URLs.

Device Verification Code:

A short-lived token representing an authorization session.

End-User Verification Code:

A short-lived token which the device displays to the end user, is entered by the user on the authorization server, and is thus used to bind the device to the user.

3. Protocol

3.1. Device Authorization Request

This specification defines a new OAuth endpoint, the device authorization endpoint. This is separate from the OAuth authorization endpoint defined in [RFC6749] with which the user interacts with via a user-agent (i.e., a browser). By comparison, when using the device authorization endpoint, the OAuth client on the device interacts with the authorization server directly without presenting the request in a user-agent, and the end user authorizes the request on a separate device. This interaction is defined as follows.

The client initiates the authorization flow by requesting a set of verification codes from the authorization server by making an HTTP "POST" request to the device authorization endpoint.

The client constructs the request with the following parameters, sent as the body of the request, encoded with the "application/x-www-form-urlencoded" encoding algorithm defined by Section 4.10.22.6 of [HTML5]:

`client_id`
REQUIRED, if the client is not authenticating with the authorization server as described in Section 3.2.1. of [RFC6749]. The client identifier as described in Section 2.2 of [RFC6749].

`scope`
OPTIONAL. The scope of the access request as described by Section 3.3 of [RFC6749].

For example, the client makes the following HTTPS request:

```
POST /device_authorization HTTP/1.1
Host: server.example.com
Content-Type: application/x-www-form-urlencoded

client_id=459691054427
```

All requests from the device MUST use the Transport Layer Security (TLS) [RFC8446] protocol and implement the best practices of BCP 195 [RFC7525].

Parameters sent without a value MUST be treated as if they were omitted from the request. The authorization server MUST ignore unrecognized request parameters. Request and response parameters MUST NOT be included more than once.

The client authentication requirements of Section 3.2.1 of [RFC6749] apply to requests on this endpoint, which means that confidential clients (those that have established client credentials) authenticate in the same manner as when making requests to the token endpoint, and public clients provide the "client_id" parameter to identify themselves.

Due to the polling nature of this protocol (as specified in Section 3.4), care is needed to avoid overloading the capacity of the token endpoint. To avoid unneeded requests on the token endpoint, the client SHOULD only commence a device authorization request when prompted by the user, and not automatically, such as when the app starts or when the previous authorization session expires or fails.

3.2. Device Authorization Response

In response, the authorization server generates a unique device verification code and an end-user code that are valid for a limited time and includes them in the HTTP response body using the "application/json" format [RFC8259] with a 200 (OK) status code. The response contains the following parameters:

`device_code`
REQUIRED. The device verification code.

`user_code`
REQUIRED. The end-user verification code.

`verification_uri`
REQUIRED. The end-user verification URI on the authorization server. The URI should be short and easy to remember as end users will be asked to manually type it into their user-agent.

`verification_uri_complete`
OPTIONAL. A verification URI that includes the "user_code" (or other information with the same function as the "user_code"), designed for non-textual transmission.

`expires_in`
REQUIRED. The lifetime in seconds of the "device_code" and "user_code".

`interval`
OPTIONAL. The minimum amount of time in seconds that the client SHOULD wait between polling requests to the token endpoint. If no value is provided, clients MUST use 5 as the default.

For example:

```
HTTP/1.1 200 OK
Content-Type: application/json
Cache-Control: no-store

{
  "device_code": "GmRhmhcxhwAzkoEqiMEg_DnyEysNkuNhszIySk9eS",
  "user_code": "WDJB-MJHT",
  "verification_uri": "https://example.com/device",
  "verification_uri_complete":
    "https://example.com/device?user_code=WDJB-MJHT",
  "expires_in": 1800,
  "interval": 5
}
```

In the event of an error (such as an invalidly configured client), the authorization server responds in the same way as the token endpoint specified in Section 5.2 of [RFC6749].

3.3. User Interaction

After receiving a successful Authorization Response, the client displays or otherwise communicates the "user_code" and the "verification_uri" to the end user and instructs them to visit the URI in a user agent on a secondary device (for example, in a browser on their mobile phone), and enter the user code.

```
+-----+
|
|   Using a browser on another device, visit:
|   https://example.com/device
|
|   And enter the code:
|   WDJB-MJHT
|
+-----+
```

Figure 2: Example User Instruction

The authorizing user navigates to the "verification_uri" and authenticates with the authorization server in a secure TLS-protected ([RFC8446]) session. The authorization server prompts the end user to identify the device authorization session by entering the "user_code" provided by the client. The authorization server should then inform the user about the action they are undertaking and ask them to approve or deny the request. Once the user interaction is complete, the server MAY inform the user to return to their device.

During the user interaction, the device continuously polls the token endpoint with the "device_code", as detailed in Section 3.4, until the user completes the interaction, the code expires, or another error occurs. The "device_code" is not intended for the end user directly, and thus should not be displayed during the interaction to avoid confusing the end user.

Authorization servers supporting this specification MUST implement a user interaction sequence that starts with the user navigating to "verification_uri" and continues with them supplying the "user_code" at some stage during the interaction. Other than that, the exact sequence and implementation of the user interaction is up to the authorization server, for example, the authorization server may enable new users to sign up for an account during the authorization flow, or add additional security verification steps.

It is NOT RECOMMENDED for authorization servers to include the user code in the verification URI ("verification_uri"), as this increases the length and complexity of the URI that the user must type. While the user must still type the same number of characters with the "user_code" separated, once they successfully navigate to the "verification_uri", any errors in entering the code can be highlighted by the authorization server to improve the user experience. The next section documents user interaction with "verification_uri_complete", which is designed to carry both pieces of information.

3.3.1. Non-textual Verification URI Optimization

When "verification_uri_complete" is included in the Authorization Response (Section 3.2), clients MAY present this URI in a non-textual manner using any method that results in the browser being opened with the URI, such as with QR (Quick Response) codes or NFC (Near Field Communication), to save the user typing the URI.

For usability reasons, it is RECOMMENDED for clients to still display the textual verification URI ("verification_uri") for users not able to use such a shortcut. Clients MUST still display the "user_code", as the authorization server will require the user to confirm it to disambiguate devices, or as a remote phishing mitigation (See Section 5.4).

If the user starts the user interaction by browsing to "verification_uri_complete", then the user interaction described in Section 3.3 is still followed, but with the optimization that the user does not need to type the "user_code". The server SHOULD display the "user_code" to the user and ask them to verify that it matches the "user_code" being displayed on the device, to confirm they are authorizing the correct device. As before, in addition to taking steps to confirm the identity of the device, the user should also be afforded the choice to approve or deny the authorization request.

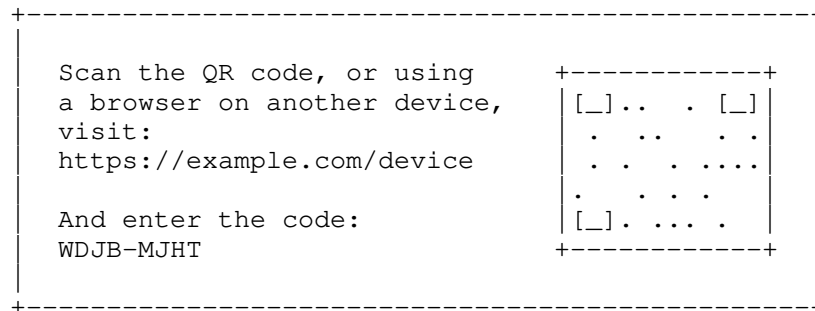


Figure 3: Example User Instruction with QR Code Representation of the Complete Verification URI

3.4. Device Access Token Request

After displaying instructions to the user, the client makes an Access Token Request to the token endpoint (as defined by Section 3.2 of [RFC6749]) with a "grant_type" of "urn:ietf:params:oauth:grant-type:device_code". This is an extension grant type (as defined by Section 4.5 of [RFC6749]) created by this specification, with the following parameters:

grant_type

REQUIRED. Value MUST be set to "urn:ietf:params:oauth:grant-type:device_code".

device_code

REQUIRED. The device verification code, "device_code" from the Device Authorization Response, defined in Section 3.2.

client_id

REQUIRED, if the client is not authenticating with the authorization server as described in Section 3.2.1. of [RFC6749]. The client identifier as described in Section 2.2 of [RFC6749].

For example, the client makes the following HTTPS request (line breaks are for display purposes only):

```
POST /token HTTP/1.1
Host: server.example.com
Content-Type: application/x-www-form-urlencoded

grant_type=urn%3Aietf%3Aparams%3Aoauth%3Agrant-type%3Adevice_code
&device_code=GmRhmhcxhwAzkoEqiMEg_DnyEysNkuNhszIySk9eS
&client_id=459691054427
```

If the client was issued client credentials (or assigned other authentication requirements), the client MUST authenticate with the authorization server as described in Section 3.2.1 of [RFC6749]. Note that there are security implications of statically distributed client credentials, see Section 5.6.

The response to this request is defined in Section 3.5. Unlike other OAuth grant types, it is expected for the client to try the Access Token Request repeatedly in a polling fashion, based on the error code in the response.

3.5. Device Access Token Response

If the user has approved the grant, the token endpoint responds with a success response defined in Section 5.1 of [RFC6749]; otherwise it responds with an error, as defined in Section 5.2 of [RFC6749].

In addition to the error codes defined in Section 5.2 of [RFC6749], the following error codes are specified for use with the device authorization grant in token endpoint responses:

authorization_pending

The authorization request is still pending as the end user hasn't yet completed the user interaction steps (Section 3.3). The client SHOULD repeat the Access Token Request to the token endpoint (a process known as polling). Before each new request the client MUST wait at least the number of seconds specified by the "interval" parameter of the Device Authorization Response (see Section 3.2), or 5 seconds if none was provided, and respect any increase in the polling interval required by the "slow_down" error.

slow_down

A variant of "authorization_pending", the authorization request is still pending and polling should continue, but the interval MUST be increased by 5 seconds for this and all subsequent requests.

access_denied

The end user denied the authorization request.

expired_token

The "device_code" has expired and the device authorization session has concluded. The client MAY commence a new Device Authorization Request but SHOULD wait for user interaction before restarting to avoid unnecessary polling.

The "authorization_pending" and "slow_down" error codes define particularly unique behavior, as they indicate that the OAuth client should continue to poll the token endpoint by repeating the token request (implementing the precise behavior defined above). If the client receives an error response with any other error code, it MUST stop polling and SHOULD react accordingly, for example, by displaying an error to the user.

On encountering a connection timeout, clients MUST unilaterally reduce their polling frequency before retrying. The use of an exponential backoff algorithm to achieve this, such as by doubling the polling interval on each such connection timeout, is RECOMMENDED.

The assumption of this specification is that the separate device the user is authorizing the request on does not have a way to communicate back to device with the OAuth client. This protocol only requires a one-way channel in order to maximise the viability of the protocol in restricted environments, like an application running on a TV that is only capable of outbound requests. If a return channel were to exist for the chosen user interaction interface, then the device MAY wait until notified on that channel that the user has completed the action before initiating the token request (as an alternative to polling). Such behavior is, however, outside the scope of this specification.

4. Discovery Metadata

Support for this specification MAY be declared in the OAuth 2.0 Authorization Server Metadata [RFC8414] by including the value "urn:ietf:params:oauth:grant-type:device_code" in the "grant_types_supported" parameter, and by adding the following new parameter:

device_authorization_endpoint

OPTIONAL. URL of the authorization server's device authorization endpoint defined in Section 3.1.

5. Security Considerations

5.1. User Code Brute Forcing

Since the user code is typed by the user, shorter codes are more desirable for usability reasons. This means the entropy is typically less than would be used for the device code or other OAuth bearer token types where the code length does not impact usability. It is therefore recommended that the server rate-limit user code attempts.

The user code SHOULD have enough entropy that when combined with rate limiting and other mitigations makes a brute-force attack infeasible. For example, it's generally held that 128-bit symmetric keys for encryption are seen as good enough today because an attacker has to put in 2^{96} work to have a 2^{-32} chance of guessing correctly via brute force. The rate limiting and finite lifetime on the user code places an artificial limit on the amount of work an attacker can "do", so if, for instance, one uses a 8-character base-20 user code (with roughly 34.5 bits of entropy), the rate-limiting interval and validity period would need to only allow 5 attempts in order to get the same 2^{-32} probability of success by random guessing.

A successful brute forcing of the user code would enable the attacker to authenticate with their own credentials and make an authorization grant to the device. This is the opposite scenario to an OAuth bearer token being brute forced, whereby the attacker gains control of the victim's authorization grant. Such attacks may not always make economic sense, for example for a video app the device owner may then be able to purchase movies using the attacker's account, though a privacy risk would still remain and thus is important to protect against. Furthermore, some uses of the device flow give the granting account the ability to perform actions such as controlling the device, which needs to be protected.

The precise length of the user code and the entropy contained within is at the discretion of the authorization server, which needs to consider the sensitivity of their specific protected resources, the practicality of the code length from a usability standpoint, and any mitigations that are in place such as rate-limiting, when determining the user code format.

5.2. Device Code Brute Forcing

An attacker who guesses the device code would be able to potentially obtain the authorization code once the user completes the flow. As the device code is not displayed to the user and thus there are no usability considerations on the length, a very high entropy code SHOULD be used.

5.3. Device Trustworthiness

Unlike other native application OAuth 2.0 flows, the device requesting the authorization is not the same as the device that the user grants access from. Thus, signals from the approving user's session and device are not relevant to the trustworthiness of the client device.

Note that if an authorization server used with this flow is malicious, then it could man-in-the-middle the backchannel flow to another authorization server. In this scenario, the man-in-the-middle is not completely hidden from sight, as the end user would end up on the authorization page of the wrong service, giving them an opportunity to notice that the URL in the browser's address bar is wrong. For this to be possible, the device manufacturer must either directly be the attacker, shipping a device intended to perform the man-in-the-middle attack, or be using an authorization server that is controlled by an attacker, possibly because the attacker compromised the authorization server used by the device. In part, the person purchasing the device is counting on it and its business partners to be trustworthy.

5.4. Remote Phishing

It is possible for the device flow to be initiated on a device in an attacker's possession. For example, an attacker might send an email instructing the target user to visit the verification URL and enter the user code. To mitigate such an attack, it is RECOMMENDED to inform the user that they are authorizing a device during the user interaction step (see Section 3.3), and to confirm that the device is in their possession. The authorization server SHOULD display information about the device so that the person can notice if a software client was attempting to impersonating a hardware device.

For authorization servers that support the option specified in Section 3.3.1 for the client to append the user code to the authorization URI, it is particularly important to confirm that the device is in the user's possession, as the user no longer has to type the code manually. One possibility is to display the code during the authorization flow and asking the user to verify that the same code is being displayed on the device they are setting up.

The user code needs to have a long enough lifetime to be useable (allowing the user to retrieve their secondary device, navigate to the verification URI, login, etc.), but should be sufficiently short to limit the usability of a code obtained for phishing. This doesn't prevent a phisher presenting a fresh token, particularly in the case

they are interacting with the user in real time, but it does limit the viability of codes sent over email or SMS.

5.5. Session Spying

While the device is pending authorization, it may be possible for a malicious user to physically spy on the device user interface (by viewing the screen on which it's displayed, for example) and hijack the session by completing the authorization faster than the user that initiated it. Devices SHOULD take into account the operating environment when considering how to communicate the code to the user to reduce the chances it will be observed by a malicious user.

5.6. Non-confidential Clients

Device clients are generally incapable of maintaining the confidentiality of their credentials, as users in possession of the device can reverse engineer it and extract the credentials. Therefore, unless additional measures are taken, they should be treated as public clients (as defined by Section 2.1 of OAuth 2.0) susceptible to impersonation. The security considerations of Section 5.3.1 of [RFC6819] and Sections 8.5 and 8.6 of [RFC8252] apply to such clients.

The user may also be able to obtain the device_code and/or other OAuth bearer tokens issued to their client, which would allow them to use their own authorization grant directly by impersonating the client. Given that the user in possession of the client credentials can already impersonate the client and create a new authorization grant (with a new device_code), this doesn't represent a separate impersonation vector.

5.7. Non-Visual Code Transmission

There is no requirement that the user code be displayed by the device visually. Other methods of one-way communication can potentially be used, such as text-to-speech audio, or Bluetooth Low Energy. To mitigate an attack in which a malicious user can bootstrap their credentials on a device not in their control, it is RECOMMENDED that any chosen communication channel only be accessible by people in close proximity. E.g., users who can see, or hear the device.

6. Usability Considerations

This section is a non-normative discussion of usability considerations.

6.1. User Code Recommendations

For many users, their nearest Internet-connected device will be their mobile phone, and typically these devices offer input methods that are more time consuming than a computer keyboard to change the case or input numbers. To improve usability (improving entry speed, and reducing retries), these limitations should be taken into account when selecting the user-code character set.

One way to improve input speed is to restrict the character set to case-insensitive A-Z characters, with no digits. These characters can typically be entered on a mobile keyboard without using modifier keys. Further removing vowels to avoid randomly creating words results in the base-20 character set: "BCDFGHJKLMNPQRSTVWXZ". Dashes or other punctuation may be included for readability.

An example user code following this guideline containing 8 significant characters and dashes added for end-user readability, with a resulting entropy of 20^8 : "WDJB-MJHT".

Pure numeric codes are also a good choice for usability, especially for clients targeting locales where A-Z character keyboards are not used, though their length needs to be longer to maintain a high entropy.

An example numeric user code containing 9 significant digits and dashes added for end-user readability, with an entropy of 10^9 : "019-450-730".

When processing the inputted user code, the server should strip dashes and other punctuation it added for readability (making the inclusion of that punctuation by the user optional). For codes using only characters in the A-Z range as with the base-20 charset defined above, the user's input should be upper-cased before comparison to account for the fact that the user may input the equivalent lower-case characters. Further stripping of all characters outside the `user_code` charset is recommended to reduce instances where an errantly typed character (like a space character) invalidates otherwise valid input.

It is RECOMMENDED to avoid character sets that contain two or more characters that can easily be confused with each other like "0" and "O", or "1", "l" and "I". Furthermore, the extent practical, where a character set contains one character that may be confused with characters outside the character set the character outside the set MAY be substituted with the one in the character set that it is commonly confused with (for example, "O" for "0" when using a numerical 0-9 character set).

6.2. Non-Browser User Interaction

Devices and authorization servers MAY negotiate an alternative code transmission and user interaction method in addition to the one described in Section 3.3. Such an alternative user interaction flow could obviate the need for a browser and manual input of the code, for example, by using Bluetooth to transmit the code to the authorization server's companion app. Such interaction methods can utilize this protocol, as ultimately, the user just needs to identify the authorization session to the authorization server; however, user interaction other than via the verification URI is outside the scope of this specification.

7. IANA Considerations

7.1. OAuth Parameters Registration

This specification registers the following values in the IANA "OAuth Parameters" registry [IANA.OAuth.Parameters] established by [RFC6749].

7.1.1. Registry Contents

- o Parameter name: device_code
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification Document: Section 3.1 of [[this specification]]

7.2. OAuth URI Registration

This specification registers the following values in the IANA "OAuth URI" registry [IANA.OAuth.Parameters] established by [RFC6755].

7.2.1. Registry Contents

- o URN: urn:ietf:params:oauth:grant-type:device_code
- o Common Name: Device flow grant type for OAuth 2.0
- o Change controller: IESG
- o Specification Document: Section 3.1 of [[this specification]]

7.3. OAuth Extensions Error Registration

This specification registers the following values in the IANA "OAuth Extensions Error Registry" registry [IANA.OAuth.Parameters] established by [RFC6749].

7.3.1. Registry Contents

- o Error name: `authorization_pending`
- o Error usage location: Token endpoint response
- o Related protocol extension: [[this specification]]
- o Change controller: IETF
- o Specification Document: Section 3.5 of [[this specification]]

- o Error name: `access_denied`
- o Error usage location: Token endpoint response
- o Related protocol extension: [[this specification]]
- o Change controller: IETF
- o Specification Document: Section 3.5 of [[this specification]]

- o Error name: `slow_down`
- o Error usage location: Token endpoint response
- o Related protocol extension: [[this specification]]
- o Change controller: IETF
- o Specification Document: Section 3.5 of [[this specification]]

- o Error name: `expired_token`
- o Error usage location: Token endpoint response
- o Related protocol extension: [[this specification]]
- o Change controller: IETF
- o Specification Document: Section 3.5 of [[this specification]]

7.4. OAuth 2.0 Authorization Server Metadata

This specification registers the following values in the IANA "OAuth 2.0 Authorization Server Metadata" registry [IANA.OAuth.Parameters] established by [RFC8414].

7.4.1. Registry Contents

- o Metadata name: `device_authorization_endpoint`
- o Metadata Description: The Device Authorization Endpoint.
- o Change controller: IESG
- o Specification Document: Section 4 of [[this specification]]

8. Normative References

[HTML5] IANA, "HTML5",
<<https://www.w3.org/TR/2014/REC-html5-20141028/>>.

[IANA.OAuth.Parameters]
IANA, "OAuth Parameters",
<<http://www.iana.org/assignments/oauth-parameters>>.

- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.
- [RFC6755] Campbell, B. and H. Tschofenig, "An IETF URN Sub-Namespace for OAuth", RFC 6755, DOI 10.17487/RFC6755, October 2012, <<https://www.rfc-editor.org/info/rfc6755>>.
- [RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<https://www.rfc-editor.org/info/rfc6819>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8252] Denniss, W. and J. Bradley, "OAuth 2.0 for Native Apps", BCP 212, RFC 8252, DOI 10.17487/RFC8252, October 2017, <<https://www.rfc-editor.org/info/rfc8252>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC8414] Jones, M., Sakimura, N., and J. Bradley, "OAuth 2.0 Authorization Server Metadata", RFC 8414, DOI 10.17487/RFC8414, June 2018, <<https://www.rfc-editor.org/info/rfc8414>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

Appendix A. Acknowledgements

The starting point for this document was the Internet-Draft draft-recordon-oauth-v2-device, authored by David Recordon and Brent Goldman, which itself was based on content in draft versions of the OAuth 2.0 protocol specification removed prior to publication due to

a then lack of sufficient deployment expertise. Thank you to the OAuth working group members who contributed to those earlier drafts.

This document was produced in the OAuth working group under the chairpersonship of Rifaat Shekh-Yusef and Hannes Tschofenig with Benjamin Kaduk, Kathleen Moriarty, and Eric Rescorla serving as Security Area Directors.

The following individuals contributed ideas, feedback, and wording that shaped and formed the final specification:

Alissa Cooper, Ben Campbell, Brian Campbell, Roshni Chandrashekhar, Eric Fazendin, Benjamin Kaduk, Jamshid Khosravian, Torsten Lodderstedt, James Manger, Dan McNulty, Breno de Medeiros, Simon Moffatt, Stein Myrseth, Emond Papegaaij, Justin Richer, Adam Roach, Nat Sakimura, Andrew Sciberras, Marius Scurtescu, Filip Skokan, Ken Wang, and Steven E. Wright.

Appendix B. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-15

- o Renamed and dropped most usage of the term "flow"
- o Documented error responses on the authorization endpoint
- o Documented client authentication for the authorization endpoint

-14

- o Added more normative text on polling behavior.
- o Added discussion on risk of user retrieving their own device_code.
- o Editorial improvements.

-13

- o Added a longer discussion about entropy, proposed by Benjamin Kaduk.
- o Added device_code to OAuth IANA registry.
- o Expanded explanation of "case insensitive".
- o Added security section on Device Code Brute Forcing.
- o application/x-www-form-urlencoded normativly referenced.
- o Editorial improvements.

-12

- o Set a default polling interval to 5s explicitly.

- o Defined the `slow_down` behavior that it should increase the current interval by 5s.
- o `expires_in` now REQUIRED
- o Other changes in response to review feedback.

-11

- o Updated reference to OAuth 2.0 Authorization Server Metadata.

-10

- o Added a missing definition of `access_denied` for use on the token endpoint.
- o Corrected text documenting which error code should be returned for expired tokens (it's `"expired_token"`, not `"invalid_grant"`).
- o Corrected section reference to RFC 8252 (the section numbers had changed after the initial reference was made).
- o Fixed line length of one diagram (was causing `xml2rfc` warnings).
- o Added line breaks so the URN `grant_type` is presented on an unbroken line.
- o Typos fixed and other stylistic improvements.

-09

- o Addressed review comments by Security Area Director Eric Rescorla about the potential of a confused deputy attack.

-08

- o Expanded the User Code Brute Forcing section to include more detail on this attack.

-07

- o Replaced the `"user_code"` URI parameter optimization with `verification_uri_complete` following the IETF99 working group discussion.
- o Added security consideration about spying.
- o Required that `device_code` not be shown.
- o Added text regarding a minimum polling interval.

-06

- o Clarified usage of the `"user_code"` URI parameter optimization following the IETF98 working group discussion.

-05

- o response_type parameter removed from authorization request.
- o Added option for clients to include the user_code on the verification URI.
- o Clarified token expiry, and other nits.

-04

- o Security & Usability sections. OAuth Discovery Metadata.

-03

- o device_code is now a URN. Added IANA Considerations

-02

- o Added token request & response specification.

-01

- o Applied spelling and grammar corrections and added the Document History appendix.

-00

- o Initial working group draft based on draft-recordon-oauth-v2-device.

Authors' Addresses

William Denniss
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
USA

Email: wdenniss@google.com
URI: <http://wdenniss.com/device-flow>

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

Hannes Tschofenig
ARM Limited
Austria

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2018

M. Jones
Microsoft
N. Sakimura
NRI
J. Bradley
Ping Identity
March 4, 2018

OAuth 2.0 Authorization Server Metadata
draft-ietf-oauth-discovery-10

Abstract

This specification defines a metadata format that an OAuth 2.0 client can use to obtain the information needed to interact with an OAuth 2.0 authorization server, including its endpoint locations and authorization server capabilities.

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 <https://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 September 5, 2018.

Copyright Notice

Copyright (c) 2018 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 (<https://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

- 1. Introduction 2
 - 1.1. Requirements Notation and Conventions 3
 - 1.2. Terminology 3
- 2. Authorization Server Metadata 4
 - 2.1. Signed Authorization Server Metadata 7
- 3. Obtaining Authorization Server Metadata 8
 - 3.1. Authorization Server Metadata Request 9
 - 3.2. Authorization Server Metadata Response 9
 - 3.3. Authorization Server Metadata Validation 10
- 4. String Operations 11
- 5. Compatibility Notes 11
- 6. Security Considerations 12
 - 6.1. TLS Requirements 12
 - 6.2. Impersonation Attacks 12
 - 6.3. Publishing Metadata in a Standard Format 13
 - 6.4. Protected Resources 13
- 7. IANA Considerations 13
 - 7.1. OAuth Authorization Server Metadata Registry 14
 - 7.1.1. Registration Template 15
 - 7.1.2. Initial Registry Contents 15
 - 7.2. Updated Registration Instructions 18
 - 7.3. Well-Known URI Registry 19
 - 7.3.1. Registry Contents 19
- 8. References 19
 - 8.1. Normative References 19
 - 8.2. Informative References 21
- Appendix A. Acknowledgements 22
- Appendix B. Document History 22
- Authors' Addresses 25

1. Introduction

This specification generalizes the metadata format defined by "OpenID Connect Discovery 1.0" [OpenID.Discovery] in a way that is compatible with OpenID Connect Discovery, while being applicable to a wider set of OAuth 2.0 use cases. This is intentionally parallel to the way that the "OAuth 2.0 Dynamic Client Registration Protocol" [RFC7591] specification generalized the dynamic client registration mechanisms defined by "OpenID Connect Dynamic Client Registration 1.0" [OpenID.Registration] in a way that was compatible with it.

The metadata for an authorization server is retrieved from a well-known location as a JSON [RFC7159] document, which declares its

endpoint locations and authorization server capabilities. This process is described in Section 3.

This metadata can either be communicated in a self-asserted fashion by the server origin via HTTPS or as a set of signed metadata values represented as claims in a JSON Web Token (JWT) [JWT]. In the JWT case, the issuer is vouching for the validity of the data about the authorization server. This is analogous to the role that the Software Statement plays in OAuth Dynamic Client Registration [RFC7591].

The means by which the client chooses an authorization server is out of scope. In some cases, its issuer identifier may be manually configured into the client. In other cases, it may be dynamically discovered, for instance, through the use of WebFinger [RFC7033], as described in Section 2 of "OpenID Connect Discovery 1.0" [OpenID.Discovery].

1.1. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

All uses of JSON Web Signature (JWS) [JWS] and JSON Web Encryption (JWE) [JWE] data structures in this specification utilize the JWS Compact Serialization or the JWE Compact Serialization; the JWS JSON Serialization and the JWE JSON Serialization are not used.

1.2. Terminology

This specification uses the terms "Access Token", "Authorization Code", "Authorization Endpoint", "Authorization Grant", "Authorization Server", "Client", "Client Authentication", "Client Identifier", "Client Secret", "Grant Type", "Protected Resource", "Redirection URI", "Refresh Token", "Resource Owner", "Resource Server", "Response Type", and "Token Endpoint" defined by OAuth 2.0 [RFC6749], the terms "Claim Name", "Claim Value", and "JSON Web Token (JWT)" defined by JSON Web Token (JWT) [JWT], and the term "Response Mode" defined by OAuth 2.0 Multiple Response Type Encoding Practices [OAuth.Responses].

2. Authorization Server Metadata

Authorization servers can have metadata describing their configuration. The following authorization server metadata values are used by this specification and are registered in the IANA "OAuth Authorization Server Metadata" registry established in Section 7.1:

issuer

REQUIRED. The authorization server's issuer identifier, which is a URL that uses the "https" scheme and has no query or fragment components. Authorization server metadata is published at a ".well-known" RFC 5785 [RFC5785] location derived from this issuer identifier, as described in Section 3. The issuer identifier is used to prevent authorization server mix-up attacks, as described in "OAuth 2.0 Mix-Up Mitigation" [I-D.ietf-oauth-mix-up-mitigation].

authorization_endpoint

URL of the authorization server's authorization endpoint [RFC6749]. This is REQUIRED unless no grant types are supported that use the authorization endpoint.

token_endpoint

URL of the authorization server's token endpoint [RFC6749]. This is REQUIRED unless only the implicit grant type is supported.

jwtks_uri

OPTIONAL. URL of the authorization server's JWK Set [JWK] document. The referenced document contains the signing key(s) the client uses to validate signatures from the authorization server. This URL MUST use the "https" scheme. The JWK Set MAY also contain the server's encryption key(s), which are used by clients to encrypt requests to the server. When both signing and encryption keys are made available, a "use" (public key use) parameter value is REQUIRED for all keys in the referenced JWK Set to indicate each key's intended usage.

registration_endpoint

OPTIONAL. URL of the authorization server's OAuth 2.0 Dynamic Client Registration endpoint [RFC7591].

scopes_supported

RECOMMENDED. JSON array containing a list of the OAuth 2.0 [RFC6749] "scope" values that this authorization server supports. Servers MAY choose not to advertise some supported scope values even when this parameter is used.

response_types_supported

REQUIRED. JSON array containing a list of the OAuth 2.0 "response_type" values that this authorization server supports. The array values used are the same as those used with the "response_types" parameter defined by "OAuth 2.0 Dynamic Client Registration Protocol" [RFC7591].

response_modes_supported

OPTIONAL. JSON array containing a list of the OAuth 2.0 "response_mode" values that this authorization server supports, as specified in OAuth 2.0 Multiple Response Type Encoding Practices [OAuth.Responses]. If omitted, the default is "["query", "fragment"]". The response mode value "form_post" is also defined in OAuth 2.0 Form Post Response Mode [OAuth.Post].

grant_types_supported

OPTIONAL. JSON array containing a list of the OAuth 2.0 grant type values that this authorization server supports. The array values used are the same as those used with the "grant_types" parameter defined by "OAuth 2.0 Dynamic Client Registration Protocol" [RFC7591]. If omitted, the default value is "["authorization_code", "implicit"]".

token_endpoint_auth_methods_supported

OPTIONAL. JSON array containing a list of client authentication methods supported by this token endpoint. Client authentication method values are used in the "token_endpoint_auth_method" parameter defined in Section 2 of [RFC7591]. If omitted, the default is "client_secret_basic" -- the HTTP Basic Authentication Scheme specified in Section 2.3.1 of OAuth 2.0 [RFC6749].

token_endpoint_auth_signing_alg_values_supported

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the token endpoint for the signature on the JWT [JWT] used to authenticate the client at the token endpoint for the "private_key_jwt" and "client_secret_jwt" authentication methods. This metadata entry MUST be present if either of these authentication methods are specified in the "token_endpoint_auth_methods_supported" entry. No default algorithms are implied if this entry is omitted. Servers SHOULD support "RS256". The value "none" MUST NOT be used.

service_documentation

OPTIONAL. URL of a page containing human-readable information that developers might want or need to know when using the authorization server. In particular, if the authorization server does not support Dynamic Client Registration, then information on how to register clients needs to be provided in this documentation.

ui_locales_supported

OPTIONAL. Languages and scripts supported for the user interface, represented as a JSON array of BCP47 [RFC5646] language tag values. If omitted, the set of supported languages and scripts is unspecified.

op_policy_uri

OPTIONAL. URL that the authorization server provides to the person registering the client to read about the authorization server's requirements on how the client can use the data provided by the authorization server. The registration process SHOULD display this URL to the person registering the client if it is given. As described in Section 5, despite the identifier "op_policy_uri", appearing to be OpenID-specific, its usage in this specification is actually referring to a general OAuth 2.0 feature that is not specific to OpenID Connect.

op_tos_uri

OPTIONAL. URL that the authorization server provides to the person registering the client to read about the authorization server's terms of service. The registration process SHOULD display this URL to the person registering the client if it is given. As described in Section 5, despite the identifier "op_tos_uri", appearing to be OpenID-specific, its usage in this specification is actually referring to a general OAuth 2.0 feature that is not specific to OpenID Connect.

revocation_endpoint

OPTIONAL. URL of the authorization server's OAuth 2.0 revocation endpoint [RFC7009].

revocation_endpoint_auth_methods_supported

OPTIONAL. JSON array containing a list of client authentication methods supported by this revocation endpoint. The valid client authentication method values are those registered in the IANA "OAuth Token Endpoint Authentication Methods" registry [IANA.OAuth.Parameters]. If omitted, the default is "client_secret_basic" -- the HTTP Basic Authentication Scheme specified in Section 2.3.1 of OAuth 2.0 [RFC6749].

revocation_endpoint_auth_signing_alg_values_supported

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the revocation endpoint for the signature on the JWT [JWT] used to authenticate the client at the revocation endpoint for the "private_key_jwt" and "client_secret_jwt" authentication methods. This metadata entry MUST be present if either of these authentication methods are specified in the "revocation_endpoint_auth_methods_supported"

entry. No default algorithms are implied if this entry is omitted. The value "none" MUST NOT be used.

`introspection_endpoint`

OPTIONAL. URL of the authorization server's OAuth 2.0 introspection endpoint [RFC7662].

`introspection_endpoint_auth_methods_supported`

OPTIONAL. JSON array containing a list of client authentication methods supported by this introspection endpoint. The valid client authentication method values are those registered in the IANA "OAuth Token Endpoint Authentication Methods" registry [IANA.OAuth.Parameters] or those registered in the IANA "OAuth Access Token Types" registry [IANA.OAuth.Parameters]. (These values are and will remain distinct, due to Section 7.2.) If omitted, the set of supported authentication methods MUST be determined by other means.

`introspection_endpoint_auth_signing_alg_values_supported`

OPTIONAL. JSON array containing a list of the JWS signing algorithms ("alg" values) supported by the introspection endpoint for the signature on the JWT [JWT] used to authenticate the client at the introspection endpoint for the "private_key_jwt" and "client_secret_jwt" authentication methods. This metadata entry MUST be present if either of these authentication methods are specified in the "introspection_endpoint_auth_methods_supported" entry. No default algorithms are implied if this entry is omitted. The value "none" MUST NOT be used.

`code_challenge_methods_supported`

OPTIONAL. JSON array containing a list of PKCE [RFC7636] code challenge methods supported by this authorization server. Code challenge method values are used in the "code_challenge_method" parameter defined in Section 4.3 of [RFC7636]. The valid code challenge method values are those registered in the IANA "PKCE Code Challenge Methods" registry [IANA.OAuth.Parameters]. If omitted, the authorization server does not support PKCE.

Additional authorization server metadata parameters MAY also be used. Some are defined by other specifications, such as OpenID Connect Discovery 1.0 [OpenID.Discovery].

2.1. Signed Authorization Server Metadata

In addition to JSON elements, metadata values MAY also be provided as a "signed_metadata" value, which is a JSON Web Token (JWT) [JWT] that asserts metadata values about the authorization server as a bundle. A set of claims that can be used in signed metadata are defined in

Section 2. The signed metadata MUST be digitally signed or MACed using JSON Web Signature (JWS) [JWS] and MUST contain an "iss" (issuer) claim denoting the party attesting to the claims in the signed metadata. Consumers of the metadata MAY ignore the signed metadata if they do not support this feature. If the consumer of the metadata supports signed metadata, metadata values conveyed in the signed metadata MUST take precedence over the corresponding values conveyed using plain JSON elements.

Signed metadata is included in the authorization server metadata JSON object using this OPTIONAL member:

`signed_metadata`

A JWT containing metadata values about the authorization server as claims. This is a string value consisting of the entire signed JWT. A "signed_metadata" metadata value SHOULD NOT appear as a claim in the JWT.

3. Obtaining Authorization Server Metadata

Authorization servers supporting metadata MUST make a JSON document containing metadata as specified in Section 2 available at a path formed by inserting a well-known URI string into the authorization server's issuer identifier between the host component and the path component, if any. By default, the well-known URI string used is `"/.well-known/oauth-authorization-server"`. This path MUST use the "https" scheme. The syntax and semantics of ".well-known" are defined in RFC 5785 [RFC5785]. The well-known URI suffix used MUST be registered in the IANA "Well-Known URIs" registry [IANA.well-known].

Different applications utilizing OAuth authorization servers in application-specific ways may define and register different well-known URI suffixes used to publish authorization server metadata as used by those applications. For instance, if the Example application uses an OAuth authorization server in an Example-specific way, and there are Example-specific metadata values that it needs to publish, then it might register and use the "example-configuration" URI suffix and publish the metadata document at the path formed by inserting `"/.well-known/example-configuration"` between the host and path components of the authorization server's issuer identifier. Alternatively, many such applications will use the default well-known URI string `"/.well-known/oauth-authorization-server"`, which is the right choice for general-purpose OAuth authorization servers, and not register an application-specific one.

An OAuth 2.0 application using this specification MUST specify what well-known URI suffix it will use for this purpose. The same

authorization server MAY choose to publish its metadata at multiple well-known locations derived from its issuer identifier, for example, publishing metadata at both `"/.well-known/example-configuration"` and `"/.well-known/oauth-authorization-server"`.

Some OAuth applications will choose to use the well-known URI suffix `"openid-configuration"`. As described in Section 5, despite the identifier `"/.well-known/openid-configuration"`, appearing to be OpenID-specific, its usage in this specification is actually referring to a general OAuth 2.0 feature that is not specific to OpenID Connect.

3.1. Authorization Server Metadata Request

An authorization server metadata document MUST be queried using an HTTP `"GET"` request at the previously specified path.

The client would make the following request when the issuer identifier is `"https://example.com"` and the well-known URI suffix is `"oauth-authorization-server"` to obtain the metadata, since the issuer identifier contains no path component:

```
GET /.well-known/oauth-authorization-server HTTP/1.1
Host: example.com
```

If the issuer identifier value contains a path component, any terminating `"/"` MUST be removed before inserting `"/.well-known/"` and the well-known URI suffix between the host component and the path component. The client would make the following request when the issuer identifier is `"https://example.com/issuer1"` and the well-known URI suffix is `"oauth-authorization-server"` to obtain the metadata, since the issuer identifier contains a path component:

```
GET /.well-known/oauth-authorization-server/issuer1 HTTP/1.1
Host: example.com
```

Using path components enables supporting multiple issuers per host. This is required in some multi-tenant hosting configurations. This use of `".well-known"` is for supporting multiple issuers per host; unlike its use in RFC 5785 [RFC5785], it does not provide general information about the host.

3.2. Authorization Server Metadata Response

The response is a set of claims about the authorization server's configuration, including all necessary endpoints and public key location information. A successful response MUST use the 200 OK HTTP status code and return a JSON object using the `"application/json"`

content type that contains a set of claims as its members that are a subset of the metadata values defined in Section 2. Other claims MAY also be returned.

Claims that return multiple values are represented as JSON arrays. Claims with zero elements MUST be omitted from the response.

An error response uses the applicable HTTP status code value.

The following is a non-normative example response:

```
HTTP/1.1 200 OK
Content-Type: application/json

{
  "issuer":
    "https://server.example.com",
  "authorization_endpoint":
    "https://server.example.com/authorize",
  "token_endpoint":
    "https://server.example.com/token",
  "token_endpoint_auth_methods_supported":
    ["client_secret_basic", "private_key_jwt"],
  "token_endpoint_auth_signing_alg_values_supported":
    ["RS256", "ES256"],
  "userinfo_endpoint":
    "https://server.example.com/userinfo",
  "jwks_uri":
    "https://server.example.com/jwks.json",
  "registration_endpoint":
    "https://server.example.com/register",
  "scopes_supported":
    ["openid", "profile", "email", "address",
     "phone", "offline_access"],
  "response_types_supported":
    ["code", "code token"],
  "service_documentation":
    "http://server.example.com/service_documentation.html",
  "ui_locales_supported":
    ["en-US", "en-GB", "en-CA", "fr-FR", "fr-CA"]
}
```

3.3. Authorization Server Metadata Validation

The "issuer" value returned MUST be identical to the authorization server's issuer identifier value into which the well-known URI string was inserted to create the URL used to retrieve the metadata. If

these values are not identical, the data contained in the response MUST NOT be used.

4. String Operations

Processing some OAuth 2.0 messages requires comparing values in the messages to known values. For example, the member names in the metadata response might be compared to specific member names such as "issuer". Comparing Unicode [UNICODE] strings, however, has significant security implications.

Therefore, comparisons between JSON strings and other Unicode strings MUST be performed as specified below:

1. Remove any JSON applied escaping to produce an array of Unicode code points.
2. Unicode Normalization [USA15] MUST NOT be applied at any point to either the JSON string or to the string it is to be compared against.
3. Comparisons between the two strings MUST be performed as a Unicode code point to code point equality comparison.

Note that this is the same equality comparison procedure described in Section 8.3 of [RFC7159].

5. Compatibility Notes

The identifiers `"/.well-known/openid-configuration"`, `"op_policy_uri"`, and `"op_tos_uri"` contain strings referring to the OpenID Connect [OpenID.Core] family of specifications that were originally defined by "OpenID Connect Discovery 1.0" [OpenID.Discovery]. Despite the reuse of these identifiers that appear to be OpenID-specific, their usage in this specification is actually referring to general OAuth 2.0 features that are not specific to OpenID Connect.

The algorithm for transforming the issuer identifier to an authorization server metadata location defined in Section 3 is equivalent to the corresponding transformation defined in Section 4 of "OpenID Connect Discovery 1.0" [OpenID.Discovery], provided that the issuer identifier contains no path component. However, they are different when there is a path component, because OpenID Connect Discovery 1.0 specifies that the well-known URI string is appended to the issuer identifier (e.g., `"https://example.com/issuer1/.well-known/openid-configuration"`), whereas this specification specifies that the well-known URI string is inserted before the path component

of the issuer identifier (e.g., "https://example.com/.well-known/openid-configuration/issuer1").

Going forward, OAuth authorization server metadata locations should use the transformation defined in this specification. However, when deployed in legacy environments in which the OpenID Connect Discovery 1.0 transformation is already used, it may be necessary during a transition period to publish metadata for issuer identifiers containing a path component at both locations. During this transition period, applications should first apply the transformation defined in this specification and attempt to retrieve the authorization server metadata from the resulting location; only if the retrieval from that location fails should they fall back to attempting to retrieve it from the alternate location obtained using the transformation defined by OpenID Connect Discovery 1.0. This backwards-compatibility behavior should only be necessary when the well-known URI suffix employed by the application is "openid-configuration".

6. Security Considerations

6.1. TLS Requirements

Implementations MUST support TLS. Which version(s) ought to be implemented will vary over time and depend on the widespread deployment and known security vulnerabilities at the time of implementation. The authorization server MUST support TLS version 1.2 [RFC5246] and MAY support additional transport-layer security mechanisms meeting its security requirements. When using TLS, the client MUST perform a TLS/SSL server certificate check, per RFC 6125 [RFC6125]. Implementation security considerations can be found in Recommendations for Secure Use of TLS and DTLS [BCP195].

To protect against information disclosure and tampering, confidentiality protection MUST be applied using TLS with a ciphersuite that provides confidentiality and integrity protection.

6.2. Impersonation Attacks

TLS certificate checking MUST be performed by the client, as described in Section 6.1, when making an authorization server metadata request. Checking that the server certificate is valid for the issuer identifier URL prevents man-in-middle and DNS-based attacks. These attacks could cause a client to be tricked into using an attacker's keys and endpoints, which would enable impersonation of the legitimate authorization server. If an attacker can accomplish this, they can access the resources that the affected client has access to using the authorization server that they are impersonating.

An attacker may also attempt to impersonate an authorization server by publishing a metadata document that contains an "issuer" claim using the issuer identifier URL of the authorization server being impersonated, but with its own endpoints and signing keys. This would enable it to impersonate that authorization server, if accepted by the client. To prevent this, the client MUST ensure that the issuer identifier URL it is using as the prefix for the metadata request exactly matches the value of the "issuer" metadata value in the authorization server metadata document received by the client.

6.3. Publishing Metadata in a Standard Format

Publishing information about the authorization server in a standard format makes it easier for both legitimate clients and attackers to use the authorization server. Whether an authorization server publishes its metadata in an ad-hoc manner or in the standard format defined by this specification, the same defenses against attacks that might be mounted that use this information should be applied.

6.4. Protected Resources

Secure determination of appropriate protected resources to use with an authorization server for all use cases is out of scope of this specification. This specification assumes that the client has a means of determining appropriate protected resources to use with an authorization server and that the client is using the correct metadata for each authorization server. Implementers need to be aware that if an inappropriate protected resource is used by the client, that an attacker may be able to act as a man-in-the-middle proxy to a valid protected resource without it being detected by the authorization server or the client.

The ways to determine the appropriate protected resources to use with an authorization server are in general, application-dependent. For instance, some authorization servers are used with a fixed protected resource or set of protected resources, the locations of which may be well known, or which could be published as metadata values by the authorization server. In other cases, the set of resources that can be used with an authorization server can be dynamically changed by administrative actions. Many other means of determining appropriate associations between authorization servers and protected resources are also possible.

7. IANA Considerations

The following registration procedure is used for the registry established by this specification.

Values are registered on a Specification Required [RFC8126] basis after a two-week review period on the `oauth-ext-review@ietf.org` mailing list, on the advice of one or more Designated Experts. However, to allow for the allocation of values prior to publication, the Designated Experts may approve registration once they are satisfied that such a specification will be published.

Registration requests sent to the mailing list for review should use an appropriate subject (e.g., "Request to register OAuth Authorization Server Metadata: example").

Within the review period, the Designated Experts will either approve or deny the registration request, communicating this decision to the review list and IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the request successful. Registration requests that are undetermined for a period longer than 21 days can be brought to the IESG's attention (using the `iesg@ietf.org` mailing list) for resolution.

Criteria that should be applied by the Designated Experts includes determining whether the proposed registration duplicates existing functionality, determining whether it is likely to be of general applicability or whether it is useful only for a single application, and whether the registration makes sense.

IANA must only accept registry updates from the Designated Experts and should direct all requests for registration to the review mailing list.

It is suggested that multiple Designated Experts be appointed who are able to represent the perspectives of different applications using this specification, in order to enable broadly-informed review of registration decisions. In cases where a registration decision could be perceived as creating a conflict of interest for a particular Expert, that Expert should defer to the judgment of the other Experts.

7.1. OAuth Authorization Server Metadata Registry

This specification establishes the IANA "OAuth Authorization Server Metadata" registry for OAuth 2.0 authorization server metadata names. The registry records the authorization server metadata member and a reference to the specification that defines it.

The Designated Experts must either:

(a) require that metadata names and values being registered use only printable ASCII characters excluding double quote (") and backslash

('\'') (the Unicode characters with code points U+0021, U+0023 through U+005B, and U+005D through U+007E), or

(b) if new metadata members or values are defined that use other code points, require that their definitions specify the exact Unicode code point sequences used to represent them. Furthermore, proposed registrations that use Unicode code points that can only be represented in JSON strings as escaped characters must not be accepted.

7.1.1. Registration Template

Metadata Name:

The name requested (e.g., "issuer"). This name is case-sensitive. Names may not match other registered names in a case-insensitive manner (one that would cause a match if the Unicode toLowerCase() operation were applied to both strings) unless the Designated Experts state that there is a compelling reason to allow an exception.

Metadata Description:

Brief description of the metadata (e.g., "Issuer identifier URL").

Change Controller:

For Standards Track RFCs, list the "IESG". For others, give the name of the responsible party. Other details (e.g., postal address, email address, home page URI) may also be included.

Specification Document(s):

Reference to the document or documents that specify the parameter, preferably including URIs that can be used to retrieve copies of the documents. An indication of the relevant sections may also be included but is not required.

7.1.2. Initial Registry Contents

- o Metadata Name: "issuer"
- o Metadata Description: Authorization server's issuer identifier URL
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "authorization_endpoint"
- o Metadata Description: URL of the authorization server's authorization endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "token_endpoint"

- o Metadata Description: URL of the authorization server's token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "jwks_uri"
- o Metadata Description: URL of the authorization server's JWK Set document
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "registration_endpoint"
- o Metadata Description: URL of the authorization server's OAuth 2.0 Dynamic Client Registration Endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "scopes_supported"
- o Metadata Description: JSON array containing a list of the OAuth 2.0 "scope" values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "response_types_supported"
- o Metadata Description: JSON array containing a list of the OAuth 2.0 "response_type" values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "response_modes_supported"
- o Metadata Description: JSON array containing a list of the OAuth 2.0 "response_mode" values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "grant_types_supported"
- o Metadata Description: JSON array containing a list of the OAuth 2.0 grant type values that this authorization server supports
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "token_endpoint_auth_methods_supported"
- o Metadata Description: JSON array containing a list of client authentication methods supported by this token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "token_endpoint_auth_signing_alg_values_supported"

- o Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the token endpoint for the signature on the JWT used to authenticate the client at the token endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "service_documentation"
- o Metadata Description: URL of a page containing human-readable information that developers might want or need to know when using the authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "ui_locales_supported"
- o Metadata Description: Languages and scripts supported for the user interface, represented as a JSON array of BCP47 language tag values
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "op_policy_uri"
- o Metadata Description: URL that the authorization server provides to the person registering the client to read about the authorization server's requirements on how the client can use the data provided by the authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "op_tos_uri"
- o Metadata Description: URL that the authorization server provides to the person registering the client to read about the authorization server's terms of service
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "revocation_endpoint"
- o Metadata Description: URL of the authorization server's OAuth 2.0 revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "revocation_endpoint_auth_methods_supported"
- o Metadata Description: JSON array containing a list of client authentication methods supported by this revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name:
"revocation_endpoint_auth_signing_alg_values_supported"
- o Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the revocation endpoint for the signature on the JWT used to authenticate the client at the revocation endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "introspection_endpoint"
- o Metadata Description: URL of the authorization server's OAuth 2.0 introspection endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "introspection_endpoint_auth_methods_supported"
- o Metadata Description: JSON array containing a list of client authentication methods supported by this introspection endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name:
"introspection_endpoint_auth_signing_alg_values_supported"
- o Metadata Description: JSON array containing a list of the JWS signing algorithms supported by the introspection endpoint for the signature on the JWT used to authenticate the client at the introspection endpoint
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

- o Metadata Name: "code_challenge_methods_supported"
- o Metadata Description: PKCE code challenge methods supported by this authorization server
- o Change Controller: IESG
- o Specification Document(s): Section 2 of [[this specification]]

7.2. Updated Registration Instructions

This specification adds to the instructions for the Designated Experts of the following IANA registries, both of which are in the "OAuth Parameters" registry [IANA.OAuth.Parameters]:

- o OAuth Access Token Types
- o OAuth Token Endpoint Authentication Methods

IANA has added a link to this specification in the Reference sections of these registries. [[RFC Editor: The above sentence is written in the past tense as it would appear in the final specification, even

though these links won't actually be created until after the IESG has requested publication of the specification. Please delete this note after the links are in place.]]

For these registries, the designated experts must reject registration requests in one registry for values already occurring in the other registry. This is necessary because the "introspection_endpoint_auth_methods_supported" parameter allows for the use of values from either registry. That way, because the values in the two registries will continue to be mutually exclusive, no ambiguities will arise.

7.3. Well-Known URI Registry

This specification registers the well-known URI defined in Section 3 in the IANA "Well-Known URIs" registry [IANA.well-known] established by RFC 5785 [RFC5785].

7.3.1. Registry Contents

- o URI suffix: "oauth-authorization-server"
- o Change controller: IESG
- o Specification document: Section 3 of [[this specification]]
- o Related information: (none)

8. References

8.1. Normative References

- [BCP195] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/bcp195>>.
- [IANA.OAuth.Parameters] IANA, "OAuth Parameters", <<http://www.iana.org/assignments/oauth-parameters>>.
- [JWE] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<http://tools.ietf.org/html/rfc7516>>.
- [JWK] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<http://tools.ietf.org/html/rfc7517>>.

- [JWS] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<http://tools.ietf.org/html/rfc7515>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<http://tools.ietf.org/html/rfc7519>>.
- [OAuth.Post] Jones, M. and B. Campbell, "OAuth 2.0 Form Post Response Mode", April 2015, <http://openid.net/specs/oauth-v2-form-post-response-mode-1_0.html>.
- [OAuth.Responses] de Medeiros, B., Ed., Scurtescu, M., Tarjan, P., and M. Jones, "OAuth 2.0 Multiple Response Type Encoding Practices", February 2014, <http://openid.net/specs/oauth-v2-multiple-response-types-1_0.html>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<https://www.rfc-editor.org/info/rfc5646>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<https://www.rfc-editor.org/info/rfc5785>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.

- [RFC7009] Lodderstedt, T., Ed., Dronia, S., and M. Scurtescu, "OAuth 2.0 Token Revocation", RFC 7009, DOI 10.17487/RFC7009, August 2013, <<https://www.rfc-editor.org/info/rfc7009>>.
- [RFC7033] Jones, P., Salgueiro, G., Jones, M., and J. Smarr, "WebFinger", RFC 7033, DOI 10.17487/RFC7033, September 2013, <<https://www.rfc-editor.org/info/rfc7033>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<https://www.rfc-editor.org/info/rfc7159>>.
- [RFC7591] Richer, J., Ed., Jones, M., Bradley, J., Machulak, M., and P. Hunt, "OAuth 2.0 Dynamic Client Registration Protocol", RFC 7591, DOI 10.17487/RFC7591, July 2015, <<https://www.rfc-editor.org/info/rfc7591>>.
- [RFC7636] Sakimura, N., Ed., Bradley, J., and N. Agarwal, "Proof Key for Code Exchange by OAuth Public Clients", RFC 7636, DOI 10.17487/RFC7636, September 2015, <<https://www.rfc-editor.org/info/rfc7636>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<https://www.rfc-editor.org/info/rfc7662>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [UNICODE] The Unicode Consortium, "The Unicode Standard", <<http://www.unicode.org/versions/latest/>>.
- [USA15] Davis, M. and K. Whistler, "Unicode Normalization Forms", Unicode Standard Annex 15, June 2015, <<http://www.unicode.org/reports/tr15/>>.

8.2. Informative References

- [I-D.ietf-oauth-mix-up-mitigation]
Jones, M., Bradley, J., and N. Sakimura, "OAuth 2.0 Mix-Up Mitigation", draft-ietf-oauth-mix-up-mitigation-01 (work in progress), July 2016.

[IANA.well-known]

IANA, "Well-Known URIs",
<<http://www.iana.org/assignments/well-known-uris>>.

[OpenID.Core]

Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and
C. Mortimore, "OpenID Connect Core 1.0", November 2014,
<http://openid.net/specs/openid-connect-core-1_0.html>.

[OpenID.Discovery]

Sakimura, N., Bradley, J., Jones, M., and E. Jay, "OpenID
Connect Discovery 1.0", November 2014,
<[http://openid.net/specs/
openid-connect-discovery-1_0.html](http://openid.net/specs/openid-connect-discovery-1_0.html)>.

[OpenID.Registration]

Sakimura, N., Bradley, J., and M. Jones, "OpenID Connect
Dynamic Client Registration 1.0", November 2014,
<[http://openid.net/specs/
openid-connect-registration-1_0.html](http://openid.net/specs/openid-connect-registration-1_0.html)>.

Appendix A. Acknowledgements

This specification is based on the OpenID Connect Discovery 1.0 specification, which was produced by the OpenID Connect working group of the OpenID Foundation. This specification standardizes the de facto usage of the metadata format defined by OpenID Connect Discovery to publish OAuth authorization server metadata.

The authors would like to thank the following people for their reviews of this specification: Shwetha Bhandari, Ben Campbell, Brian Campbell, Brian Carpenter, William Denniss, Vladimir Dzhuvinov, Donald Eastlake, Samuel Erdtman, George Fletcher, Dick Hardt, Phil Hunt, Alexey Melnikov, Tony Nadalin, Mark Nottingham, Eric Rescorla, Justin Richer, Adam Roach, Hannes Tschofenig, and Hans Zandbelt.

Appendix B. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-10

- o Clarified the meaning of "case-insensitive", as suggested by Alexey Melnikov.

-09

- o Revised the transformation between the issuer identifier and the authorization server metadata location to conform to BCP 190, as suggested by Adam Roach.
- o Defined the characters allowed in registered metadata names and values, as suggested by Alexey Melnikov.
- o Changed to using the RFC 8174 boilerplate instead of the RFC 2119 boilerplate, as suggested by Ben Campbell.
- o Acknowledged additional reviewers.

-08

- o Changed the "authorization_endpoint" to be REQUIRED only when grant types are supported that use the authorization endpoint.
- o Added the statement, to provide historical context, that this specification standardizes the de facto usage of the metadata format defined by OpenID Connect Discovery to publish OAuth authorization server metadata.
- o Applied clarifications suggested by Mark Nottingham about when application-specific well-known suffixes are and are not appropriate.
- o Acknowledged additional reviewers.

-07

- o Applied clarifications suggested by EKR.

-06

- o Incorporated resolutions to working group last call comments.

-05

- o Removed the "protected_resources" element and the reference to draft-jones-oauth-resource-metadata.

-04

- o Added the ability to list protected resources with the "protected_resources" element.
- o Added ability to provide signed metadata with the "signed_metadata" element.

- o Removed "Discovery" from the name, since this is now just about authorization server metadata.

-03

- o Changed term "issuer URL" to "issuer identifier" for terminology consistency, paralleling the same terminology consistency change in the mix-up mitigation spec.

-02

- o Changed the title to OAuth 2.0 Authorization Server Discovery Metadata.
- o Made "jwks_uri" and "registration_endpoint" OPTIONAL.
- o Defined the well-known URI string `"/.well-known/oauth-authorization-server"`.
- o Added security considerations about publishing authorization server discovery metadata in a standard format.
- o Added security considerations about protected resources.
- o Added more information to the "grant_types_supported" and "response_types_supported" definitions.
- o Referenced the working group Mix-Up Mitigation draft.
- o Changed some example metadata values.
- o Acknowledged individuals for their contributions to the specification.

-01

- o Removed WebFinger discovery.
- o Clarified the relationship between the issuer identifier URL and the well-known URI path relative to it at which the discovery metadata document is located.

-00

- o Created the initial working group version based on draft-jones-oauth-discovery-01, with no normative changes.

Authors' Addresses

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

Nat Sakimura
Nomura Research Institute, Ltd.

Email: n-sakimura@nri.co.jp
URI: <http://nat.sakimura.org/>

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 12, 2019

N. Sakimura
Nomura Research Institute
J. Bradley
Yubico
June 10, 2019

The OAuth 2.0 Authorization Framework: JWT Secured Authorization Request
(JAR)
draft-ietf-oauth-jwsreq-19

Abstract

The authorization request in OAuth 2.0 described in RFC 6749 utilizes query parameter serialization, which means that Authorization Request parameters are encoded in the URI of the request and sent through user agents such as web browsers. While it is easy to implement, it means that (a) the communication through the user agents are not integrity protected and thus the parameters can be tainted, and (b) the source of the communication is not authenticated. Because of these weaknesses, several attacks to the protocol have now been put forward.

This document introduces the ability to send request parameters in a JSON Web Token (JWT) instead, which allows the request to be signed with JSON Web Signature (JWS) and encrypted with JSON Web Encryption (JWE) so that the integrity, source authentication and confidentiality property of the Authorization Request is attained. The request can be sent by value or by reference.

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 <https://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 December 12, 2019.

Copyright Notice

Copyright (c) 2019 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 (https://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

- 1. Introduction 3
 - 1.1. Requirements Language 5
- 2. Terminology 5
 - 2.1. Request Object 5
 - 2.2. Request Object URI 6
- 3. Symbols and abbreviated terms 6
- 4. Request Object 6
- 5. Authorization Request 8
 - 5.1. Passing a Request Object by Value 9
 - 5.2. Passing a Request Object by Reference 9
 - 5.2.1. URI Referencing the Request Object 11
 - 5.2.2. Request using the "request_uri" Request Parameter . . 11
 - 5.2.3. Authorization Server Fetches Request Object 11
- 6. Validating JWT-Based Requests 12
 - 6.1. Encrypted Request Object 12
 - 6.2. JWS Signed Request Object 13
 - 6.3. Request Parameter Assembly and Validation 13
- 7. Authorization Server Response 13
- 8. TLS Requirements 13
- 9. IANA Considerations 14
- 10. Security Considerations 14
 - 10.1. Choice of Algorithms 14
 - 10.2. Request Source Authentication 15
 - 10.3. Explicit Endpoints 15
 - 10.4. Risks Associated with request_uri 16
 - 10.4.1. DDoS Attack on the Authorization Server 16
 - 10.4.2. Request URI Rewrite 16
- 11. TLS security considerations 17
- 12. Privacy Considerations 17
 - 12.1. Collection limitation 17
 - 12.2. Disclosure Limitation 18

12.2.1. Request Disclosure	18
12.2.2. Tracking using Request Object URI	18
13. Acknowledgements	18
14. Revision History	19
15. References	24
15.1. Normative References	25
15.2. Informative References	26
Authors' Addresses	27

1. Introduction

The Authorization Request in OAuth 2.0 [RFC6749] utilizes query parameter serialization and is typically sent through user agents such as web browsers.

For example, the parameters "response_type", "client_id", "state", and "redirect_uri" are encoded in the URI of the request:

```
GET /authorize?response_type=code&client_id=s6BhdRkqt3&state=xyz
&redirect_uri=https%3A%2F%2Fclient%2Eexample%2Ecom%2Fcb HTTP/1.1
Host: server.example.com
```

While it is easy to implement, the encoding in the URI does not allow application layer security with confidentiality and integrity protection to be used. While TLS is used to offer communication security between the Client and the user-agent as well as the user-agent and the Authorization Server, TLS sessions are terminated in the user-agent. In addition, TLS sessions may be terminated prematurely at some middlebox (such as a load balancer).

As the result, the Authorization Request of [RFC6749] has shortcomings in that:

- (a) the communication through the user agents are not integrity protected and thus the parameters can be tainted (integrity protection failure)
- (b) the source of the communication is not authenticated (source authentication failure)
- (c) the communication through the user agents can be monitored (containment / confidentiality failure).

Due to these inherent weaknesses, several attacks against the protocol, such as Redirection URI rewriting and Mix-up attack [FETT], have been identified.

The use of application layer security mitigates these issues.

The use of application layer security allows requests to be prepared by a third party so that a client application cannot request more permissions than previously agreed. This offers an additional degree of privacy protection.

Furthermore, the request by reference allows the reduction of over-the-wire overhead.

The JWT [RFC7519] encoding has been chosen because of

- (1) its close relationship with JSON, which is used as OAuth's response format
- (2) its developer friendliness due to its textual nature
- (3) its relative compactness compared to XML
- (4) its development status that it is an RFC and so is its associated signing and encryption methods as [RFC7515] and [RFC7516]
- (5) the relative ease of JWS and JWE compared to XML Signature and Encryption.

The parameters "request" and "request_uri" are introduced as additional authorization request parameters for the OAuth 2.0 [RFC6749] flows. The "request" parameter is a JSON Web Token (JWT) [RFC7519] whose JWT Claims Set holds the JSON encoded OAuth 2.0 authorization request parameters. This JWT is integrity protected and source authenticated using JWS.

The JWT [RFC7519] can be passed to the authorization endpoint by reference, in which case the parameter "request_uri" is used instead of the "request".

Using JWT [RFC7519] as the request encoding instead of query parameters has several advantages:

- (a) (integrity protection) The request can be signed so that the integrity of the request can be checked.
- (b) (source authentication) The request can be signed so that the signer can be authenticated.
- (c) (confidentiality protection) The request can be encrypted so that end-to-end confidentiality can be provided even if the TLS connection is terminated at one point or another.

- (d) (collection minimization) The request can be signed by a third party attesting that the authorization request is compliant with a certain policy. For example, a request can be pre-examined by a third party that all the personal data requested is strictly necessary to perform the process that the end-user asked for, and statically signed by that third party. The authorization server then examines the signature and shows the conformance status to the end-user, who would have some assurance as to the legitimacy of the request when authorizing it. In some cases, it may even be desirable to skip the authorization dialogue under such circumstances.

There are a few cases that request by reference is useful such as:

1. When it is desirable to reduce the size of transmitted request. The use of application layer security increases the size of the request, particularly when public key cryptography is used.
2. When the client does not want to do the crypto. The Authorization Server may provide an endpoint to accept the Authorization Request through direct communication with the Client so that the Client is authenticated and the channel is TLS protected.

This capability is in use by OpenID Connect [OpenID.Core].

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 RFC 2119 [RFC2119].

2. Terminology

For the purposes of this specification, the following terms and definitions in addition to what is defined in OAuth 2.0 Framework [RFC6749], JSON Web Signature [RFC7515], and JSON Web Encryption [RFC7519] apply.

2.1. Request Object

JWT [RFC7519] that holds an OAuth 2.0 authorization request as JWT Claims Set

2.2. Request Object URI

Absolute URI from which the Request Object (Section 2.1) can be obtained

3. Symbols and abbreviated terms

The following abbreviations are common to this specification.

JSON Javascript Object Notation

JWT JSON Web Token

JWS JSON Web Signature

JWE JSON Web Encryption

URI Uniform Resource Identifier

URL Uniform Resource Locator

4. Request Object

A Request Object (Section 2.1) is used to provide authorization request parameters for an OAuth 2.0 authorization request. It MUST contain all the OAuth 2.0 [RFC6749] authorization request parameters including extension parameters. The parameters are represented as the JWT claims. Parameter names and string values MUST be included as JSON strings. Since Request Objects are handled across domains and potentially outside of a closed ecosystem, per section 8.1 of [RFC8259], these JSON strings MUST be encoded using UTF-8 [RFC3629]. Numerical values MUST be included as JSON numbers. It MAY include any extension parameters. This JSON [RFC7159] constitutes the JWT Claims Set defined in JWT [RFC7519]. The JWT Claims Set is then signed or signed and encrypted.

To sign, JSON Web Signature (JWS) [RFC7515] is used. The result is a JWS signed JWT [RFC7519]. If signed, the Authorization Request Object SHOULD contain the Claims "iss" (issuer) and "aud" (audience) as members, with their semantics being the same as defined in the JWT [RFC7519] specification. The value of "aud" should be the value of the Authorization Server (AS) "issuer" as defined in RFC8414 [RFC8414].

To encrypt, JWE [RFC7516] is used. When both signature and encryption are being applied, the JWT MUST be signed then encrypted as advised in the section 11.2 of [RFC7519]. The result is a Nested JWT, as defined in [RFC7519].

The following RSA public key, represented in JWK format, can be used to validate the Request Object signature in this and subsequent Request Object examples (with line wraps within values for display purposes only):

```
{
  "kty": "RSA",
  "kid": "k2bdc",
  "n": "y9Lqv4fCp6Ei-u2-ZCKq83YvbFEk6JMs_pSj76eMkddWRuWX2aBKGHAtK1E5P
7_vn__PCKZWePt3vGkB6ePgzaFu08NmKemwE5bQI0e6kIChtt_6KzT5OaaXDF
I6qCLJmk51Cc4VYFagqgevMncYrzaW_50mZ1yGSFIQzLYP8bijAHGVjdEFgZa
ZEN91sn_GdWLaJpHrB3RO1S50E45wxrlg9xMncVb8qDPuXZarvghLL0HzOuYR
adBJVoWZowDNTpKpk2RklZ7QaB07XDv3uR7s_sf2g-bAjSYxYUGsqkNA9b3xV
W53am_UZZ3tZbFTIh557JICWKH1Wj5uzeJXaw",
  "e": "AQAB"
}
```

5. Authorization Request

The client constructs the authorization request URI by adding one of the following parameters but not both to the query component of the authorization endpoint URI using the "application/x-www-form-urlencoded" format:

`request` The Request Object (Section 2.1) that holds authorization request parameters stated in section 4 of OAuth 2.0 [RFC6749].

`request_uri` The absolute URI as defined by RFC3986 [RFC3986] that points to the Request Object (Section 2.1) that holds authorization request parameters stated in section 4 of OAuth 2.0 [RFC6749].

The client directs the resource owner to the constructed URI using an HTTP redirection response, or by other means available to it via the user-agent.

For example, the client directs the end user's user-agent to make the following HTTPS request:

```
GET /authz?request=eyJhbGciOiJIbGciLCJ0eSI6ImF1dG8iLCJ0aWQiOiJk2bdcIiwiaWF0Ijoi2019-06-01T12:00:00Z\"
Host: server.example.com
```

The value for the request parameter is abbreviated for brevity.

The authorization request object MUST be one of the following:

- (a) JWS signed

(b) JWS signed and JWE encrypted

The client MAY send the parameters included in the request object duplicated in the query parameters as well for the backward compatibility etc. However, the authorization server supporting this specification MUST only use the parameters included in the request object.

5.1. Passing a Request Object by Value

The Client sends the Authorization Request as a Request Object to the Authorization Endpoint as the "request" parameter value.

The following is an example of an Authorization Request using the "request" parameter (with line wraps within values for display purposes only):

```
https://server.example.com/authorize?
request=eyJhbGciOiJSUzI1NiIsImtpZCI6ImN5YmRjIn0.ew0KICJpc3MiOiA
icZzCAGRsa3F0MyIsdQogImFlZCI6ICJodHRwczovL3NlcnZlci5leGFtcGxlLmN
vbSIIsdQogInJlc3BvbmlX3R5cGUiOiAiY29kZSBpZF90b2t1biIsdQogImNsaWV
udF9pZCI6ICJzNkJoZlJrcXQzIiwNCiAicmVkaXJlY3RfdXJpIjogImh0dHBzOi8
vY2xpZW50LmV4YW1wbGUub3JnL2NiIiwNCiAic2NvcGUiOiAib3BlbmlkIiwNCiA
ic3RhdGUiOiAiYWYwaWZqc2xka2oiLA0KICJub25jZSI6ICJuLTB0bnl9XekEyTwo
iLA0KICJtYXhfYWdlIjogODY0MDAsdQogImNsYWltcyI6IA0KICB7DQogICAidXN
lcmluZm8iOiANCiAgICB7DQogICAgICJnaXZlbnl9uYW11IjogeyJlc3NlbnRpYWw
iOiB0cnVlfSwNCiAgICAgIm5pY2tuYW11IjogbnVsbCwNCiAgICAgImVtYWlsIjo
geyJlc3NlbnRpYWwiOiB0cnVlfSwNCiAgICAgImVtYWlsX3Zlcm1maWVkiJogeyJ
lc3NlbnRpYWwiOiB0cnVlfSwNCiAgICAgInBpY3RlcmUiOiBudWxsDQogICAgfSw
NCiAgICJpZF90b2t1biI6IA0KICAgIHsNCiAgICAgImdlbmRlciI6IG51bGwsDQo
gICAgICJiaXJ0aGRhdGUiOiB7ImVzc2VudG1hbCI6IHRYdWV9LA0KICAgICAgIWN
yIjogeyJ2YWxlZXMiOiBbInVyb29tbW9uOmlhcDpzaWx2ZXIiXX0
NCiAgICB9DQogIH0NCn0.nwwnNsk1-ZkbnmvsF6zTHm8CHERFMGQPhos-EJcaH4H
h-sMgk8ePrGhw_trPYs8KQxsn6R9Emo_wHwajyFKzuMXZFSZ3p6Mb8dkxtVyjoy2
GIzvuJT_u7PkY2t8QU9hjBcHs68PkgjDVTTrGluRTx0GxFbuPbj96tVuJ1lpTnmFC
UR6IEOXKYr7iGOCRB3btfJhM0_AKQUfqKnRlrRsc8Kol-cSLWoYE915QqholImz
jT_cMnNIznW9E7CDyWXTsO70xnB4SkG6pXfLSjLLlxmPGiyon_-Te111V8uE83I1
zCYIb_NMXvtTIVcljpspnTSD7xMbpL-2QgwUsAlMGzw
```

5.2. Passing a Request Object by Reference

The "request_uri" Authorization Request parameter enables OAuth authorization requests to be passed by reference, rather than by value. This parameter is used identically to the "request" parameter, other than that the Request Object value is retrieved from the resource identified by the specified URI rather than passed by value.

5.2.1. URI Referencing the Request Object

The Client stores the Request Object resource either locally or remotely at a URI the Authorization Server can access. Such facility may be provided by the authorization server or a third party. For example, the authorization server may provide a URL to which the client POSTs the request object and obtains the Request URI. This URI is the Request Object URI, "request_uri".

It is possible for the Request Object to include values that are to be revealed only to the Authorization Server. As such, the "request_uri" MUST have appropriate entropy for its lifetime. For the guidance, refer to 5.1.4.2.2 of [RFC6819]. It is RECOMMENDED that it be removed after a reasonable timeout unless access control measures are taken.

The following is an example of a Request Object URI value (with line wraps within values for display purposes only):

```
https://tfp.example.org/request.jwt#
GkurKxf5T0Y-mnPFCHqWOMiZi4VS138cQO_V7PZHAdM
```

5.2.2. Request using the "request_uri" Request Parameter

The Client sends the Authorization Request to the Authorization Endpoint.

The following is an example of an Authorization Request using the "request_uri" parameter (with line wraps within values for display purposes only):

```
https://server.example.com/authorize?
response_type=code%20id_token
&client_id=s6BhdRkqt3
&request_uri=https%3A%2F%2Ftfp.example.org%2Frequest.jwt
%23GkurKxf5T0Y-mnPFCHqWOMiZi4VS138cQO_V7PZHAdM
&state=af0ifjsldkj
```

5.2.3. Authorization Server Fetches Request Object

Upon receipt of the Request, the Authorization Server MUST send an HTTP "GET" request to the "request_uri" to retrieve the referenced Request Object, unless it is stored in a way so that it can retrieve it through other mechanism securely, and parse it to recreate the Authorization Request parameters.

6.2. JWS Signed Request Object

The Authorization Server MUST perform the signature validation of the JSON Web Signature [RFC7515] signed request object. For this, the "alg" Header Parameter in its JOSE Header MUST match the value of the pre-registered algorithm. The signature MUST be validated against the appropriate key for that "client_id" and algorithm.

If signature validation fails, the Authorization Server MUST return an "invalid_request_object" error.

6.3. Request Parameter Assembly and Validation

The Authorization Server MUST extract the set of Authorization Request parameters from the Request Object value. The Authorization Server MUST only use the parameters in the Request Object even if the same parameter is provided in the query parameter. The Authorization Server then validates the request as specified in OAuth 2.0 [RFC6749].

If the validation fails, then the Authorization Server MUST return an error as specified in OAuth 2.0 [RFC6749].

7. Authorization Server Response

Authorization Server Response is created and sent to the client as in Section 4 of OAuth 2.0 [RFC6749] .

In addition, this document uses these additional error values:

`invalid_request_uri` The "request_uri" in the Authorization Request returns an error or contains invalid data.

`invalid_request_object` The request parameter contains an invalid Request Object.

`request_not_supported` The Authorization Server does not support the use of the "request" parameter.

`request_uri_not_supported` The Authorization Server does not support the use of the "request_uri" parameter.

8. TLS Requirements

Client implementations supporting the Request Object URI method MUST support TLS following Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) [BCP195].

To protect against information disclosure and tampering, confidentiality protection MUST be applied using TLS with a cipher suite that provides confidentiality and integrity protection.

HTTP clients MUST also verify the TLS server certificate, using subjectAltName dNSName identities as described in [RFC6125], to avoid man-in-the-middle attacks. The rules and guidelines defined in [RFC6125] apply here, with the following considerations:

- o Support for DNS-ID identifier type (that is, the dNSName identity in the subjectAltName extension) is REQUIRED. Certification authorities which issue server certificates MUST support the DNS-ID identifier type, and the DNS-ID identifier type MUST be present in server certificates.
- o DNS names in server certificates MAY contain the wildcard character "*".
- o Clients MUST NOT use CN-ID identifiers; a CN field may be present in the server certificate's subject name, but MUST NOT be used for authentication within the rules described in [BCP195].
- o SRV-ID and URI-ID as described in Section 6.5 of [RFC6125] MUST NOT be used for comparison.

9. IANA Considerations

This specification requests no actions by IANA.

10. Security Considerations

In addition to the all the security considerations discussed in OAuth 2.0 [RFC6819], the security considerations in [RFC7515], [RFC7516], and [RFC7518] needs to be considered. Also, there are several academic papers such as [BASIN] that provide useful insight into the security properties of protocols like OAuth.

In consideration of the above, this document advises taking the following security considerations into account.

10.1. Choice of Algorithms

When sending the authorization request object through "request" parameter, it MUST either be signed using JWS [RFC7515] or encrypted using JWE [RFC7516] with then considered appropriate algorithm.

10.2. Request Source Authentication

The source of the Authorization Request MUST always be verified. There are several ways to do it in this specification.

- (a) Verifying the JWS Signature of the Request Object.
- (b) Verifying that the symmetric key for the JWE encryption is the correct one if the JWE is using symmetric encryption.
- (c) Verifying the TLS Server Identity of the Request Object URI. In this case, the Authorization Server MUST know out-of-band that the Client uses Request Object URI and only the Client is covered by the TLS certificate. In general, it is not a reliable method.
- (d) Authorization Server is providing an endpoint that provides a Request Object URI in exchange for a Request Object. In this case, the Authorization Server MUST perform Client Authentication to accept the Request Object and bind the Client Identifier to the Request Object URI it is providing. Since Request Object URI can be replayed, the lifetime of the Request Object URI MUST be short and preferably one-time use. The entropy of the Request Object URI MUST be sufficiently large. The adequate shortness of the validity and the entropy of the Request Object URI depends on the risk calculation based on the value of the resource being protected. A general guidance for the validity time would be less than a minute and the Request Object URI is to include a cryptographic random value of 128bit or more at the time of the writing of this specification.
- (e) A third party, such as a Trust Framework Provider, provides an endpoint that provides a Request Object URI in exchange for a Request Object. The same requirements as (b) above apply. In addition, the Authorization Server MUST know out-of-band that the Client utilizes the Trust Framework Operator.

10.3. Explicit Endpoints

Although this specification does not require them, research such as [BASIN] points out that it is a good practice to explicitly state the intended interaction endpoints and the message position in the sequence in a tamper evident manner so that the intent of the initiator is unambiguous. The endpoints that come into question in this specification are :

- (a) Protected Resources ("protected_resources")

- (b) Authorization Endpoint ("authorization_endpoint")
- (c) Redirection URI ("redirect_uri")
- (d) Token Endpoint ("token_endpoint")

Further, if dynamic discovery is used, then the discovery related endpoints also come into question.

In [RFC6749], while Redirection URI is included, others are not included in the Authorization Request. As the result, the same applies to Authorization Request Object.

The lack of the link among those endpoints are cited as the cause of Cross-Phase Attacks introduced in [FETT]. An extension specification should be created as a measure to address the risk.

10.4. Risks Associated with request_uri

The introduction of "request_uri" introduces several attack possibilities.

10.4.1. DDoS Attack on the Authorization Server

A set of malicious client can launch a DoS attack to the authorization server by pointing the "request_uri" to a uri that returns extremely large content or extremely slow to respond. Under such an attack, the server may use up its resource and start failing.

Similarly, a malicious client can specify the "request_uri" value that itself points to an authorization request URI that uses "request_uri" to cause the recursive lookup.

To prevent such attack to succeed, the server should (a) check that the value of "request_uri" parameter does not point to an unexpected location, (b) check the content type of the response is "application/json" (c) implement a time-out for obtaining the content of "request_uri", and (d) do not perform recursive GET on the "request_uri".

10.4.2. Request URI Rewrite

The value of "request_uri" is not signed thus it can be tampered by Man-in-the-browser attacker. Several attack possibilities rise because of this, e.g., (a) attacker may create another file that the rewritten URI points to making it possible to request extra scope (b) attacker launches a DoS attack to a victim site by setting the value of "request_uri" to be that of the victim.

To prevent such attack to succeed, the server should (a) check that the value of "request_uri" parameter does not point to an unexpected location, (b) check the content type of the response is "application/jwt" (c) implement a time-out for obtaining the content of "request_uri".

11. TLS security considerations

Current security considerations can be found in Recommendations for Secure Use of TLS and DTLS [BCP195]. This supersedes the TLS version recommendations in OAuth 2.0 [RFC6749].

12. Privacy Considerations

When the Client is being granted access to a protected resource containing personal data, both the Client and the Authorization Server need to adhere to Privacy Principles. RFC 6973 Privacy Considerations for Internet Protocols [RFC6973] gives excellent guidance on the enhancement of protocol design and implementation. The provision listed in it should be followed.

Most of the provision would apply to The OAuth 2.0 Authorization Framework [RFC6749] and The OAuth 2.0 Authorization Framework: Bearer Token Usage [RFC6750] and are not specific to this specification. In what follows, only the specific provisions to this specification are noted.

12.1. Collection limitation

When the Client is being granted access to a protected resource containing personal data, the Client SHOULD limit the collection of personal data to that which is within the bounds of applicable law and strictly necessary for the specified purpose(s).

It is often hard for the user to find out if the personal data asked for is strictly necessary. A Trust Framework Provider can help the user by examining the Client request and comparing to the proposed processing by the Client and certifying the request. After the certification, the Client, when making an Authorization Request, can submit Authorization Request to the Trust Framework Provider to obtain the Request Object URI.

Upon receiving such Request Object URI in the Authorization Request, the Authorization Server first verifies that the authority portion of the Request Object URI is a legitimate one for the Trust Framework Provider. Then, the Authorization Server issues HTTP GET request to the Request Object URI. Upon connecting, the Authorization Server MUST verify the server identity represented in the TLS certificate is

legitimate for the Request Object URI. Then, the Authorization Server can obtain the Request Object, which includes the "client_id" representing the Client.

The Consent screen MUST indicate the Client and SHOULD indicate that the request has been vetted by the Trust Framework Operator for the adherence to the Collection Limitation principle.

12.2. Disclosure Limitation

12.2.1. Request Disclosure

This specification allows extension parameters. These may include potentially sensitive information. Since URI query parameter may leak through various means but most notably through referrer and browser history, if the authorization request contains a potentially sensitive parameter, the Client SHOULD JWE [RFC7516] encrypt the request object.

Where Request Object URI method is being used, if the request object contains personally identifiable or sensitive information, the "request_uri" SHOULD be used only once, have a short validity period, and MUST have large enough entropy deemed necessary with applicable security policy unless the Request Object itself is JWE [RFC7516] Encrypted. The adequate shortness of the validity and the entropy of the Request Object URI depends on the risk calculation based on the value of the resource being protected. A general guidance for the validity time would be less than a minute and the Request Object URI is to include a cryptographic random value of 128bit or more at the time of the writing of this specification.

12.2.2. Tracking using Request Object URI

Even if the protected resource does not include a personally identifiable information, it is sometimes possible to identify the user through the Request Object URI if persistent per-user Request Object URI is used. A third party may observe it through browser history etc. and start correlating the user's activity using it. In a way, it is a data disclosure as well and should be avoided.

Therefore, per-user Request Object URI should be avoided.

13. Acknowledgements

The following people contributed to the creation of this document in the OAuth WG. (Affiliations at the time of the contribution are used.)

Sergey Beryozkin, Brian Campbell (Ping Identity), Vladimir Dzhuvinov (Connect2id), Michael B. Jones (Microsoft), Torsten Lodderstedt (YES) Jim Manico, Axel Nenker (Deutsche Telecom), Hannes Tschofenig (ARM), Ben Campbell, Kathleen Moriarty (as AD), and Steve Kent (as SECDIR).

The following people contributed to creating this document through the OpenID Connect Core 1.0 [OpenID.Core].

Brian Campbell (Ping Identity), George Fletcher (AOL), Ryo Itou (Mixi), Edmund Jay (Illumila), Michael B. Jones (Microsoft), Breno de Medeiros (Google), Hideki Nara (TACT), Justin Richer (MITRE).

In addition, the following people contributed to this and previous versions through the OAuth Working Group.

Dirk Balfanz (Google), James H. Manger (Telstra), John Panzer (Google), David Recordon (Facebook), Marius Scurtescu (Google), Luke Shepard (Facebook).

14. Revision History

Note to the RFC Editor: Please remove this section from the final RFC.

-19

- o AD cooments
- o Section 5.2.1. s/Requiest URI/Request URI/
- o Section 8 s/[BCP195] ./[BCP195]./
- o Section 10.3. s/sited/cited/
- o Section 11. Typo. s/Curent/Current/

-17

- o #78 Typos in content-type

-16

- o Treated remaining Ben Campbell comments.

-15

- o Removed further duplication

-14

- o #71 Reiterate dynamic params are included.
- o #70 Made clear that AS must return error.
- o #69 Inconsistency of the need to sign.
- o Fixed Mime-type.
- o #67 Inconsistence in requiring HTTPS in request uri.
- o #66 Dropped ISO 29100 reference.
- o #25 Removed Encrypt only option.
- o #59 Same with #25.

-13

- o add TLS Security Consideration section
- o replace RFC7525 reference with BCP195
- o moved front tag in FETT reference to fix XML structure
- o changes reference from SoK to FETT

-12

- o fixes #62 - Alexey Melnikov Discuss
- o fixes #48 - OPSDIR Review : General - delete semicolons after list items
- o fixes #58 - DP Comments for the Last Call
- o fixes #57 - GENART - Remove "non-normative ... " from examples.
- o fixes #45 - OPSDIR Review : Introduction - are attacks discovered or already opened
- o fixes #49 - OPSDIR Review : Introduction - Inconsistent colons after initial sentence of list items.
- o fixes #53 - OPSDIR Review : 6.2 JWS Signed Request Object - Clarify JOSE Header

- o fixes #42 - OPSDIR Review : Introduction - readability of 'and' is confusing
- o fixes #50 - OPSDIR Review : Section 4 Request Object - Clarify 'signed, encrypted, or signed and encrypted'
- o fixes #39 - OPSDIR Review : Abstract - Explain/Clarify JWS and JWE
- o fixed #50 - OPSDIR Review : Section 4 Request Object - Clarify 'signed, encrypted, or signed and encrypted'
- o fixes #43 - OPSDIR Review : Introduction - 'properties' sounds awkward and are not exactly 'properties'
- o fixes #56 - OPSDIR Review : 12 Acknowledgements - 'contribution is' => 'contribution are'
- o fixes #55 - OPSDIR Review : 11.2.2 Privacy Considerations - ' It is in a way' => 'In a way, it is'
- o fixes #54 - OPSDIR Review : 11 Privacy Considerations - 'and not specific' => 'and are not specific'
- o fixes #51 - OPSDIR Review : Section 4 Request Object - 'It is fine' => 'It is recommended'
- o fixes #47 - OPSDIR Review : Introduction - 'over- the- wire' => 'over-the-wire'
- o fixes #46 - OPSDIR Review : Introduction - 'It allows' => 'The use of application security' for
- o fixes #44 - OPSDIR Review : Introduction - 'has' => 'have'
- o fixes #41 - OPSDIR Review : Introduction - missing 'is' before 'typically sent'
- o fixes #38 - OPSDIR Review : Section 11 - Delete 'freely accessible' regarding ISO 29100

-11

- o s/bing/being/
- o Added history for -10

-10

- o #20: KM1 -- some wording that is awkward in the TLS section.
- o #21: KM2 - the additional attacks against OAuth 2.0 should also have a pointer
- o #22: KM3 -- Nit: in the first line of 10.4:
- o #23: KM4 -- Mention RFC6973 in Section 11 in addition to ISO 29100
- o #24: SECDIR review: Section 4 -- Confusing requirements for sign+encrypt
- o #25: SECDIR review: Section 6 -- authentication and integrity need not be provided if the requestor encrypts the token?
- o #26: SECDIR Review: Section 10 -- why no reference for JWS algorithms?
- o #27: SECDIR Review: Section 10.2 - how to do the agreement between client and server "a priori"?
- o #28: SECDIR Review: Section 10.3 - Indication on "large entropy" and "short lifetime" should be indicated
- o #29: SECDIR Review: Section 10.3 - Typo
- o #30: SECDIR Review: Section 10.4 - typos and missing articles
- o #31: SECDIR Review: Section 10.4 - Clearer statement on the lack of endpoint identifiers needed
- o #32: SECDIR Review: Section 11 - ISO29100 needs to be moved to normative reference
- o #33: SECDIR Review: Section 11 - Better English and Entropy language needed
- o #34: Section 4: Typo
- o #35: More Acknowledgment
- o #36: DP - More precise qualification on Encryption needed.

-09

- o Minor Editorial Nits.
- o Section 10.4 added.

- o Explicit reference to Security consideration (10.2) added in section 5 and section 5.2.
- o , (add yourself) removed from the acknowledgment.

-08

- o Applied changes proposed by Hannes on 2016-06-29 on IETF OAuth list recorded as <https://bitbucket.org/Nat/oauth-jwsreq/issues/12/>.
- o TLS requirements added.
- o Security Consideration reinforced.
- o Privacy Consideration added.
- o Introduction improved.

-07

- o Changed the abbrev to OAuth JAR from oauth-jar.
- o Clarified sig and enc methods.
- o Better English.
- o Removed claims from one of the example.
- o Re-worded the URI construction.
- o Changed the example to use request instead of request_uri.
- o Clarified that Request Object parameters take precedence regardless of request or request_uri parameters were used.
- o Generalized the language in 4.2.1 to convey the intent more clearly.
- o Changed "Server" to "Authorization Server" as a clarification.
- o Stopped talking about request_object_signing_alg.
- o IANA considerations now reflect the current status.
- o Added Brian Campbell to the contributors list. Made the lists alphabetic order based on the last names. Clarified that the affiliation is at the time of the contribution.

- o Added "older versions of " to the reference to IE uri length limitations.
- o Stopped talking about signed or unsigned JWS etc.
- o 1.Introduction improved.

-06

- o Added explanation on the 512 chars URL restriction.
- o Updated Acknowledgements.

-05

- o More alignment with OpenID Connect.

-04

- o Fixed typos in examples. (request_url -> request_uri, cliend_id -> client_id)
- o Aligned the error messages with the OAuth IANA registry.
- o Added another rationale for having request object.

-03

- o Fixed the non-normative description about the advantage of static signature.
- o Changed the requirement for the parameter values in the request itself and the request object from 'MUST MATCH' to 'Req Obj takes precedence.

-02

- o Now that they are RFCs, replaced JWS, JWE, etc. with RFC numbers.

-01

- o Copy Edits.

15. References

15.1. Normative References

- [BCP195] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, May 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<https://www.rfc-editor.org/info/rfc6750>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<https://www.rfc-editor.org/info/rfc7159>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.

- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<https://www.rfc-editor.org/info/rfc7516>>.
- [RFC7518] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <<https://www.rfc-editor.org/info/rfc7518>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC8141] Saint-Andre, P. and J. Klensin, "Uniform Resource Names (URNs)", RFC 8141, DOI 10.17487/RFC8141, April 2017, <<https://www.rfc-editor.org/info/rfc8141>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC8414] Jones, M., Sakimura, N., and J. Bradley, "OAuth 2.0 Authorization Server Metadata", RFC 8414, DOI 10.17487/RFC8414, June 2018, <<https://www.rfc-editor.org/info/rfc8414>>.

15.2. Informative References

- [BASIN] Basin, D., Cremers, C., and S. Meier, "Provably Repairing the ISO/IEC 9798 Standard for Entity Authentication", Journal of Computer Security - Security and Trust Principles Volume 21 Issue 6, Pages 817-846, November 2013, <<https://www.cs.ox.ac.uk/people/cas.cremers/downloads/papers/BCM2012-iso9798.pdf>>.
- [FETT] Fett, D., Kusters, R., and G. Schmitz, "A Comprehensive Formal Security Analysis of OAuth 2.0", CCS '16 Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security Pages 1204-1215, October 2016, <<https://infsec.uni-trier.de/people/publications/paper/FettKuestersSchmitz-CCS-2016.pdf>>.

[OpenID.Core]

Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", OpenID Foundation Standards, February 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.

[RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<https://www.rfc-editor.org/info/rfc6819>>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

Authors' Addresses

Nat Sakimura
Nomura Research Institute
Otemachi Financial City Grand Cube, 1-9-2 Otemachi
Chiyoda-ku, Tokyo 100-0004
Japan

Phone: +81-3-5533-2111
Email: n-sakimura@nri.co.jp
URI: <http://nat.sakimura.org/>

John Bradley
Yubico
Casilla 177, Sucursal Talagante
Talagante, RM
Chile

Phone: +1.202.630.5272
Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2017

M. Jones
Microsoft
J. Bradley
Ping Identity
N. Sakimura
NRI
July 6, 2016

OAuth 2.0 Mix-Up Mitigation
draft-ietf-oauth-mix-up-mitigation-01

Abstract

This specification defines an extension to The OAuth 2.0 Authorization Framework that enables the authorization server to dynamically provide the client using it with additional information about the current protocol interaction that can be validated by the client and that enables the client to dynamically provide the authorization server with additional information about the current protocol interaction that can be validated by the authorization server. This additional information can be used by the client and the authorization server to prevent classes of attacks in which the client might otherwise be tricked into using inconsistent sets of metadata from multiple authorization servers, including potentially using a token endpoint that does not belong to the same authorization server as the authorization endpoint used. Recent research publications refer to these as "IdP Mix-Up" and "Malicious Endpoint" attacks.

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 7, 2017.

Copyright Notice

Copyright (c) 2016 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

1.	Introduction	3
1.1.	Requirements Notation and Conventions	4
1.2.	Terminology	4
2.	The OAuth Issuer Identifier	4
3.	Mitigation Data Returned in Authorization Response	5
3.1.	Mitigation Data Returned in Authorization Response Parameters	5
3.1.1.	Example Authorization Response using Response Parameters	5
3.2.	Mitigation Data Returned in JWT	6
3.2.1.	Example Authorization Response using JWT	6
4.	Validating the Authorization Response	7
5.	Mitigation Data Sent to the Token Endpoint	8
5.1.	Example Token Request	8
6.	Validating the Token Request	9
7.	Security Considerations	9
7.1.	IdP Mix-Up and Malicious Endpoint Attacks	9
7.2.	Duplicate Information Attacks	9
7.3.	Cut-and-Paste Attacks	10
8.	IANA Considerations	11
8.1.	OAuth Parameters Registration	11
8.1.1.	Registry Contents	11
9.	References	11
9.1.	Normative References	11
9.2.	Informative References	12
Appendix A.	Implementation Notes	13
Appendix B.	Acknowledgements	13
Appendix C.	Open Issues	14
Appendix D.	Document History	14
Authors' Addresses	14

1. Introduction

OAuth 2.0 [RFC6749] clients use multiple authorization server endpoints when using some OAuth response types. For instance, when using the "code" response type, the client uses both the authorization endpoint and the token endpoint. It is important that endpoints belonging to the same authorization server always be used together. Otherwise, information produced by one authorization server could mistakenly be sent by the client to different authorization server, resulting in some of the attacks described in Section 7. Recent research publications refer to these specific attacks as "IdP Mix-Up" [arXiv.1601.01229v2] and "Malicious Endpoint" [arXiv.1508.04324v2] attacks.

The client obviously cannot be confused into using endpoints from multiple authorization servers in an authorization flow if the client is configured to use only a single authorization server. However, the client can potentially be tricked into mixing endpoints if it is configured to use more than one authorization server, whether the configuration is dynamic or static. The client may be confused if it has no way to determine whether the set of endpoints belongs to the same authorization server. Or, a client may be confused simply because it is receiving authorization responses from more than one authorization server at the same redirection endpoint and the client is insufficiently able to determine that the response received is associated with the correct authorization server.

This specification enables the authorization server to dynamically provide the client using it with additional information about the current protocol interaction that can be validated by the client and that enables the client to dynamically provide the authorization server with additional information about the current protocol interaction that can be validated by the authorization server. This enables them to abort interactions in which endpoints from multiple authorization servers would otherwise be used.

The mitigation data provided by the authorization server to the client is an issuer identifier, which is used to identify the authorization server, and a client ID, which is used to verify that the response is from the correct authorization server and is intended for this client. The issuer identifier is defined in Section 2 of [OAuth.Discovery]. If supported by the authorization server, the issuer identifier can also be used to obtain a consistent set of metadata describing the authorization server configuration, as also described in [OAuth.Discovery].

This mitigation data is returned to the client in the authorization response. The syntax for returning the mitigation data from the

authorization server is dependent upon the OAuth response type being used. The syntax used with the existing response types registered in the IANA "OAuth Authorization Endpoint Response Types" registry [IANA.OAuth.Parameters] as of the time of this writing is defined by this specification. Two of these response types are defined by RFC 6749 [RFC6749]; the rest are defined by [OAuth.Responses].

The mitigation data provided by the client to the authorization server is the existing "state" value defined by RFC 6749 [RFC6749], but adding also sending it from the client to the token endpoint. This is used by the authorization server to verify that the authorization code and state both belong to the same protocol interaction.

1.1. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Terminology

This specification uses the terms "Access Token", "Authorization Code", "Authorization Endpoint", "Authorization Grant", "Authorization Server", "Client", "Client Authentication", "Client Identifier", "Client Secret", "Grant Type", "Protected Resource", "Redirection URI", "Refresh Token", "Resource Owner", "Resource Server", "Response Type", and "Token Endpoint" defined by OAuth 2.0 [RFC6749], the terms "Claim Name", "Claim Value", and "JSON Web Token (JWT)" defined by JSON Web Token (JWT) [JWT].

2. The OAuth Issuer Identifier

The OAuth issuer identifier serves as a concrete identifier for the authorization server. As defined in [OAuth.Discovery], the issuer identifier is a URL that uses the "https" scheme and has no query or fragment components. Also as specified there, this is the location where ".well-known" RFC 5785 [RFC5785] resources containing information about the authorization server are published. In particular, when discovery is supported, the authorization server's metadata is retrieved as a JSON document [RFC7159] from a path derived from this URL. This metadata document contains a consistent set of metadata describing the authorization server configuration.

Implementations supporting this specification MAY also support discovery or they MAY simply use the issuer identifier as a concrete identifier for the authorization server. This specification does not

rely upon the authorization server publishing or the client retrieving a discovery metadata document.

3. Mitigation Data Returned in Authorization Response

Mitigating the attacks relies on the authorization server returning additional data about the interaction and the client checking that data. The mitigation data returned is the client ID and the issuer identifier. The syntax for returning the mitigation data from the authorization server is dependent upon the OAuth response type being used.

3.1. Mitigation Data Returned in Authorization Response Parameters

Some OAuth response types do not already return the issuer identifier and client ID in the authorization response. When this is the case, the mitigation data is returned as additional OAuth response parameters.

These new response parameters are defined for this purpose:

client_id

Client that this response is intended for. It MUST contain the OAuth 2.0 client ID of the client as its value.

iss

Issuer identifier for the authorization server issuing the response. The "iss" value is a case-sensitive URL using the "https" scheme that contains scheme, host, and optionally, port number and path components and no query or fragment components.

As of the time of this writing, these are the existing response types that are registered in the IANA "OAuth Authorization Endpoint Response Types" registry [IANA.OAuth.Parameters] that do not already return the issuer identifier and client ID in the authorization response: "code", "code token", "none", and "token". Therefore, the client ID and issuer are returned using the new authorization response parameters when using these response types. To avoid duplication, as discussed in Section 7.2, it is NOT RECOMMENDED to also return them in this manner when the response type already returns these values in the authorization response.

3.1.1. Example Authorization Response using Response Parameters

The following example authorization response is to a request that used the "code" response type. It uses the "iss" and "client_id" response parameters to return the mitigation information to the client.

The example successful authorization response follows (with line breaks within lines for display purposes only):

```
HTTP/1.1 302 Found
Location: https://client.example.org/cb?
  code=Qcb0Orvlzh30vLlMPRsbm-diHiMwcLyZvnlarpZv-Jxf_11jnpEX3Tgfvk
  &state=nrsz6AnHzPSVBYRVTXV6ZTXQeg_eih7hdpewHNXmZ8
  &iss=https://server.example.com
  &client_id=5d9e8a36-569d-4c40-8d6b-6e279ac1c5f1
```

3.2. Mitigation Data Returned in JWT

As of the time of this writing, these are the existing response types that are registered in the IANA "OAuth Authorization Endpoint Response Types" registry [IANA.OAuth.Parameters] that already return the issuer identifier and client ID in the authorization response: "code id_token", "code id_token token", "id_token", and "id_token token". All of these return these values as the "iss" (issuer) claim value and as an "aud" (audience) claim value in a signed ID Token, which is a JSON Web Token [JWT], as specified in "OpenID Connect Core 1.0" [OpenID.Core]. When using these response types, the client MUST use the client ID and issuer values returned in the ID Token for validating the mitigation data.

3.2.1. Example Authorization Response using JWT

The following example authorization response is to a request that used the "id_token token" response type. It uses the "iss" and "aud" claims in the ID Token to return the mitigation information to the client.

The example successful authorization response follows (with line breaks within lines for display purposes only):

```
HTTP/1.1 302 Found
Location: https://client.example.org/cb#
  access_token=jHkWEduXMU1BwAsC4vtUsZwnNvTIXe10z9K3vx5KF0Y
  &token_type=Bearer
  &id_token=eyJraWQiaWZlZG53IiwiaXNjaWkiOiJ1MyNTYifQ.
ewogImZlcyI6ICJodHRwczovL3NlcnZlci5leGFtcGxlLmNvbSIsCiAic3ViIjog
IjI0ODI4OTc2MTAwMSIsCiAiYXVkiJogInM2QmhhkUmtxdDMiLAogIm5vbmNlIjog
Im4tMFM2Xld6QTJNaIIsCiAiZXhwIjogMTMxMTI4MTk3MCwKICJpYXQiOiAxMzEx
Mjg0OTcwLAogImF0X2hhc2giOiAiNzdRbVVQdGpQZnpXdeEYyQW5wSzlsUSIKfQ.
kdqTmftlaXg5WBYYBr1wkhkqCGZPc0k8vTiV5g2jj67jQ7XkrDamYx2bOkZLdZrp
MPIzkdYB1nZI_G8vQGQuamRhJcEIt2lkb1GPZ-yhEhdkAiZIZLu38rChalDS2Mh0
glE_rke5XXRhmqqoEFFdZiFdnO3p61-7y5lco840EAZvARSINQaOWIzvioRfs4zw
IFOaT33Vpxfqr8HDyh3lzo9eBW2dSQuCa071z0ENWChWoPliK1JCo_Bk9eDg2uwo
2ZwhsvHzj6TMQ01YOTzufSlSmXIKfj1Os3nftQeR697_hA-nMZyAdL8_NRfaC37
XnAbW8WB9wCfECp7cuNuOg
  &state=af0ifjsldkj
```

Decoding the ID Token in the response will yield the following claims, which includes the mitigation information in the "iss" and "aud" claims:

```
{
  "iss": "https://server.example.com",
  "sub": "248289761001",
  "aud": "s6BhdRkqt3",
  "nonce": "n-0S6_WzA2Mj",
  "exp": 1311281970,
  "iat": 1311280970,
  "at_hash": "77QmUPtjPzfWtF2AnpK9RQ"
}
```

4. Validating the Authorization Response

Upon receiving the mitigation data in an authorization response, the client MUST validate that the response was intended for it and that the authorization server metadata that it obtained at client registration time is consistent with the authorization server metadata contained in the metadata referenced by the issuer identifier.

The client MUST validate the authorization server configuration as follows:

1. Compare the issuer identifier for the authorization server that the client received when it registered at the authorization

server that it made the request to with the issuer value returned in the "iss" response parameter or the "iss" claim in the ID Token, depending upon the response type being used. If they do not exactly match, the client MUST NOT proceed with the authorization.

2. Verify that the response is intended for this client by confirming that the client's client identifier for the authorization server the request was made to matches the value of the "client_id" response parameter or that the client's client identifier is an audience value of the ID Token, depending upon the response type being used. If not, the client MUST NOT proceed with the authorization.

5. Mitigation Data Sent to the Token Endpoint

Mitigating the attacks also relies on the client sending additional data about the interaction to the token endpoint, for response types that use it, and the authorization server checking that data. The mitigation data sent is the same state value that is sent in the authorization request and returned in the authorization response. This specification defines the new "state" token request parameter for passing this additional information.

As of the time of this writing, these are the existing response types that are registered in the IANA "OAuth Authorization Endpoint Response Types" registry [IANA.OAuth.Parameters] that use the token endpoint: "code", "code id_token", "code id_token token", and "code token". The state value is to be sent in the "state" token request parameter when using these response types, and any new response types registered that use the token endpoint.

5.1. Example Token Request

The following example token request is part of a protocol interaction that used the "code" response type. It uses the "state" request parameter to send mitigation information to the authorization server.

The example of token request follows (with line breaks within lines for display purposes only):

```
POST /token HTTP/1.1
Host: server.example.com
Content-Type: application/x-www-form-urlencoded
Authorization: Basic czZCaGRSa3F0MzpnWDFmQmF0M2JW

grant_type=authorization_code
&code=Splx10BeZQQYbYS6WxSbIA
&redirect_uri=https%3A%2F%2Fclient.example.org%2Fcb
&state=ZSGXNBavNc-B3kU3DeJnZoWWOzYxsbvj7jp-S0x_z8U
```

6. Validating the Token Request

When the authorization server receives a token request at the token endpoint that contains a value in the "state" parameter, it MUST validate that the state value received exactly matches the state value previously received in the corresponding authorization request. If the recorded state value and the state value received do not exactly match, the authorization server MUST NOT proceed with the authorization.

7. Security Considerations

7.1. IdP Mix-Up and Malicious Endpoint Attacks

The attacks mitigated by this extension are described in detail in "A Comprehensive Formal Security Analysis of OAuth 2.0" [arXiv.1601.01229v2] and "On the security of modern Single Sign-On Protocols: Second-Order Vulnerabilities in OpenID Connect" [arXiv.1508.04324v2]. To mitigate these attacks, clients configured to use more than one authorization server should use authorization servers that return issuer and client ID information and should validate that a consistent set of authorization server endpoints are being used when using response types that utilize multiple endpoints.

When registering, clients SHOULD NOT allow multiple authorization servers to return the same issuer value, and MUST NOT allow multiple authorization servers to return the same issuer and client ID value pair.

7.2. Duplicate Information Attacks

If a protocol is defined to return the same information in multiple locations, this can create an additional attack surface. Knowing that the information is supposed to be the same, recipients will often be lazy and use the information from only one of the locations,

not validating that all the supposedly duplicate instances are the same. This can enable attackers to create illegal protocol messages that have different values in the multiple locations and those illegal messages will not be detected or rejected by these lazy recipients.

For this reason, if an OAuth profile is being used that returns the mitigation information defined by this specification in one location, it SHOULD NOT also be returned in another. In particular, if a JWT containing the client ID and issuer values is being returned in the authorization response, they SHOULD NOT also be returned as individual authorization response parameters.

7.3. Cut-and-Paste Attacks

OAuth authorization responses are sent as redirects to redirection URIs, with the response parameters typically passed as URI query parameters or fragment values. A "cut-and-paste" attack is performed by the attacker creating what appears to be a legitimate authorization response, but that substitutes some of the response parameter values with values of the attacker's choosing. Sometimes this is done by copying or "cutting" some values out of a legitimate response and replacing or "pasting" some of these values into a different response, the original version of which may have also been legitimate, creating a combination of response values that are not legitimate and that may cause behaviors sought by the attacker. The Code Substitution threat described in Section 4.4.1.13 of [RFC6819] is one example of the use of a cut-and-paste attack.

A concern with returning the mitigation information as new individual authorization response parameters whose values are not cryptographically bound together is that cut-and-paste attacks against their values will not be detected. A security analysis has not been done of the effects of the new attacks that the use of cut-and-paste against these new values will enable.

To prevent replay of the state in another browser instance by an attacker, the state value MUST be tied to the browser instance in a way that cannot be forged by an attacker. Section 4 of [I-D.bradley-oauth-jwt-encoded-state] provides several examples of how a client can accomplish this.

In the replay attack, the attacker can set cookies in the browser. Using an unsigned cookie to bind state to the browser is not sufficient.

8. IANA Considerations

8.1. OAuth Parameters Registration

This specification registers the following parameters in the IANA "OAuth Parameters" registry [IANA.OAuth.Parameters] established by RFC 6749 [RFC6749].

8.1.1. Registry Contents

- o Parameter name: "client_id"
- o Parameter usage location: Authorization Response
- o Change controller: IESG
- o Specification document(s): Section 3.1 of [[this specification]]
- o Related information: None

- o Parameter name: "iss"
- o Parameter usage location: Authorization Response
- o Change controller: IESG
- o Specification document(s): Section 3.1 of [[this specification]]
- o Related information: None

- o Parameter name: "state"
- o Parameter usage location: Token Request
- o Change controller: IESG
- o Specification document(s): Section 5 of [[this specification]]
- o Related information: None

9. References

9.1. Normative References

- [IANA.OAuth.Parameters]
IANA, "OAuth Parameters",
<<http://www.iana.org/assignments/oauth-parameters>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015,
<<http://tools.ietf.org/html/rfc7519>>.
- [OAuth.Discovery]
Jones, M., Sakimura, N., and J. Bradley, "OAuth 2.0 Discovery", draft-ietf-oauth-discovery-02 (work in progress), March 2016, <<http://tools.ietf.org/html/draft-ietf-oauth-discovery-02>>.

[OAuth.Responses]

de Medeiros, B., Ed., Scurtescu, M., Tarjan, P., and M. Jones, "OAuth 2.0 Multiple Response Type Encoding Practices", February 2014, <http://openid.net/specs/oauth-v2-multiple-response-types-1_0.html>.

[OpenID.Core]

Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", November 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<http://www.rfc-editor.org/info/rfc3629>>.

[RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<http://www.rfc-editor.org/info/rfc5785>>.

[RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.

[RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<http://www.rfc-editor.org/info/rfc6819>>.

[RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.

[RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<http://www.rfc-editor.org/info/rfc7662>>.

9.2. Informative References

[arXiv.1508.04324v2]

Mladenov, V., Mainka, C., and J. Schwenk, "On the security of modern Single Sign-On Protocols: Second-Order Vulnerabilities in OpenID Connect", arXiv 1508.04324v2, January 2016, <<http://arxiv.org/abs/1508.04324v2/>>.

[arXiv.1601.01229v2]

Fett, D., Kuesters, R., and G. Schmitz, "A Comprehensive Formal Security Analysis of OAuth 2.0", arXiv 1601.01229v2, January 2016, <<http://arxiv.org/abs/1601.01229v2/>>.

[I-D.bradley-oauth-jwt-encoded-state]

Bradley, J., Lodderstedt, T., and H. Zandbelt, "Encoding claims in the OAuth 2 state parameter using a JWT", draft-bradley-oauth-jwt-encoded-state-05 (work in progress), December 2015.

[RFC7591] Richer, J., Ed., Jones, M., Bradley, J., Machulak, M., and P. Hunt, "OAuth 2.0 Dynamic Client Registration Protocol", RFC 7591, DOI 10.17487/RFC7591, July 2015, <<http://www.rfc-editor.org/info/rfc7591>>.

Appendix A. Implementation Notes

The authorization server can compare the two state values either by recording the complete state value between the authorization request and the token request, possibly in the same data structure in which the authorization code issued was recorded, or by recording only a cryptographic hash of the state value, possibly resulting in substantial size savings.

Appendix B. Acknowledgements

Alfred Albrecht, John Bradley, Brian Campbell, Joerg Connotte, William Denniss, Sebastian Ebling, Florian Feldmann, Daniel Fett, Roland Hedberg, Phil Hunt, Ralf Kuesters, Torsten Lodderstedt, Christian Mainka, Vladislav Mladenov, Anthony Nadalin, Justin Richer, Nat Sakimura, Antonio Sanso, Guido Schmitz, Joerg Schwenk, Hannes Tschofenig, and Hans Zandbelt all contributed to the discussions that led to the creation of this specification.

This specification is partially based on the OpenID Connect Core 1.0 specification, which was produced by the OpenID Connect working group of the OpenID Foundation.

Appendix C. Open Issues

- o We need to do a security analysis of the cut-and-paste attacks that may be enabled when mitigation information is returned to the client using individual authorization response parameters.

Appendix D. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-01

- o Changed terms "issuer URL" and "configuration information location" to "issuer identifier" so that consistent terminology is used for this.

-00

- o Created the initial working group draft from draft-jones-oauth-mix-up-mitigation-01 with no normative changes and adding Nat Sakimura as an editor.

Authors' Addresses

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Nat Sakimura
Nomura Research Institute, Ltd.

Email: n-sakimura@nri.co.jp
URI: <http://nat.sakimura.org/>

OAuth Working Group
Internet-Draft
Updates: 6749 (if approved)
Intended status: Best Current Practice
Expires: December 11, 2017

W. Denniss
Google
J. Bradley
Ping Identity
June 9, 2017

OAuth 2.0 for Native Apps
draft-ietf-oauth-native-apps-12

Abstract

OAuth 2.0 authorization requests from native apps should only be made through external user-agents, primarily the user's browser. This specification details the security and usability reasons why this is the case, and how native apps and authorization servers can implement this best practice.

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 December 11, 2017.

Copyright Notice

Copyright (c) 2017 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

1. Introduction	2
2. Notational Conventions	3
3. Terminology	3
4. Overview	4
4.1. Authorization Flow for Native Apps Using the Browser	5
5. Using Inter-app URI Communication for OAuth	6
6. Initiating the Authorization Request from a Native App	6
7. Receiving the Authorization Response in a Native App	7
7.1. Private-use URI Scheme Redirection	8
7.2. Claimed HTTPS URI Redirection	9
7.3. Loopback Interface Redirection	9
8. Security Considerations	10
8.1. Protecting the Authorization Code	10
8.2. OAuth Implicit Grant Authorization Flow	11
8.3. Loopback Redirect Considerations	11
8.4. Registration of Native App Clients	11
8.5. Client Authentication	12
8.6. Client Impersonation	12
8.7. Fake External User-Agent	13
8.8. Malicious External User-Agent	13
8.9. Cross-App Request Forgery Protections	14
8.10. Authorization Server Mix-Up Mitigation	14
8.11. Non-Browser External User-Agents	14
8.12. Embedded User-Agents	14
9. IANA Considerations	15
10. References	16
10.1. Normative References	16
10.2. Informative References	16
Appendix A. Server Support Checklist	17
Appendix B. Operating System Specific Implementation Details	17
B.1. iOS Implementation Details	18
B.2. Android Implementation Details	18
B.3. Windows Implementation Details	19
B.4. macOS Implementation Details	19
B.5. Linux Implementation Details	20
Appendix C. Acknowledgements	20
Authors' Addresses	20

1. Introduction

The OAuth 2.0 [RFC6749] authorization framework documents two approaches in Section 9 for native apps to interact with the

authorization endpoint: an embedded user-agent, and an external user-agent.

This best current practice requires that only external user-agents like the browser are used for OAuth by native apps. It documents how native apps can implement authorization flows using the browser as the preferred external user-agent, and the requirements for authorization servers to support such usage.

This practice is also known as the AppAuth pattern, in reference to open source libraries [AppAuth] that implement it.

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in Key words for use in RFCs to Indicate Requirement Levels [RFC2119]. If these words are used without being spelled in uppercase then they are to be interpreted with their normal natural language meanings.

3. Terminology

In addition to the terms defined in referenced specifications, this document uses the following terms:

"native app" An app or application that is installed by the user to their device, as distinct from a web app that runs in the browser context only. Apps implemented using web-based technology but distributed as a native app, so-called hybrid apps, are considered equivalent to native apps for the purpose of this specification.

"app" In this document, "app" means a "native app" unless further specified.

"app store" An ecommerce store where users can download and purchase apps.

"OAuth" In this document, OAuth refers to the OAuth 2.0 Authorization Framework [RFC6749].

"external user-agent" A user-agent capable of handling the authorization request that is a separate entity or security domain to the native app making the request (such as a browser), such that the app cannot access the cookie storage, nor inspect or modify page content.

"embedded user-agent" A user-agent hosted inside the native app itself (such as via a web-view), with which the app has control over to the extent it is capable of accessing the cookie storage and/or modifying the page content.

"browser" The default application launched by the operating system to handle "http" and "https" scheme URI content.

"in-app browser tab" A programmatic instantiation of the browser that is displayed inside a host app, but retains the full security properties and authentication state of the browser. Has different platform-specific product names, such as SFSafariViewController on iOS, and Custom Tabs on Android.

"inter-app communication" Communication between two apps on a device.

"claimed HTTPS URI" Some platforms allow apps to claim a HTTPS scheme URI after proving ownership of the domain name. URIs claimed in such a way are then opened in the app instead of the browser.

"private-use URI scheme" A private-use URI scheme defined by the app and registered with the operating system. URI requests to such schemes trigger the app which registered it to be launched to handle the request.

"web-view" A web browser UI (user interface) component that can be embedded in apps to render web pages, used to create embedded user-agents.

"reverse domain name notation" A naming convention based on the domain name system, but where the domain components are reversed, for example "app.example.com" becomes "com.example.app".

4. Overview

The best current practice for authorizing users in native apps is to perform the OAuth authorization request in an external user-agent (typically the browser), rather than an embedded user-agent (such as one implemented with web-views).

Previously it was common for native apps to use embedded user-agents (commonly implemented with web-views) for OAuth authorization requests. That approach has many drawbacks, including the host app being able to copy user credentials and cookies, and the user needing to authenticate from scratch in each app. See Section 8.12 for a deeper analysis of using embedded user-agents for OAuth.

Native app authorization requests that use the browser are more secure and can take advantage of the user's authentication state. Being able to use the existing authentication session in the browser enables single sign-on, as users don't need to authenticate to the authorization server each time they use a new app (unless required by authorization server policy).

Supporting authorization flows between a native app and the browser is possible without changing the OAuth protocol itself, as the authorization request and response are already defined in terms of URIs, which encompasses URIs that can be used for inter-app communication. Some OAuth server implementations that assume all clients are confidential web-clients will need to add an understanding of public native app clients and the types of redirect URIs they use to support this best practice.

4.1. Authorization Flow for Native Apps Using the Browser




Figure 1: Native App Authorization via External User-agent

Figure 1 illustrates the interaction of the native app with a browser external user-agent to authorize the user.

- (1) The client app opens a browser tab with the authorization request.
- (2) Authorization endpoint receives the authorization request, authenticates the user and obtains authorization. Authenticating the user may involve chaining to other authentication systems.
- (3) Authorization server issues an authorization code to the redirect URI.
- (4) Client receives the authorization code from the redirect URI.
- (5) Client app presents the authorization code at the token endpoint.
- (6) Token endpoint validates the authorization code and issues the tokens requested.

5. Using Inter-app URI Communication for OAuth

Just as URIs are used for OAuth 2.0 [RFC6749] on the web to initiate the authorization request and return the authorization response to the requesting website, URIs can be used by native apps to initiate the authorization request in the device's browser and return the response to the requesting native app.

By adopting the same methods used on the web for OAuth, benefits seen in the web context like the usability of a single sign-on session and the security of a separate authentication context are likewise gained in the native app context. Re-using the same approach also reduces the implementation complexity and increases interoperability by relying on standards-based web flows that are not specific to a particular platform.

To conform to this best practice, native apps MUST use an external user-agent to perform OAuth authentication requests. This is achieved by opening the authorization request in the browser (detailed in Section 6), and using a redirect URI that will return the authorization response back to the native app, as defined in Section 7.

6. Initiating the Authorization Request from a Native App

Native apps needing user authorization create an authorization request URI with the authorization code grant type per Section 4.1 of OAuth 2.0 [RFC6749], using a redirect URI capable of being received by the native app.

The function of the redirect URI for a native app authorization request is similar to that of a web-based authorization request. Rather than returning the authorization response to the OAuth client's server, the redirect URI used by a native app returns the response to the app. Several options for a redirect URI that will return the authorization response to the native app in different platforms are documented in Section 7. Any redirect URI that allows the app to receive the URI and inspect its parameters is viable.

Public native app clients MUST implement the Proof Key for Code Exchange (PKCE [RFC7636]) extension to OAuth, and authorization servers MUST support PKCE for such clients, for the reasons detailed in Section 8.1.

After constructing the authorization request URI, the app uses platform-specific APIs to open the URI in an external user-agent. Typically the external user-agent used is the default browser, that is, the application configured for handling "http" and "https" scheme URIs on the system, but different browser selection criteria and other categories of external user-agents MAY be used.

This best practice focuses on the browser as the RECOMMENDED external user-agent for native apps. An external user-agent designed specifically for processing authorization requests capable of processing the request and redirect URIs in the same way MAY also be used. Other external user-agents, such as a native app provided by the authorization server may meet the criteria set out in this best practice, including using the same redirection URI properties, but their use is out of scope for this specification.

Some platforms support a browser feature known as in-app browser tabs, where an app can present a tab of the browser within the app context without switching apps, but still retain key benefits of the browser such as a shared authentication state and security context. On platforms where they are supported, it is RECOMMENDED for usability reasons that apps use in-app browser tabs for the authorization request.

7. Receiving the Authorization Response in a Native App

There are several redirect URI options available to native apps for receiving the authorization response from the browser, the availability and user experience of which varies by platform.

To fully support this best practice, authorization servers MUST offer at least the following three redirect URI options to native apps. Native apps MAY use whichever redirect option suits their needs best, taking into account platform specific implementation details.

7.1. Private-use URI Scheme Redirection

Many mobile and desktop computing platforms support inter-app communication via URIs by allowing apps to register private-use URI schemes (sometimes colloquially referred to as custom URL schemes) like "com.example.app". When the browser or another app attempts to load a URI with a custom scheme, the app that registered it is launched to handle the request.

To perform an OAuth 2.0 authorization request with a private-use URI scheme redirect, the native app launches the browser with a standard authorization request, but one where the redirection URI utilizes a custom URI scheme it registered with the operating system.

When choosing a URI scheme to associate with the app, apps MUST use a URI scheme based on a domain name under their control, expressed in reverse order, as recommended by Section 3.8 of [RFC7595] for private-use URI schemes.

For example, an app that controls the domain name "app.example.com" can use "com.example.app" as their scheme. Some authorization servers assign client identifiers based on domain names, for example "client1234.usercontent.example.net", which can also be used as the domain name for the scheme when reversed in the same manner. A scheme such as "myapp" however would not meet this requirement, as it is not based on a domain name.

Care must be taken when there are multiple apps by the same publisher that each scheme is unique within that group. On platforms that use app identifiers that are also based on reverse order domain names, those can be reused as the private-use URI scheme for the OAuth redirect to help avoid this problem.

Following the requirements of [RFC3986] Section 3.2, as there is no naming authority for private-use URI scheme redirects, only a single slash ("/") appears after the scheme component. A complete example of a redirect URI utilizing a private-use URI scheme:

```
com.example.app:/oauth2redirect/example-provider
```

When the authentication server completes the request, it redirects to the client's redirection URI as it would normally. As the redirection URI uses a custom scheme it results in the operating system launching the native app, passing in the URI as a launch parameter. The native app then processes the authorization response like normal.

7.2. Claimed HTTPS URI Redirection

Some operating systems allow apps to claim HTTPS scheme [RFC7230] URIs in domains they control. When the browser encounters a claimed URI, instead of the page being loaded in the browser, the native app is launched with the URI supplied as a launch parameter.

Such URIs can be used as redirect URIs by native apps. They are indistinguishable to the authorization server from a regular web-based client redirect URI. An example is:

```
https://app.example.com/oauth2redirect/example-provider
```

As the redirect URI alone is not enough to distinguish public native app clients from confidential web clients, it is REQUIRED in Section 8.4 that the client type be recorded during client registration to enable the server to determine the client type and act accordingly.

App-claimed HTTPS redirect URIs have some advantages compared to other native app redirect options in that the identity of the destination app is guaranteed to the authorization server by the operating system. For this reason, native apps SHOULD use them over the other options where possible.

7.3. Loopback Interface Redirection

Native apps that are able to open a port on the loopback network interface without needing special permissions (typically, those on desktop operating systems) can use the loopback interface to receive the OAuth redirect.

Loopback redirect URIs use the HTTP scheme and are constructed with the loopback IP literal and whatever port the client is listening on. That is, "http://127.0.0.1:{port}/{path}" for IPv4, and "http://[::1]:{port}/{path}" for IPv6. An example redirect using the IPv4 loopback interface with a randomly assigned port:

```
http://127.0.0.1:50719/oauth2redirect/example-provider
```

An example redirect using the IPv6 loopback interface with a randomly assigned port:

```
http://[::1]:61023/oauth2redirect/example-provider
```

The authorization server MUST allow any port to be specified at the time of the request for loopback IP redirect URIs, to accommodate

clients that obtain an available ephemeral port from the operating system at the time of the request.

Clients SHOULD NOT assume the device supports a particular version of the Internet Protocol. It is RECOMMENDED that clients attempt to bind to the loopback interface using both IPv4 and IPv6, and use whichever is available.

8. Security Considerations

8.1. Protecting the Authorization Code

The redirect URI options documented in Section 7 share the benefit that only a native app on the same device can receive the authorization code which limits the attack surface, however code interception by a different native app running on the same device may be possible.

A limitation of using private-use URI schemes for redirect URIs is that multiple apps can typically register the same scheme, which makes it indeterminate as to which app will receive the Authorization Code. Section 1 of PKCE [RFC7636] details how this limitation can be used to execute a code interception attack.

Loopback IP based redirect URIs may be susceptible to interception by other apps accessing the same loopback interface on some operating systems.

App-claimed HTTPS redirects are less susceptible to URI interception due to the presence of the URI authority, but they are still public clients and the URI is sent using the operating system's URI dispatch handler with unknown security properties.

The Proof Key for Code Exchange by OAuth Public Clients (PKCE [RFC7636]) standard was created specifically to mitigate against this attack. It is a proof of possession extension to OAuth 2.0 that protects the code grant from being used if it is intercepted. It achieves this by having the client generate a secret verifier, a hash of which it passes in the initial authorization request, and which it must present in full when redeeming the authorization code grant. An app that intercepted the authorization code would not be in possession of this secret, rendering the code useless.

Section 6 requires that both clients and servers use PKCE for public native app clients. Authorization servers SHOULD reject authorization requests from native apps that don't use PKCE by returning an error message as defined in Section 4.4.1 of PKCE [RFC7636].

8.2. OAuth Implicit Grant Authorization Flow

The OAuth 2.0 implicit grant authorization flow as defined in Section 4.2 of OAuth 2.0 [RFC6749] generally works with the practice of performing the authorization request in the browser, and receiving the authorization response via URI-based inter-app communication. However, as the implicit flow cannot be protected by PKCE [RFC7636] (which is a required in Section 8.1), the use of the Implicit Flow with native apps is NOT RECOMMENDED.

Tokens granted via the implicit flow also cannot be refreshed without user interaction, making the authorization code grant flow - which can issue refresh tokens - the more practical option for native app authorizations that require refreshing.

8.3. Loopback Redirect Considerations

Loopback interface redirect URIs use the "http" scheme (i.e., without TLS). This is acceptable for loopback interface redirect URIs as the HTTP request never leaves the device.

Clients should open the network port only when starting the authorization request, and close it once the response is returned.

Clients should listen on the loopback network interface only, to avoid interference by other network actors.

While redirect URIs using localhost (i.e., "http://localhost:{port}/") function similarly to loopback IP redirects described in Section 7.3, the use of "localhost" is NOT RECOMMENDED. Specifying a redirect URI with the loopback IP literal rather than localhost avoids inadvertently listening on network interfaces other than the loopback interface. It is also less susceptible to client side firewalls, and misconfigured host name resolution on the user's device.

8.4. Registration of Native App Clients

Native apps, except when using a mechanism like Dynamic Client Registration [RFC7591] to provision per-instance secrets, are classified as public clients, as defined by Section 2.1 of OAuth 2.0 [RFC6749] and MUST be registered with the authorization server as such. Authorization servers MUST record the client type in the client registration details in order to identify and process requests accordingly.

Authorization servers MUST require clients to register their complete redirect URI (including the path component), and reject authorization

requests that specify a redirect URI that doesn't exactly match the one that was registered, with the exception of loopback redirects, where an exact match is required except for the port URI component.

For private-use URI scheme based redirects, authorization servers SHOULD enforce the requirement in Section 7.1 that clients use reverse domain name based schemes. At a minimum, any scheme that doesn't contain a period character ("."), SHOULD be rejected.

In addition to the collision resistant properties, requiring a URI scheme based on a domain name that is under the control of the app can help to prove ownership in the event of a dispute where two apps claim the same private-use URI scheme (where one app is acting maliciously). For example, if two apps claimed "com.example.app", the owner of "example.com" could petition the app store operator to remove the counterfeit app. Such a petition is harder to prove if a generic URI scheme was used.

Authorization servers MAY request the inclusion of other platform-specific information, such as the app package or bundle name, or other information used to associate the app that may be useful for verifying the calling app's identity, on operating systems that support such functions.

8.5. Client Authentication

Secrets that are statically included as part of an app distributed to multiple users should not be treated as confidential secrets, as one user may inspect their copy and learn the shared secret. For this reason, and those stated in Section 5.3.1 of [RFC6819], it is NOT RECOMMENDED for authorization servers to require client authentication of public native apps clients using a shared secret, as this serves little value beyond client identification which is already provided by the "client_id" request parameter.

Authorization servers that still require a statically included shared secret for native app clients MUST treat the client as a public client (as defined by Section 2.1 of OAuth 2.0 [RFC6749]), and not accept the secret as proof of the client's identity. Without additional measures, such clients are subject to client impersonation (see Section 8.6).

8.6. Client Impersonation

As stated in Section 10.2 of OAuth 2.0 [RFC6749], the authorization server SHOULD NOT process authorization requests automatically without user consent or interaction, except when the identity of the client can be assured. This includes the case where the user has

previously approved an authorization request for a given client id - unless the identity of the client can be proven, the request SHOULD be processed as if no previous request had been approved.

Measures such as claimed HTTPS redirects MAY be accepted by authorization servers as identity proof. Some operating systems may offer alternative platform-specific identity features which MAY be accepted, as appropriate.

8.7. Fake External User-Agent

The native app which is initiating the authorization request has a large degree of control over the user interface and can potentially present a fake external user-agent, that is, an embedded user-agent made to appear as an external user agent.

The advantage when all good actors are using external user-agents is that it is possible for security experts to detect bad actors, as anyone faking an external user-agent is provably bad. If good and bad actors alike are using embedded user-agents, bad actors don't need to fake anything, making them harder to detect. Once malicious apps are detected, it may be possible to use this knowledge to blacklist the apps signatures in malware scanning software, take removal action in the case of apps distributed by app stores, and other steps to reduce the impact and spread of the malicious app.

Authorization servers can also directly protect against fake external user-agents by requiring an authentication factor only available to true external user-agents.

Users who are particularly concerned about their security when using in-app browser tabs may also take the additional step of opening the request in the full browser from the in-app browser tab, and complete the authorization there, as most implementations of the in-app browser tab pattern offer such functionality.

8.8. Malicious External User-Agent

If a malicious app is able to configure itself as the default handler for "https" scheme URIs in the operating system, it will be able to intercept authorization requests that use the default browser and abuse this position of trust for malicious ends such as phishing the user.

Many operating systems mitigate this issue by requiring an explicit user action to change the default handler for HTTP URIs. This attack is not confined to OAuth for Native Apps, a malicious app configured

in this way would present a general and ongoing risk to the user beyond OAuth usage.

8.9. Cross-App Request Forgery Protections

Section 5.3.5 of [RFC6819] recommends using the "state" parameter to link client requests and responses to prevent CSRF (Cross Site Request Forgery) attacks.

To mitigate CSRF style attacks using inter-app URI communication, it is similarly RECOMMENDED that native apps include a high entropy secure random number in the "state" parameter of the authorization request, and reject any incoming authorization responses without a state value that matches a pending outgoing authorization request.

8.10. Authorization Server Mix-Up Mitigation

To protect against a compromised or malicious authorization server attacking another authorization server used by the same app, it is REQUIRED that a unique redirect URI is used for each authorization server used by the app (for example, by varying the path component), and that authorization responses are rejected if the redirect URI they were received on doesn't match the redirect URI in a outgoing authorization request.

The native app MUST store the redirect URI used in the authorization request with the authorization session data (i.e., along with "state" and other related data), and MUST verify that the URI on which the authorization response was received exactly matches it.

The requirements of Section 8.4 that authorization servers reject requests with URIs that don't match what was registered are also required to prevent such attacks.

8.11. Non-Browser External User-Agents

This best practice recommends a particular type of external user-agent, the user's browser. Other external user-agent patterns may also be viable for secure and usable OAuth. This document makes no comment on those patterns.

8.12. Embedded User-Agents

OAuth 2.0 [RFC6749] Section 9 documents two approaches for native apps to interact with the authorization endpoint. This best current practice requires that native apps MUST NOT use embedded user-agents to perform authorization requests, and allows that authorization endpoints MAY take steps to detect and block authorization requests

in embedded user-agents. The security considerations for these requirements are detailed herein.

Embedded user-agents are an alternative method for authorizing native apps. These embedded user agents are unsafe for use by third-parties to the authorization server by definition, as the app that hosts the embedded user-agent can access the user's full authentication credential, not just the OAuth authorization grant that was intended for the app.

In typical web-view based implementations of embedded user-agents, the host application can: log every keystroke entered in the form to capture usernames and passwords; automatically submit forms and bypass user-consent; copy session cookies and use them to perform authenticated actions as the user.

Even when used by trusted apps belonging to the same party as the authorization server, embedded user-agents violate the principle of least privilege by having access to more powerful credentials than they need, potentially increasing the attack surface.

Encouraging users to enter credentials in an embedded user-agent without the usual address bar and visible certificate validation features that browsers have makes it impossible for the user to know if they are signing in to the legitimate site, and even when they are, it trains them that it's OK to enter credentials without validating the site first.

Aside from the security concerns, embedded user-agents do not share the authentication state with other apps or the browser, requiring the user to login for every authorization request which is often considered an inferior user experience.

9. IANA Considerations

[RFC Editor: please do NOT remove this section.]

This document has no IANA actions.

Section 7.1 specifies how private-use URI schemes are used for inter-app communication in OAuth protocol flows. This document requires in Section 7.1 that such schemes are based on domain names owned or assigned to the app, as recommended in Section 3.8 of [RFC7595]. Per Section 6 of [RFC7595], registration of domain based URI schemes with IANA is not required.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", BCP 35, RFC 7595, DOI 10.17487/RFC7595, June 2015, <<http://www.rfc-editor.org/info/rfc7595>>.
- [RFC7636] Sakimura, N., Ed., Bradley, J., and N. Agarwal, "Proof Key for Code Exchange by OAuth Public Clients", RFC 7636, DOI 10.17487/RFC7636, September 2015, <<http://www.rfc-editor.org/info/rfc7636>>.

10.2. Informative References

- [RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<http://www.rfc-editor.org/info/rfc6819>>.
- [RFC7591] Richer, J., Ed., Jones, M., Bradley, J., Machulak, M., and P. Hunt, "OAuth 2.0 Dynamic Client Registration Protocol", RFC 7591, DOI 10.17487/RFC7591, July 2015, <<http://www.rfc-editor.org/info/rfc7591>>.
- [AppAuth] Denniss, W., Wright, S., McGinniss, I., Ravikumar, R., and others, "AppAuth", May 22, <<https://appauth.io>>.

[AppAuth.iOSmacOS]

Wright, S., Denniss, W., and others, "AppAuth for iOS and macOS", February 2016, <<https://github.com/openid/AppAuth-iOS>>.

[AppAuth.Android]

McGinniss, I., Denniss, W., and others, "AppAuth for Android", February 2016, <<https://github.com/openid/AppAuth-Android>>.

[SamplesForWindows]

Denniss, W., "OAuth for Apps: Samples for Windows", July 2016, <<https://github.com/googlesamples/oauth-apps-for-windows>>.

Appendix A. Server Support Checklist

OAuth servers that support native apps must:

1. Support private-use URI scheme redirect URIs. This is required to support mobile operating systems. See Section 7.1.
2. Support HTTPS scheme redirect URIs for use with public native app clients. This is used by apps on advanced mobile operating systems that allow app-claimed URIs. See Section 7.2.
3. Support loopback IP redirect URIs. This is required to support desktop operating systems. See Section 7.3.
4. Not assume native app clients can keep a secret. If secrets are distributed to multiple installs of the same native app, they should not be treated as confidential. See Section 8.5.
5. Support PKCE [RFC7636]. Required to protect authorization code grants sent to public clients over inter-app communication channels. See Section 8.1

Appendix B. Operating System Specific Implementation Details

This document primarily defines best practices in a generic manner, referencing techniques commonly available in a variety of environments. This non-normative section documents operating system specific implementation details of the best practice.

The implementation details herein are considered accurate at the time of publishing but will likely change over time. It is hoped that such change won't invalidate the generic principles in the rest of

the document, and those principles should take precedence in the event of a conflict.

B.1. iOS Implementation Details

Apps can initiate an authorization request in the browser without the user leaving the app, through the `SFSafariViewController` class which implements the in-app browser tab pattern. Safari can be used to handle requests on old versions of iOS without `SFSafariViewController`.

To receive the authorization response, both private-use URI scheme redirects (referred to as Custom URL Schemes) and claimed HTTPS links (known as Universal Links) are viable choices, and function the same whether the request is loaded in `SFSafariViewController` or the Safari app. Apps can claim Custom URI schemes with the `"CFBundleURLTypes"` key in the application's property list file `"Info.plist"`, and HTTPS links using the Universal Links feature with an entitlement file and an association file on the domain.

Universal Links are the preferred choice on iOS 9 and above due to the ownership proof that is provided by the operating system.

A complete open source sample is included in the AppAuth for iOS and macOS [`AppAuth.iOSmacOS`] library.

B.2. Android Implementation Details

Apps can initiate an authorization request in the browser without the user leaving the app, through the Android Custom Tab feature which implements the in-app browser tab pattern. The user's default browser can be used to handle requests when no browser supports Custom Tabs.

Android browser vendors should support the Custom Tabs protocol (by providing an implementation of the `"CustomTabsService"` class), to provide the in-app browser tab user experience optimization to their users. Chrome is one such browser that implements Custom Tabs.

To receive the authorization response, private-use URI schemes are broadly supported through Android Implicit Intents. Claimed HTTPS redirect URIs through Android App Links are available on Android 6.0 and above. Both types of redirect URIs are registered in the application's manifest.

A complete open source sample is included in the AppAuth for Android [`AppAuth.Android`] library.

B.3. Windows Implementation Details

Both traditional and Universal Windows Platform (UWP) apps can perform authorization requests in the user's browser. Traditional apps typically use a loopback redirect to receive the authorization response, and listening on the loopback interface is allowed by default firewall rules. When creating the loopback network socket, apps SHOULD set the "SO_EXCLUSIVEADDRUSE" socket option to prevent other apps binding to the same socket.

UWP apps can use private-use URI scheme redirects to receive the authorization response from the browser, which will bring the app to the foreground. Known on the platform as "URI Activation", the URI scheme is limited to 39 characters in length, and may include the "." character, making short reverse domain name based schemes (as recommended in Section 7.1) possible.

UWP apps can alternatively use the Web Authentication Broker API in SSO (Single Sign-on) mode, which is an external user agent designed for authorization flows. Cookies are shared between invocations of the broker but not the user's preferred browser, meaning the user will need to sign-in again even if they have an active session in their browser, but the session created in the broker will be available to subsequent apps that use the broker. Personalisations the user has made to their browser, such as configuring a password manager may not be available in the broker. To qualify as an external user-agent, the broker MUST be used in SSO mode.

To use the Web Authentication Broker in SSO mode, the redirect URI must be of the form "msapp://{appSID}" where "appSID" is the app's SID, which can be found in the app's registration information. While Windows enforces the URI authority on such redirects, ensuring only the app with the matching SID can receive the response on Windows, the URI scheme could be claimed by apps on other platforms without the same authority present, thus this redirect type should be treated similar to private-use URI scheme redirects for security purposes.

An open source sample demonstrating these patterns is available [SamplesForWindows].

B.4. macOS Implementation Details

Apps can initiate an authorization request in the user's default browser using platform APIs for opening URIs in the browser.

To receive the authorization response, private-use URI schemes are a good redirect URI choice on macOS, as the user is returned right back to the app they launched the request from. These are registered in

the application's bundle information property list using the "CFBundleURLSchemes" key. Loopback IP redirects are another viable option, and listening on the loopback interface is allowed by default firewall rules.

A complete open source sample is included in the AppAuth for iOS and macOS [AppAuth.iOSmacOS] library.

B.5. Linux Implementation Details

Opening the Authorization Request in the user's default browser requires a distro-specific command, "xdg-open" is one such tool.

The loopback redirect is the recommended redirect choice for desktop apps on Linux to receive the authorization response. Apps SHOULD NOT set the "SO_REUSEPORT" or "SO_REUSEADDR" socket options, to prevent other apps binding to the same socket.

Appendix C. Acknowledgements

The author would like to acknowledge the work of Marius Scurtescu, and Ben Wiley Sittler whose design for using private-use URI schemes in native OAuth 2.0 clients at Google formed the basis of Section 7.1.

The following individuals contributed ideas, feedback, and wording that shaped and formed the final specification:

Andy Zmolek, Steven E Wright, Brian Campbell, Nat Sakimura, Eric Sachs, Paul Madsen, Iain McGinniss, Rahul Ravikumar, Breno de Medeiros, Hannes Tschofenig, Ashish Jain, Erik Wahlstrom, Bill Fisher, Sudhi Umarji, Michael B. Jones, Vittorio Bertocci, Dick Hardt, David Waite, Ignacio Fiorentino, Kathleen Moriarty, and Elwyn Davies.

Authors' Addresses

William Denniss
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
USA

Email: wdenniss@google.com
URI: <http://wdenniss.com/appauth>

John Bradley
Ping Identity

Phone: +1 202-630-5272

Email: ve7jtb@ve7jtb.com

URI: <http://www.thread-safe.com/p/appauth.html>

OAuth
Internet-Draft
Intended status: Informational
Expires: January 9, 2017

P. Hunt, Ed.
Oracle Corporation
J. Richer

W. Mills

P. Mishra
Oracle Corporation
H. Tschofenig
ARM Limited
July 8, 2016

OAuth 2.0 Proof-of-Possession (PoP) Security Architecture
draft-ietf-oauth-pop-architecture-08.txt

Abstract

The OAuth 2.0 bearer token specification, as defined in RFC 6750, allows any party in possession of a bearer token (a "bearer") to get access to the associated resources (without demonstrating possession of a cryptographic key). To prevent misuse, bearer tokens must be protected from disclosure in transit and at rest.

Some scenarios demand additional security protection whereby a client needs to demonstrate possession of cryptographic keying material when accessing a protected resource. This document motivates the development of the OAuth 2.0 proof-of-possession security mechanism.

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 9, 2017.

Copyright Notice

Copyright (c) 2016 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

1. Introduction	3
2. Terminology	3
3. Use Cases	3
3.1. Preventing Access Token Re-Use by the Resource Server	4
3.2. TLS and DTLS Channel Binding Support	4
3.3. Access to a Non-TLS Protected Resource	4
3.4. Offering Application Layer End-to-End Security	5
4. Security and Privacy Threats	5
5. Requirements	6
6. Threat Mitigation	10
6.1. Confidentiality Protection	11
6.2. Sender Constraint	11
6.3. Key Confirmation	12
6.4. Summary	13
7. Architecture	14
7.1. Client and Authorization Server Interaction	15
7.1.1. Symmetric Keys	15
7.1.2. Asymmetric Keys	16
7.2. Client and Resource Server Interaction	17
7.3. Resource and Authorization Server Interaction (Token Introspection)	18
8. Security Considerations	19
9. IANA Considerations	19
10. Acknowledgments	19
11. References	20
11.1. Normative References	20
11.2. Informative References	21
Authors' Addresses	22

1. Introduction

The OAuth 2.0 protocol family ([RFC6749], [RFC6750], and [RFC6819]) offer a single token type known as the "bearer" token to access protected resources. RFC 6750 [RFC6750] specifies the bearer token mechanism and defines it as follows:

"A security token with the property that any party in possession of the token (a "bearer") can use the token in any way that any other party in possession of it can. Using a bearer token does not require a bearer to prove possession of cryptographic key material."

The bearer token meets the security needs of a number of use cases the OAuth 2.0 protocol had originally been designed for. There are, however, other scenarios that require stronger security properties and ask for active participation of the OAuth client in form of cryptographic computations when presenting an access token to a resource server.

This document outlines additional use cases requiring stronger security protection in Section 3, identifies threats in Section 4, proposes different ways to mitigate those threats in Section 6, outlines an architecture for a solution that builds on top of the existing OAuth 2.0 framework in Section 7, and concludes with a requirements list in Section 5.

2. Terminology

The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in this specification are to be interpreted as described in [RFC2119], with the important qualification that, unless otherwise stated, these terms apply to the design of the protocol, not its implementation or application.

3. Use Cases

The main use case that motivates improvement upon "bearer" token security is the desire of resource servers to obtain additional assurance that the client is indeed authorized to present an access token. The expectation is that the use of additional credentials (symmetric or asymmetric keying material) will encourage developers to take additional precautions when transferring and storing access token in combination with these credentials.

Additional use cases listed below provide further requirements for the solution development. Note that a single solution does not necessarily need to offer support for all use cases.

3.1. Preventing Access Token Re-Use by the Resource Server

In a scenario where a resource server receives a valid access token, the resource server then re-uses it with other resource server. The reason for re-use may be malicious or may well be legitimate. In a legitimate case, the intent is to support chaining of computations whereby a resource server needs to consult other third party resource servers to complete a requested operation. In both cases it may be assumed that the scope and audience of the access token is sufficiently defined that to allow such a re-use. For example, imagine a case where a company operates email services as well as picture sharing services and that company had decided to issue access tokens with a scope and audience that allows access to both services.

With this use case the desire is to prevent such access token re-use. This also implies that the legitimate use cases require additional enhancements for request chaining.

3.2. TLS and DTLS Channel Binding Support

In this use case we consider the scenario where an OAuth 2.0 request to a protected resource is secured using TLS or DTLS (see [RFC4347]), but the client and the resource server demand that the underlying TLS/DTLS exchange is bound to additional application layer security to prevent cases where the TLS/DTLS connection is terminated at a TLS/DTLS intermediary, which splits the TLS/DTLS connection into two separate connections.

In this use case additional information should be conveyed to the resource server to ensure that no entity entity has tampered with the TLS/DTLS connection.

3.3. Access to a Non-TLS Protected Resource

This use case is for a web client that needs to access a resource that makes data available (such as videos) without offering integrity and confidentiality protection using TLS. Still, the initial resource request using OAuth, which includes the access token, must be protected against various threats (e.g., token replay, token modification).

While it is possible to utilize bearer tokens in this scenario with TLS protection when the request to the protected resource is made, as described in [RFC6750], there may be the desire to avoid using TLS

between the client and the resource server at all. In such a case the bearer token approach is not possible since it relies on TLS for ensuring integrity and confidentiality protection of the access token exchange since otherwise replay attacks are possible: First, an eavesdropper may steal an access token and present it at a different resource server. Second, an eavesdropper may steal an access token and replay it against the same resource server at a later point in time. In both cases, if the attack is successful, the adversary gets access to the resource owners data or may perform an operation selected by the adversary (e.g., sending a message). Note that the adversary may obtain the access token (if the recommendations in [RFC6749] and [RFC6750] are not followed) using a number of ways, including eavesdropping the communication on the wireless link.

Consequently, the important assumption in this use case is that a resource server does not have TLS support and the security solution should work in such a scenario. Furthermore, it may not be necessary to provide authentication of the resource server towards the client.

3.4. Offering Application Layer End-to-End Security

In Web deployments resource servers are often placed behind load balancers, which are deployed by the same organization that operates the resource servers. These load balancers may terminate the TLS connection setup and HTTP traffic is transmitted without TLS protection from the load balancer to the resource server. With application layer security in addition to the underlying TLS security it is possible to allow application servers to perform cryptographic verification on an end-to-end basis.

The key aspect in this use case is therefore to offer end-to-end security in the presence of load balancers via application layer security. Enterprise networks also deploy proxies that inspect traffic and thereby break TLS.

4. Security and Privacy Threats

The following list presents several common threats against protocols utilizing some form of token. This list of threats is based on NIST Special Publication 800-63 [NIST800-63]. We exclude a discussion of threats related to any form of identity proofing and authentication of the resource owner to the authorization server since these procedures are not part of the OAuth 2.0 protocol specification itself.

Token manufacture/modification:

An attacker may generate a bogus token or modify the token content (such as authentication or attribute statements) of an existing token, causing resource server to grant inappropriate access to the client. For example, an attacker may modify the token to extend the validity period. A client, which MAY be a normal client or MAY be assumed to be constrained (see [RFC7252]), may modify the token to have access to information that they should not be able to view.

Token disclosure:

Tokens may contain personal data, such as real name, age or birthday, payment information, etc.

Token redirect:

An attacker uses the token generated for consumption by the resource server to obtain access to another resource server.

Token reuse:

An attacker attempts to use a token that has already been used once with a resource server. The attacker may be an eavesdropper who observes the communication exchange or, worse, one of the communication end points. A client may, for example, leak access tokens because it cannot keep secrets confidential. A client may also reuse access tokens for some other resource servers. Finally, a resource server may use a token it had obtained from a client and use it with another resource server that the client interacts with. A resource server, offering relatively unimportant application services, may attempt to use an access token obtained from a client to access a high-value service, such as a payment service, on behalf of the client using the same access token.

Token repudiation:

Token repudiation refers to a property whereby a resource server is given an assurance that the authorization server cannot deny to have created a token for the client.

5. Requirements

RFC 4962 [RFC4962] gives useful guidelines for designers of authentication and key management protocols. While RFC 4962 was written with the AAA framework used for network access authentication in mind the offered suggestions are useful for the design of other key management systems as well. The following requirements list

applies OAuth 2.0 terminology to the requirements outlined in RFC 4962.

These requirements include

Cryptographic Algorithm Independent:

The key management protocol MUST be cryptographic algorithm independent.

Strong, fresh session keys:

Session keys MUST be strong and fresh. Each session deserves an independent session key, i.e., one that is generated specifically for the intended use. In context of OAuth this means that keying material is created in such a way that can only be used by the combination of a client instance, protected resource, and authorization scope.

Limit Key Scope:

Following the principle of least privilege, parties MUST NOT have access to keying material that is not needed to perform their role. Any protocol that is used to establish session keys MUST specify the scope for session keys, clearly identifying the parties to whom the session key is available.

Replay Detection Mechanism:

The key management protocol exchanges MUST be replay protected. Replay protection allows a protocol message recipient to discard any message that was recorded during a previous legitimate dialogue and presented as though it belonged to the current dialogue.

Authenticate All Parties:

Each party in the key management protocol MUST be authenticated to the other parties with whom they communicate. Authentication mechanisms MUST maintain the confidentiality of any secret values used in the authentication process. Secrets MUST NOT be sent to another party without confidentiality protection.

Authorization:

Client and resource server authorization MUST be performed. These entities MUST demonstrate possession of the appropriate keying material, without disclosing it. Authorization is REQUIRED

whenever a client interacts with an authorization server.
Authorization checking prevents an elevation of privilege attack.

Keying Material Confidentiality and Integrity:

While preserving algorithm independence, confidentiality and integrity of all keying material MUST be maintained.

Confirm Cryptographic Algorithm Selection:

The selection of the "best" cryptographic algorithms SHOULD be securely confirmed. The mechanism SHOULD detect attempted roll-back attacks.

Uniquely Named Keys:

Key management proposals require a robust key naming scheme, particularly where key caching is supported. The key name provides a way to refer to a key in a protocol so that it is clear to all parties which key is being referenced. Objects that cannot be named cannot be managed. All keys MUST be uniquely named, and the key name MUST NOT directly or indirectly disclose the keying material.

Prevent the Domino Effect:

Compromise of a single client MUST NOT compromise keying material held by any other client within the system, including session keys and long-term keys. Likewise, compromise of a single resource server MUST NOT compromise keying material held by any other Resource Server within the system. In the context of a key hierarchy, this means that the compromise of one node in the key hierarchy must not disclose the information necessary to compromise other branches in the key hierarchy. Obviously, the compromise of the root of the key hierarchy will compromise all of the keys; however, a compromise in one branch MUST NOT result in the compromise of other branches. There are many implications of this requirement; however, two implications deserve highlighting. First, the scope of the keying material must be defined and understood by all parties that communicate with a party that holds that keying material. Second, a party that holds keying material in a key hierarchy must not share that keying material with parties that are associated with other branches in the key hierarchy.

Bind Key to its Context:

Keying material MUST be bound to the appropriate context. The context includes the following.

- * The manner in which the keying material is expected to be used.
- * The other parties that are expected to have access to the keying material.
- * The expected lifetime of the keying material. Lifetime of a child key SHOULD NOT be greater than the lifetime of its parent in the key hierarchy.

Any party with legitimate access to keying material can determine its context. In addition, the protocol MUST ensure that all parties with legitimate access to keying material have the same context for the keying material. This requires that the parties are properly identified and authenticated, so that all of the parties that have access to the keying material can be determined. The context will include the client and the resource server identities in more than one form.

Authorization Restriction:

If client authorization is restricted, then the client SHOULD be made aware of the restriction.

Client Identity Confidentiality:

A client has identity confidentiality when any party other than the resource server and the authorization server cannot sufficiently identify the client within the anonymity set. In comparison to anonymity and pseudonymity, identity confidentiality is concerned with eavesdroppers and intermediaries. A key management protocol SHOULD provide this property.

Resource Owner Identity Confidentiality:

Resource servers SHOULD be prevented from knowing the real or pseudonymous identity of the resource owner, since the authorization server is the only entity involved in verifying the resource owner's identity.

Collusion:

Resource servers that collude can be prevented from using information related to the resource owner to track the individual. That is, two different resource servers can be prevented from determining that the same resource owner has authenticated to both

of them. Authorization servers MUST bind different keying material to access tokens used for resource servers from different origins (or similar concepts in the app world).

AS-to-RS Relationship Anonymity:

For solutions using asymmetric key cryptography the client MAY conceal information about the resource server it wants to interact with. The authorization server MAY reject such an attempt since it may not be able to enforce access control decisions.

Channel Binding:

A solution MUST enable support for channel bindings. The concept of channel binding, as defined in [RFC5056], allows applications to establish that the two end-points of a secure channel at one network layer are the same as at a higher layer by binding authentication at the higher layer to the channel at the lower layer.

There are performance concerns with the use of asymmetric cryptography. Although symmetric key cryptography offers better performance asymmetric cryptography offers additional security properties. A solution MUST therefore offer the capability to support both symmetric as well as asymmetric keys.

There are threats that relate to the experience of the software developer as well as operational practices. Verifying the servers identity in TLS is discussed at length in [RFC6125].

A number of the threats listed in Section 4 demand protection of the access token content and a standardized solution, for example, in the form of a JSON-based format, is available with the JWT [RFC7519].

6. Threat Mitigation

A large range of threats can be mitigated by protecting the content of the token, for example using a digital signature or a keyed message digest. Alternatively, the content of the token could be passed by reference rather than by value (requiring a separate message exchange to resolve the reference to the token content).

To simplify discussion in the following example we assume that the token itself cannot be modified by the client, either due to cryptographic protection (such as signature or encryption) or use of a reference value with sufficient entropy and associated secure lookup. The token remains opaque to the client. These are characteristics shared with bearer tokens and more information on

best practices can be found in [RFC6819] and in the security considerations section of [RFC6750].

To deal with token redirect it is important for the authorization server to include the identifier of the intended recipient - the resource server. A resource server must not be allowed to accept access tokens that are not meant for its consumption.

To provide protection against token disclosure two approaches are possible, namely (a) not to include sensitive information inside the token or (b) to ensure confidentiality protection. The latter approach requires at least the communication interaction between the client and the authorization server as well as the interaction between the client and the resource server to experience confidentiality protection. As an example, TLS with a ciphersuite that offers confidentiality protection has to be applied as per [RFC7525]. Encrypting the token content itself is another alternative. In our scenario the authorization server would, for example, encrypt the token content with a symmetric key shared with the resource server.

To deal with token reuse more choices are available.

6.1. Confidentiality Protection

In this approach confidentiality protection of the exchange is provided on the communication interfaces between the client and the resource server, and between the client and the authorization server. No eavesdropper on the wire is able to observe the token exchange. Consequently, a replay by a third party is not possible. An authorization server wants to ensure that it only hands out tokens to clients it has authenticated first and who are authorized. For this purpose, authentication of the client to the authorization server will be a requirement to ensure adequate protection against a range of attacks. This is, however, true for the description in Section 6.2 and Section 6.3 as well. Furthermore, the client has to make sure it does not distribute (or leak) the access token to entities other than the intended the resource server. For that purpose the client will have to authenticate the resource server before transmitting the access token.

6.2. Sender Constraint

Instead of providing confidentiality protection, the authorization server could also put the identifier of the client into the protected token with the following semantic: 'This token is only valid when presented by a client with the following identifier.' When the access token is then presented to the resource server how does it

know that it was provided by the client? It has to authenticate the client! There are many choices for authenticating the client to the resource server, for example by using client certificates in TLS [RFC5246], or pre-shared secrets within TLS [RFC4279]. The choice of the preferred authentication mechanism and credential type may depend on a number of factors, including

- o security properties
- o available infrastructure
- o library support
- o credential cost (financial)
- o performance
- o integration into the existing IT infrastructure
- o operational overhead for configuration and distribution of credentials

This long list hints to the challenge of selecting at least one mandatory-to-implement client authentication mechanism.

6.3. Key Confirmation

A variation of the mechanism of sender authentication, described in Section 6.2, is to replace authentication with the proof-of-possession of a specific (session) key, i.e., key confirmation. In this model the resource server would not authenticate the client itself but would rather verify whether the client knows the session key associated with a specific access token. Examples of this approach can be found with the OAuth 1.0 MAC token [RFC5849], and Kerberos [RFC4120] when utilizing the AP_REQ/AP_REP exchange (see also [I-D.hardjono-oauth-kerberos] for a comparison between Kerberos and OAuth).

To illustrate key confirmation, the first example is borrowed from Kerberos and use symmetric key cryptography. Assume that the authorization server shares a long-term secret with the resource server, called $K(\text{Authorization Server-Resource Server})$. This secret would be established between them out-of-band. When the client requests an access token the authorization server creates a fresh and unique session key K_s and places it into the token encrypted with the long term key $K(\text{Authorization Server-Resource Server})$. Additionally, the authorization server attaches K_s to the response message to the client (in addition to the access token itself) over a

confidentiality protected channel. When the client sends a request to the resource server it has to use Ks to compute a keyed message digest for the request (in whatever form or whatever layer). The resource server, when receiving the message, retrieves the access token, verifies it and extracts K(Authorization Server-Resource Server) to obtain Ks. This key Ks is then used to verify the keyed message digest of the request message.

Note that in this example one could imagine that the mechanism to protect the token itself is based on a symmetric key based mechanism to avoid any form of public key infrastructure but this aspect is not further elaborated in the scenario.

A similar mechanism can also be designed using asymmetric cryptography. When the client requests an access token the authorization server creates an ephemeral public / privacy key pair (PK/SK) and places the public key PK into the protected token. When the authorization server returns the access token to the client it also provides the PK/SK key pair over a confidentiality protected channel. When the client sends a request to the resource server it has to use the privacy key SK to sign the request. The resource server, when receiving the message, retrieves the access token, verifies it and extracts the public key PK. It uses this ephemeral public key to verify the attached signature.

6.4. Summary

As a high level message, there are various ways the threats can be mitigated. While the details of each solution are somewhat different, they all accomplish the goal of mitigating the threats.

The three approaches are:

Confidentiality Protection:

The weak point with this approach, which is briefly described in Section 6.1, is that the client has to be careful to whom it discloses the access token. What can be done with the token entirely depends on what rights the token entitles the presenter and what constraints it contains. A token could encode the identifier of the client but there are scenarios where the client is not authenticated to the resource server or where the identifier of the client rather represents an application class rather than a single application instance. As such, it is possible that certain deployments choose a rather liberal approach to security and that everyone who is in possession of the access token is granted access to the data.

Sender Constraint:

The weak point with this approach, which is briefly described in Section 6.2, is to setup the authentication infrastructure such that clients can be authenticated towards resource servers. Additionally, the authorization server must encode the identifier of the client in the token for later verification by the resource server. Depending on the chosen layer for providing client-side authentication there may be additional challenges due to Web server load balancing, lack of API access to identity information, etc.

Key Confirmation:

The weak point with this approach, see Section 6.3, is the increased complexity: a complete key distribution protocol has to be defined.

In all cases above it has to be ensured that the client is able to keep the credentials secret.

7. Architecture

The proof-of-possession security concept assumes that the authorization server acts as a trusted third party that binds keys to access tokens. These keys are then used by the client to demonstrate the possession of the secret to the resource server when accessing the resource. The resource server, when receiving an access token, needs to verify that the key used by the client matches the one included in the access token.

There are slight differences between the use of symmetric keys and asymmetric keys when they are bound to the access token and the subsequent interaction between the client and the authorization server when demonstrating possession of these keys. Figure 1 shows the symmetric key procedure and Figure 2 illustrates how asymmetric keys are used. While symmetric cryptography provides better performance properties the use of asymmetric cryptography allows the client to keep the private key locally and never expose it to any other party.

For example, with the JSON Web Token (JWT) [RFC7519] a standardized format for access tokens is available. The necessary elements to bind symmetric or asymmetric keys to a JWT are described in [I-D.ietf-oauth-proof-of-possession].

Note: The negotiation of cryptographic algorithms between the client and the authorization server is not shown in the examples below and

assumed to be present in a protocol solution to meet the requirements for crypto-agility.

7.1. Client and Authorization Server Interaction

7.1.1. Symmetric Keys




Figure 1: Interaction between the Client and the Authorization Server (Symmetric Keys).

In order to request an access token the client interacts with the authorization server as part of the a normal grant exchange, as shown in Figure 1. However, it needs to include additional information elements for use with the PoP security mechanism, as depicted in message (I). In message (II) the authorization server then returns the requested access token. In addition to the access token itself, the symmetric key is communicated to the client. This symmetric key is a unique and fresh session key with sufficient entropy for the given lifetime. Furthermore, information within the access token ties it to this specific symmetric key.

Note: For this security mechanism to work the client as well as the resource server need to have access to the session key. While the key transport mechanism from the authorization server to the client has been explained in the previous paragraph there are three ways for communicating this session key from the authorization server to the resource server, namely

Embedding the symmetric key inside the access token itself. This requires that the symmetric key is confidentiality protected.

The resource server queries the authorization server for the symmetric key. This is an approach envisioned by the token introspection endpoint [RFC7662].

The authorization server and the resource server both have access to the same back-end database. Smaller, tightly coupled systems might prefer such a deployment strategy.

7.1.2. Asymmetric Keys

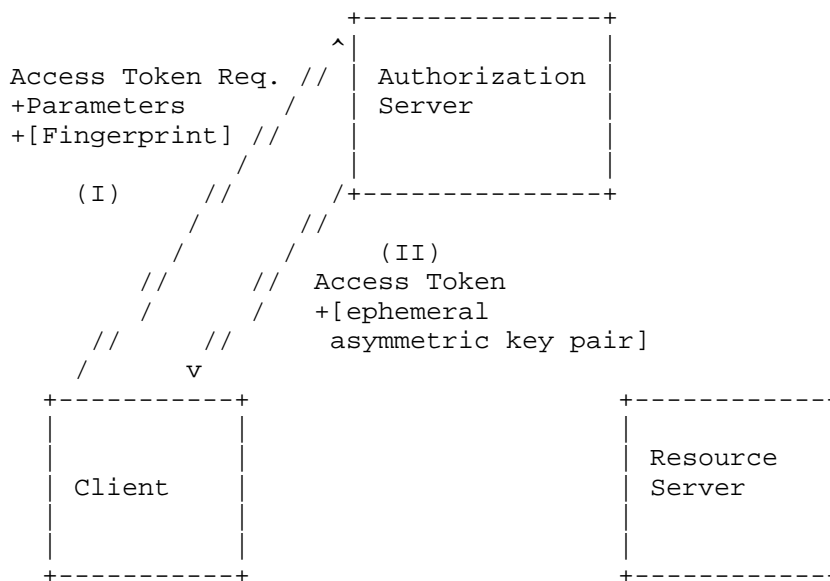


Figure 2: Interaction between the Client and the Authorization Server (Asymmetric Keys).

The use of asymmetric keys is slightly different since the client or the server could be involved in the generation of the ephemeral key pair. This exchange is shown in Figure 1. If the client generates the key pair it either includes a fingerprint of the public key or the public key in the request to the authorization server. The authorization server would include this fingerprint or public key in the confirmation claim inside the access token and thereby bind the asymmetric key pair to the token. If the client did not provide a fingerprint or a public key in the request then the authorization server is asked to create an ephemeral asymmetric key pair, binds the fingerprint of the public key to the access token, and returns the

asymmetric key pair (public and private key) to the client. Note that there is a strong preference for generating the private/public key pair locally at the client rather than at the server.

7.2. Client and Resource Server Interaction

The specification describing the interaction between the client and the authorization server, as shown in Figure 1 and in Figure 2, can be found in [I-D.ietf-oauth-pop-key-distribution].

Once the client has obtained the necessary access token and keying material it can start to interact with the resource server. To demonstrate possession of the key bound to the access token it needs to apply this key to the request by computing a keyed message digest (i.e., a symmetric key-based cryptographic primitive) or a digital signature (i.e., an asymmetric cryptographic computation). When the resource server receives the request it verifies it and decides whether access to the protected resource can be granted. This exchange is shown in Figure 3.




Figure 3: Client Demonstrates PoP.

The specification describing the ability to sign the HTTP request from the client to the resource server can be found in [I-D.ietf-oauth-signed-http-request].

7.3. Resource and Authorization Server Interaction (Token Introspection)

So far the examples talked about access tokens that are passed by value and allow the resource server to make authorization decisions immediately after verifying the request from the client. In some deployments a real-time interaction between the authorization server and the resource server is envisioned that lowers the need to pass self-contained access tokens around. In that case the access token merely serves as a handle or a reference to state stored at the authorization server. As a consequence, the resource server cannot autonomously make an authorization decision when receiving a request

from a client but has to consult the authorization server. This can, for example, be done using the token introspection endpoint (see [RFC7662]). Figure 4 shows the protocol interaction graphically. Despite the additional token exchange previous descriptions about associating symmetric and asymmetric keys to the access token are still applicable to this scenario.




Figure 4: Token Introspection and Access Token Handles.

8. Security Considerations

The purpose of this document is to provide use cases, requirements, and motivation for developing an OAuth security solution extending Bearer Tokens. As such, this document is only about security.

9. IANA Considerations

This document does not require actions by IANA.

10. Acknowledgments

This document is the result of conference calls late 2012/early 2013 and in design team conference calls February 2013 of the IETF OAuth working group. The following persons (in addition to the OAuth WG chairs, Hannes Tschofenig, and Derek Atkins) provided their input during these calls: Bill Mills, Justin Richer, Phil Hunt, Prateek Mishra, Mike Jones, George Fletcher, Leif Johansson, Lucy Lynch, John

Bradley, Tony Nadalin, Klaas Wierenga, Thomas Hardjono, Brian Campbell

In the appendix of this document we reuse content from [RFC4962] and the authors would like thank Russ Housely and Bernard Aboba for their work on RFC 4962.

We would like to thank Reddy Tirumaleswar for his review.

11. References

11.1. Normative References

- [I-D.ietf-oauth-pop-key-distribution]
Bradley, J., Hunt, P., Jones, M., and H. Tschofenig, "OAuth 2.0 Proof-of-Possession: Authorization Server to Client Key Distribution", draft-ietf-oauth-pop-key-distribution-02 (work in progress), October 2015.
- [I-D.ietf-oauth-proof-of-possession]
Jones, M., Bradley, J., and H. Tschofenig, "Proof-of-Possession Key Semantics for JSON Web Tokens (JWTs)", draft-ietf-oauth-proof-of-possession-11 (work in progress), December 2015.
- [I-D.ietf-oauth-signed-http-request]
Richer, J., Bradley, J., and H. Tschofenig, "A Method for Signing HTTP Requests for OAuth", draft-ietf-oauth-signed-http-request-02 (work in progress), February 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<http://www.rfc-editor.org/info/rfc7519>>.

- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<http://www.rfc-editor.org/info/rfc7662>>.

11.2. Informative References

- [I-D.hardjono-oauth-kerberos]
Hardjono, T., "OAuth 2.0 support for the Kerberos V5 Authentication Protocol", draft-hardjono-oauth-kerberos-01 (work in progress), December 2010.
- [NIST800-63]
Burr, W., Dodson, D., Perlner, R., Polk, T., Gupta, S., and E. Nabbus, "NIST Special Publication 800-63-1, INFORMATION SECURITY", December 2008.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, DOI 10.17487/RFC4120, July 2005, <<http://www.rfc-editor.org/info/rfc4120>>.
- [RFC4279] Eronen, P., Ed. and H. Tschofenig, Ed., "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)", RFC 4279, DOI 10.17487/RFC4279, December 2005, <<http://www.rfc-editor.org/info/rfc4279>>.
- [RFC4347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, DOI 10.17487/RFC4347, April 2006, <<http://www.rfc-editor.org/info/rfc4347>>.
- [RFC4962] Housley, R. and B. Aboba, "Guidance for Authentication, Authorization, and Accounting (AAA) Key Management", BCP 132, RFC 4962, DOI 10.17487/RFC4962, July 2007, <<http://www.rfc-editor.org/info/rfc4962>>.
- [RFC5056] Williams, N., "On the Use of Channel Bindings to Secure Channels", RFC 5056, DOI 10.17487/RFC5056, November 2007, <<http://www.rfc-editor.org/info/rfc5056>>.
- [RFC5849] Hammer-Lahav, E., Ed., "The OAuth 1.0 Protocol", RFC 5849, DOI 10.17487/RFC5849, April 2010, <<http://www.rfc-editor.org/info/rfc5849>>.

- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<http://www.rfc-editor.org/info/rfc6125>>.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<http://www.rfc-editor.org/info/rfc6750>>.
- [RFC6819] Lodderstedt, T., Ed., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, DOI 10.17487/RFC6819, January 2013, <<http://www.rfc-editor.org/info/rfc6819>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.

Authors' Addresses

Phil Hunt (editor)
Oracle Corporation

Email: phil.hunt@yahoo.com

Justin Richer

Email: ietf@justin.richer.org

William Mills

Email: wmills@yahoo-inc.com

Prateek Mishra
Oracle Corporation

Email: prateek.mishra@oracle.com

Hannes Tschofenig
ARM Limited
Hall in Tirol 6060
Austria

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 28, 2019

J. Bradley
Ping Identity
P. Hunt
Oracle Corporation
M. Jones
Microsoft
H. Tschofenig
Arm Ltd.
M. Meszaros
GITDA
March 27, 2019

OAuth 2.0 Proof-of-Possession: Authorization Server to Client Key
Distribution
draft-ietf-oauth-pop-key-distribution-07

Abstract

RFC 6750 specified the bearer token concept for securing access to protected resources. Bearer tokens need to be protected in transit as well as at rest. When a client requests access to a protected resource it hands-over the bearer token to the resource server.

The OAuth 2.0 Proof-of-Possession security concept extends bearer token security and requires the client to demonstrate possession of a key when accessing a protected resource.

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 <https://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 September 28, 2019.

Copyright Notice

Copyright (c) 2019 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 (https://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

1. Introduction 2
2. Terminology 4
3. Processing Instructions 4
4. Examples 5
4.1. Symmetric Key Transport 5
4.1.1. Client-to-AS Request 5
4.1.2. Client-to-AS Response 6
4.2. Asymmetric Key Transport 9
4.2.1. Client-to-AS Request 9
4.2.2. Client-to-AS Response 10
5. Security Considerations 11
6. IANA Considerations 13
6.1. OAuth Access Token Types 13
6.2. OAuth Parameters Registration 13
6.3. OAuth Extensions Error Registration 13
7. Acknowledgements 13
8. References 14
8.1. Normative References 14
8.2. Informative References 15
Authors' Addresses 16

1. Introduction

The work on proof-of-possession tokens, an extended token security mechanisms for OAuth 2.0, is motivated in [22]. This document defines the ability for the client request and to obtain PoP tokens from the authorization server. After successfully completing the exchange the client is in possession of a PoP token and the keying material bound to it. Clients that access protected resources then need to demonstrate knowledge of the secret key that is bound to the PoP token.

To best describe the scope of this specification, the OAuth 2.0 protocol exchange sequence is shown in Figure 1. The extension defined in this document piggybacks on the message exchange marked with (C) and (D). To demonstrate possession of the private/secret key to the resource server protocol mechanisms outside the scope of this document are used.

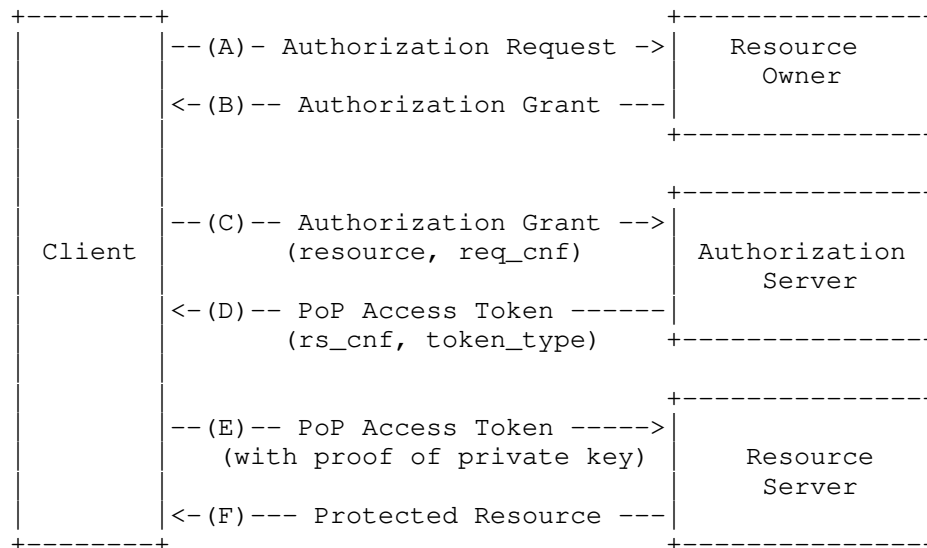


Figure 1: Augmented OAuth 2.0 Protocol Flow

In OAuth 2.0 [2] access tokens can be obtained via authorization grants and using refresh tokens. The core OAuth specification defines four authorization grants, see Section 1.3 of [2], and [19] adds an assertion-based authorization grant to that list. The token endpoint, which is described in Section 3.2 of [2], is used with every authorization grant except for the implicit grant type. In the implicit grant type the access token is issued directly.

This specification extends the functionality of the token endpoint, i.e., the protocol exchange between the client and the authorization server, to allow keying material to be bound to an access token. Two types of keying material can be bound to an access token, namely symmetric keys and asymmetric keys. Conveying symmetric keys from the authorization server to the client is described in Section 4.1 and the procedure for dealing with asymmetric keys is described in Section 4.2.

This document describes how the client requests and obtains a PoP access token from the authorization server for use with HTTPS-based

transport. The use of alternative transports, such as Constrained Application Protocol (CoAP), is described in [24].

2. Terminology

The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in this specification are to be interpreted as described in [1].

Session Key:

In the context of this specification 'session key' refers to fresh and unique keying material established between the client and the resource server. This session key has a lifetime that corresponds to the lifetime of the access token, is generated by the authorization server and bound to the access token.

This document uses the following abbreviations:

JWA: JSON Web Algorithms[7]

JWT: JSON Web Token[9]

JWS: JSON Web Signature[6]

JWK: JSON Web Key[5]

JWE: JSON Web Encryption[8]

CWT: CBOR Web Token[13]

COSE: CBOR Object Signing and Encryption[14]

3. Processing Instructions

Step (0): As an initial step the client typically determines the resource server it wants to interact with. This may, for example, happen as part of a discovery procedure or via manual configuration.

Step (1): The client starts the OAuth 2.0 protocol interaction based on the selected grant type.

Step (2): When the client interacts with the token endpoint to obtain an access token it MUST use the resource identifier parameter, defined in [16], or the audience parameter, defined in [15], when symmetric PoP tokens are used. For asymmetric PoP tokens the use of resource indicators and audience is optional but

RECOMMENDED. The parameters 'audience' and 'resource' both allow the client to express the location of the target service and the difference between the two is described in [15]. As a summary, 'audience' allows expressing a logical name while 'resource' contains an absolute URI. More details about the 'resource' parameter can be found in [16].

Step (3): The authorization server parses the request from the server and determines the suitable response based on OAuth 2.0 and the PoP token credential procedures.

Note that PoP access tokens may be encoded in a variety of ways:

JWT The access token may be encoded using the JSON Web Token (JWT) format [9]. The proof-of-possession token functionality is described in [10]. A JWT encoded PoP token MUST be protected against modification by either using a digital signature or a keyed message digest, as described in [6]. The JWT may also be encrypted using [8].

CWT [13] defines an alternative token format based on CBOR. The proof-of-possession token functionality is defined in [12]. A CWT encoded PoP token MUST be protected against modification by either using a digital signature or a keyed message digest, as described in [12].

If the access token is only a reference then a look-up by the resource server is needed, as described in the token introspection specification [23].

Note that the OAuth 2.0 framework nor this specification does not mandate a specific PoP token format but using a standardized format will improve interoperability and will lead to better code re-use.

Application layer interactions between the client and the resource server are beyond the scope of this document.

4. Examples

This section provides a number of examples.

4.1. Symmetric Key Transport

4.1.1. Client-to-AS Request

The client starts with a request to the authorization server indicating that it is interested to obtain a token for `https://resource.example.com`

```
POST /token HTTP/1.1
Host: authz.example.com
Authorization: Basic czZCaGRSa3F0MzpnWDFmQmF0M2JW
Content-Type: application/x-www-form-urlencoded;charset=UTF-8

grant_type=authorization_code
&code=Sp1xl0BeZQQYbYS6WxSbIA
&scope=calendar%20contacts
&redirect_uri=https%3A%2F%2Fclient%2Eexample%2Ecom%2Fcb
&resource=https%3A%2F%2Fresource.example.com
```

Example Request to the Authorization Server

4.1.2. Client-to-AS Response

If the access token request has been successfully verified by the authorization server and the client is authorized to obtain a PoP token for the indicated resource server, the authorization server issues an access token and optionally a refresh token.

Figure 2 shows a response containing a token and a "cnf" parameter with a symmetric proof-of-possession key both encoded in a JSON-based serialization format. The "cnf" parameter contains the RFC 7517 [5] encoded key element.

```
HTTP/1.1 200 OK
Content-Type: application/json
Cache-Control: no-store

{
  "access_token":"SlAV32hkKG ...
  (remainder of JWT omitted for brevity;
  JWT contains JWK in the cnf claim)",
  "token_type":"pop",
  "expires_in":3600,
  "refresh_token":"8xLOxBtZp8",
  "cnf":{
    {"keys":
      [
        {"kty":"oct",
         "alg":"A128KW",
         "k":"GawgguFyGrWKav7AX4VKUg"}
      ]
    }
  }
}
```

Figure 2: Example: Response from the Authorization Server (Symmetric Variant)

Note that the cnf payload in Figure 2 is not encrypted at the application layer since Transport Layer Security is used between the AS and the client and the content of the cnf payload is consumed by the client itself. Alternatively, a JWE could be used to encrypt the key distribution, as shown in Figure 3.


```
{
  "access_token":"SlAV32hkKG ...
    (remainder of JWT omitted for brevity;
    JWT contains JWK in the cnf claim)",
  "token_type":"pop",
  "expires_in":3600,
  "refresh_token":"8xLOxBtZp8",
  "cnf":{
    "jwe":
      "eyJhbGciOiJSU0EtT0FFUCIsImVuYyI6IkJkExMjhdQkMtSFMyNTYifQ.
      (remainder of JWE omitted for brevity)"
    }
  }
}
```

Figure 3: Example: Encrypted Symmetric Key

The content of the 'access_token' in JWT format contains the 'cnf' (confirmation) claim. The confirmation claim is defined in [10]. The digital signature or the keyed message digest offering integrity protection is not shown in this example but has to be present in a real deployment to mitigate a number of security threats.

The JWK in the key element of the response from the authorization server, as shown in Figure 2, contains the same session key as the JWK inside the access token, as shown in Figure 4. It is, in this example, protected by TLS and transmitted from the authorization server to the client (for processing by the client).

```
{
  "iss": "https://server.example.com",
  "sub": "24400320",
  "aud": "s6BhdRkqt3",
  "exp": 1311281970,
  "iat": 1311280970,
  "cnf":{
    "jwe":
      "eyJhbGciOiJSU0EtT0FFUCIsImVuYyI6IkJkExMjhdQkMtSFMyNTYifQ.
      (remainder of JWE omitted for brevity)"
    }
  }
}
```

Figure 4: Example: Access Token in JWT Format

Note: When the JWK inside the access token contains a symmetric key it must be confidentiality protected using a JWE to maintain the security goals of the PoP architecture since content is meant for consumption by the selected resource server only. The details are described in [22].

4.2. Asymmetric Key Transport

4.2.1. Client-to-AS Request

This example illustrates the case where an asymmetric key shall be bound to an access token. The client makes the following HTTPS request shown in Figure 5. Extra line breaks are for display purposes only.

```
POST /token HTTP/1.1
Host: server.example.com
Authorization: Basic czZCaGRSa3F0MzpnWDFmQmF0M2JW
Content-Type: application/x-www-form-urlencoded; charset=UTF-8

grant_type=authorization_code
&code=Sp1xl0BeZQQYbYS6WxSbIA
&redirect_uri=https%3A%2F%2Fclient%2Eexample%2Ecom%2Fcb
&token_type=pop
&req_cnf=eyJhbGciOiJSU0ExXzUi ...
(remainder of JWK omitted for brevity)
```

Figure 5: Example Request to the Authorization Server (Asymmetric Key Variant)

As shown in Figure 6 the content of the 'req_cnf' parameter contains the ECC public key the client would like to associate with the access token (in JSON format).

```
"jwk":{
  "kty": "EC",
  "use": "sig",
  "crv": "P-256",
  "x": "18wHLeIgW9wVN6VD1Txgpqy2LszYkMf6J8njVAibvhM",
  "y": "-V4dS4UaLMgP_4fY4j8ir7cl1TX1FdAgcx55o7TkcSA"
}
```

Figure 6: Client Providing Public Key to Authorization Server

4.2.2. Client-to-AS Response

If the access token request is valid and authorized, the authorization server issues an access token and optionally a refresh token. The authorization server also places information about the public key used by the client into the access token to create the binding between the two. The new token type "pop" is placed into the 'token_type' parameter.

An example of a successful response is shown in Figure 7.

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache

{
  "access_token":"2YotnFZFE....jrlzCsicMWpAA",
  "token_type":"pop",
  "expires_in":3600,
  "refresh_token":"tGzv3JOkF0XG5Qx2TlKWIA"
}
```

Figure 7: Example: Response from the Authorization Server (Asymmetric Variant)

The content of the 'access_token' field contains an encoded JWT, as shown in Figure 8. The digital signature covering the access token offering authenticity and integrity protection is not shown below (but must be present).

```
{
  "iss": "https://authz.example.com",
  "aud": "https://resource.example.com",
  "exp": "1361398824",
  "nbf": "1360189224",
  "cnf": {
    "jwk": {
      "kty": "EC",
      "crv": "P-256",
      "x": "usWxHK2PmfHkXPS54m0kTcGJ90UiglWiGahtagnv8",
      "y": "IBOL+C3BttVivg+lSreASjpkttcsz+lrb7btKLv8EX4"
    }
  }
}
```

Figure 8: Example: Access Token Structure (Asymmetric Variant)

Note: In this example there is no need for the authorization server to convey further keying material to the client since the client is already in possession of the private key (as well as the public key).

5. Security Considerations

[22] describes the architecture for the OAuth 2.0 proof-of-possession security architecture, including use cases, threats, and requirements. This requirements describes one solution component of that architecture, namely the mechanism for the client to interact with the authorization server to either obtain a symmetric key from the authorization server, to obtain an asymmetric key pair, or to offer a public key to the authorization. In any case, these keys are then bound to the access token by the authorization server.

To summarize the main security recommendations: A large range of threats can be mitigated by protecting the contents of the access token by using a digital signature or a keyed message digest. Consequently, the token integrity protection MUST be applied to prevent the token from being modified, particularly since it contains a reference to the symmetric key or the asymmetric key. If the access token contains the symmetric key (see Section 2.2 of [10] for a description about how symmetric keys can be securely conveyed within the access token) this symmetric key MUST be encrypted by the authorization server with a long-term key shared with the resource server.

To deal with token redirect, it is important for the authorization server to include the identity of the intended recipient (the audience), typically a single resource server (or a list of resource servers), in the token. Using a single shared secret with multiple

authorization server to simplify key management is NOT RECOMMENDED since the benefit from using the proof-of-possession concept is significantly reduced.

Token replay is also not possible since an eavesdropper will also have to obtain the corresponding private key or shared secret that is bound to the access token. Nevertheless, it is good practice to limit the lifetime of the access token and therefore the lifetime of associated key.

The authorization server MUST offer confidentiality protection for any interactions with the client. This step is extremely important since the client will obtain the session key from the authorization server for use with a specific access token. Not using confidentiality protection exposes this secret (and the access token) to an eavesdropper thereby making the OAuth 2.0 proof-of-possession security model completely insecure. OAuth 2.0 [2] relies on TLS to offer confidentiality protection and additional protection can be applied using the JWK [5] offered security mechanism, which would add an additional layer of protection on top of TLS for cases where the keying material is conveyed, for example, to a hardware security module. Which version(s) of TLS ought to be implemented will vary over time, and depend on the widespread deployment and known security vulnerabilities at the time of implementation. At the time of this writing, TLS version 1.2 [4] is the most recent version. The client MUST validate the TLS certificate chain when making requests to protected resources, including checking the validity of the certificate.

Similarly to the security recommendations for the bearer token specification [17] developers MUST ensure that the ephemeral credentials (i.e., the private key or the session key) is not leaked to third parties. An adversary in possession of the ephemeral credentials bound to the access token will be able to impersonate the client. Be aware that this is a real risk with many smart phone app and Web development environments.

Clients can at any time request a new proof-of-possession capable access token. Using a refresh token to regularly request new access tokens that are bound to fresh and unique keys is important. Keeping the lifetime of the access token short allows the authorization server to use shorter key sizes, which translate to a performance benefit for the client and for the resource server. Shorter keys also lead to shorter messages (particularly with asymmetric keying material).

When authorization servers bind symmetric keys to access tokens then they SHOULD scope these access tokens to a specific permissions.

6. IANA Considerations

6.1. OAuth Access Token Types

This specification registers the following error in the IANA "OAuth Access Token Types" [25] established by [17].

- o Name: pop
- o Change controller: IESG
- o Specification document(s): [[this specification]]

6.2. OAuth Parameters Registration

This specification registers the following value in the IANA "OAuth Parameters" registry [25] established by [2].

- o Parameter name: cnf_req
- o Parameter usage location: authorization request, token request
- o Change controller: IESG
- o Specification document(s): [[this specification]]

- o Parameter name: cnf
- o Parameter usage location: authorization response, token response
- o Change controller: IESG
- o Specification document(s): [[this specification]]

- o Parameter name: rs_cnf
- o Parameter usage location: token response
- o Change controller: IESG
- o Specification document(s): [[this specification]]

6.3. OAuth Extensions Error Registration

This specification registers the following error in the IANA "OAuth Extensions Error Registry" [25] established by [2].

- o Error name: invalid_token_type
- o Error usage location: implicit grant error response, token error response
- o Related protocol extension: token_type parameter
- o Change controller: IESG
- o Specification document(s): [[this specification]]

7. Acknowledgements

We would like to thank Chuck Mortimore and James Manger for their review comments.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [2] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.
- [3] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [4] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [5] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.
- [6] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [7] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <<https://www.rfc-editor.org/info/rfc7518>>.
- [8] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<https://www.rfc-editor.org/info/rfc7516>>.
- [9] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [10] Jones, M., Bradley, J., and H. Tschofenig, "Proof-of-Possession Key Semantics for JSON Web Tokens (JWTs)", RFC 7800, DOI 10.17487/RFC7800, April 2016, <<https://www.rfc-editor.org/info/rfc7800>>.

- [11] Jones, M. and N. Sakimura, "JSON Web Key (JWK) Thumbprint", RFC 7638, DOI 10.17487/RFC7638, September 2015, <<https://www.rfc-editor.org/info/rfc7638>>.
- [12] Jones, M., Seitz, L., Selander, G., Erdtman, S., and H. Tschofenig, "Proof-of-Possession Key Semantics for CBOR Web Tokens (CWTs)", draft-ietf-ace-cwt-proof-of-possession-06 (work in progress), February 2019.
- [13] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/info/rfc8392>>.
- [14] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.
- [15] Jones, M., Nadalin, A., Campbell, B., Bradley, J., and C. Mortimore, "OAuth 2.0 Token Exchange", draft-ietf-oauth-token-exchange-16 (work in progress), October 2018.
- [16] Campbell, B., Bradley, J., and H. Tschofenig, "Resource Indicators for OAuth 2.0", draft-ietf-oauth-resource-indicators-02 (work in progress), January 2019.

8.2. Informative References

- [17] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<https://www.rfc-editor.org/info/rfc6750>>.
- [18] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [19] Campbell, B., Mortimore, C., Jones, M., and Y. Goland, "Assertion Framework for OAuth 2.0 Client Authentication and Authorization Grants", RFC 7521, DOI 10.17487/RFC7521, May 2015, <<https://www.rfc-editor.org/info/rfc7521>>.
- [20] Sakimura, N., Ed., Bradley, J., and N. Agarwal, "Proof Key for Code Exchange by OAuth Public Clients", RFC 7636, DOI 10.17487/RFC7636, September 2015, <<https://www.rfc-editor.org/info/rfc7636>>.

- [21] Richer, J., Ed., Jones, M., Bradley, J., Machulak, M., and P. Hunt, "OAuth 2.0 Dynamic Client Registration Protocol", RFC 7591, DOI 10.17487/RFC7591, July 2015, <<https://www.rfc-editor.org/info/rfc7591>>.
- [22] Hunt, P., Richer, J., Mills, W., Mishra, P., and H. Tschofenig, "OAuth 2.0 Proof-of-Possession (PoP) Security Architecture", draft-ietf-oauth-pop-architecture-08 (work in progress), July 2016.
- [23] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<https://www.rfc-editor.org/info/rfc7662>>.
- [24] Seitz, L., Selander, G., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "Authentication and Authorization for Constrained Environments (ACE) using the OAuth 2.0 Framework (ACE-OAuth)", draft-ietf-ace-oauth-authz-24 (work in progress), March 2019.
- [25] IANA, "OAuth Parameters", October 2018.
- [26] IANA, "JSON Web Token Claims", June 2018.

Authors' Addresses

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Phil Hunt
Oracle Corporation

Email: phil.hunt@yahoo.com
URI: <http://www.independentid.com>

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

Hannes Tschofenig
Arm Ltd.
Absam 6067
Austria

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Mihaly Meszaros
GITDA
Debrecen 4033
Hungary

Email: bakfitty@gmail.com
URI: <https://github.com/misi>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 20, 2016

M. Jones
Microsoft
J. Bradley
Ping Identity
H. Tschofenig
ARM Limited
December 18, 2015

Proof-of-Possession Key Semantics for JSON Web Tokens (JWTs)
draft-ietf-oauth-proof-of-possession-11

Abstract

This specification defines how to declare in a JSON Web Token (JWT) that the presenter of the JWT possesses a particular proof-of-possession key and that the recipient can cryptographically confirm proof-of-possession of the key by the presenter. Being able to prove possession of a key is also sometimes described as the presenter being a holder-of-key.

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 June 20, 2016.

Copyright Notice

Copyright (c) 2015 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

1. Introduction	3
1.1. Notational Conventions	5
2. Terminology	5
3. Representations for Proof-of-Possession Keys	6
3.1. Confirmation Claim	6
3.2. Representation of an Asymmetric Proof-of-Possession Key	7
3.3. Representation of an Encrypted Symmetric Proof-of-Possession Key	8
3.4. Representation of a Key ID for a Proof-of-Possession Key	9
3.5. Representation of a URL for a Proof-of-Possession Key	9
3.6. Specifics Intentionally Not Specified	10
4. Security Considerations	10
5. Privacy Considerations	11
6. IANA Considerations	11
6.1. JSON Web Token Claims Registration	12
6.1.1. Registry Contents	12
6.2. JWT Confirmation Methods Registry	12
6.2.1. Registration Template	12
6.2.2. Initial Registry Contents	13
7. References	13
7.1. Normative References	13
7.2. Informative References	14
Appendix A. Acknowledgements	15
Appendix B. Document History	15
Authors' Addresses	17

1. Introduction

This specification defines how a JSON Web Token [JWT] can declare that the presenter of the JWT possesses a particular proof-of-possession (PoP) key and that the recipient can cryptographically confirm proof-of-possession of the key by the presenter. Proof-of-possession of a key is also sometimes described as the presenter being a holder-of-key. The [I-D.ietf-oauth-pop-architecture] specification describes key confirmation, among other confirmation mechanisms. This specification defines how to communicate key confirmation key information in JWTs.

Envision the following two use cases. The first use case employs a symmetric proof-of-possession key and the second use case employs an asymmetric proof-of-possession key.




Figure 1: Proof-of-Possession with a Symmetric Key

In the case illustrated in Figure 1, either the presenter generates a symmetric key and privately sends it to the issuer (1) or the issuer generates a symmetric key and privately sends it to the presenter (1). The issuer generates a JWT with an encrypted copy of this symmetric key in the confirmation claim. This symmetric key is

encrypted with a key known only to the issuer and the recipient, which was previously established in step (0). The entire JWT is integrity protected by the issuer. The JWT is then (2) sent to the presenter. Now, the presenter is in possession of the symmetric key as well as the JWT (which includes the confirmation claim). When the presenter (3) presents the JWT to the recipient, it also needs to demonstrate possession of the symmetric key; the presenter, for example, (4) uses the symmetric key in a challenge/response protocol with the recipient. The recipient is then able to verify that it is interacting with the genuine presenter by decrypting the key in the confirmation claim of the JWT. By doing this, the recipient obtains the symmetric key, which it then uses to verify cryptographically protected messages exchanged with the presenter (4). This symmetric key mechanism described above is conceptually similar to the use of Kerberos tickets.

Note that for simplicity, the diagram above and associated text describe the direct use of symmetric keys without the use of derived keys. A more secure practice is to derive the symmetric keys actually used from secrets exchanged, such as the key exchanged in step (0), using a Key Derivation Function (KDF) and use the derived keys, rather than directly using the secrets exchanged.

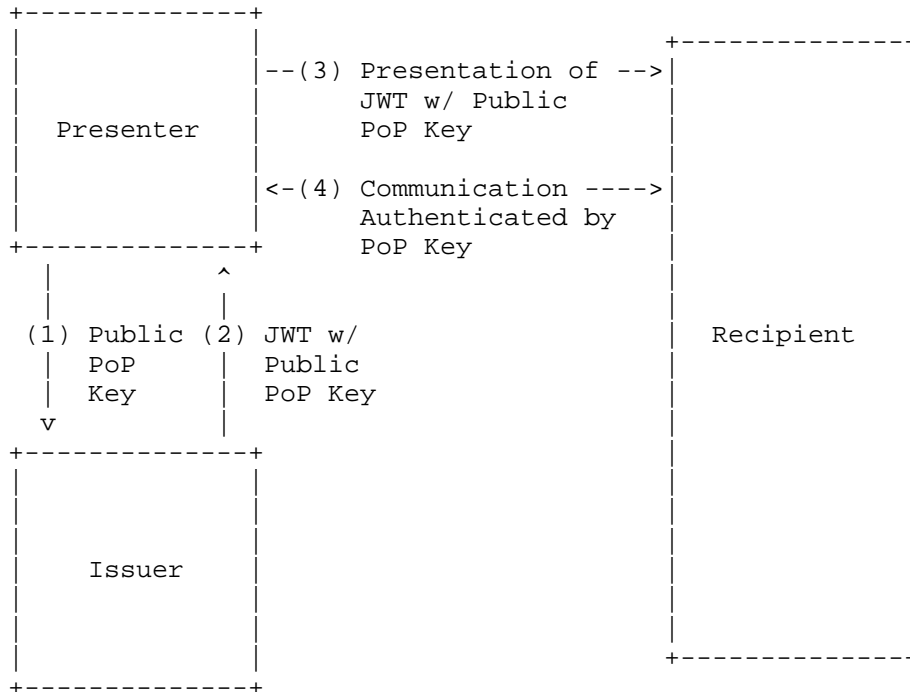


Figure 2: Proof-of-Possession with an Asymmetric Key

In the case illustrated in Figure 2, the presenter generates a public/private key pair and (1) sends the public key to the issuer, which creates a JWT that contains the public key (or an identifier for it) in the confirmation claim. The entire JWT is integrity protected using a digital signature to protect it against modifications. The JWT is then (2) sent to the presenter. When the presenter (3) presents the JWT to the recipient, it also needs to demonstrate possession of the private key. The presenter, for example, (4) uses the private key in a TLS exchange with the recipient or (4) signs a nonce with the private key. The recipient is able to verify that it is interacting with the genuine presenter by extracting the public key from the confirmation claim of the JWT (after verifying the digital signature of the JWT) and utilizing it with the private key in the TLS exchange or by checking the nonce signature.

In both cases, the JWT may contain other claims that are needed by the application.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Unless otherwise noted, all the protocol parameter names and values are case sensitive.

2. Terminology

This specification uses terms defined in the JSON Web Token [JWT], JSON Web Key [JWK], and JSON Web Encryption [JWE] specifications.

These terms are defined by this specification:

Issuer

Party that creates the JWT and binds the proof-of-possession key to it.

Presenter

Party that proves possession of a private key (for asymmetric key cryptography) or secret key (for symmetric key cryptography) to a recipient.

Recipient

Party that receives the JWT containing the proof-of-possession key information from the presenter.

3. Representations for Proof-of-Possession Keys

By including a "cnf" (confirmation) claim in a JWT, the issuer of the JWT declares that the presenter possesses a particular key, and that the recipient can cryptographically confirm that the presenter has possession of that key. The value of the "cnf" claim is a JSON object and the members of that object identify the proof-of-possession key.

The presenter can be identified in one of several ways by the JWT, depending upon the application requirements. If the JWT contains a "sub" (subject) claim [JWT], the presenter is normally the subject identified by the JWT. (In some applications, the subject identifier will be relative to the issuer identified by the "iss" (issuer) claim [JWT].) If the JWT contains no "sub" (subject) claim, the presenter is normally the issuer identified by the JWT using the "iss" (issuer) claim. The case in which the presenter is the subject of the JWT is analogous to SAML 2.0 [OASIS.saml-core-2.0-os] SubjectConfirmation usage. At least one of the "sub" and "iss" claims MUST be present in the JWT. Some use cases may require that both be present.

Another means used by some applications to identify the presenter is an explicit claim, such as the "azp" (authorized party) claim defined by OpenID Connect [OpenID.Core]. Ultimately, the means of identifying the presenter is application-specific, as is the means of confirming possession of the key that is communicated.

3.1. Confirmation Claim

The "cnf" (confirmation) claim is used in the JWT to contain members used to identify the proof-of-possession key. Other members of the "cnf" object may be defined because a proof-of-possession key may not be the only means of confirming the authenticity of the token. This is analogous to the SAML 2.0 [OASIS.saml-core-2.0-os] SubjectConfirmation element, in which a number of different subject confirmation methods can be included, including proof-of-possession key information.

The set of confirmation members that a JWT must contain to be considered valid is context dependent and is outside the scope of this specification. Specific applications of JWTs will require implementations to understand and process some confirmation members in particular ways. However, in the absence of such requirements,

all confirmation members that are not understood by implementations MUST be ignored.

This specification establishes the IANA "JWT Confirmation Methods" registry for these members in Section 6.2 and registers the members defined by this specification. Other specifications can register other members used for confirmation, including other members for conveying proof-of-possession keys, possibly using different key representations.

The "cnf" claim value MUST represent only a single proof-of-possession key; thus, at most one of the "jwk", "jwe", and "jku" confirmation values defined below may be present. Note that if an application needs to represent multiple proof-of-possession keys in the same JWT, one way for it to achieve this is to use other claim names, in addition to "cnf", to hold the additional proof-of-possession key information. These claims could use the same syntax and semantics as the "cnf" claim. Those claims would be defined by applications or other specifications and could be registered in the IANA "JSON Web Token Claims" registry [IANA.JWT.Claims].

3.2. Representation of an Asymmetric Proof-of-Possession Key

When the key held by the presenter is an asymmetric private key, the "jwk" member is a JSON Web Key [JWK] representing the corresponding asymmetric public key. The following example demonstrates such a declaration in the JWT Claims Set of a JWT:

```
{
  "iss": "https://server.example.com",
  "aud": "https://client.example.org",
  "exp": 1361398824,
  "cnf": {
    "jwk": {
      "kty": "EC",
      "use": "sig",
      "crv": "P-256",
      "x": "18wHLeIgW9wVN6VD1Txgppy2LszYkMf6J8njVAibvhM",
      "y": "-V4dS4UaLMgP_4fY4j8ir7cl1TXlFdAgcx55o7TkcSA"
    }
  }
}
```

The JWK MUST contain the required key members for a JWK of that key type and MAY contain other JWK members, including the "kid" (key ID) member.

The "jwk" member MAY also be used for a JWK representing a symmetric

key, provided that the JWT is encrypted so that the key is not revealed to unintended parties. If the JWT is not encrypted, the symmetric key MUST be encrypted as described below.

3.3. Representation of an Encrypted Symmetric Proof-of-Possession Key

When the key held by the presenter is a symmetric key, the "jwe" member is an encrypted JSON Web Key [JWK] encrypted to a key known to the recipient using the JWE Compact Serialization containing the symmetric key. The rules for encrypting a JWK are found in Section 7 of the JSON Web Key [JWK] specification.

The following example illustrates a symmetric key that could subsequently be encrypted for use in the "jwe" member:

```
{
  "kty": "oct",
  "alg": "HS256",
  "k": "ZoRSOrFzN_FzUA5XKMYoVHyzff5oRJxl-IXRtztJ6uE"
}
```

The UTF-8 [RFC3629] encoding of this JWK is used as the JWE Plaintext when encrypting the key.

The following example is a JWE Header that could be used when encrypting this key:

```
{
  "alg": "RSA-OAEP",
  "enc": "A128CBC-HS256"
}
```

The following example JWT Claims Set of a JWT illustrates the use of an encrypted symmetric key as the "jwe" member value:

```
{
  "iss": "https://server.example.com",
  "sub": "24400320",
  "aud": "s6BhdRkqt3",
  "nonce": "n-0S6_WzA2Mj",
  "exp": 1311281970,
  "iat": 1311280970,
  "cnf": {
    "jwe":
      "eyJhbGciOiJSU0EtT0FFUCIsImVuYyI6IkJkZXNjbDQkMtsFMjNTYifQ.
      (remainder of JWE omitted for brevity)"
  }
}
```

3.4. Representation of a Key ID for a Proof-of-Possession Key

The proof-of-possession key can also be identified by the use of a Key ID instead of communicating the actual key, provided the recipient is able to obtain the identified key using the Key ID. In this case, the issuer of a JWT declares that the presenter possesses a particular key and that the recipient can cryptographically confirm proof-of-possession of the key by the presenter by including a "cnf" (confirmation) claim in the JWT whose value is a JSON object, with the JSON object containing a "kid" (key ID) member identifying the key.

The following example demonstrates such a declaration in the JWT Claims Set of a JWT:

```
{
  "iss": "https://server.example.com",
  "aud": "https://client.example.org",
  "exp": 1361398824,
  "cnf": {
    "kid": "dfdl1aa97-6d8d-4575-a0fe-34b96de2bfad"
  }
}
```

The content of the "kid" value is application specific. For instance, some applications may choose to use a JWK Thumbprint [JWK.Thumbprint] value as the "kid" value.

3.5. Representation of a URL for a Proof-of-Possession Key

The proof-of-possession key can be passed by reference instead of being passed by value. This is done using the "jku" (JWK Set URL) member. Its value is a URI [RFC3986] that refers to a resource for a set of JSON-encoded public keys represented as a JWK Set [JWK], one of which is the proof-of-possession key. If there are multiple keys in the referenced JWK Set document, a "kid" member MUST also be included, with the referenced key's JWK also containing the same "kid" value.

The protocol used to acquire the resource MUST provide integrity protection. An HTTP GET request to retrieve the JWK Set MUST use Transport Layer Security (TLS) [RFC5246] and the identity of the server MUST be validated, as per Section 6 of RFC 6125 [RFC6125].

The following example demonstrates such a declaration in the JWT Claims Set of a JWT:

```
{
  "iss": "https://server.example.com",
  "sub": "17760704",
  "aud": "https://client.example.org",
  "exp": 1440804813,
  "cnf": {
    "jku": "https://keys.example.net/pop-keys.json",
    "kid": "2015-08-28"
  }
}
```

3.6. Specifics Intentionally Not Specified

Proof-of-possession is typically demonstrated by having the presenter sign a value determined by the recipient using the key possessed by the presenter. This value is sometimes called a "nonce" or a "challenge".

The means of communicating the nonce and the nature of its contents are intentionally not described in this specification, as different protocols will communicate this information in different ways. Likewise, the means of communicating the signed nonce is also not specified, as this is also protocol-specific.

Note that another means of proving possession of the key when it is a symmetric key is to encrypt the key to the recipient. The means of obtaining a key for the recipient is likewise protocol-specific.

For examples using the mechanisms defined in this specification, see [I-D.ietf-oauth-pop-architecture].

4. Security Considerations

All of the security considerations that are discussed in [JWT] also apply here. In addition, proof-of-possession introduces its own unique security issues. Possessing a key is only valuable if it is kept secret. Appropriate means must be used to ensure that unintended parties do not learn private key or symmetric key values.

Applications utilizing proof-of-possession should also utilize audience restriction, as described in Section 4.1.3 of [JWT], as it provides different protections. Proof-of-possession can be used by recipients to reject messages from unauthorized senders. Audience restriction can be used by recipients to reject messages intended for different recipients.

A recipient might not understand the "cnf" claim. Applications that

require the proof-of-possession keys communicated with it to be understood and processed must ensure that the parts of this specification that they use are implemented.

Proof-of-possession via encrypted symmetric secrets is subject to replay attacks. This attack can be avoided when a signed nonce or challenge is used, since the recipient can use a distinct nonce or challenge for each interaction. Replay can also be avoided if a sub-key is derived from a shared secret that is specific to the instance of the PoP demonstration.

Similarly to other information included in a JWT, it is necessary to apply data origin authentication and integrity protection (via a keyed message digest or a digital signature). Data origin authentication ensures that the recipient of the JWT learns about the entity that created the JWT, since this will be important for any policy decisions. Integrity protection prevents an adversary from changing any elements conveyed within the JWT payload. Special care has to be applied when carrying symmetric keys inside the JWT, since those not only require integrity protection, but also confidentiality protection.

5. Privacy Considerations

A proof-of-possession key can be used as a correlation handle if the same key is used with multiple parties. Thus, for privacy reasons, it is recommended that different proof-of-possession keys be used when interacting with different parties.

6. IANA Considerations

The following registration procedure is used for all the registries established by this specification.

Values are registered on a Specification Required [RFC5226] basis after a three-week review period on the `oauth-pop-reg-review@ietf.org` mailing list, on the advice of one or more Designated Experts. However, to allow for the allocation of values prior to publication, the Designated Experts may approve registration once they are satisfied that such a specification will be published. [[Note to the RFC Editor: The name of the mailing list should be determined in consultation with the IESG and IANA. Suggested name: `oauth-pop-reg-review@ietf.org`.]]

Registration requests sent to the mailing list for review should use an appropriate subject (e.g., "Request to register JWT Confirmation

Method: example"). Registration requests that are undetermined for a period longer than 21 days can be brought to the IESG's attention (using the `iesg@ietf.org` mailing list) for resolution.

Criteria that should be applied by the Designated Experts include determining whether the proposed registration duplicates existing functionality, determining whether it is likely to be of general applicability or whether it is useful only for a single application, evaluating the security properties of the item being registered, and whether the registration makes sense.

It is suggested that multiple Designated Experts be appointed who are able to represent the perspectives of different applications using this specification, in order to enable broadly-informed review of registration decisions. In cases where a registration decision could be perceived as creating a conflict of interest for a particular Expert, that Expert should defer to the judgment of the other Experts.

6.1. JSON Web Token Claims Registration

This specification registers the "cnf" claim in the IANA "JSON Web Token Claims" registry [IANA.JWT.Claims] established by [JWT].

6.1.1. Registry Contents

- o Claim Name: "cnf"
- o Claim Description: Confirmation
- o Change Controller: IESG
- o Specification Document(s): Section 3.1 of [[this document]]

6.2. JWT Confirmation Methods Registry

This specification establishes the IANA "JWT Confirmation Methods" registry for JWT "cnf" member values. The registry records the confirmation method member and a reference to the specification that defines it.

6.2.1. Registration Template

Confirmation Method Value:

The name requested (e.g., "kid"). Because a core goal of this specification is for the resulting representations to be compact, it is RECOMMENDED that the name be short -- not to exceed 8 characters without a compelling reason to do so. This name is case-sensitive. Names may not match other registered names in a case-insensitive manner unless the Designated Experts state that there is a compelling reason to allow an exception.

Confirmation Method Description:

Brief description of the confirmation method (e.g., "Key Identifier").

Change Controller:

For Standards Track RFCs, list the "IESG". For others, give the name of the responsible party. Other details (e.g., postal address, email address, home page URI) may also be included.

Specification Document(s):

Reference to the document or documents that specify the parameter, preferably including URIs that can be used to retrieve copies of the documents. An indication of the relevant sections may also be included but is not required.

6.2.2. Initial Registry Contents

- o Confirmation Method Value: "jwk"
- o Confirmation Method Description: JSON Web Key Representing Public Key
- o Change Controller: IESG
- o Specification Document(s): Section 3.2 of [[this document]]

- o Confirmation Method Value: "jwe"
- o Confirmation Method Description: Encrypted JSON Web Key
- o Change Controller: IESG
- o Specification Document(s): Section 3.3 of [[this document]]

- o Confirmation Method Value: "kid"
- o Confirmation Method Description: Key Identifier
- o Change Controller: IESG
- o Specification Document(s): Section 3.4 of [[this document]]

- o Confirmation Method Value: "jku"
- o Confirmation Method Description: JWK Set URL
- o Change Controller: IESG
- o Specification Document(s): Section 3.5 of [[this document]]

7. References**7.1. Normative References**

[IANA.JWT.Claims]

IANA, "JSON Web Token Claims",
<<http://www.iana.org/assignments/jwt>>.

[JWE]

Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)",

- RFC 7516, DOI 10.17487/RFC7156, May 2015,
<<http://www.rfc-editor.org/info/rfc7516>>.
- [JWK] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/
RFC7157, May 2015,
<<http://www.rfc-editor.org/info/rfc7517>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token
(JWT)", RFC 7519, DOI 10.17487/RFC7159, May 2015,
<<http://www.rfc-editor.org/info/rfc7519>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/
RFC2119, March 1997,
<<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO
10646", STD 63, RFC 3629, DOI 10.17487/RFC3629,
November 2003, <<http://www.rfc-editor.org/info/rfc3629>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform
Resource Identifier (URI): Generic Syntax", STD 66,
RFC 3986, DOI 10.17487/RFC3986, January 2005,
<<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an
IANA Considerations Section in RFCs", BCP 26, RFC 5226,
DOI 10.17487/RFC5226, May 2008,
<<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security
(TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/
RFC5246, August 2008,
<<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and
Verification of Domain-Based Application Service Identity
within Internet Public Key Infrastructure Using X.509
(PKIX) Certificates in the Context of Transport Layer
Security (TLS)", RFC 6125, DOI 10.17487/RFC6125,
March 2011, <<http://www.rfc-editor.org/info/rfc6125>>.

7.2. Informative References

- [I-D.ietf-oauth-pop-architecture]
Hunt, P., Richer, J., Mills, W., Mishra, P., and H.
Tschofenig, "OAuth 2.0 Proof-of-Possession (PoP) Security
Architecture", draft-ietf-oauth-pop-architecture-05 (work

in progress), October 2015.

[JWK.Thumbprint]

Jones, M. and N. Sakimura, "JSON Web Key (JWK) Thumbprint", RFC 7638, DOI 10.17487/RFC7638, September 2015, <<http://www.rfc-editor.org/info/rfc7638>>.

[OASIS.saml-core-2.0-os]

Cantor, S., Kemp, J., Philpott, R., and E. Maler, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-core-2.0-os, March 2005.

[OpenID.Core]

Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", November 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.

Appendix A. Acknowledgements

The authors wish to thank Brian Campbell, Stephen Farrell, Barry Leiba, Kepeng Li, Chris Lonvick, James Manger, Kathleen Moriarty, Justin Richer, and Nat Sakimura for their reviews of the specification.

Appendix B. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-11

- o Addressed Sec-Dir review comments by Chris Lonvick and ballot comments by Stephen Farrell.

-10

- o Addressed ballot comments by Barry Leiba.

-09

- o Removed erroneous quotation marks around numeric "exp" claim values in examples.

-08

- o Added security consideration about also utilizing audience restriction.

-07

- o Addressed review comments by Hannes Tschofenig, Kathleen Moriarty, and Justin Richer. Changes were:
- o Clarified that symmetric proof-of-possession keys can be generated by either the presenter or the issuer.
- o Clarified that confirmation members that are not understood must be ignored unless otherwise specified by the application.

-06

- o Added diagrams to the introduction.

-05

- o Addressed review comments by Kepeng Li.

-04

- o Allowed the use of "jwk" for symmetric keys when the JWT is encrypted.
- o Added the "jku" (JWK Set URL) member.
- o Added privacy considerations.
- o Reordered sections so that the "cnf" (confirmation) claim is defined before it is used.
- o Noted that applications can define new claim names, in addition to "cnf", to represent additional proof-of-possession keys, using the same representation as "cnf".
- o Applied wording clarifications suggested by Nat Sakimura.

-03

- o Separated the "jwk" and "jwe" confirmation members; the former represents a public key as a JWK and the latter represents a symmetric key as a JWE encrypted JWK.
- o Changed the title to indicate that a proof-of-possession key is being communicated.

- o Updated language that formerly assumed that the issuer was an OAuth 2.0 authorization server.
- o Described ways that applications can choose to identify the presenter, including use of the "iss", "sub", and "azp" claims.
- o Harmonized the registry language with that used in JWT [RFC 7519].
- o Addressed other issues identified during working group last call.
- o Referenced the JWT and JOSE RFCs.

-02

- o Defined the terms Issuer, Presenter, and Recipient and updated their usage within the document.
- o Added a description of a use case using an asymmetric proof-of-possession key to the introduction.
- o Added the "kid" (key ID) confirmation method.
- o These changes address the open issues identified in the previous draft.

-01

- o Updated references.

-00

- o Created the initial working group draft from draft-jones-oauth-proof-of-possession-02.

Authors' Addresses

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Hannes Tschofenig
ARM Limited
Austria

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 9, 2017

J. Richer, Ed.
J. Bradley
Ping Identity
H. Tschofenig
ARM Limited
August 08, 2016

A Method for Signing HTTP Requests for OAuth
draft-ietf-oauth-signed-http-request-03

Abstract

This document a method for offering data origin authentication and integrity protection of HTTP requests. To convey the relevant data items in the request a JSON-based encapsulation is used and the JSON Web Signature (JWS) technique is re-used. JWS offers integrity protection using symmetric as well as asymmetric cryptography.

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 February 9, 2017.

Copyright Notice

Copyright (c) 2016 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

1. Introduction	2
2. Terminology	3
3. Generating a JSON Object from an HTTP Request	3
3.1. Calculating the query parameter list and hash	4
3.2. Calculating the header list and hash	5
4. Sending the signed object	6
4.1. HTTP Authorization header	6
4.2. HTTP Form body	6
4.3. HTTP Query parameter	7
5. Validating the request	7
5.1. Validating the query parameter list and hash	7
5.2. Validating the header list and hash	8
6. IANA Considerations	9
6.1. The 'pop' OAuth Access Token Type	9
6.2. JSON Web Signature and Encryption Type Values Registration	9
7. Security Considerations	9
7.1. Offering Confidentiality Protection for Access to Protected Resources	9
7.2. Plaintext Storage of Credentials	10
7.3. Entropy of Keys	10
7.4. Denial of Service	10
7.5. Validating the integrity of HTTP message	11
8. Privacy Considerations	12
9. Acknowledgements	12
10. Normative References	12
Authors' Addresses	13

1. Introduction

In order to prove possession of an access token and its associated key, an OAuth 2.0 client needs to compute some cryptographic function and present the results to the protected resource as a signature. The protected resource then needs to verify the signature and compare that to the expected keys associated with the access token. This is in addition to the normal token protections provided by a bearer token [RFC6750] and transport layer security (TLS).

Furthermore, it is desirable to bind the signature to the HTTP request. Ideally, this should be done without replicating the information already present in the HTTP request more than required. However, many HTTP application frameworks insert extra headers, query

parameters, and otherwise manipulate the HTTP request on its way from the web server into the application code itself. It is the goal of this draft to have a signature protection mechanism that is sufficiently robust against such deployment constraints while still providing sufficient security benefits.

The key required for this signature calculation is distributed via mechanisms described in companion documents (see [I-D.ietf-oauth-pop-key-distribution] and [I-D.ietf-oauth-pop-architecture]). The JSON Web Signature (JWS) specification [RFC7515] is used for computing a digital signature (which uses asymmetric cryptography) or a keyed message digest (in case of symmetric cryptography).

The mechanism described in this document assumes that a client is in possession of an access token and associated key. That client then creates a JSON object including the access token, signs the JSON object using JWS, and issues an request to a resource server for access to a protected resource using the signed object as its authorization. The protected resource validates the JWS signature and parses the JSON object to obtain token information.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Other terms such as "client", "authorization server", "access token", and "protected resource" are inherited from OAuth 2.0 [RFC6749].

We use the term 'sign' (or 'signature') to denote both a keyed message digest and a digital signature operation.

3. Generating a JSON Object from an HTTP Request

This specification uses JSON Web Signatures [RFC7515] to protect the access token and, optionally, parts of the request.

This section describes how to generate a JSON [RFC7159] object from the HTTP request. Each value below is included as a member of the JSON object at the top level.

at REQUIRED. The access token value. This string is assumed to have no particular format or structure and remains opaque to the client.

- ts RECOMMENDED. The timestamp. This integer provides replay protection of the signed JSON object. Its value MUST be a number containing an integer value representing number of whole integer seconds from midnight, January 1, 1970 GMT.
- m OPTIONAL. The HTTP Method used to make this request. This MUST be the uppercase HTTP verb as a JSON string.
- u OPTIONAL. The HTTP URL host component as a JSON string. This MAY include the port separated from the host by a colon in host:port format.
- p OPTIONAL. The HTTP URL path component of the request as an HTTP string.
- q OPTIONAL. The hashed HTTP URL query parameter map of the request as a two-part JSON array. The first part of this array is a JSON array listing all query parameters that were used in the calculation of the hash in the order that they were added to the hashed value as described below. The second part of this array is a JSON string containing the Base64URL encoded hash itself, calculated as described below.
- h OPTIONAL. The hashed HTTP request headers as a two-part JSON array. The first part of this array is a JSON array listing all headers that were used in the calculation of the hash in the order that they were added to the hashed value as described below. The second part of this array is a JSON string containing the Base64URL encoded hash itself, calculated as described below.
- b OPTIONAL. The base64URL encoded hash of the HTTP Request body, calculated as the SHA256 of the byte array of the body

All hashes SHALL be calculated using the SHA256 algorithm. [[Note to WG: do we want crypto agility here? If so how do we signal this]]

The JSON object is signed using the algorithm appropriate to the associated access token key, usually communicated as part of key distribution [I-D.ietf-oauth-pop-key-distribution].

3.1. Calculating the query parameter list and hash

To generate the query parameter list and hash, the client creates two data objects: an ordered list of strings to hold the query parameter names and a string buffer to hold the data to be hashed.

The client iterates through all query parameters in whatever order it chooses and for each query parameter it does the following:

1. Adds the name of the query parameter to the end of the list.
2. Percent-encodes the name and value of the parameter as specified in [RFC3986]. Note that if the name and value have already been percent-encoded for transit, they are not re-encoded for this step.
3. Encodes the name and value of the query parameter as "name=value" and appends it to the string buffer separated by the ampersand "&" character.

Repeated parameter names are processed separately with no special handling. Parameters MAY be skipped by the client if they are not required (or desired) to be covered by the signature.

The client then calculates the hash over the resulting string buffer. The list and the hash result are added to a list as the value of the "q" member.

For example, the query parameter set of "b=bar", "a=foo", "c=duck" is concatenated into the string:

```
b=bar&a=foo&c=duck
```

When added to the JSON structure using this process, the results are:

```
"q": [{"b", "a", "c"}, "u4LgkGUWhP9MsKrEjA4dizI1lDXluDku6ZqCeyuR-JY"]
```

3.2. Calculating the header list and hash

To generate the header list and hash, the client creates two data objects: an ordered list of strings to hold the header names and a string buffer to hold the data to be hashed.

The client iterates through all query parameters in whatever order it chooses and for each query parameter it does the following:

1. Lowercases the header name.
2. Adds the name of the header to the end of the list.
3. Encodes the name and value of the header as "name: value" and appends it to the string buffer separated by a newline "\n" character.

Repeated header names are processed separately with no special handling. Headers MAY be skipped by the client if they are not required (or desired) to be covered by the signature.

The client then calculates the hash over the resulting string buffer. The list and the hash result are added to a list as the value of the "h" member.

For example, the headers "Content-Type: application/json" and "Etag: 742-3u8f34-3r2nv3" are concatenated into the string:

```
content-type: application/json
etag: 742-3u8f34-3r2nv3

"h": [{"content-type", "etag"},
      "bZA981YJBrPlIzOvplbu3e7ueREXXr38vSkxIBYOaxI"]
```

4. Sending the signed object

In order to send the signed object to the protected resource, the client includes it in one of the following three places.

4.1. HTTP Authorization header

The client SHOULD send the signed object to the protected resource in the Authorization header. The value of the signed object in JWS compact form is appended to the Authorization header as a PoP value. This is the preferred method. Note that if this method is used, the Authorization header MUST NOT be included in the protected elements of the signed object.

```
GET /resource/foo
Authorization: PoP eyJ....omitted for brevity...
```

4.2. HTTP Form body

If the client is sending the request as a form-encoded HTTP message with parameters in the body, the client MAY send the signed object as part of that form body. The value of the signed object in JWS compact form is sent as the form parameter `pop_access_token`. Note that if this method is used, the body hash cannot be included in the protected elements of the signed object.

```
POST /resource
Content-type: application/www-form-encoded

pop_access_token=eyJ....omitted for brevity...
```

4.3. HTTP Query parameter

If neither the Authorization header nor the form-encoded body parameter are available to the client, the client MAY send the signed object as a query parameter. The value of the signed object in JWS compact form is sent as the query parameter `pop_access_token`. Note that if this method is used, the `pop_access_token` parameter MUST NOT be included in the protected elements of the signed object.

```
GET /resource?pop_access_token=eyJ....
```

5. Validating the request

Just like with a bearer token [RFC6750], while the access token value included in the signed object is opaque to the client, it MUST be understood by the protected resource in order to fulfill the request. Also like a bearer token, the protected resource traditionally has several methods at its disposal for understanding the access token. It can look up the token locally (such as in a database), it can parse a structured token (such as JWT [RFC7519]), or it can use a service to look up token information (such as introspection [RFC7662]). Whatever method is used to look up token information, the protected resource MUST have access to the key associated with the access token, as this key is required to validate the signature of the incoming request. Validation of the signature is done using normal JWS validation for the signature and key type.

Additionally, in order to trust any of the hashed components of the HTTP request, the protected resource MUST re-create and verify a hash for each component as described below. This process is a mirror of the process used to create the hashes in the first place, with a mind toward the fact that order may have changed and that elements may have been added or deleted. The protected resource MUST similarly compare the replicated values included in various JSON fields with the corresponding actual values from the request. Failure to do so will allow an attacker to modify the underlying request while at the same time having the application layer verify the signature correctly.

5.1. Validating the query parameter list and hash

The client has at its disposal a map that indexes the query parameter names to the values given. The client creates a string buffer for calculating the hash. The client then iterates through the "list" portion of the "p" parameter. For each item in the list (in the order of the list) it does the following:

1. Fetch the value of the parameter from the HTTP request query parameter map. If a parameter is found in the list of signed parameters but not in the map, the validation fails.
2. Percent-encodes the name and value of the parameter as specified in [RFC3986]. Note that if the name and value have already been percent-encoded for transit, they are not re-encoded for this step.
3. Encode the parameter as "name=value" and concatenate it to the end of the string buffer, separated by an ampersand character.

The client calculates the hash of the string buffer and base64url encodes it. The protected resource compares that string to the string passed in as the hash. If the two match, the hash validates, and all named parameters and their values are considered covered by the signature.

There MAY be additional query parameters that are not listed in the list and are therefore not covered by the signature. The client MUST decide whether or not to accept a request with these uncovered parameters.

5.2. Validating the header list and hash

The client has at its disposal a map that indexes the header names to the values given. The client creates a string buffer for calculating the hash. The client then iterates through the "list" portion of the "h" parameter. For each item in the list (in the order of the list) it does the following:

1. Fetch the value of the header from the HTTP request header map. If a header is found in the list of signed parameters but not in the map, the validation fails.
2. Encode the parameter as "name: value" and concatenate it to the end of the string buffer, separated by a newline character.

The client calculates the hash of the string buffer and base64url encodes it. The protected resource compares that string to the string passed in as the hash. If the two match, the hash validates, and all named headers and their values are considered covered by the signature.

There MAY be additional headers that are not listed in the list and are therefore not covered by the signature. The client MUST decide whether or not to accept a request with these uncovered headers.

6. IANA Considerations

6.1. The 'pop' OAuth Access Token Type

Section 11.1 of [RFC6749] defines the OAuth Access Token Type Registry and this document adds another token type to this registry.

Type name: pop

Additional Token Endpoint Response Parameters: (none)

HTTP Authentication Scheme(s): Proof-of-possession access token for use with OAuth 2.0

Change controller: IETF

Specification document(s): [[this document]]

6.2. JSON Web Signature and Encryption Type Values Registration

This specification registers the "pop" type value in the IANA JSON Web Signature and Encryption Type Values registry [RFC7515]:

- o "typ" Header Parameter Value: "pop"
- o Abbreviation for MIME Type: None
- o Change Controller: IETF
- o Specification Document(s): [[this document]]

7. Security Considerations

7.1. Offering Confidentiality Protection for Access to Protected Resources

This specification can be used with and without Transport Layer Security (TLS).

Without TLS this protocol provides a mechanism for verifying the integrity of requests, it provides no confidentiality protection. Consequently, eavesdroppers will have full access to communication content and any further messages exchanged between the client and the resource server. This could be problematic when data is exchanged that requires care, such as personal data.

When TLS is used then confidentiality of the transmission can be ensured between endpoints, including both the request and the

response. The use of TLS in combination with the signed HTTP request mechanism is highly recommended to ensure the confidentiality of the data returned from the protected resource.

7.2. Plaintext Storage of Credentials

The mechanism described in this document works in a similar way to many three-party authentication and key exchange mechanisms. In order to compute the signature over the HTTP request, the client must have access to a key bound to the access token in plaintext form. If an attacker were to gain access to these stored secrets at the client or (in case of symmetric keys) at the resource server they would be able to perform any action on behalf of any client just as if they had stolen a bearer token.

It is therefore paramount to the security of the protocol that the private keys associated with the access tokens are protected from unauthorized access.

7.3. Entropy of Keys

Unless TLS is used between the client and the resource server, eavesdroppers will have full access to requests sent by the client. They will thus be able to mount off-line brute-force attacks to attempt recovery of the session key or private key used to compute the keyed message digest or digital signature, respectively.

This specification assumes that the key used herein has been distributed via other mechanisms, such as [I-D.ietf-oauth-pop-key-distribution]. Hence, it is the responsibility of the authorization server and or the client to be careful when generating fresh and unique keys with sufficient entropy to resist such attacks for at least the length of time that the session keys (and the access tokens) are valid.

For example, if the key bound to the access token is valid for one day, authorization servers must ensure that it is not possible to mount a brute force attack that recovers that key in less than one day. Of course, servers are urged to err on the side of caution, and use the longest key length possible within reason.

7.4. Denial of Service

This specification includes a number of features which may make resource exhaustion attacks against resource servers possible. For example, a resource server may need to process the incoming request, verify the access token, perform signature verification, and might (in certain circumstances) have to consult back-end databases or the

authorization server before granting access to the protected resource. Many of these actions are shared with bearer tokens, but the additional cryptographic overhead of validating the signed request needs to be taken into consideration with deployment of this specification.

An attacker may exploit this to perform a denial of service attack by sending a large number of invalid requests to the server. The computational overhead of verifying the keyed message digest alone is not likely sufficient to mount a denial of service attack. To help combat this, it is RECOMMENDED that the protected resource validate the access token (contained in the "at" member of the signed structure) before performing any cryptographic verification calculations.

7.5. Validating the integrity of HTTP message

This specification provides flexibility for selectively validating the integrity of the HTTP request, including header fields, query parameters, and message bodies. Since all components of the HTTP request are only optionally validated by this method, and even some components may be validated only in part (e.g., some headers but not others) it is up to protected resource developers to verify that any vital parameters in a request are actually covered by the signature. Failure to do so could allow an attacker to inject vital parameters or headers into the request, outside of the protection of the signature.

The application verifying this signature MUST NOT assume that any particular parameter is appropriately covered by the signature unless it is included in the signed structure and the hash is verified. Any applications that are sensitive of header or query parameter order MUST verify the order of the parameters on their own. The application MUST also compare the values in the JSON container with the actual parameters received with the HTTP request (using a direct comparison or a hash calculation, as appropriate). Failure to make this comparison will render the signature mechanism useless for protecting these elements.

The behavior of repeated query parameters or repeated HTTP headers is undefined by this specification. If a header or query parameter is repeated on either the outgoing request from the client or the incoming request to the protected resource, that query parameter or header name MUST NOT be covered by the hash and signature.

This specification records the order in which query parameters and headers are hashed, but it does not guarantee that order is preserved between the client and protected resource. If the order of

parameters or headers are significant to the underlying application, it MUST confirm their order on its own, apart from the signature and HTTP message validation.

8. Privacy Considerations

This specification addresses machine to machine communications and raises no privacy considerations beyond existing OAuth transactions.

9. Acknowledgements

The authors thank the OAuth Working Group for input into this work.

10. Normative References

[I-D.ietf-oauth-pop-architecture]

Hunt, P., Richer, J., Mills, W., Mishra, P., and H. Tschofenig, "OAuth 2.0 Proof-of-Possession (PoP) Security Architecture", draft-ietf-oauth-pop-architecture-08 (work in progress), July 2016.

[I-D.ietf-oauth-pop-key-distribution]

Bradley, J., Hunt, P., Jones, M., and H. Tschofenig, "OAuth 2.0 Proof-of-Possession: Authorization Server to Client Key Distribution", draft-ietf-oauth-pop-key-distribution-02 (work in progress), October 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.

[RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<http://www.rfc-editor.org/info/rfc6749>>.

[RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<http://www.rfc-editor.org/info/rfc6750>>.

- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<http://www.rfc-editor.org/info/rfc7515>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<http://www.rfc-editor.org/info/rfc7519>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<http://www.rfc-editor.org/info/rfc7662>>.

Authors' Addresses

Justin Richer (editor)

Email: ietf@justin.richer.org

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com
URI: <http://www.thread-safe.com/>

Hannes Tschofenig
ARM Limited
Austria

Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 22, 2019

M. Jones
A. Nadalin
Microsoft
B. Campbell, Ed.
J. Bradley
Ping Identity
C. Mortimore
Salesforce
October 19, 2018

OAuth 2.0 Token Exchange
draft-ietf-oauth-token-exchange-16

Abstract

This specification defines a protocol for an HTTP- and JSON- based Security Token Service (STS) by defining how to request and obtain security tokens from OAuth 2.0 authorization servers, including security tokens employing impersonation and delegation.

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 <https://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 April 22, 2019.

Copyright Notice

Copyright (c) 2018 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 (<https://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

1.	Introduction	3
1.1.	Delegation vs. Impersonation Semantics	4
1.2.	Requirements Notation and Conventions	5
1.3.	Terminology	6
2.	Token Exchange Request and Response	6
2.1.	Request	6
2.1.1.	Relationship Between Resource, Audience and Scope	8
2.2.	Response	9
2.2.1.	Successful Response	9
2.2.2.	Error Response	10
2.3.	Example Token Exchange	11
3.	Token Type Identifiers	13
4.	JSON Web Token Claims and Introspection Response Parameters	14
4.1.	"act" (Actor) Claim	14
4.2.	"scope" (Scopes) Claim	16
4.3.	"client_id" (Client Identifier) Claim	17
4.4.	"may_act" (May Act For) Claim	17
5.	Security Considerations	18
6.	Privacy Considerations	19
7.	IANA Considerations	19
7.1.	OAuth URI Registration	19
7.1.1.	Registry Contents	19
7.2.	OAuth Parameters Registration	20
7.2.1.	Registry Contents	20
7.3.	OAuth Access Token Type Registration	21
7.3.1.	Registry Contents	21
7.4.	JSON Web Token Claims Registration	21
7.4.1.	Registry Contents	21
7.5.	OAuth Token Introspection Response Registration	22
7.5.1.	Registry Contents	22
7.6.	OAuth Extensions Error Registration	22
7.6.1.	Registry Contents	22
8.	References	22
8.1.	Normative References	22
8.2.	Informative References	23
Appendix A.	Additional Token Exchange Examples	24
A.1.	Impersonation Token Exchange Example	24
A.1.1.	Token Exchange Request	24
A.1.2.	Subject Token Claims	25
A.1.3.	Token Exchange Response	25
A.1.4.	Issued Token Claims	26

- A.2. Delegation Token Exchange Example 26
 - A.2.1. Token Exchange Request 26
 - A.2.2. Subject Token Claims 27
 - A.2.3. Actor Token Claims 28
 - A.2.4. Token Exchange Response 28
 - A.2.5. Issued Token Claims 28
- Appendix B. Acknowledgements 29
- Appendix C. Document History 29
- Authors' Addresses 33

1. Introduction

A security token is a set of information that facilitates the sharing of identity and security information in heterogeneous environments or across security domains. Examples of security tokens include JSON Web Tokens (JWTs) [JWT] and SAML 2.0 Assertions [OASIS.saml-core-2.0-os]. Security tokens are typically signed to achieve integrity and sometimes also encrypted to achieve confidentiality. Security tokens are also sometimes described as Assertions, such as in [RFC7521].

A Security Token Service (STS) is a service capable of validating and issuing security tokens, which enables clients to obtain appropriate access credentials for resources in heterogeneous environments or across security domains. Web Service clients have used WS-Trust [WS-Trust] as the protocol to interact with an STS for token exchange. While WS-Trust uses XML and SOAP, the trend in modern Web development has been towards RESTful patterns and JSON. The OAuth 2.0 Authorization Framework [RFC6749] and OAuth 2.0 Bearer Tokens [RFC6750] have emerged as popular standards for authorizing third-party applications access to HTTP and RESTful resources. The conventional OAuth 2.0 interaction involves the exchange of some representation of resource owner authorization for an access token, which has proven to be an extremely useful pattern in practice, however, its input and output are somewhat too constrained as is to fully accommodate a security token exchange framework.

This specification defines a protocol extending OAuth 2.0 that enables clients to request and obtain security tokens from authorization servers acting in the role of an STS. Similar to OAuth 2.0, this specification focuses on client developer simplicity and requires only an HTTP client and JSON parser, which are nearly universally available in modern development environments. The STS protocol defined in this specification is not itself RESTful (an STS doesn't lend itself particularly well to a REST approach) but does utilize communication patterns and data formats that should be familiar to developers accustomed to working with RESTful systems.

A new grant type for a token exchange request and the associated specific parameters for such a request to the token endpoint are defined by this specification. A token exchange response is a normal OAuth 2.0 response from the token endpoint with a few additional parameters defined herein to provide information to the client.

The entity that makes the request to exchange tokens is considered the client in the context of the token exchange interaction. However, that does not restrict usage of this profile to traditional OAuth clients. An OAuth resource server, for example, might assume the role of the client during token exchange in order to trade an access token, which it received in a protected resource request, for a new token that is appropriate to include in a call to a backend service. The new token might be an access token that is more narrowly scoped for the downstream service or it could be an entirely different kind of token.

The scope of this specification is limited to the definition of a basic request and response protocol for an STS-style token exchange utilizing OAuth 2.0. Although a few new JWT claims are defined that enable delegation semantics to be expressed, the specific syntax, semantics and security characteristics of the tokens themselves (both those presented to the authorization server and those obtained by the client) are explicitly out of scope and no requirements are placed on the trust model in which an implementation might be deployed. Additional profiles may provide more detailed requirements around the specific nature of the parties and trust involved, such as whether signing and/or encryption of tokens is needed or if proof-of-possession style tokens will be required or issued; however, such details will often be policy decisions made with respect to the specific needs of individual deployments and will be configured or implemented accordingly.

The security tokens obtained may be used in a number of contexts, the specifics of which are also beyond the scope of this specification.

1.1. Delegation vs. Impersonation Semantics

When principal A impersonates principal B, A is given all the rights that B has within some defined rights context and is indistinguishable from B in that context. Thus, when principal A impersonates principal B, then in so far as any entity receiving such a token is concerned, they are actually dealing with B. It is true that some members of the identity system might have awareness that impersonation is going on, but it is not a requirement. For all intents and purposes, when A is impersonating B, A is B.

Delegation semantics are different than impersonation semantics, though the two are closely related. With delegation semantics, principal A still has its own identity separate from B and it is explicitly understood that while B may have delegated some of its rights to A, any actions taken are being taken by A representing B. In a sense, A is an agent for B.

Delegation and impersonation are not inclusive of all situations. When a principal is acting directly on its own behalf, for example, neither delegation nor impersonation are in play. They are, however, the more common semantics operating for token exchange and, as such, are given more direct treatment in this specification.

Delegation semantics are typically expressed in a token by including information about both the primary subject of the token as well as the actor to whom that subject has delegated some of its rights. Such a token is sometimes referred to as a composite token because it is composed of information about multiple subjects. Typically, in the request, the "subject_token" represents the identity of the party on behalf of whom the token is being requested while the "actor_token" represents the identity of the party to whom the access rights of the issued token are being delegated. A composite token issued by the authorization server will contain information about both parties. When and if a composite token is issued is at the discretion of the authorization server and applicable policy and configuration.

The specifics of representing a composite token and even whether or not such a token will be issued depend on the details of the implementation and the kind of token. The representations of composite tokens that are not JWTs are beyond the scope of this specification. The "actor_token" request parameter, however, does provide a means for providing information about the desired actor and the JWT "act" claim can provide a representation of a chain of delegation.

1.2. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.3. Terminology

This specification uses the terms "access token type", "authorization server", "client", "client identifier", "resource server", "token endpoint", "token request", and "token response" defined by OAuth 2.0 [RFC6749], and the terms "Base64url Encoding", "Claim", and "JWT Claims Set" defined by JSON Web Token (JWT) [JWT].

2. Token Exchange Request and Response

2.1. Request

A client requests a security token by making a token request to the authorization server's token endpoint using the extension grant type mechanism defined in Section 4.5 of OAuth 2.0 [RFC6749].

Client authentication to the authorization server is done using the normal mechanisms provided by OAuth 2.0. Section 2.3.1 of The OAuth 2.0 Authorization Framework [RFC6749] defines password-based authentication of the client, however, client authentication is extensible and other mechanisms are possible. For example, [RFC7523] defines client authentication using JSON Web Tokens (JWTs) [JWT]. The supported methods of client authentication and whether or not to allow unauthenticated or unidentified clients are deployment decisions that are at the discretion of the authorization server.

The client makes a token exchange request to the token endpoint with an extension grant type by including the following parameters using the "application/x-www-form-urlencoded" format with a character encoding of UTF-8 in the HTTP request entity-body:

grant_type

REQUIRED. The value "urn:ietf:params:oauth:grant-type:token-exchange" indicates that a token exchange is being performed.

resource

OPTIONAL. Indicates the location of the target service or resource where the client intends to use the requested security token. This enables the authorization server to apply policy as appropriate for the target, such as determining the type and content of the token to be issued or if and how the token is to be encrypted. In many cases, a client will not have knowledge of the logical organization of the systems with which it interacts and will only know the location of the service where it intends to use the token. The "resource" parameter allows the client to indicate to the authorization server where it intends to use the issued token by providing the location, typically as an https URL, in the token exchange request in the same form that will be used to

access that resource. The authorization server will typically have the capability to map from a resource URI value to an appropriate policy. The value of the "resource" parameter MUST be an absolute URI, as specified by Section 4.3 of [RFC3986], which MAY include a query component and MUST NOT include a fragment component. Multiple "resource" parameters may be used to indicate that the issued token is intended to be used at the multiple resources listed.

audience

OPTIONAL. The logical name of the target service where the client intends to use the requested security token. This serves a purpose similar to the "resource" parameter, but with the client providing a logical name rather than a location. Interpretation of the name requires that the value be something that both the client and the authorization server understand. An OAuth client identifier, a SAML entity identifier [OASIS.saml-core-2.0-os], an OpenID Connect Issuer Identifier [OpenID.Core], or a URI are examples of things that might be used as "audience" parameter values. Multiple "audience" parameters may be used to indicate that the issued token is intended to be used at the multiple audiences listed. The "audience" and "resource" parameters may be used together to indicate multiple target services with a mix of logical names and locations.

scope

OPTIONAL. A list of space-delimited, case-sensitive strings, as defined in Section 3.3 of [RFC6749], that allow the client to specify the desired scope of the requested security token in the context of the service or resource where the token will be used. The values and associated semantics of scope are service specific and expected to be described in the relevant service documentation.

requested_token_type

OPTIONAL. An identifier, as described in Section 3, for the type of the requested security token. If the requested type is unspecified, the issued token type is at the discretion of the authorization server and may be dictated by knowledge of the requirements of the service or resource indicated by the "resource" or "audience" parameter.

subject_token

REQUIRED. A security token that represents the identity of the party on behalf of whom the request is being made. Typically, the subject of this token will be the subject of the security token issued in response to this request.

subject_token_type

REQUIRED. An identifier, as described in Section 3, that indicates the type of the security token in the "subject_token" parameter.

actor_token

OPTIONAL. A security token that represents the identity of the acting party. Typically, this will be the party that is authorized to use the requested security token and act on behalf of the subject.

actor_token_type

An identifier, as described in Section 3, that indicates the type of the security token in the "actor_token" parameter. This is REQUIRED when the "actor_token" parameter is present in the request but MUST NOT be included otherwise.

In processing the request, the authorization sever MUST validate the subject token as appropriate for the indicated token type and, if the actor token is present, also validate it according to its token type. The validity criteria and details of any particular token are beyond the scope of this document and are specific to the respective type of token and its content.

In the absence of one-time-use or other semantics specific to the token type, the act of performing a token exchange has no impact on the validity of the subject token or actor token. Furthermore, the validity of the subject token or actor token have no impact on the validity of the issued token after the exchange has occurred.

2.1.1. Relationship Between Resource, Audience and Scope

When requesting a token, the client can indicate the desired target service(s) where it intends to use that token by way of the "audience" and "resource" parameters, as well as indicating the desired scope of the requested token using the "scope" parameter. The semantics of such a request are that the client is asking for a token with the requested scope that is usable at all the requested target services. Effectively, the requested access rights of the token are the cartesian product of all the scopes at all the target services.

An authorization server may be unwilling or unable to fulfill any token request but the likelihood of an unfulfillable request is significantly higher when very broad access rights are being solicited. As such, in the absence of specific knowledge about the relationship of systems in a deployment, clients should exercise discretion in the breadth of the access requested, particularly the

number of target services. An authorization server can use the "invalid_target" error code, defined in Section 2.2.2, to inform a client that it requested access to too many target services simultaneously.

2.2. Response

The authorization server responds to a token exchange request with a normal OAuth 2.0 response from the token endpoint, as specified in Section 5 of [RFC6749]. Additional details and explanation are provided in the following subsections.

2.2.1. Successful Response

If the request is valid and meets all policy and other criteria of the authorization server, a successful token response is constructed by adding the following parameters to the entity-body of the HTTP response using the "application/json" media type, as specified by [RFC7159], and an HTTP 200 status code. The parameters are serialized into a JavaScript Object Notation (JSON) structure by adding each parameter at the top level. Parameter names and string values are included as JSON strings. Numerical values are included as JSON numbers. The order of parameters does not matter and can vary.

access_token

REQUIRED. The security token issued by the authorization server in response to the token exchange request. The "access_token" parameter from Section 5.1 of [RFC6749] is used here to carry the requested token, which allows this token exchange protocol to use the existing OAuth 2.0 request and response constructs defined for the token endpoint. The identifier "access_token" is used for historical reasons and the issued token need not be an OAuth access token.

issued_token_type

REQUIRED. An identifier, as described in Section 3, for the representation of the issued security token.

token_type

REQUIRED. A case-insensitive value specifying the method of using the access token issued, as specified in Section 7.1 of [RFC6749]. It provides the client with information about how to utilize the access token to access protected resources. For example, a value of "Bearer", as specified in [RFC6750], indicates that the security token is a bearer token and the client can simply present it as is without any additional proof of eligibility beyond the contents of the token itself. Note that the meaning of this

parameter is different from the meaning of the "issued_token_type" parameter, which declares the representation of the issued security token; the term "token type" is typically used with this meaning, as it is in all "*_token_type" parameters in this specification. If the issued token is not an access token or usable as an access token, then the "token_type" value "N_A" is used to indicate that an OAuth 2.0 "token_type" identifier is not applicable in that context.

expires_in

RECOMMENDED. The validity lifetime, in seconds, of the token issued by the authorization server. Oftentimes the client will not have the inclination or capability to inspect the content of the token and this parameter provides a consistent and token type agnostic indication of how long the token can be expected to be valid. For example, the value 1800 denotes that the token will expire in thirty minutes from the time the response was generated.

scope

OPTIONAL, if the scope of the issued security token is identical to the scope requested by the client; otherwise, REQUIRED.

refresh_token

OPTIONAL. A refresh token will typically not be issued when the exchange is of one temporary credential (the subject_token) for a different temporary credential (the issued token) for use in some other context. A refresh token can be issued in cases where the client of the token exchange needs the ability to access a resource even when the original credential is no longer valid (e.g., user-not-present or offline scenarios where there is no longer any user entertaining an active session with the client). Profiles or deployments of this specification should clearly document the conditions under which a client should expect a refresh token in response to "urn:ietf:params:oauth:grant-type:token-exchange" grant type requests.

2.2.2. Error Response

If the request itself is not valid or if either the "subject_token" or "actor_token" are invalid for any reason, or are unacceptable based on policy, the authorization server MUST construct an error response, as specified in Section 5.2 of [RFC6749]. The value of the "error" parameter MUST be the "invalid_request" error code.

If the authorization server is unwilling or unable to issue a token for all the target services indicated by the "resource" or "audience" parameters, the "invalid_target" error code SHOULD be used in the error response.

The authorization server MAY include additional information regarding the reasons for the error using the "error_description" and/or "error_uri" parameters.

Other error codes may also be used, as appropriate.

2.3. Example Token Exchange

The following example demonstrates a hypothetical token exchange in which an OAuth resource server assumes the role of the client during token exchange in order to trade an access token that it received in a protected resource request for a token that it will use to call to a backend service (extra line breaks and indentation in the examples are for display purposes only).

The resource server receives the following request containing an OAuth access token in the Authorization request header, as specified in Section 2.1 of [RFC6750].

```
GET /resource HTTP/1.1
Host: frontend.example.com
Authorization: Bearer accVkjCjyb4BWCxGsndESCJQbdFMogUC5PbRDqceLTC
```

Figure 1: Protected Resource Request

The resource server assumes the role of the client for the token exchange and the access token from the request above is sent to the authorization server using a request as specified in Section 2.1. The value of the "subject_token" parameter carries the access token and the value of the "subject_token_type" parameter indicates that it is an OAuth 2.0 access token. The resource server, acting in the role of the client, uses its identifier and secret to authenticate to the authorization server using the HTTP Basic authentication scheme. The "resource" parameter indicates the location of the backend service, `https://backend.example.com/api`, where the issued token will be used.

```

POST /as/token.oauth2 HTTP/1.1
Host: as.example.com
Authorization: Basic cnMwODpsb25nLXNlY3VyZS1yYW5kb20tc2VjcmV0
Content-Type: application/x-www-form-urlencoded

grant_type=urn%3Aietf%3Aparams%3Aoauth%3Agrant-type%3Atoken-exchange
&resource=https%3A%2F%2Fbackend.example.com%2Fapi%20
&subject_token=accVkjcJyb4BWCxGsndESCJQbdfMogUC5PbRDqceLTC
&subject_token_type=
urn%3Aietf%3Aparams%3Aoauth%3Atoken-type%3Aaccess_token

```

Figure 2: Token Exchange Request

The authorization server validates the client credentials and the "subject_token" presented in the token exchange request. From the "resource" parameter, the authorization server is able to determine the appropriate policy to apply to the request and issues a token suitable for use at `https://backend.example.com`. The "access_token" parameter of the response contains the new token, which is itself a bearer OAuth access token that is valid for one minute. The token happens to be a JWT; however, its structure and format are opaque to the client so the "issued_token_type" indicates only that it is an access token.

```

HTTP/1.1 200 OK
Content-Type: application/json
Cache-Control: no-cache, no-store

{
  "access_token": "eyJhbGciOiJIUzI1NiIsImtpZCI6IjllciJ9.eyJhdWQiOiJodHRwczovL2JhY2t1bmQuZXhhbXBsZS5jb20iLCJpc3MiOiJodHRwczovL2FzLmV4YW1wbGUuY29tIiwiaXNjaXhwIjoxNDQxOTE3NTkzLCJpYXQiOiJEONDE5MTclMzMsInN1YiI6ImJjQGV4YW1wbGUuY29tIiwic2NvcGUiOiJhcGkifQ.K4Ik-igqOKi_4CnBu4dG3-gGUObfgv-rJhgXVDNOWW_MHgVwddhgVLLQf_bm3xlpQM6wHrLbMaZC4LicsQC23g",
  "issued_token_type":
    "urn:ietf:params:oauth:token-type:access_token",
  "token_type": "Bearer",
  "expires_in": 60
}

```

Figure 3: Token Exchange Response

The resource server can then use the newly acquired access token in making a request to the backend server.

Indicates that the token is a base64url-encoded SAML 2.0 [OASIS.saml-core-2.0-os] assertion.

The value "urn:ietf:params:oauth:token-type:jwt", which is defined in Section 9 of [JWT], indicates that the token is a JWT.

The distinction between an access token and a JWT is subtle. An access token represents a delegated authorization decision, whereas JWT is a token format. An access token can be formatted as a JWT but doesn't necessarily have to be. And a JWT might well be an access token but not all JWTs are access tokens. The intent of this specification is that "urn:ietf:params:oauth:token-type:access_token" be an indicator that the token is a typical OAuth access token issued by the authorization server in question, opaque to the client, and usable the same manner as any other access token obtained from that authorization server. (It could well be a JWT, but the client isn't and needn't be aware of that fact.) Whereas, "urn:ietf:params:oauth:token-type:jwt" is to indicate specifically that a JWT is being requested or sent (perhaps in a cross-domain use-case where the JWT is used as an authorization grant to obtain an access token from a different authorization server as is facilitated by [RFC7523]).

4. JSON Web Token Claims and Introspection Response Parameters

It is useful to have defined mechanisms to express delegation within a token as well as to express authorization to delegate or impersonate. Although the token exchange protocol described herein can be used with any type of token, this section defines claims to express such semantics specifically for JWTs and in an OAuth 2.0 Token Introspection [RFC7662] response. Similar definitions for other types of tokens are possible but beyond the scope of this specification.

Note that the claims not established herein but used in examples and descriptions, such as "iss", "sub", "exp", etc., are defined by [JWT].

4.1. "act" (Actor) Claim

The "act" (actor) claim provides a means within a JWT to express that delegation has occurred and identify the acting party to whom authority has been delegated. The "act" claim value is a JSON object and members in the JSON object are claims that identify the actor. The claims that make up the "act" claim identify and possibly provide additional information about the actor. For example, the combination of the two claims "iss" and "sub" might be necessary to uniquely identify an actor.

However, claims within the "act" claim pertain only to the identity of the actor and are not relevant to the validity of the containing JWT in the same manner as the top-level claims. Consequently, non-identity claims (e.g., "exp", "nbf", and "aud") are not meaningful when used within an "act" claim, and therefore must not be used.

The following example illustrates the "act" (actor) claim within a JWT Claims Set. The claims of the token itself are about user@example.com while the "act" claim indicates that admin@example.com is the current actor.

```
{
  "aud": "https://consumer.example.com",
  "iss": "https://issuer.example.com",
  "exp": 1443904177,
  "nbf": 1443904077,
  "sub": "user@example.com",
  "act":
  {
    "sub": "admin@example.com"
  }
}
```

Figure 5: Actor Claim

A chain of delegation can be expressed by nesting one "act" claim within another. The outermost "act" claim represents the current actor while nested "act" claims represent prior actors. The least recent actor is the most deeply nested.

For the purpose of applying access control policy, the consumer of a token MUST only consider the token's top-level claims and the party identified as the current actor by the "act" claim. Prior actors identified by any nested "act" claims are informational only and are not to be considered in access control decisions.

The following example illustrates nested "act" (actor) claims within a JWT Claims Set. The claims of the token itself are about user@example.com while the "act" claim indicates that the system https://service16.example.com is the current actor and https://service77.example.com was a prior actor. Such a token might come about as the result of service16 receiving a token in a call from service77 and exchanging it for a token suitable to call service26 while the authorization server notes the situation in the newly issued token.

```
{
  "aud": "https://service26.example.com",
  "iss": "https://issuer.example.com",
  "exp": 1443904100,
  "nbf": 1443904000,
  "sub": "user@example.com",
  "act":
  {
    "sub": "https://service16.example.com",
    "act":
    {
      "sub": "https://service77.example.com",
    }
  }
}
```

Figure 6: Nested Actor Claim

When included as a top-level member of an OAuth token introspection response, "act" has the same semantics and format as the claim of the same name.

4.2. "scope" (Scopes) Claim

The value of the "scope" claim is a JSON string containing a space-separated list of scopes associated with the token, in the format described in Section 3.3 of OAuth 2.0 [RFC6749].

The following example illustrates the "scope" claim within a JWT Claims Set.

```
{
  "aud": "https://consumer.example.com",
  "iss": "https://issuer.example.com",
  "exp": 1443904177,
  "nbf": 1443904077,
  "sub": "dgaf4mvfs75Fci_FL3heQA",
  "scope": "email profile phone address"
}
```

Figure 7: Scopes Claim

OAuth 2.0 Token Introspection [RFC7662] already defines the "scope" parameter to convey the scopes associated with the token.

4.3. "client_id" (Client Identifier) Claim

The "client_id" claim carries the client identifier of the OAuth 2.0 [RFC6749] client that requested the token.

The following example illustrates the "client_id" claim within a JWT Claims Set indicating an OAuth 2.0 client with "s6BhdRkqt3" as its identifier.

```
{
  "aud": "https://consumer.example.com",
  "iss": "https://issuer.example.com",
  "exp": 1443904177,
  "sub": "user@example.com",
  "client_id": "s6BhdRkqt3"
}
```

Figure 8: Client Identifier Claim

OAuth 2.0 Token Introspection [RFC7662] already defines the "client_id" parameter as the client identifier for the OAuth 2.0 client that requested the token.

4.4. "may_act" (May Act For) Claim

The "may_act" claim makes a statement that one party is authorized to become the actor and act on behalf of another party. The claim value is a JSON object and members in the JSON object are claims that identify the party that is asserted as being eligible to act for the party identified by the JWT containing the claim. The claims that make up the "may_act" claim identify and possibly provide additional

information about the authorized actor. For example, the combination of the two claims "iss" and "sub" are sometimes necessary to uniquely identify an authorized actor, while the "email" claim might be used to provide additional useful information about that party.

However, claims within the "may_act" claim pertain only to the identity of that party and are not relevant to the validity of the containing JWT in the same manner as top-level claims. Consequently, claims such as "exp", "nbf", and "aud" are not meaningful when used within a "may_act" claim, and therefore should not be used.

The following example illustrates the "may_act" claim within a JWT Claims Set. The claims of the token itself are about user@example.com while the "may_act" claim indicates that admin@example.com is authorized to act on behalf of user@example.com.

```
{
  "aud": "https://consumer.example.com",
  "iss": "https://issuer.example.com",
  "exp": 1443904177,
  "nbf": 1443904077,
  "sub": "user@example.com",
  "may_act": {
    {
      "sub": "admin@example.com"
    }
  }
}
```

Figure 9: May Act For Claim

When included as a top-level member of an OAuth token introspection response, "may_act" has the same semantics and format as the claim of the same name.

5. Security Considerations

All of the normal security issues that are discussed in [JWT], especially in relationship to comparing URIs and dealing with unrecognized values, also apply here.

In addition, both delegation and impersonation introduce unique security issues. Any time one principal is delegated the rights of another principal, the potential for abuse is a concern. The use of the "scope" claim is suggested to mitigate potential for such abuse, as it restricts the contexts in which the delegated rights can be exercised.

6. Privacy Considerations

Tokens employed in the context of the functionality described herein may contain privacy-sensitive information and, to prevent disclosure of such information to unintended parties, should only be transmitted over encrypted channels, such as Transport Layer Security (TLS). In cases where it is desirable to prevent disclosure of certain information to the client, the token should be encrypted to its intended recipient. Deployments should determine the minimally necessary amount of data and only include such information in issued tokens. In some cases, data minimization may include representing only an anonymous or pseudonymous user.

7. IANA Considerations

7.1. OAuth URI Registration

This specification registers the following values in the IANA "OAuth URI" registry [IANA.OAuth.Parameters] established by [RFC6755].

7.1.1. Registry Contents

- o URN: urn:ietf:params:oauth:grant-type:token-exchange
- o Common Name: Token exchange grant type for OAuth 2.0
- o Change controller: IESG
- o Specification Document: Section 2.1 of [[this specification]]

- o URN: urn:ietf:params:oauth:token-type:access_token
- o Common Name: Token type URI for an OAuth 2.0 access token
- o Change controller: IESG
- o Specification Document: Section 3 of [[this specification]]

- o URN: urn:ietf:params:oauth:token-type:refresh_token
- o Common Name: Token type URI for an OAuth 2.0 refresh token
- o Change controller: IESG
- o Specification Document: Section 3 of [[this specification]]

- o URN: urn:ietf:params:oauth:token-type:id_token
- o Common Name: Token type URI for an ID Token
- o Change controller: IESG
- o Specification Document: Section 3 of [[this specification]]

- o URN: urn:ietf:params:oauth:token-type:saml1
- o Common Name: Token type URI for a base64url-encoded SAML 1.1 assertion
- o Change Controller: IESG
- o Specification Document: Section 3 of [[this specification]]

- o URN: urn:ietf:params:oauth:token-type:saml2
- o Common Name: Token type URI for a base64url-encoded SAML 2.0 assertion
- o Change Controller: IESG
- o Specification Document: Section 3 of [[this specification]]

7.2. OAuth Parameters Registration

This specification registers the following values in the IANA "OAuth Parameters" registry [IANA.OAuth.Parameters] established by [RFC6749].

7.2.1. Registry Contents

- o Parameter name: resource
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: audience
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: requested_token_type
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: subject_token
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: subject_token_type
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: actor_token
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: actor_token_type
- o Parameter usage location: token request
- o Change controller: IESG
- o Specification document(s): Section 2.1 of [[this specification]]

- o Parameter name: issued_token_type
- o Parameter usage location: token response
- o Change controller: IESG
- o Specification document(s): Section 2.2.1 of [[this specification]]

7.3. OAuth Access Token Type Registration

This specification registers the following access token type in the IANA "OAuth Access Token Types" registry [IANA.OAuth.Parameters] established by [RFC6749].

7.3.1. Registry Contents

- o Type name: N_A
- o Additional Token Endpoint Response Parameters: (none)
- o HTTP Authentication Scheme(s): (none)
- o Change controller: IESG
- o Specification document(s): Section 2.2.1 of [[this specification]]

7.4. JSON Web Token Claims Registration

This specification registers the following Claims in the IANA "JSON Web Token Claims" registry [IANA.JWT.Claims] established by [JWT].

7.4.1. Registry Contents

- o Claim Name: "act"
- o Claim Description: Actor
- o Change Controller: IESG
- o Specification Document(s): Section 4.1 of [[this specification]]

- o Claim Name: "scope"
- o Claim Description: Scope Values
- o Change Controller: IESG
- o Specification Document(s): Section 4.2 of [[this specification]]

- o Claim Name: "client_id"
- o Claim Description: Client Identifier
- o Change Controller: IESG
- o Specification Document(s): Section 4.3 of [[this specification]]

- o Claim Name: "may_act"
- o Claim Description: May Act For
- o Change Controller: IESG
- o Specification Document(s): Section 4.4 of [[this specification]]

7.5. OAuth Token Introspection Response Registration

This specification registers the following values in the IANA "OAuth Token Introspection Response" registry [IANA.OAuth.Parameters] established by [RFC7662].

7.5.1. Registry Contents

- o Claim Name: "act"
- o Claim Description: Actor
- o Change Controller: IESG
- o Specification Document(s): Section 4.1 of [[this specification]]

- o Claim Name: "may_act"
- o Claim Description: May Act For
- o Change Controller: IESG
- o Specification Document(s): Section 4.4 of [[this specification]]

7.6. OAuth Extensions Error Registration

This specification registers the following values in the IANA "OAuth Extensions Error" registry [IANA.OAuth.Parameters] established by [RFC6749].

7.6.1. Registry Contents

- o Error Name: "invalid_target"
- o Error Usage Location: token error response
- o Related Protocol Extension: OAuth 2.0 Token Exchange
- o Change Controller: IETF
- o Specification Document(s): Section 2.2.2 of [[this specification]]

8. References

8.1. Normative References

- [IANA.JWT.Claims]
IANA, "JSON Web Token Claims",
<<http://www.iana.org/assignments/jwt>>.
- [IANA.OAuth.Parameters]
IANA, "OAuth Parameters",
<<http://www.iana.org/assignments/oauth-parameters>>.
- [JWT] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015,
<<http://tools.ietf.org/html/rfc7519>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<https://www.rfc-editor.org/info/rfc7159>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<https://www.rfc-editor.org/info/rfc7662>>.

8.2. Informative References

- [OASIS.saml-core-1.1]
Maler, E., Mishra, P., and R. Philpott, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V1.1", OASIS Standard oasis-sstc-saml-core-1.1, September 2003.
- [OASIS.saml-core-2.0-os]
Cantor, S., Kemp, J., Philpott, R., and E. Maler, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-core-2.0-os, March 2005.
- [OpenID.Core]
Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", August 2015, <http://openid.net/specs/openid-connect-core-1_0.html>.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<https://www.rfc-editor.org/info/rfc6750>>.

- [RFC6755] Campbell, B. and H. Tschofenig, "An IETF URN Sub-Namespace for OAuth", RFC 6755, DOI 10.17487/RFC6755, October 2012, <<https://www.rfc-editor.org/info/rfc6755>>.
- [RFC7521] Campbell, B., Mortimore, C., Jones, M., and Y. Goland, "Assertion Framework for OAuth 2.0 Client Authentication and Authorization Grants", RFC 7521, DOI 10.17487/RFC7521, May 2015, <<https://www.rfc-editor.org/info/rfc7521>>.
- [RFC7523] Jones, M., Campbell, B., and C. Mortimore, "JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants", RFC 7523, DOI 10.17487/RFC7523, May 2015, <<https://www.rfc-editor.org/info/rfc7523>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [WS-Trust]
Nadalin, A., Goodner, M., Gudgin, M., Barbir, A., and H. Granqvist, "WS-Trust 1.4", February 2012, <<http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/ws-trust.html>>.

Appendix A. Additional Token Exchange Examples

Two example token exchanges are provided in the following sections illustrating impersonation and delegation, respectively (with extra line breaks and indentation for display purposes only).

A.1. Impersonation Token Exchange Example

A.1.1. Token Exchange Request

In the following token exchange request, a client is requesting a token with impersonation semantics. The client tells the authorization server that it needs a token for use at the target service with the logical name "urn:example:cooperation-context".

```

POST /as/token.oauth2 HTTP/1.1
Host: as.example.com
Content-Type: application/x-www-form-urlencoded

grant_type=urn%3Aietf%3Aparams%3Aoauth%3Agrant-type%3Atoken-exchange
&audience=urn%3Aexample%3Acooperation-context
&subject_token=eyJhbGciOiJIJFUzI1NiIsImtpZCI6IjE2In0.eyJhdWQiOiJodHRwczovL2FzLmV4YW1wbGUuY29tIiwiaXNzIjoiaHR0cHM6Ly9vcmlnaW5hbC1pc3NlZXIuZXhhbXBsZS5uZXQiLCJleHAiOjE0NDE5MTA2MDAsIm5iZiI6MTQ0MTkwOTAwMCwic3ViIjoiyMNAZXhhbXBsZS5uZXQiLCJzY29wZSI6Im9yZGVycyBwcm9maWx1IGhpc3RvcnkifQ.u0slqvbngU43EvI_itGdFJ11StrAwXlxczYfMYsaR5p4J_gBp019mxljSx
&subject_token_type=urn%3Aietf%3Aparams%3Aoauth%3Atoken-type%3Ajwt

```

Figure 10: Token Exchange Request

A.1.1.2. Subject Token Claims

The "subject_token" in the prior request is a JWT and the decoded JWT Claims Set is shown here. The JWT is intended for consumption by the authorization server within a specific time window. The subject of the JWT ("bc@example.net") is the party on behalf of whom the new token is being requested.

```

{
  "aud": "https://as.example.com",
  "iss": "https://original-issuer.example.net",
  "exp": 1441910600,
  "nbf": 1441909000,
  "sub": "bc@example.net",
  "scope": "orders profile history"
}

```

Figure 11: Subject Token Claims

A.1.1.3. Token Exchange Response

The "access_token" parameter of the token exchange response shown below contains the new token that the client requested. The other parameters of the response indicate that the token is a bearer access token that expires in an hour.

A.2.3. Actor Token Claims

The "actor_token" in the prior request is a JWT and the decoded JWT Claims Set is shown here. This JWT is also intended for consumption by the authorization server before a specific expiration time. The subject of the JWT ("admin@example.net") is the actor that will wield the security token being requested.

```
{
  "aud": "https://as.example.com",
  "iss": "https://original-issuer.example.net",
  "exp": 1441910060,
  "sub": "admin@example.net"
}
```

Figure 16: Actor Token Claims

A.2.4. Token Exchange Response

The "access_token" parameter of the token exchange response shown below contains the new token that the client requested. The other parameters of the response indicate that the token is a JWT that expires in an hour and that the access token type is not applicable since the issued token is not an access token.

```
HTTP/1.1 200 OK
Content-Type: application/json
Cache-Control: no-cache, no-store

{
  "access_token": "eyJhbGciOiJIJFUzI1NiIsImtpZCI6IjcyIn0.eyJhdWQiOiJlcm42ZXhhbXBsZTpjb29wZXhhdGlvbiljb250ZXh0IiwiaXNzIjoiaHR0cHM6Ly9hcy5leGFtcGxlLnV4bSIsImV4cCI6MTQ0MTkxMzYxMCwic2NvcGUiOiJzdGF0dXMgZmVlZCIsInN1YiI6InVzZXhhdGlvbiIsImZhbGciOiJIJFUzIiwiaWF0IjoiYWRtaW5AZXhhbXBsZS5uZXQifX0.3paKl9UySKYB5ng6_cUtQ2qlO8Rc_y7Mea7IwEXTcYbNdwG9-GlEKCFe5fw3H0hwX-MSZ49Wpcb1SiAZaOQBtw",
  "issued_token_type": "urn:ietf:params:oauth:token-type:jwt",
  "token_type": "N_A",
  "expires_in": 3600
}
```

Figure 17: Token Exchange Response

A.2.5. Issued Token Claims

The decoded JWT Claims Set of the issued token is shown below. The new JWT is issued by the authorization server and intended for consumption by a system entity known by the logical name

"urn:example:cooperation-context" any time before its expiration. The subject ("sub") of the JWT is the same as the subject of the "subject_token" used to make the request. The actor ("act") of the JWT is the same as the subject of the "actor_token" used to make the request. This indicates delegation and identifies "admin@example.net" as the current actor to whom authority has been delegated to act on behalf of "user@example.net".

```
{
  "aud": "urn:example:cooperation-context",
  "iss": "https://as.example.com",
  "exp": 1441913610,
  "scope": "status feed",
  "sub": "user@example.net",
  "act":
  {
    "sub": "admin@example.net"
  }
}
```

Figure 18: Issued Token Claims

Appendix B. Acknowledgements

This specification was developed within the OAuth Working Group, which includes dozens of active and dedicated participants. It was produced under the chairmanship of Hannes Tschofenig, Derek Atkins, and Rifaat Shekh-Yusef with Kathleen Moriarty, Stephen Farrell, Eric Rescorla, and Benjamin Kaduk serving as Security Area Directors. The following individuals contributed ideas, feedback, and wording to this specification:

Caleb Baker, Vittorio Bertocci, Thomas Broyer, William Denniss, Vladimir Dzhuvinov, Phil Hunt, Benjamin Kaduk, Jason Keglovitz, Torsten Lodderstedt, Adam Lewis, James Manger, Nov Matake, Matt Miller, Hilarie Orman, Matthew Perry, Eric Rescorla, Justin Richer, Rifaat Shekh-Yusef, Scott Tomilson, and Hannes Tschofenig.

Appendix C. Document History

[[to be removed by the RFC Editor before publication as an RFC]]

-16

- o Fixed typo and added an AD to Acknowledgements.

-15

- o Updated the nested actor claim example to (hopefully) be more straightforward.
- o Reworked Privacy Considerations to say to use TLS in transit, minimize the amount of information in the token, and encrypt the token if disclosure of its information to the client is a concern per https://mailarchive.ietf.org/arch/msg/secdir/KJhx4aq_U5uk3k6zpYP-CEHbpVM
- o Moved the Security and Privacy Considerations sections to before the IANA Considerations.

-14

- o Added text in Section 4.1 about the "act" claim stating that only the top-level claims and the current actor are to be considered in applying access control decisions.

-13

- o Updated the claim name and value syntax for scope to be consistent with the treatment of scope in RFC 7662 OAuth 2.0 Token Introspection.
- o Updated the client identifier claim name to be consistent with the treatment of client id in RFC 7662 OAuth 2.0 Token Introspection.

-12

- o Updated to use the boilerplate from RFC 8174.

-11

- o Added new WG chair and AD to the Acknowledgements.
- o Applied clarifications suggested during AD review by EKR.

-10

- o Defined token type URIs for base64url-encoded SAML 1.1 and SAML 2.0 assertions.
- o Applied editorial fixes.

-09

- o Changed "security tokens obtained could be used in a number of contexts" to "security tokens obtained may be used in a number of contexts" per a WGLC suggestion.
- o Clarified that the validity of the subject or actor token have no impact on the validity of the issued token after the exchange has occurred per a WGLC comment.

- o Changed use of `invalid_target` error code to a SHOULD per a WGLC comment.
- o Clarified text about non-identity claims within the "act" claim being meaningless per a WGLC comment.
- o Added brief Privacy Considerations section per WGLC comments.

-08

- o Use the bibxml reference for OpenID.Core rather than defining it inline.
- o Added editor role for Campbell.
- o Minor clarification of the text for `actor_token`.

-07

- o Fixed typo (desecration -> discretion).
- o Added an explanation of the relationship between scope, audience and resource in the request and added an "invalid_target" error code enabling the AS to tell the client that the requested audiences/resources were too broad.

-06

- o Drop "An STS for the REST of Us" from the title.
- o Drop "heavyweight" and "lightweight" from the abstract and introduction.
- o Clarifications on the language around `xxxxxx_token_type`.
- o Remove the `want_composite` parameter.
- o Add a short mention of proof-of-possession style tokens to the introduction and remove the respective open issue.

-05

- o Defined the JWT claim "cid" to express the OAuth 2.0 client identifier of the client that requested the token.
- o Defined and requested registration for "act" and "may_act" as Token introspection response parameters (in addition to being JWT claims).
- o Loosen up the language about `refresh_token` in the response to OPTIONAL from NOT RECOMMENDED based on feedback from real world deployment experience.
- o Add clarifying text about the distinction between JWT and access token URIs.
- o Close out (remove) some of the Open Issues bullets that have been resolved.

-04

- o Clarified that the "resource" and "audience" request parameters can be used at the same time (via <http://www.ietf.org/mail-archive/web/oauth/current/msg15335.html>).
- o Clarified subject/actor token validity after token exchange and explained a bit more about the recommendation to not issue refresh tokens (via <http://www.ietf.org/mail-archive/web/oauth/current/msg15318.html>).
- o Updated the examples appendix to use an issuer value that doesn't imply that the client issued and signed the tokens and used "Bearer" and "urn:ietf:params:oauth:token-type:access_token" in one of the responses (via <http://www.ietf.org/mail-archive/web/oauth/current/msg15335.html>).
- o Defined and registered urn:ietf:params:oauth:token-type:id_token, since some use cases perform token exchanges for ID Tokens and no URI to indicate that a token is an ID Token had previously been defined.

-03

- o Updated the document editors (adding Campbell, Bradley, and Mortimore).
- o Added to the title.
- o Added to the abstract and introduction.
- o Updated the format of the request to use application/x-www-form-urlencoded request parameters and the response to use the existing token endpoint JSON parameters defined in OAuth 2.0.
- o Changed the grant type identifier to urn:ietf:params:oauth:grant-type:token-exchange.
- o Added RFC 6755 registration requests for urn:ietf:params:oauth:token-type:refresh_token, urn:ietf:params:oauth:token-type:access_token, and urn:ietf:params:oauth:grant-type:token-exchange.
- o Added RFC 6749 registration requests for request/response parameters.
- o Removed the Implementation Considerations and the requirement to support JWTs.
- o Clarified many aspects of the text.
- o Changed "on_behalf_of" to "subject_token", "on_behalf_of_token_type" to "subject_token_type", "act_as" to "actor_token", and "act_as_token_type" to "actor_token_type".
- o Added an "audience" request parameter used to indicate the logical names of the target services at which the client intends to use the requested security token.
- o Added a "want_composite" request parameter used to indicate the desire for a composite token rather than trying to infer it from the presence/absence of token(s) in the request.

- o Added a "resource" request parameter used to indicate the URLs of resources at which the client intends to use the requested security token.
- o Specified that multiple "audience" and "resource" request parameter values may be used.
- o Defined the JWT claim "act" (actor) to express the current actor or delegation principal.
- o Defined the JWT claim "may_act" to express that one party is authorized to act on behalf of another party.
- o Defined the JWT claim "scp" (scopes) to express OAuth 2.0 scope-token values.
- o Added the "N_A" (not applicable) OAuth Access Token Type definition for use in contexts in which the token exchange syntax requires a "token_type" value, but in which the token being issued is not an access token.
- o Added examples.

-02

- o Enabled use of Security Token types other than JWTs for "act_as" and "on_behalf_of" request values.
- o Referenced the JWT and OAuth Assertions RFCs.

-01

- o Updated references.

-00

- o Created initial working group draft from draft-jones-oauth-token-exchange-01.

Authors' Addresses

Michael B. Jones
Microsoft

Email: mbj@microsoft.com
URI: <http://self-issued.info/>

Anthony Nadalin
Microsoft

Email: tonynad@microsoft.com

Brian Campbell (editor)
Ping Identity

Email: brian.d.campbell@gmail.com

John Bradley
Ping Identity

Email: ve7jtb@ve7jtb.com

Chuck Mortimore
Salesforce

Email: cmortimore@salesforce.com

OAuth Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 19, 2018

N. Sakimura
Nomura Research Institute
N. Mataka
GREE, Inc.
S. Preibisch
CA Technologies
November 15, 2017

OAuth Response Metadata
draft-sakimura-oauth-meta-08

Abstract

This specification defines an extensible metadata framework that may be inserted into the OAuth 2.0 responses to assist the clients to process those responses. It is expressed either as a link header, or query parameters. It will allow the client to learn the metadata about the particular response. For example, the client can learn where the members in the response could be used, what is the characteristics of the payload is, how it should be processed, and so on. Since they are just additional response header/query parameters, any client that does not understand this extension should not break and work normally while supporting clients can utilize the metadata to take the advantage of the extension.

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 <https://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on May 19, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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

1. Introduction	2
1.1. Requirements Notation and Conventions	3
1.2. terminology	3
2. Resource Endpoint Resonse	3
3. Token Endpoint Response	4
4. Authorization Endpoint HEAD response	5
5. Authorization Response	6
6. IANA Considerations	6
6.1. Link Type Registration	6
7. Security Considerations	7
7.1. Authorization Response Query Parameter Tampering by a Bad User	7
8. Acknowledgements	7
9. Document History	7
10. References	9
10.1. Normative References	9
10.2. Informational References	9
Authors' Addresses	10

1. Introduction

Although OAuth 2.0 [RFC6749] has been known for its REST friendliness, OAuth itself is not RESTful, as it heavily relies on out-of-band information to drive the interactions. This situation can be eased by hypertext-enabling the endpoint responses through the introduction of data structure that represents such hypertext and other metadata.

Hyper-text enabling the OAuth responses has many advantages. For example,

- o The protected resource can tell which authorization servers it supports.
- o Permissioned resource discovery: It is possible to tell the client which resource endpoint it should use. This has a privacy

advantage. The location of the resource by itself may be a sensitive information as its location may reveal information about the resource owner. Therefore, it may be sensible to tell the location only after the user consent.

- o It is possible to give a hint on the processing of the payload.
- o It will be resitant to IdP Mix-up attack.
- o It will be resitant to Code Phishing Attack.

This specification defines methods to represent such metadata in the authorization and token responses.

1.1. Requirements Notation and Conventions

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].

1.2. terminology

This specification uses the terms "Access Token", "Authorization Code", "Authorization Endpoint", "Authorization Grant", "Authorization Server", "Client", "Client Authentication", "Protected Resource", "Refresh Token", and "Token Endpoint" defined by OAuth 2.0 [RFC6749].

2. Resource Endpoint Resonse

Resource Endpoints that implement this specification returns the following link relation (rel) and the corresponsing URI value as defined in [RFC5988] in the response header. The response header can be returned in response to HEAD, GET, or POST request to the endpoint.

duri The URI of the corresponding authorization server's discovery document, from which the client can learn the server capabilities and endpoints.

auri The URI of the corresponding Authorization Endpoint URI.

A typical example of the use of these header values are in the case of the Client accessing the protected resource without a propoer credential. For example, in the case of an [RFC6750] protected resource, the unauthorized access may result in a response header that includes both WWW-Authenticate header as well as the Web Linking

header indicating either the Authorization Endpoint URI or the discovery document URI.

There is no cardinality restriction on relations put in place by [RFC5988]. Therefore, the resource can respond with multiple Authorization Endpoint URI or Discovery Document URIs from which the Client may choose the appropriate one.

```
HTTP/1.1 401 Unauthorized
WWW-Authenticate: Bearer realm="example"
Link: <https://example.com/.well-known/openid-configuration>; rel="duri",
    <https://example.net/.well-known/openid-configuration>; rel="duri",
    <https://example.com/authz/>; rel="auri",
    <https://example.net/az/>; rel="auri",
    <https://example.com/payment-upon-trial-expiry>; rel="payments"
```

3. Token Endpoint Response

Token Endpoints that implements this specification returns the following link relation (rel) and the corresponding URI value as defined in [RFC5988] in the Access Token Response defined in [RFC6749].

ruri Resource URI. The value of this parameter is the URI of the Resource Endpoint that the Access Token is supposed to be used at. If this value is present, the client MUST NOT send the Access Token to any other URI.

turi Token Endpoint URI. The value of this parameter is the URI of the Token Endpoint that the Refresh Token can be sent to obtain a new Access Token. If this value is present, then the client MUST NOT send the refresh token to any other places.

Any other rels that are registered in Link Relation Type Registry defined in [RFC5988] registry can be used.

There is no cardinality restriction on relations put in place by [RFC5988]. Therefore, the Token Endpoing can respond with multiple Resource Endpoint URI or Discovery Document URIs from which the Client may choose the appropriate one.

Following is an example of an HTTPS response.


```
HTTP/1.1 200 OK
Link: <https://example.com/userinfo>; rel="ruri",
      <https://example.net/photostream/>; rel="ruri",
      <https://example.com/payment-upon-trial-expiry>; rel="payments"
Content-Type: application/JSON; charset=utf-8

{
  "access_token": "aCeSsToKen"
}
```

4. Authorization Endpoint HEAD response

Authorization Endpoints that implements this specification returns the following link relation (rel) and the corresponding URI value as defined in [RFC5988] in the response to the HEAD request.

auri The canonical URI of the Authorization Endpoint URI.

turi Token Endpoint URI. The value of this parameter is the URI of the Token Endpoint that the Authorization Code can be sent to obtain the Access Token.

duri The URI of discovery document, from which the client can learn the server capabilities and endpoints.

ruri Resource URI. The value of this parameter is the URI of the Resource Endpoint that the Access Token can be used at. If this parameter is specified, the client MUST NOT send the Access Token to any other URIs than the value of this parameter.

There is no cardinality restriction on relations put in place by [RFC5988]. Therefore, the Authorization Endpoint can respond with multiple Endpoint URIs with a same relation type from which the Client may choose the appropriate one.

Following is an example of an HTTPS response.

```
HTTP/1.1 200 OK
Link: <https://example.com/.well-known/openid-configuration>; rel="duri",
      <https://example.net/.well-known/openid-configuration>; rel="duri",
      <https://example.com/payment-upon-trial-expiry>; rel="payments"
Content-Type: application/JSON; charset=utf-8
```

5. Authorization Response

While [RFC5988] defines a useful way of conveying link relations, it cannot be utilized for a redirect based communication such as the authorization response of OAuth 2.0. This section defines a way to return a limited set of those link relations as query parameters so that it can be conveyed over the redirection.

The authorization response of the implementation of this specification may return the following query parameter in the redirect URI.

`turi` Token Endpoint URI. The value of this parameter is the URI of the Token Endpoint that the Authorization Code can be sent to obtain the Access Token. If this parameter is specified, the client MUST check that the value of `turi` matches exactly with the pre-registered token endpoint URI of the Authorization Server that the session recovered from the state variable points to. The client MUST NOT send the code to any other URIs than the value of this parameter.

`ruri` Resource URI. The value of this parameter is the URI of the Resource Endpoint that the Access Token can be used at. If this parameter is specified, the client MUST NOT send the Access Token to any other URIs than the value of this parameter.

As long as the link relation type string does not collide with the underlying protocol parameters, they can also be specified as a query parameter. The value MUST be encoded in `application/x-www-form-urlencoded`.

The following is an example of such response. Line breaks are for display purposes only.

```
HTTP/1.1 302 Found
Location: https://client.example.com/cb?code=Splxl0BeZQQYbYS6WxSbIA
        &turi=https%3A%2F%2Fexample.com%2Ftoken
        &state=xyz
```

6. IANA Considerations

6.1. Link Type Registration

Pursuant to [RFC5988], the following link type registrations [[will be]] registered by mail to `link-relations@ietf.org`.

- o Relation Name: `turi`

- o Description: An OAuth 2.0 Token Endpoint specified in section 3.2 of [RFC6749].
- o Reference: This specification
- o Relation Name: ruri
- o Description: An OAuth 2.0 Resource Endpoint specified in section 3.2 of [RFC6750].
- o Reference: This specification

7. Security Considerations

7.1. Authorization Response Query Parameter Tampering by a Bad User

The query response parameters may be tampered by the man-in-the-browser. It can also be tampered by a malicious user. In general, anything that comes via the browser/user-agent can be tainted and untrusted.

This specification mandates the turi check so that tampering of turi by the malicious user will be detected. It does not mandate ruri check as the user can get the Access Token and send it to anywhere he wants anyways when it is returned to the browser.

However, other parameters are not protected. The Client MUST treat them tainted and implement its own check rules for each parameters.

To solve this "Tampering by bad user", either HMAC(concat(params)) need to be sent with them or have all of them inside the JWS.

8. Acknowledgements

Members of OAuth WG helped to form this specification. Notably: Hannes tschofenig, John Bradley, Justin Richer, Kaoru Maeda, Masashi Kurabayashi, Michael B. Jones, Phil Hunt, William Dennis, James Manger, (add yourselves).

9. Document History

-07

- o Added note that there is no cardinality requirements so that multiple endpoints can be returned by repeating a same rel.
- o Added resource endpoint response.

-06

- o Removed duri description from token response as it is not needed.
- o Made the processing instruction more precise.
- o Added RFC5988 defined link relation type in the example.
- o Swaped the order of the authorization response and token response. Now, token response gets explained first so that the reader will grasp the basic concept according to RFC5988 and regards the authorization response extension as a mapping of RFC5988 into query parameter form.

-05

- o Factored out JSON Meta and now using query param and Web Linking.

-04

- o Date refresh.

-03

- o Date refresh.

-02

- o Added Mike Kelly as an author.
- o xref fix.
- o Introduced "operations" as in draft-ietf-scim-api-00#section-3.5.
- o Updated the informative reference to HAL.
- o Added description to OAuth Token Endpoint hrefs.
- o Added content-type to the example.
- o Added Area and Working Group.

-01

- o Some format changes, reference fix, and typo fixes.
- o Changed 'items' to 'elements' to match the JSON terminology.

-00

o Initial Draft

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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, DOI 10.17487/RFC2616, June 1999, <<https://www.rfc-editor.org/info/rfc2616>>.
- [RFC5988] Nottingham, M., "Web Linking", RFC 5988, DOI 10.17487/RFC5988, October 2010, <<https://www.rfc-editor.org/info/rfc5988>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, DOI 10.17487/RFC6750, October 2012, <<https://www.rfc-editor.org/info/rfc6750>>.

10.2. Informational References

- [HAL] Kelly, M., "JSON Hypermedia API Language", February 2013.
- [oauth-lrdd] Mills, W., "Link Type Registrations for OAuth 2", October 2012.
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", RFC 4627, DOI 10.17487/RFC4627, July 2006, <<https://www.rfc-editor.org/info/rfc4627>>.

[RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M.,
and D. Orchard, "URI Template", RFC 6570,
DOI 10.17487/RFC6570, March 2012,
<<https://www.rfc-editor.org/info/rfc6570>>.

Authors' Addresses

Nat Sakimura
Nomura Research Institute

Email: sakimura@gmail.com

Nov Matake
GREE, Inc.

Email: nov@matake.jp
URI: <http://matake.jp>

Sascha Preibisch
CA Technologies

Email: Sascha.Preibisch@gmail.com