

Workload Identity in Multi System Environments
Internet-Draft
Intended status: Standards Track
Expires: 19 April 2026

B. Campbell
Ping Identity
J. Salowey
CyberArk
A. Schwenkschuster
SPIRL
Y. Sheffer
Intuit
16 October 2025

WIMSE Workload-to-Workload Authentication
draft-ietf-wimse-s2s-protocol-07

Abstract

The WIMSE architecture defines authentication and authorization for software workloads in a variety of runtime environments, from the most basic ones up to complex multi-service, multi-cloud, multi-tenant deployments. This document defines the simplest, atomic unit of this architecture: the protocol between two workloads that need to verify each other's identity in order to communicate securely. The scope of this protocol is a single HTTP request-and-response pair. To address the needs of different setups, we propose two protocols, one at the application level and one that makes use of trusted TLS transport. These two protocols are compatible, in the sense that a single call chain can have some calls use one protocol and some use the other. Workload A can call Workload B with mutual TLS authentication, while the next call from Workload B to Workload C would be authenticated at the application level.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://ietf-wg-wimse.github.io/draft-ietf-wimse-s2s-protocol/draft-ietf-wimse-s2s-protocol.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-wimse-s2s-protocol/>.

Discussion of this document takes place on the Workload Identity in Multi System Environments Working Group mailing list (<mailto:wimse@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/wimse/>. Subscribe at <https://www.ietf.org/mailman/listinfo/wimse/>.

Source for this draft and an issue tracker can be found at <https://github.com/ietf-wg-wimse/draft-ietf-wimse-s2s-protocol>.

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 19 April 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Extending This Protocol to Other Use Cases	4
1.2. Deployment Architecture and Message Flow	4
1.3. Workload Identifiers and Authentication Granularity	6
2. Conventions and Definitions	7
3. Application Level Workload-to-Workload Authentication	7
3.1. The Workload Identity Token	7
3.1.1. The WIT HTTP Header	11
3.1.2. Including Additional Claims	12
3.1.3. A note on iss claim and key distribution	12
3.2. Option 1: DPoP-Inspired Authentication	12
3.3. Option 2: Authentication Based on HTTP Message Signatures	17
3.4. Comparing the DPoP Inspired Option with Message Signatures	20

3.5.	Error Conditions	21
3.6.	Coexistence with JWT Bearer Tokens	22
4.	Using Mutual TLS for Workload-to-Workload Authentication . .	22
4.1.	The Workload Identity Certificate	22
4.2.	Workload Identity Certificate Validation	23
4.2.1.	Server Name Validation	23
4.3.	Client Authorization Using the Workload Identity	24
5.	Implementation Status	24
6.	Security Considerations	25
6.1.	Workload Identity	25
6.2.	Workload Identity Token and Proof of Possession	26
6.3.	Workload Identity Key Management	27
6.4.	Middle Boxes	27
6.5.	Privacy Considerations	28
7.	IANA Considerations	28
7.1.	JSON Web Token Claims	28
7.2.	Media Type Registration	28
7.2.1.	application/wit+jwt	28
7.2.2.	application/wpt+jwt	29
7.3.	Hypertext Transfer Protocol (HTTP) Field Name Registration	30
7.3.1.	Workload-Identity-Token	31
7.3.2.	Workload-Proof-Token	31
7.4.	URI Scheme Registration	31
8.	References	31
8.1.	Normative References	32
8.2.	Informative References	33
Appendix A.	Document History	34
A.1.	draft-ietf-wimse-s2s-protocol-07	34
A.2.	draft-ietf-wimse-s2s-protocol-06	34
A.3.	draft-ietf-wimse-s2s-protocol-05	35
A.4.	draft-ietf-wimse-s2s-protocol-04	35
A.5.	draft-ietf-wimse-s2s-protocol-03	35
A.6.	draft-ietf-wimse-s2s-protocol-02	35
A.7.	draft-ietf-wimse-s2s-protocol-01	36
A.8.	draft-ietf-wimse-s2s-protocol-00	36
	Acknowledgments	36
	Authors' Addresses	36

1. Introduction

This document defines authentication and authorization in the context of interaction between two workloads. This is the core component of the WIMSE architecture [I-D.ietf-wimse-arch]. For simplicity, this document focuses on HTTP-based services, and the workload-to-workload call consists of a single HTTP request and its response. We define the credentials that both workloads should possess and how they are used to protect the HTTP exchange.

There are multiple deployment styles in use today, and they result in different security properties. We propose to address them differently.

- * Many use cases have various middleboxes inserted between pairs of workloads, resulting in a transport layer that is not end-to-end encrypted. We propose to address these use cases by protecting the HTTP messages at the application level (Section 3).
- * The other commonly deployed architecture has a mutual-TLS connection between each pair of workloads. This setup can be addressed by a simpler solution (Section 4).

It is an explicit goal of this protocol that a workload deployment can include both architectures across a multi-chain call. In other words, Workload A can call Workload B with mutual TLS protection, while the next call to Workload C is protected at the application level.

For application-level protection we currently propose two alternative solutions, one inspired by DPoP [RFC9449] in Section 3.2 and one which is a profile of HTTP Message Signatures [RFC9421] in Section 3.3. The design team believes that we need to pick one of these two alternatives for standardization, once we have understood their pros and cons.

1.1. Extending This Protocol to Other Use Cases

The protocol defined here is narrowly scoped, targeting only HTTP-based request/response services. To secure workloads communicating over other transports, new protocol bindings will need to be defined. We note though that this protocol is designed to allow some level of reuse. In particular, we expect that the Workload Identity Token (WIT) construct will be reusable in other settings. The Workload Proof Token (WPT) may be adaptable with some changes to different environments.

1.2. Deployment Architecture and Message Flow

Regardless of the transport between the workloads, we assume the following logical architecture (numbers refer to the sequence of steps listed below):

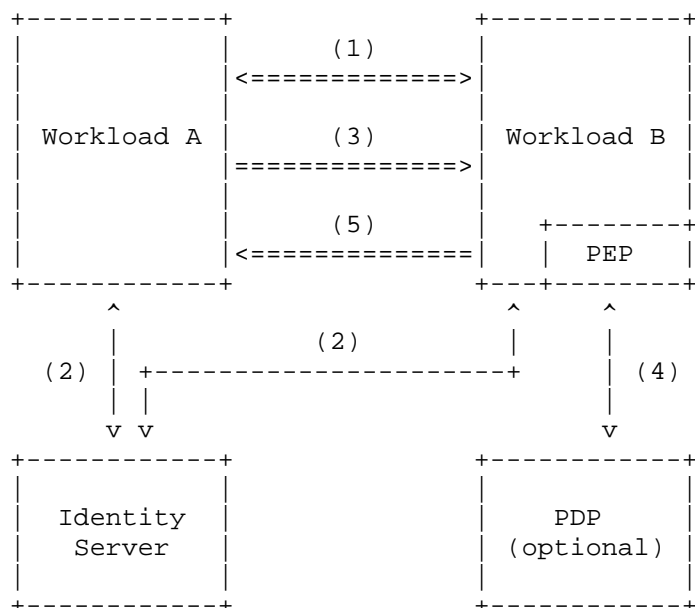


Figure 1: Sequence of Operations

The Identity Server provisions credentials to each of the workloads. At least Workload A (and possibly both) must be provisioned with a credential before the call can proceed. Details of communication with the Identity Server are out of scope of this document, however we do describe the credential received by the workload.

PEP is a Policy Enforcement Point, the component that allows the call to go through or blocks it. PDP is an optional Policy Decision Point, which may be deployed in architectures where policy management is centralized. All details of policy management and message authorization are out of scope of this document.

The high-level message flow is as follows:

1. A transport connection is set up. In the case of mutual TLS, this includes authentication of both workloads to one another. In the case of application-level security, the TLS connection is typically one-way authenticated, and workload-level authentication does not yet take place.
2. Workload A (and similarly, Workload B) obtains a credential from the Identity Server. This happens periodically, e.g. once every 24 hours.

3. Workload A makes an HTTP call into Workload B. This is a regular HTTP request, with the additional protection mechanisms defined below.
4. In the case of application-level security, Workload B authenticates Workload A (when using mutual TLS, this happened in step 1). In either case, Workload B decides whether to authorize the call. In certain architectures, Workload B may need to consult with an external server when making this decision.
5. Workload B returns a response to Workload A, which may be an error response or a regular one.

1.3. Workload Identifiers and Authentication Granularity

The specific format of workload identifiers (see [I-D.ietf-wimse-arch]) is set by local policy for each deployment, and this choice has several implications.

Prior to WIMSE, many use cases did not allow for fully granular authentication in containerized runtime platforms. For instance, with mutual TLS, there's often no clear way to map the request's external access reference (e.g., Kubernetes Ingress path, service name, or host header) to the SubjectAltName value in the server certificate. This means that the client could only verify if the server certificate is valid within a trust domain, not if it's tied to a specific workload.

To enable mutual and granular authentication between workloads, two things must be in place:

- * Each workload must know its own identifier.
- * There needs to be an explicit mapping from the external handle used to access a workload (such as an Ingress path or service DNS name) to its workload identifier.

Once these conditions are met, the methods described in this document can be used for the caller and callee to mutually authenticate.

Implementations **MUST** allow for defining this mapping between the workload's access path and the workload identifier (e.g., through callback functions). Deployments **SHOULD** use these features to establish a consistent set of identifiers within their environment.

2. Conventions and Definitions

All terminology in this document follows [I-D.ietf-wimse-arch].

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.

3. Application Level Workload-to-Workload Authentication

As noted in the Introduction, for many deployments communication between workloads cannot use end-to-end TLS. For these deployment styles, this document proposes application-level protections.

The current version of the document includes two alternatives, both using the newly introduced Workload Identity Token (Section 3.1). The first alternative (Section 3.2) is inspired by the OAuth DPoP specification. The second (Section 3.3) is based on the HTTP Message Signatures RFC. We present both alternatives and expect the working group to select one of them as this document progresses towards IETF consensus. A comparison of the two alternatives is attempted in Section 3.4.

3.1. The Workload Identity Token

The Workload Identity Token (WIT) is a JWS [RFC7515] signed JWT [RFC7519] that represents the identity of a workload. It is issued by the Identity Server and binds a public key to the workload identity. See Section 6.3 for security considerations.

A WIT MUST contain the following claims, except where noted:

- * in the JOSE header:

- alg: An identifier for a JWS asymmetric digital signature algorithm (registered algorithm identifiers are listed in the IANA JOSE Algorithms registry [IANA.JOSE.ALGS]). The value none MUST NOT be used.
- typ: the WIT is explicitly typed, as recommended in Section 3.11 of [RFC8725], using the wit+jwt media type.

- * in the JWT claims:

- iss: The issuer of the token, which is the Identity Server, represented by a URI. The iss claim is RECOMMENDED but optional, see Section 3.1.3 for more.
- sub: The subject of the token, which is the identity of the workload, represented by a URI. See [I-D.ietf-wimse-arch] for details of the Workload Identifier. And see Section 1.3 for security implications of these identifiers.
- exp: The expiration time of the token (as defined in Section 4.1.4 of [RFC7519]). WITs should be refreshed regularly, e.g. on the order of hours.
- jti: A unique identifier for the token. This claim is OPTIONAL. The jti claim is frequently useful for auditing issuance of individual WITs or to revoke them, but some token generation environments do not support it.
- cnf: A confirmation claim referencing the public key of the workload.
 - o jwk: Within the cnf claim, a jwk key MUST be present that contains the public key of the workload as defined in Section 3.2 of [RFC7800]. The workload MUST prove possession of the corresponding private key when presenting the WIT to another party, which can be accomplished by using it in conjunction with one of the methods in Section 3.2 or Section 3.3. As such, it MUST NOT be used as a bearer token and is not intended for use in the Authorization header.
 - + alg: Within the jwk object, an alg field MUST be present. Allowed values are listed in the IANA "JSON Web Signature and Encryption Algorithms" registry established by [RFC7518]. The presented proof (WPT or http-sig) MUST be produced with the algorithm specified in this field. The value none MUST NOT be used. Algorithms used in combination with symmetric keys MUST NOT be used. Also encryption algorithms MUST NOT be used as this would require additional key distribution outside of the WIT. To promote interoperability, the ES256 signing algorithm MUST be supported by general purpose implementations of this document.


```
{
  "cnf": {
    "jwk": {
      "alg": "EdDSA",
      "crv": "Ed25519",
      "kty": "OKP",
      "x": "1CXXvflN_LVVsIsYXsUvB03JmlGWeCHqQVuouCF92bg"
    }
  },
  "exp": 1745512510,
  "iat": 1745508910,
  "jti": "x-_1CTL2cca3CSE4cwb_1",
  "sub": "wimse://example.com/specific-workload"
}
```

Figure 4: Example WIT Claims

The claims indicate that the example WIT:

- * is valid until Thu Apr 24 2025 16:35:10 GMT (represented as NumericDate Section 2 of [RFC7519] value 1745512510).
- * identifies the workload to which the token was issued as wimse://example.com/specific-workload.
- * has a unique identifier of x-_1CTL2cca3CSE4cwb_1.
- * binds the public key represented by the jwk confirmation method to the workload wimse://example.com/specific-workload.
- * requires the proof to be produced with the EdDSA signature algorithm.

For elucidative purposes only, the workload's key, including the private part, is shown below in JWK [RFC7517] format:

```
{
  "kty": "OKP",
  "crv": "Ed25519",
  "x": "1CXXvflN_LVVsIsYXsUvB03JmlGWeCHqQVuouCF92bg",
  "d": "sdLX8yCYKqo_XvGBLn-ZWeKT7l1YeeQpgeCaXVxb5kY"
}
```

Figure 5: Example Workload's Key

The afore-exemplified WIT is signed with the private key of the Identity Server. The public key(s) of the Identity Server need to be known to all workloads in order to verify the signature of the WIT. The Identity Server's public key from this example is shown below in JWK [RFC7517] format:

```
{
  "kty": "EC",
  "kid": "June 5",
  "crv": "P-256",
  "x": "kXqnA2Op7hgd4zRMbw0iFcc_hDxUxhojxOFVGjE2gks",
  "y": "n__VndPMR021-59UAs0b9qDTFT-EZtT6xSNs_xFskLo"
}
```

Figure 6: Example Identity Server Key

3.1.1.1. The WIT HTTP Header

A WIT is conveyed in an HTTP header field named Workload-Identity-Token.

ABNF [RFC5234] for the value of Workload-Identity-Token header field is provided in Figure 7:

```
ALPHA = %x41-5A / %x61-7A ; A-Z / a-z
DIGIT = %x30-39 ; 0-9
base64url = 1*(ALPHA / DIGIT / "-" / "_")
JWT = base64url "." base64url "." base64url
WIT = JWT
```

Figure 7: Workload-Identity-Token Header Field ABNF

The following shows the WIT from Figure 2 in an example of a Workload-Identity-Token header field:

===== NOTE: '\ ' line wrapping per RFC 8792 =====

```
Workload-Identity-Token: eyJhbGciOiJFUzI1NiIsImtpZCI6Ikp1bmUgNSIsInR\
5cCI6IndpdCtqd3QifQ.eyJjbYiOnsiandrIjp7ImFsZyI6IkJkVFRFNBIiwiaY3J2Ijo\
iWQyNTUxOSIsImt0eSI6IkJk9LUCIsIngiaOiIxQ1hYdmZsTl9MVlZzSXNZWHNVdkIwM0pt\
bEdXZUNicVFWdW9lQ0Y5MmJnInl9LCJleHAiOiJlMTAsImhhdCI6MTc0NTUw\
ODkxMCwianRpIjoieCI6MUNUTDJjY2EzQ1NFNGN3Y19sIiwic3ViIjoieCI6Mjltc2U6Ly9l\
eGFtcGxlLmNvbS9zcGVjaWZpYy13b3JrbG9hZCJ9.6KraSQUxWdl9dxFQ3Fj6dPY0Vi8\
8OkwFwZpAIOhLeq6AbXANLLQgOp8U9UDGcBuYF3KiNv6oKQD1ZWazrMZOJw
```

Figure 8: An example Workload Identity Token HTTP Header Field

Note that per [RFC9110], header field names are case insensitive; thus, Workload-Identity-Token, workload-identity-token, WORKLOAD-IDENTITY-TOKEN, etc., are all valid and equivalent header field names. However, case is significant in the header field value.

3.1.2. Including Additional Claims

The WIT contains JSON structures and therefore can be trivially extended by adding more claims beyond those defined in the current specification. This, however, could result in interoperability issues, which the following rules are addressing.

- * To ensure interoperability in WIMSE environments, the use of Private claim names (Sec. 4.3 of [RFC7519]) is NOT RECOMMENDED.
- * In closed environments, deployers MAY freely add claims to the WIT. Such claims SHOULD be collision-resistant, such as example.com/myclaim.
- * A recipient that does not understand such claims MUST ignore them, as per Sec. 4 of [RFC7519].
- * Outside of closed environments, new claims MUST be registered with IANA [IANA.JWT.CLAIMS] before they can be used.

3.1.3. A note on iss claim and key distribution

It is RECOMMENDED that the WIT carries an iss claim. This specification itself does not make use of a potential iss claim but also carries the trust domain in the workload identifier (see [I-D.ietf-wimse-arch] for a definition of the identifier and related rules). Implementations MAY include the iss claim in the form of a https URL to facilitate key distribution via mechanisms like the jwks_uri from [RFC8414] but alternative key distribution methods may make use of the trust domain included in the workload identifier which is carried in the mandatory sub claim.

3.2. Option 1: DPoP-Inspired Authentication

This option, inspired by the OAuth DPoP specification [RFC9449], uses a DPoP-like mechanism to authenticate the calling workload in the context of the request. The Workload Identity Token (Section 3.1) is sent in the request as described in Section 3.1.1. An additional JWT, the Workload Proof Token (WPT), is signed by the private key corresponding to the public key in the WIT. The WPT is sent in the Workload-Proof-Token header field of the request. The ABNF syntax of the Workload-Proof-Token header field is:

WPT = JWT

Figure 9: Workload-Proof-Token Header Field ABNF

where the JWT projection is defined in Figure 7.

A WPT MUST contain the following:

- * in the JOSE header:
 - alg: An identifier for an appropriate JWS asymmetric digital signature algorithm corresponding to the confirmation key in the associated WIT. The value MUST match the alg value of the jwk in the cnf claim of the WIT. See Section 3.1 for valid values and restrictions.
 - typ: the WPT is explicitly typed, as recommended in Section 3.11 of [RFC8725], using the application/wpt+jwt media type.
- * in the JWT claims:
 - aud: The audience SHOULD contain the HTTP target URI (Section 7.1 of [RFC9110]) of the request to which the WPT is attached, without query or fragment parts. However, there may be some normalization, rewriting or other process that requires the audience to be set to a deployment-specific value. See also Section 1.3 for more details.
 - exp: The expiration time of the WPT (as defined in Section 4.1.4 of [RFC7519]). WPT lifetimes MUST be short, e.g., on the order of minutes or seconds.
 - jti: An identifier for the token. The value MUST be unique, at least within the scope of the sender.
 - wth: Hash of the Workload Identity Token, defined in Section 3.1. The value is the base64url encoding of the SHA-256 hash of the ASCII encoding of the WIT's value.
 - ath: Hash of the OAuth access token, if present in the request, which might convey end-user identity and/or authorization context of the request. The value, as per Section 4.1 of [RFC9449], is the base64url encoding of the SHA-256 hash of the ASCII encoding of the access token's value.

- tth: Hash of the Txn-Token [I-D.ietf-oauth-transaction-tokens], if present in the request, which might convey end-user identity and/or authorization context of the request. The value MUST be the result of a base64url encoding (as defined in Section 2 of [RFC7515]) of the SHA-256 hash of the ASCII encoding of the associated token's value.
- oth: Hash(es) of other token(s) in the request that convey end-user identity and/or authorization context of the request. The value is a JSON object with a key-value pair for each such token. For each, in the absence of an application profile specifying details, the key corresponds to the header field name containing the token, and the value is the base64url encoding of the SHA-256 hash of the ASCII bytes of the header field value with any leading or trailing spaces removed. The header field name MUST be normalized by converting it to all lower case. Header fields occurring multiple times in the request are not supported by default. An application profile may specify different behavior for a key, such as using a different hash algorithm or means of locating the token in the request.

To clarify: the ath, tth and oth claims are each mandatory if the respective tokens are included in the request.

The rules for using non-standard claims in WPTs are similar to the rules for WITs, Section 3.1.2.

An example WPT might look like the following:

===== NOTE: '\ ' line wrapping per RFC 8792 =====

```
eyJhbGciOiJIJFZERTQSIzInR5cCI6IndwdCtqd3QifQ.eyJhdGgiOiJDTDR3amZwUm10Z\
ilizFlJYllMblY5ZDVyTUFSR3dLWUUXMHdVd3pDMGpJIiwiYXVkJjoiaHR0cHM6Ly93b\
3JrbG9hZC5leGFtcGxlLmNvbS9wYXR0IiwizXhwIjoxNzQlNTEwMDE2LCJqdGkiOiJfX\
2J3YzRFU0MzYWNjMkxUQzEtX3giLCJ3dGgiOiJBYVlVZkMzNEQxZGkyRnhRTHBpSupKN\
lNnOFZaNm84T0Nkd1NmOUlUb0xnIn0.PI7d9AcYhLoEgPgbJvcM132lkBKnM1NXU-5hZ\
UzVTIy2dRJaTLFs6e5NYv5gg6AqcV7NvkYCwfgMgWasWQzCA
```

Figure 10: Example Workload Proof Token (WPT)

The decoded JOSE header of the WPT from the example above is shown here:

```
{
  "alg": "EdDSA",
  "typ": "wpt+jwt"
}
```

Figure 11: Example WPT JOSE Header

The decoded JWT claims of the WPT from the example above are shown here:

```
{
  "ath": "CL4wjfpRmNf-bdYIbYLnV9d5rMARGwKYEl0wUwzC0jI",
  "aud": "https://workload.example.com/path",
  "exp": 1740755048,
  "jti": "0c740386caldcad37de1b5f9de1b0705",
  "wth": "aA0W_oFJK7qV7zYhcmzR1K0XVCHjd2x6c4sOQLvE90Y"
}
```

Figure 12: Example WPT Claims

An example of an HTTP request with both the WIT and WPT from prior examples is shown below:

===== NOTE: '\ ' line wrapping per RFC 8792 =====

```
POST /path HTTP/1.1
Host: workload.example.com
Content-Type: application/json
Authorization: Bearer 16_mAd0GiwaZokU26_0902100
Workload-Identity-Token: eyJhbGciOiJFUzI1NiIsImtpZCI6Ikp1bmUgNSIsInR\
5cCI6IndpdCtqd3QifQ.eyJjbmYiOnsiaandrIjp7ImFsZyI6IkwkRFRNBIiw\
iY3J2Ijo\
RWQyNTUxOSIsImt0eSI6IkwkRFRNBIiw\
bEdXZUNlcmVudW91Q0Y5MmJnInl9LCJleHAiOiJlMTI1MTAsImhhdCI6MTc\
0NTUw\
ODkxMCwianRpIjoieCI6IjMUNUTDJjY2EzQ1NFNGN3Yl9sIiwic3ViIjo\
id2ltc2U6Ly9l\
eGFTcGx1LmNvbS9zGVjwZpYy13b3JrbG9hZCJ9.6KraSQUxWd19dx\
FQ3Fj6dPY0Vi8\
80kwFwZpAIOhLeq6AbXAnLLQgOp8U9UDGcBuYF3KiNv6oKQD1ZWAzr\
MZOJw
Workload-Proof-Token: eyJhbGciOiJFZERTQSI6ImtpZCI6Ikp1bmUgNSIsInR\
5cCI6IndwdCtqd3QifQ.eyJ\
hdGgiOiJlMTI1MTAsImhhdCI6MTc\
0NTUw\
ODkxMCwianRpIjoieCI6IjMUNUTDJjY2EzQ1NFNGN3Yl9sIiwic3ViIjo\
id2ltc2U6Ly9l\
eGFTcGx1LmNvbS9zGVjwZpYy13b3JrbG9hZCJ9.6KraSQUxWd19dx\
FQ3Fj6dPY0Vi8\
80kwFwZpAIOhLeq6AbXAnLLQgOp8U9UDGcBuYF3KiNv6oKQD1ZWAzr\
MZOJw
{"do_stuff":"please"}
```

Figure 13: Example HTTP Request with WIT and WPT

To validate the WPT in the request, the recipient MUST ensure the following:

- * There is exactly one Workload-Proof-Token header field in the request.

- * The Workload-Proof-Token header field value is a single and well-formed JWT.
- * The signature algorithm in the alg JOSE header string-equal matches the alg attribute of the jwk in the cnf claim of the WIT.
- * The WPT signature is valid using the public key from the confirmation claim of the WIT.
- * The typ JOSE header parameter of the WPT conveys a media type of wpt+jwt.
- * The aud claim of the WPT matches the target URI, or an acceptable alias or normalization thereof, of the HTTP request in which the WPT was received, ignoring any query and fragment parts. See also Section 1.3 for implementation advice on this verification check.
- * The exp claim is present and conveys a time that has not passed. WPTs with an expiration time unreasonably far in the future SHOULD be rejected.
- * The wth claim is present and matches the hash of the token value conveyed in the Workload-Identity-Token header.
- * It is RECOMMENDED to check that the value of the jti claim has not been used before in the time window in which the respective WPT would be considered valid.
- * If presented in conjunction with an OAuth access token, the value of the ath claim matches the hash of that token's value.
- * If presented in conjunction with a Txn-Token, the value of the tth claim matches the hash of that token's value.
- * If presented in conjunction with a token conveying end-user identity or authorization context, the value of the oth claim matches the hash of that token's value.
- * If the oth claim is present, verify the hashes of all tokens listed in the oth claim per the default behavior defined in Section 3.2 or as specified by an application specific profile. If the oth claim contains entries that are not understood by the recipient, the WPT MUST be rejected. Conversely, additional tokens not covered by the oth claim MUST NOT be used by the recipient to make authorization decisions.

3.3. Option 2: Authentication Based on HTTP Message Signatures

This option uses the Workload Identity Token (Section 3.1) and the private key associated with its public key, to sign the request and optionally, the response. See Section 6.3 for security considerations. This section defines a profile of the Message Signatures specification [RFC9421].

The request is signed as per [RFC9421]. The following derived components MUST be signed:

- * @method
- * @request-target

In addition, the following request headers MUST be signed when they exist:

- * Content-Type
- * Content-Digest
- * Authorization
- * Txn-Token [I-D.ietf-oauth-transaction-tokens]
- * Workload-Identity-Token

If the response is signed, the following components MUST be signed:

- * @status
- * @method:req
- * @request-target:req
- * Content-Type if it exists
- * Content-Digest if it exists
- * Workload-Identity-Token

To ensure the message is fully integrity-protected, if the request or response includes a message body, the sender MUST include (and the receiver MUST verify) a Content-Digest header.

For both requests and responses, the following signature parameters MUST be included:

- * created
- * expires - expiration MUST be short, e.g. on the order of minutes. The WIMSE architecture will provide separate mechanisms in support of long-lived compute processes.
- * nonce
- * tag - the value for implementations of this specification is wimse-workload-to-workload

The following signature parameters in the Signature-Input header MUST NOT be used:

- * keyid - The signing key is sent along with the message in the WIT. Additionally specifying the key identity would add confusion.
- * alg - The signature algorithm is specified in the jwk section of the cnf claim in the WIT. See Section 3.1 and Sec. 3.3.7 of [RFC9421] for details.

It is RECOMMENDED to include only one signature with the HTTP message. If multiple ones are included, then the signature label included in both the Signature-Input and Signature headers SHOULD be wimse.

A sender MUST ensure that each nonce it generates is unique, at least among messages sent to the same recipient. To detect message replays, a recipient SHOULD reject a message (request or response) if a nonce generated by a certain peer is seen more than once.

Implementors need to be aware that the WIT is extracted from the message before the message signature is validated. Recipients of signed HTTP messages MUST validate the signature and content of the WIT before validating the HTTP message signature. They MUST ensure that the message is not processed further before it has been fully validated.

Either client or server MAY send an Accept-Signature header, but is not required to do so. When this header is sent, it MUST include the header components listed above.

Following is a non-normative example of a signed request and a signed response, where the caller is using the keys specified in Figure 5 (TODO: it is actually using a different key but that'll need to be fixed later).

===== NOTE: '\ ' line wrapping per RFC 8792 =====

Response:

```
HTTP/1.1 404 Not Found
Connection: close
Content-Digest: sha-256=:47DEQpj8HBSa+/TImW+5JCeuQeRkm5NMpJWZG3hSuFU\
=:
Content-Type: text/plain
Signature: wimse=:WAjxzuiuCiYRqCzetetDwaTS7Ka9yMwB+dAHVJPw3VkuH+c8c4A\
5BKrCsPlD/ymy+7PgWxl3y3mVdaD4ww7WqDA=:
Signature-Input: wimse=("@status" "workload-identity-token" "content\
-type" "content-digest" "@method";req "@request-target";req);created\
=1754558248;expires=1754558550;nonce="abcd2222";tag="wimse-workload-\
to-workload"
Workload-Identity-Token: eyJhbGciOiJFZERTQSIsImtpZCI6ImIzZmVlcilrZXk\
iLCJ0eXAiOiJ3aWlzc3Q3QifQ.eyJjb2YiOiJmYXNjaW50eSI6Ik9LUCIsIngi\
OiJneJhU0pFLWc5dzFyYmdKaU5wc3RHYjJhUGs1MGs1b0pVRWJMRHVzYXl\
eHAiOiJ3NTQ1NTg1NTAsImIhdCI6MTc1NDU1ODI1MCwiaXNzIjoiaHR0cHM6Ly9leG\
FTcGx1LmNvbS9pc3NlZXIiLCJqdGkiOiJ3aXQtMTc1NDU1ODI1ODQwMjY4MTAwMCIs\
InNlYiI6IndpbXNlOi8vZXhhbXBsZS5jb20vc3ZjQjI9.yJ4gGQ1DQBX0E9Xp5fI\
g6Ucpf9\LI_gHnhURZxqu6atT-3hpbFTgw4rd-6knM7-HClok4b6N2ViZaEDcz6IMCg
```

No ice cream today.

Figure 16: Signed Response

3.4. Comparing the DPOp Inspired Option with Message Signatures

The two workload protection options have different strengths and weaknesses regarding implementation complexity, extensibility, and security. Here is a summary of the main differences between Section 3.2 and Section 3.3.

- * The DPOp-inspired solution is less HTTP-specific, making it easier to adapt for other protocols beyond HTTP. This flexibility is particularly valuable for asynchronous communication scenarios, such as event-driven systems.
- * Message Signatures, on the other hand, benefit from an existing HTTP-specific RFC with some established implementations. This existing groundwork means that this option could be simpler to deploy, to the extent such implementations are available and easily integrated.

- * Given that the WIT (Workload Identity Token) is a type of JWT, the DPoP-inspired approach that also uses JWT is less complex and technology-intensive than Message Signatures. In contrast, Message Signatures introduce an additional layer of technology, potentially increasing the complexity of the overall system.
- * Message Signatures offer superior integrity protection, particularly by mitigating message modification by middleboxes. See also Section 6.4.
- * A key advantage of Message Signatures is that they support response signing. This opens up the possibility for future decisions about whether to make response signing mandatory, allowing for flexibility in the specification and/or in specific deployment scenarios.
- * In general, Message Signatures provide greater flexibility compared to the DPoP-inspired approach. Future versions of this draft (and subsequent implementations) can decide whether specific aspects of message signing, such as coverage of particular fields, should be mandatory or optional. Covering more fields will constrain the proof so it cannot be easily reused in another context, which is often a security improvement. The DPoP inspired approach could be designed to include extensibility to sign other fields, but this would make it closer to trying to reinvent Message Signatures.

3.5. Error Conditions

Errors may occur during the processing of the message signature or WPT. If the signature verification fails for any reason, such as an invalid signature, an expired validity time window, or a malformed data structure, an error is returned. Typically, this will be in response to an API call, so an HTTP status code such as 400 (Bad Request) is appropriate. This response could include more details as per [RFC9457], such as an indicator that the wrong key material or algorithm was used. The use of HTTP status code 401 is NOT RECOMMENDED for this purpose because it requires a WWW-Authenticate with acceptable http auth mechanisms in the error response and an associated Authorization header in the subsequent request. The use of these headers for the WIT or WPT is not compatible with this specification.

3.6. Coexistence with JWT Bearer Tokens

The WIT and WPT define new HTTP headers. They can therefore be presented along with existing headers used for JWT bearer tokens. This property allows for transition from mechanisms using identity tokens based on bearer JWTs to proof of possession based WITs. A workload may implement a policy that accepts both bearer tokens and WITs during a transition period. This policy may be configurable per-caller to allow the workload to reject bearer tokens from callers that support WITs. Once a deployment fully supports WITs, then the use of bearer tokens for identity can be disabled through policy. Implementations should be careful when implementing such a transition strategy, since the decision which token to prefer is made when the caller's identity has still not been authenticated, and needs to be revalidated following the authentication step.

The WIT can also coexist with tokens used to establish security context, such as transaction tokens [I-D.ietf-oauth-transaction-tokens]. In this case a workload's authorization policy may take into account both the sending workload's identity and the information in the context token. For example, the identity in the WIT may be used to establish which API calls can be made and information in the context token may be used to determine which specific resources can be accessed.

4. Using Mutual TLS for Workload-to-Workload Authentication

As noted in the introduction, for many deployments, transport-level protection of application traffic using TLS is ideal.

4.1. The Workload Identity Certificate

The Workload Identity Certificate is an X.509 certificate. The workload identity MUST be encoded in a SubjectAltName extension of type URI. There MUST be only one SubjectAltName extension of type URI in a Workload Identity Certificate. If the workload will act as a TLS server for clients that do not understand workload identities it is RECOMMENDED that the Workload Identity Certificate contain a SubjectAltName of type DNSName with the appropriate DNS names for the server. The certificate MAY contain SubjectAltName extensions of other types.

4.2. Workload Identity Certificate Validation

Workload Identity Certificates may be used to authenticate both the server and client side of the connections. When validating a Workload Identity Certificate, the relying party MUST use the trust anchors configured for the trust domain in the workload identity to validate the peer's certificate. Other PKIX [RFC5280] path validation rules apply. Workloads acting as TLS clients and servers MUST validate that the trust domain portion of the Workload Identity Certificate matches the expected trust domain for the other side of the connection.

Servers wishing to use the Workload Identity Certificate for authorizing the client MUST require client certificate authentication in the TLS handshake. Other methods of post handshake authentication are not specified by this document.

Workload Identity Certificates used by TLS servers SHOULD have the id-kp-serverAuth extended key usage [RFC5280] field set and Workload Identity Certificates used by TLS clients SHOULD have the id-kp-clientAuth extended key usage field set. A certificate that is used for both client and server connections may have both fields set. This specification does not make any other requirements beyond [RFC5280] on the contents of Workload Identity Certificates or on the certification authorities that issue workload certificates.

4.2.1. Server Name Validation

If the WIMSE client uses a hostname to connect to the server and the server certificate contain a DNS SAN the client MUST perform standard host name validation (Section 6.3 of [RFC9525]) unless it is configured with the additional information necessary to perform alternate validation of the peer's workload identity. If the client did not perform standard host name validation then the WIMSE client SHOULD further use the workload identifier to validate the server. The host portion of the workload identifier is NOT treated as a host name as specified in section 6.4 of [RFC9525] but rather as a trust domain. The server identity is encoded in the path portion of the workload identifier in a deployment specific way. Validating the workload identity could be a simple match on the trust domain and path portions of the identifier or validation may be based on the specific details on how the identifier is constructed. The path portion of the WIMSE identifier MUST always be considered in the scope of the trust domain. In most cases it is preferable to validate the entire workload identifier, see Section 1.3 for additional implementation advice.

4.3. Client Authorization Using the Workload Identity

The server application retrieves the workload identifier from the client certificate subjectAltName, which in turn is obtained from the TLS layer. The identifier is used in authorization, accounting and auditing. For example, the full workload identifier may be matched against ACLs to authorize actions requested by the peer and the identifier may be included in log messages to associate actions to the client workload for audit purposes. A deployment may specify other authorization policies based on the specific details of how the workload identifier is constructed. The path portion of the workload identifier MUST always be considered in the scope of the trust domain. See Section 1.3 on additional security implications of workload identifiers.

5. Implementation Status

// Note to RFC Editor: please remove this section, as well as the // reference to RFC 7942, before publication. This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to RFC 7942, "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

SPIFFE (Standard)

- * Organization: CNCF

- * Maturity:

- Workload Identity Certificate: fully compatible with the X509-SVID and widely used.

- Workload Identity Token: beta
- * Coverage: Workload Identity Certificate, WIT
- * Contact: SPIFFE sig-spec community
(<https://github.com/spiffe/spiffe/tree/main/community/sig-spec>)

SPIRL

- * Organization: SPIRL
- * Maturity:
 - Workload Identity Certificate: production
 - Workload Identity Token/Workload Proof Token: alpha
- * Coverage: Workload Identity Certificate, WIT, WPT
- * Contact: arndt@spirl.com

Teleport - Machine & Workload Identity

- * Organization: Teleport
- * Maturity:
 - Workload Identity Certificate: production
 - Workload Identity Token/Workload Proof Token: research
- * Coverage: Workload Identity Certificate
- * Contact: noah@goteleport.com

6. Security Considerations

6.1. Workload Identity

The Workload Identifier is scoped within an issuer and therefore any sub-components (path portion of Identifier) are only unique within a trust domain defined by the issuer. Using a Workload Identifier without taking into account the trust domain could allow one domain to issue tokens to spoof identities in another domain. Additionally, the trust domain must be tied to an authorized issuer cryptographic trust anchor through some mechanism such as a JWKS or X.509 certificate chain. The association of an issuer, trust domain and a cryptographic trust anchor MUST be communicated securely out of band.

6.2. Workload Identity Token and Proof of Possession

The Workload Identity Token (WIT) is bound to a secret cryptographic key and is always presented with a proof of possession as described in Section 3.1. The WIT is a general purpose token that can be presented in multiple contexts. The WIT and its PoP are only used in the application-level options, and both are not used in MTLS. The WIT MUST NOT be used as a bearer token. While this helps reduce the sensitivity of the token it is still possible that a token and its proof of possession may be captured and replayed within the PoP's lifetime. The following are some mitigations for the capture and reuse of the proof of possession (PoP):

- * Preventing Eavesdropping and Interception with TLS

An attacker observing or intercepting the communication channel can view the token and its proof of possession and attempt to replay it to gain an advantage. In order to prevent this the token and proof of possession MUST be sent over a secure, server authenticated TLS connection unless a secure channel is provided by some other mechanisms. Host name validation according to Section 4.2.1 MUST be performed by the client.

- * Limiting Proof of Possession Lifespan

The proof of possession MUST be time limited. A PoP should only be valid over the time necessary for it to be successfully used for the purpose it is needed. This will typically be on the order of minutes. PoPs received outside their validity time MUST be rejected.

- * Limiting Proof of Possession Scope

In order to reduce the risk of theft and replay the PoP should have a limited scope. For example, a PoP may be targeted for use with a specific workload and even a specific transaction to reduce the impact of a stolen PoP. In some cases a workload may wish to reuse a PoP for a period of time or have it accepted by multiple target workloads. A careful analysis is warranted to understand the impacts to the system if a PoP is disclosed allowing it to be presented by an attacker along with a captured WIT.

- * Replay Protection

A proof of possession includes the jti claim that MUST uniquely identify it, within the scope of a particular sender. This claim SHOULD be used by the receiver to perform basic replay protection against tokens it has already seen. Depending upon the design of the system it may be difficult to synchronize the replay cache across all

token validators. If an attacker can somehow influence the identity of the validator (e.g. which cluster member receives the message) then replay protection would not be effective.

- * Binding to TLS Endpoint

The POP MAY be bound to a transport layer sender such as the client identity of a TLS session or TLS channel binding parameters. The mechanisms for binding are outside the scope of this specification.

6.3. Workload Identity Key Management

Both the Workload Identity Token and the Workload Identity Certificate carry a public key. The corresponding private key:

- * MUST be kept private
- * MUST be individual for each Workload Identifier (see [I-D.ietf-wimse-arch])
- * MUST NOT be used once the credential is expired
- * SHOULD be re-generated for each new Workload Identity Token or Certificate.

6.4. Middle Boxes

In some deployments the Workload Identity Token and proof of possession may pass through multiple systems. The communication between the systems is over TLS, but the token and PoP are available in the clear at each intermediary. While the intermediary cannot modify the token or the information within the PoP they can attempt to capture and replay the token or modify the data not protected by the PoP.

Mitigations listed in Section 3 can be used to provide some protection from middle boxes. However we note that the DPoP-inspired solution (Section 3.2) does not protect major portions of the request and response and therefore does not provide protection from an actively malicious middle box. Deployments should perform analysis on their situation to determine if it is appropriate to trust and allow traffic to pass through a middle box.

6.5. Privacy Considerations

WITs and the proofs of possession may contain private information such as user names or other identities. Care should be taken to prevent the disclosure of this information. The use of TLS helps protect the privacy of WITs and proofs of possession.

WITs and certificates with workload identifiers are typically associated with a workload and not a specific user, however in some deployments the workload may be associated directly to a user. While these are exceptional cases a deployment should evaluate if the disclosure of WITs or certificates can be used to track a user.

7. IANA Considerations

7.1. JSON Web Token Claims

IANA is requested to add the following entries to the "JSON Web Token Claims" registry [IANA.JWT.CLAIMS]:

Claim Name	Claim Description	Change Controller	Reference
tth	Transaction Token hash	IETF	RFC XXX, Section 3.2
wth	Workload Identity Token hash	IETF	RFC XXX, Section 3.2
oth	Other Tokens hashes	IETF	RFC XXX, Section 3.2

Table 1

7.2. Media Type Registration

IANA is requested to register the following entries to the "Media Types" registry [IANA.MEDIA.TYPES]:

- * application/wit+jwt, per Section 7.2.1.
- * application/wpt+jwt, per Section 7.2.2.

7.2.1. application/wit+jwt

Type name: application

Subtype name: wit+jwt

Required parameters: N/A

Optional parameters: N/A

Encoding considerations: Encoding considerations are identical to those specified for the "application/jwt" media type. See [RFC7519].

Security considerations: See the Security Considerations section of RFC XXX.

Interoperability considerations: N/A

Published specification: RFC XXX, Section 3.1.

Applications that use this media type: Identity servers that vend Workload Identity Tokens, and Workloads that use these tokens to authenticate to each other.

Fragment identifier considerations: N/A

Additional information:

Deprecated alias names for this type: N/A

Magic number(s): N/A

File extension(s): None

Macintosh file type code(s): N/A

Person & email address to contact for further information:

See the Authors' Addresses section of RFC XXX.

Intended usage: COMMON

Restrictions on usage: N/A

Author: See the Authors' Addresses section of RFC XXX.

Change controller: Internet Engineering Task Force (iesg@ietf.org).

7.2.2. application/wpt+jwt

Type name: application

Subtype name: wpt+jwt

Required parameters: N/A

Optional parameters: N/A

Encoding considerations: Encoding considerations are identical to those specified for the "application/jwt" media type. See [RFC7519].

Security considerations: See the Security Considerations section of RFC XXX.

Interoperability considerations: N/A

Published specification: RFC XXX, Section 3.2.

Applications that use this media type: Workloads that use these tokens to integrity-protect messages in the WIMSE workload-to-workload protocol.

Fragment identifier considerations: N/A

Additional information:

Deprecated alias names for this type: N/A

Magic number(s): N/A

File extension(s): None

Macintosh file type code(s): N/A

Person & email address to contact for further information:

See the Authors' Addresses section of RFC XXX.

Intended usage: COMMON

Restrictions on usage: N/A

Author: See the Authors' Addresses section of RFC XXX.

Change controller: Internet Engineering Task Force (iesg@ietf.org).

7.3. Hypertext Transfer Protocol (HTTP) Field Name Registration

IANA is requested to register the following entries to the "Hypertext Transfer Protocol (HTTP) Field Name Registry" [IANA.HTTP.FIELDS]:

- * Workload-Identity-Token, per Section 7.3.1.

- * Workload-Proof-Token, per Section 7.3.2.

7.3.1. Workload-Identity-Token

- * Field Name: Workload-Identity-Token

- * Status: permanent

- * Structured Type: N/A

- * Specification Document: RFC XXX, Section 3.1.1

- * Comments: see reference above for an ABNF syntax of this field

7.3.2. Workload-Proof-Token

- * Field Name: Workload-Proof-Token

- * Status: permanent

- * Structured Type: N/A

- * Specification Document: RFC XXX, Section 3.2

- * Comments: see reference above for an ABNF syntax of this field

7.4. URI Scheme Registration

IANA is requested to register the "wimse" scheme to the "URI Schemes" registry [IANA.URI.SCHEMES]:

- * Scheme name: wimse

- * Status: permanent

- * Applications/protocols that use this scheme name: the WIMSE workload-to-workload authentication protocol.

- * Contact: IETF Chair chair@ietf.org (mailto:chair@ietf.org)

- * Change controller: IESG iesg@ietf.org (mailto:iesg@ietf.org)

- * References: Section 3 of this document (RFC XXX).

8. References

8.1. 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/rfc/rfc2119>>.
- [RFC5234] 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/rfc/rfc5234>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/rfc/rfc5280>>.
- [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/rfc/rfc7515>>.
- [RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/rfc/rfc7517>>.
- [RFC7518] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <<https://www.rfc-editor.org/rfc/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/rfc/rfc7519>>.
- [RFC7800] 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/rfc/rfc7800>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/rfc/rfc7942>>.
- [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/rfc/rfc8174>>.

- [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/rfc/rfc8414>>.
- [RFC8725] Sheffer, Y., Hardt, D., and M. Jones, "JSON Web Token Best Current Practices", BCP 225, RFC 8725, DOI 10.17487/RFC8725, February 2020, <<https://www.rfc-editor.org/rfc/rfc8725>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/rfc/rfc9110>>.
- [RFC9421] Backman, A., Ed., Richer, J., Ed., and M. Sporny, "HTTP Message Signatures", RFC 9421, DOI 10.17487/RFC9421, February 2024, <<https://www.rfc-editor.org/rfc/rfc9421>>.
- [RFC9525] Saint-Andre, P. and R. Salz, "Service Identity in TLS", RFC 9525, DOI 10.17487/RFC9525, November 2023, <<https://www.rfc-editor.org/rfc/rfc9525>>.

8.2. Informative References

- [I-D.ietf-oauth-transaction-tokens]
Tulshibagwale, A., Fletcher, G., and P. Kasselmann,
"Transaction Tokens", Work in Progress, Internet-Draft,
draft-ietf-oauth-transaction-tokens-06, 28 July 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-oauth-transaction-tokens-06>>.
- [I-D.ietf-wimse-arch]
Salowey, J. A., Rosomakho, Y., and H. Tschofenig,
"Workload Identity in a Multi System Environment (WIMSE)
Architecture", Work in Progress, Internet-Draft, draft-
ietf-wimse-arch-06, 30 September 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-wimse-arch-06>>.
- [IANA.HTTP.FIELDS]
IANA, "Hypertext Transfer Protocol (HTTP) Field Name
Registry", <<https://www.iana.org/assignments/http-fields>>.
- [IANA.JOSE.ALGS]
IANA, "JSON Web Signature and Encryption Algorithms",
<<https://www.iana.org/assignments/jose>>.

[IANA.JWT.CLAIMS]
IANA, "JSON Web Token Claims",
<<https://www.iana.org/assignments/jwt>>.

[IANA.MEDIA.TYPES]
IANA, "Media Types",
<<https://www.iana.org/assignments/media-types>>.

[IANA.URI.SCHEMES]
IANA, "Uniform Resource Identifier (URI) Schemes",
<<https://www.iana.org/assignments/uri-schemes>>.

[RFC9449] Fett, D., Campbell, B., Bradley, J., Lodderstedt, T.,
Jones, M., and D. Waite, "OAuth 2.0 Demonstrating Proof of
Possession (DPoP)", RFC 9449, DOI 10.17487/RFC9449,
September 2023, <<https://www.rfc-editor.org/rfc/rfc9449>>.

[RFC9457] Nottingham, M., Wilde, E., and S. Dalal, "Problem Details
for HTTP APIs", RFC 9457, DOI 10.17487/RFC9457, July 2023,
<<https://www.rfc-editor.org/rfc/rfc9457>>.

Appendix A. Document History

// RFC Editor: please remove before publication.

A.1. draft-ietf-wimse-s2s-protocol-07

- * Rework the WPT's oth claim.
- * update the media types.
- * Discuss extensibility of WIT and WPT.
- * Clarify error handling, specifically why not HTTP 401.
- * Correct the code examples.
- * Add registration request content for a wimse URI scheme.
- * New section on key management.
- * Use of the Accept-Signature header.

A.2. draft-ietf-wimse-s2s-protocol-06

- * Explicit definition of the Workload Identity Certificate.

- * Definition of the validation of workload identifiers as part of workload authentication. Still work in progress.

A.3. draft-ietf-wimse-s2s-protocol-05

- * Removed the entire Workload Identity section which is now covered in the Architecture document.
- * Content-Digest is mandatory with HTTP-Sig.
- * Some wording on extending the protocol beyond HTTP.
- * IANA considerations.

A.4. draft-ietf-wimse-s2s-protocol-04

- * Require `cnf.jwk.alg` in WIT which restricts signature algorithm of WPT or HTTP-Sig.
- * Replay protection as a SHOULD for both WPT and HTTP-Sig.
- * Consolidate terminology with the Architecture draft.

A.5. draft-ietf-wimse-s2s-protocol-03

- * Consistently use "workload".
- * Implement comments from the SPIFFE community.
- * Make iss claim in WIT optional and add wording about its relation to key distribution.
- * Remove iss claim from WPT.
- * Make jti claim in WIT optional.
- * Error handling for the application level methods.

A.6. draft-ietf-wimse-s2s-protocol-02

- * Coexistence with bearer tokens.
- * Improve the architecture diagram.
- * Some more ABNF.
- * Clarified identifiers and URIs.

- * Moved an author to acknowledgments.

A.7. draft-ietf-wimse-s2s-protocol-01

- * Addressed multiple comments from Pieter.
- * Clarified WIMSE identity concepts, specifically "trust domain" and "workload identifier".
- * Much more detail around mTLS, including some normative language.
- * WIT (the identity token) is now included in the WPT proof of possession.
- * Added a section comparing the DPoP-inspired app-level security option to the Message Signature-based alternative.

A.8. draft-ietf-wimse-s2s-protocol-00

- * Initial WG draft, an exact copy of draft-sheffer-wimse-s2s-protocol-00
- * Added this document history section

Acknowledgments

The authors would like to thank Pieter Kasselmann for his detailed comments.

We thank Daniel Feldman for his contributions to earlier versions of this document.

Authors' Addresses

Brian Campbell
Ping Identity
Email: bcampbell@pingidentity.com

Joe Salowey
CyberArk
Email: joe@salowey.net

Arndt Schwenkschuster
SPIRL
Email: arndts.ietf@gmail.com

Yaron Sheffer
Intuit
Email: yaronf.ietf@gmail.com