Workload Identity in Multi System Environments
Internet-Draft
Intended status: Standards Track
Expires: 4 January 2026

Y. Rosomakho
Zscaler
J. Salowey
CyberArk

3 July 2025

Workload Identifier
draft-rosomakho-wimse-identifier-00

Abstract

This document defines a canonical identifier for workloads, referred to as the Workload Identifier. A Workload Identifier is a URI that uniquely identifies a workload within the context of a specific trust domain. This identifier can be embedded in digital credentials, including X.509 certificates and security tokens, to support authentication, authorization, and policy enforcement across diverse systems. The Workload Identifier format ensures interoperability, facilitates secure identity federation, and enables consistent identity semantics.

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://yaroslavros.github.io/wimse-identifier/draft-rosomakho-wimse-identifier.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-rosomakho-wimse-identifier/.

Discussion of this document takes place on the Workload Identity in Multi System Environments 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/yaroslavros/wimse-identifier.

Status of This Memo

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

This Internet-Draft will expire on 4 January 2026.

Copyright Notice

Table of Contents

1. Introduction 3
2. Conventions and Definitions 4
3. Terminology 4
4. Workload Identifier Specification 4
  4.1. URI Requirements 5
  4.2. Scheme Specific Portion 5
  4.3. Trust Domain Association 6
  4.4. Stability and Uniqueness 6
5. Usage in Credentials and Tokens 7
  5.1. X.509 Certificates 7
  5.2. JSON Web Tokens (JWT) 7
  5.3. Interpretation by Consumers 7
6. Security Considerations 7
  6.1. Identifier Authenticity 8
  6.2. Trust Domain Validation 8
  6.3. Identifier Reuse and Collision 8
  6.4. Information Disclosure 8
  6.5. Wildcard and Prefix Matching 8
7. IANA Considerations 9
8. References 9

8.1. Normative References 9
8.2. Informative References 9
Acknowledgments 10
Authors' Addresses 10

1. Introduction

In modern distributed systems, workloads such as services, applications, or containerised tasks require cryptographically verifiable identities to support secure communication, access control, and auditability. As systems scale across trust domains, administrative boundaries, and heterogeneous platforms, the need for a consistent and interoperable identifier format becomes critical.

This document defines the Workload Identifier, a URI-based [URI] identifier intended to uniquely represent a workload within the context of an issuing authority. The identifier is designed to be stable, globally unique within a given trust domain, and suitable for use in digital credentials such as X.509 certificates , JSON Web Tokens (JWTs, [JWT]), and other security artifacts.

The Workload Identifier format is simple yet expressive. It enables organisations to define trust boundaries, delegate identity management, and reason about workloads in a uniform way across service meshes, cloud environments, and on-premises infrastructure. This specification is intended to be generic and reusable beyond the context of any single system or architecture, including but not limited to the Workload Identity in Multi-System Environments (WIMSE) architecture [ARCH].

The primary goals of this specification are:

* To define the syntax and semantics of a Workload Identifier.

* To establish requirements for issuers and consumers of such identifiers.

* To promote interoperability across different identity systems and domains.

This document does not prescribe how identifiers are issued or verified. Instead, it focuses on the identifier' s format, uniqueness guarantees, and its relationship to trust domains.

2. Conventions and Definitions

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. Terminology

The following terms are used throughout this document:

Workload: An independently addressable and executable software entity. This may include microservices, containers, virtual machines, serverless functions, or similar components that initiate or receive network communications.

Workload Identifier: A URI-based identifier that uniquely represents a workload within a specific trust domain. It is intended to be included in security credentials and interpreted within the scope of an issuing authority.

Trust Domain: A security boundary defined and controlled by a single administrative authority. A trust domain establishes its own rules for identity issuance, validation, and policy enforcement.

Issuer: An entity responsible for assigning and validating Workload Identifiers.

Consumer: An entity that evaluates, verifies or uses a Workload Identifier, typically as part of authentication or authorisation decisions. This includes relying parties, verifiers, and policy enforcement points.

4. Workload Identifier Specification

A Workload Identifier is a URI [URI] that uniquely identifies a workload. It encodes both the trust domain and a workload-specific path, enabling unambiguous identification of workloads across administrative and organisational boundaries.

The identifier is designed to be stable and suitable for inclusion in digital credentials such as X.509 certificates and security tokens. This section defines the format, structure, and associated requirements for Workload Identifiers.

4.1. URI Requirements

A Workload Identifier MUST be an absolute URI, as defined in Section 4.3 of [URI]. In addition the URI MUST include an authority that identifies the trust domain within which the identifier is scoped. The scheme and scheme specific part are not defined by this specification. The URI format allows different schemes (e.g., spiffe as defined in [SPIFFE-ID], wimse) depending on deployment requirements. Example identifiers:

```
spiffe://incubation.example.org/ns/experimental/analytics/ingest
wimse://trust.corp.example.com/workload/af3e86cb-7013-4e33-b717-11c4edd25679
```

(Note that the wimse scheme is used as an example and is not defined in this document).

4.2. Scheme Specific Portion

The format and semantics scheme specific part of the URI that follows the identity is determined by the issuer in the trust domain. What the identity refers to is also determined by the issuer. For example a workload identity may refer to a specific instance of a running piece of software or it may refer just to a specific software version running in a particular environment, or it may refer to the role that the software performs within the system. The scheme specific part of the URI may just be an opaque unique identifier used to look up the additional identity information in another system. Some examples of these concepts are given below:

* Opaque identifier

```
spiffe://prod.trust.domain/89a6ec51-f877-44c0-9501-b213597f2d1d
```

* Application role

```
spiffe://prod.trust.domain/ns/prod-01/sa/foo-service
```

* Specific instance of application role

```
spiffe://prod.trust.domain/ns/prod-01/sa/foo-service/iid-
      1f814646-87b5-4e26-bb55-1d13caccdd8d
```

* Specific code for an application role

```
spiffe://prod.trust.domaain/foo-servce#@sha256:
     c4dbb1a06030e142cb0ed4be61421967618289a19c0c7760bdd745ac67779ca7
```

Other concepts may be defined within the trust domain depending on what is important in the system and what information is available when the identity is issued. A trust domain should define the scheme specific portion of the URI to meet their auditing and authorization needs.

4.3. Trust Domain Association

The authority component of the URI defines the trust domain which is responsible for issuing, validating, and managing Workload Identifiers within its scope. The trust domain SHOULD be a fully qualified domain name belonging to the organization defining the trust domain to help provide uniqueness for the trust domain identifier. While IP addresses are allowed as host names in the URI encoding rules, they MUST NOT be used to represent trust domains except in the case where they are needed for compatibility with legacy naming schemes.

Workload Identifiers are interpreted in the context of the trust domain that issued the credential. Identifiers with identical path components but different trust domains represent different workloads.

Issuers within a trust domain MUST ensure uniqueness of all Workload Identifiers they assign.

4.4. Stability and Uniqueness

Workload Identifiers are intended to be stable over time. An identifier assigned to a specific workload should not be reassigned to a different workload unless explicitly intended by the policies of the trust domain.

Workload Identifiers are globally unique when the trust domain is globally unique. This is typically achieved by using a fully qualified domain name (FQDN) under organisational control.

For example, the following contains identifiers of two distinct globally unique Workload Identifiers

spiffe://dev.example.com/ns/default/database/backend
spiffe://prod.example.com/ns/default/database/backend

5. Usage in Credentials and Tokens

Workload Identifiers are designed to be embedded in cryptographic credentials and security tokens that are used to assert the identity of workloads during authentication, authorisation, and auditing. This section describes how such identifiers may be represented in commonly used formats.

5.1. X.509 Certificates

Workload Identifier included in an X.509 are encoded in the subject alternative name extension as a URI using the uniformResourceIdentifier field, as defined in Section 4.2.1.6 of [X509-PROFILE].

For example,

```
X509v3 extensions:
    X509v3 Subject Alternative Name:
        URI:spiffe://example.org/ns/default/analytics/ingest
```

Consumers MUST NOT attempt to interpret or derive workload identity from other certificate fields such as the Common Name.

5.2. JSON Web Tokens (JWT)

When a Workload Identifier is included in a JWT, it MUST appear in the "sub" (Subject) claim, as defined in Section 4.1.2 of [JWT].

5.3. Interpretation by Consumers

Consumers of credentials and tokens MUST validate that the Workload Identifier is consistent with the expected trust domain and issuing authority. Consumers SHOULD NOT make assumptions about internal structure or semantics of the identifier beyond the URI format defined in this specification.

For authorisation decisions, consumers may map Workload Identifiers to policies or roles. However, such mappings are out of scope for this specification.

6. Security Considerations

The Workload Identifier is intended to be used as a stable, verifiable identity for workloads. Its use in cryptographic credentials means it must be protected against spoofing, ambiguity, and misinterpretation. This section outlines security considerations for issuers, consumers, and system designers.

6.1. Identifier Authenticity

Workload Identifiers MUST only be considered valid when presented in a credential or token that has been cryptographically verified. An identifier received outside such a context, such as a plaintext string in a request, MUST NOT be treated as authenticated.

Consumers MUST verify the signature, issuer, and validity of the credential or token before considering Workload Identifier as authenticated.

6.2. Trust Domain Validation

Consumers MUST validate that the trust domain in the Workload Identifier matches an expected or explicitly trusted domain. Failure to do so may allow identifiers from unauthorised domains to be accepted as legitimate.

Where appropriate, consumers should maintain an allowlist of trusted domains or trusted issuing authorities.

6.3. Identifier Reuse and Collision

Issuers SHOULD ensure that Workload Identifiers are not reused across different workloads unless such reuse is intentional and well-scoped. Reassignment of identifiers to unrelated entities can result in privilege escalation or confusion in audit trails.

Consumers SHOULD assume that identifiers are permanent within their domain of interpretation and treat unexpected reuse with suspicion.

6.4. Information Disclosure

Because Workload Identifiers may encode topological or semantic information, they may inadvertently reveal deployment details. Issuers and system designers should take care not to expose sensitive naming conventions in externally visible identifiers.

Where possible, identifier paths SHOULD be minimally descriptive and avoid exposing internal implementation details unless necessary for interoperation.

6.5. Wildcard and Prefix Matching

Consumers SHOULD NOT interpret Workload Identifiers using wildcard or prefix matching unless explicitly specified by policy. For example, treating all identifiers under prefix of spiffe://example.org/ns/db/ as equivalent may lead to incorrect authorisation.

7. IANA Considerations

This document has no IANA actions.

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>.

[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>.

[URI] 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/rfc/rfc3986>.

8.2. Informative References

[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-04, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-wimse-arch-04>.

[JWT] 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>.

[SPIFFE-ID] "The SPIFFE Identity and Verifiable Identity Document", January 2025, <https://github.com/spiffe/spiffe/blob/main/standards/SPIFFE-ID.md>.

[X509-PROFILE] 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>.

Acknowledgments

Authors would like to thank Evan Gilman for his review of the initial text of this document and his guidance.

Authors' Addresses

Yaroslav Rosomakho
Zscaler
Email: yaroslavros@gmail.com

Joe Salowey
CyberArk
Email: joe@salowey.net