WIMSE
Internet-Draft
Intended status: Informational
Expires: 5 January 2026

M. Novak
J.P. Morgan Chase
Y. Deshpande
arm
H. Birkholz
Fraunhofer SIT

4 July 2025

WIMSE Extensions for Trustworthy Workload Identity
draft-ccc-wimse-twi-extensions-00

Abstract

This document contains a gap analysis that is the output of the Confidential Computing Consortium identifying areas in the IETF WIMSE WG work where the current WIMSE architecture should be extended to accommodate workloads running in Confidential Computing environments. This document contains a high-level outline for these extensions.

About This Document

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

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ccc-wimse-twi-extensions/.

information can be found at https://confidentialcomputing.io/about/committees/.

Source for this draft and an issue tracker can be found at https://github.com/confidential-computing/twi-wimse.

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 5 January 2026.

Table of Contents

1. Introduction

Until recently, there were few scenarios requiring data-in-use protection. This is starting to change. Regulatory bodies worldwide are increasingly requiring data-in-use protection and privacy enhancing technologies. Outside of regulatory requirements, companies are exploring:

* Multiparty computation

* Machine learning training & inferencing

* Addressing the risks of computing in cases where the operator of the workload does not fully trust the hosting environments, such as public clouds, high-risk locations and IoT deployments

* Entrusting confidential data to SaaS providers

* Insider threats, and other reasons for protecting data-in-use

Correspondingly, there is an increased push to harmonize management and governance of human and non-human identities. Modern workloads may operate on their own behalf with their own credentials, or as agents on behalf of other entities with delegated credentials. Entities interested in strong assurances around the security of their deployed workloads, for regulatory, contractual or peace of mind reasons, are facing large and challenging tasks of upgrading their existing computing system infrastructure to meet these requirements. Current ways of issuing and managing workload identities, as well as those required for effective protection of data-in-use, are subject to multiple architectural challenges; chief among them:

1. Lack of workload isolation against the hardware and the operating system owners/administrators, as well as peer workload instances

2. Lack of strong binding between a workload credential and the workload instance to which that credential had been issued

3. Lack of verifiable composition of the workload, and inability to associate a credential with a set of decisions leading up to its issuance

It is important to highlight that these shortcomings are related: lack of process isolation eases credential exfiltration and leads to credential leakage and reuse.

Confidential Computing can close these architectural gaps due to its unique features (i.e., verifiable composition, strong workload isolation) and broad availability (i.e., support by all major hardware vendors). Multiple emerging regulations will mean that customers will be looking to these features and capabilities to satisfy them.

Confidential Computing-assisted mechanisms have to align with the emerging Workload Identity solution ecosystem. Correspondingly, the evolution of the Workload Identity ecosystem should remain in

alignment with the expectations of the owners and operators of Confidential Computing workloads. To address this set of requirements, this document defines and elaborates on the concept of Trustworthy Workload Identity (TWI) in the following Sections.

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.

This document uses terms and concepts defined by the WIMSE and RATS architectures, as well as the terms defined by the Trustworthy Workload Identity Special Interest Group at the Confidential Computing Consortium. For a complete glossary, see Section 4 of [RFC9334] , [I-D.draft-ietf-wimse-arch] & [TWISIGCharter].

The definitions of terms like Workload Identity, Workload Credential and Workload Provenance match those specified by the TWI SIG Charter [TWISIGCharter].

Workload: [I-D.draft-ietf-wimse-arch] defines ’Workload’ as "an instance of software executing for a specific purpose". Here we restrict that definition to the portions of the deployed software and its configuration that are subject to Remote Attestation.

Workload Identifier: a stable construct around which Relying Parties can form long-lived Workload authorization policies.

Workload Identity: the definition of Workload Identity is identical to the definition of the same term by [I-D.draft-ietf-wimse-arch]: "a combination of three basic building blocks: trust domain, Workload Identifier and identity credentials.

Workload Credential: an ephemeral identity document containing the Workload Identifier and a number of additional claims, that can be short-lived or long-lived, and that is used to represent and prove Workload Identity to a Relying Party.

Workload Provenance: a unique linkage between a Workload Credential and the trusted entities (such as a vendor, developer, or issuer) responsible for the production and/or the remote attestation of the corresponding Workload.

3. Gap Analysis

The following shortcomings were identified by performing a gap analysis of the current WIMSE architecture [I-D.draft-ietf-wimse-arch] against the requirements underlying Trustworthy Workload Identity (TWI). The gap analysis provides the basis for identifying extensions necessary to meet the level of trustworthiness required by Confidential Computing environments.

* Protection of Credentials at runtime and insufficient Runtime Attestation:
  - Without data-in-use protection, there is always a risk that Credentials in memory could be exposed if a Workload’s execution environment is compromised.
  - Confidential Computing does not trust its hosting environment and relies on each Workload attesting itself; in particular, the "Agent" must be fully trusted and included in Remote Attestation.
* Token replay and misuse:
  - Token binding, even if used, is not sufficient unless the secrets underpinning the Workload Credentials are protected from leakage.
* Lack of verifiable workload composition accessible through Credential Provenance:
  - When a Relying Party receives a Credential, or when an auditor examines a log of decisions by a Relying Party, it is unable to peform additional checks on the security properties of the Workload or the process involved in creating it, beyond the claims communicated inside the Credential.

4. Alignment or Synergy with WIMSE Architecture

WIMSE defines an architecture for utilizing Workload Identity in multi-system environments.

1. The WIMSE Architecture explains the core concept of Workload Identity in-line with the concept of identity in the TWI world.
2. WIMSE model has a CA/Credential issuer that is responsible for provisioning identity Credentials to the Workload. TWI requirement is roughly similar in terms of issuing credentials.

3. WIMSE Architecture defines Trust Domain, which is the authority that identifies domain within which the identifier is scoped. TWI Architecture is aligned with this basic building block.

5. Extensions to the WIMSE Architecture

Trustworthy Workload Identity as detailed in this document proposes the following extensions to the core WIMSE Architecture.

5.1. Workload Isolation and the Role of the Hosting Environment

Under Confidential Computing, the confidentiality and integrity of a Workload is isolated from the hosting environment and other Workloads. The underlying assumption in the WIMSE architecture cannot be that the hosting environment can be fully trusted to establish the identity of a hosted Workload, as that is incompatible with the assumptions of Confidential Computing. Confidential Workloads perform their own Remote Attestation and Workload isolation ensures that the hosting environment cannot interfere with this process. WIMSE Arctecture and all its associated protocols and schemes MUST accommodate both confidential and non-confidential Workloads.

5.2. Remote Attestation

Under Confidential Computing, Remote Attestation plays a fundamental role in assessing the trustworthiness of the Workload and ensuring that the Workload is running on a platform that provides required security guarantees. Specifically, when performing Remote Attestation, the Workload cannot rely on an untrusted Agent, therefore including the Agent, if any, in the Remote Attestation process is of paramount importance.

5.3. Provenance

Provenance is a consideration that is foundational for Confidential Computing workloads.

6. Integration of Confidential Computing into WIMSE

6.1. Secure Key Storage & Cryptographic Operations

With TEEs, a Workload’s private keys and sensitive cryptographic operations (such as signing or validating tokens) can be isolated from the hosting environment, reducing the risk of key leakage even if the hosting environment is compromised. (For instance, the WIMSE token -- be it a WIT or an X.509 certificate -- can be generated and signed within a TEE, ensuring that the proof-of-possession mechanism

remains intact.)

6.2. Enhanced Bootstrapping with Attestation

Strengthening the initial bootstrapping process: a TEE can enable hardware-based Remote Attestation, proving that a Workload is running in a secure, isolated environment. This attestation could be used as an additional factor during Workload Credential provisioning, ensuring that only Workloads running in a TEE and matching the Workload Credential issuer's attestation policies receive valid Credentials.

6.3. Protected Credential Exchange

For the Credential exchange patterns defined in the WIMSE Credential Exchange draft, Confidential Computing can provide a Trusted Execution Environment in which the exchange logic runs. This ensures that the process of exchanging or re-provisioning credentials is protected against tampering and eavesdropping.

6.4. Mitigating Runtime Compromise

Executing the Workload inside a Trusted Execution Environment can lower the risk that runtime attacks (such as memory scraping or side-channel attacks) can expose critical identity or authentication tokens. For example, a Trusted Execution Environment can be used to securely generate and verify proofs of possession that are important within the WIMSE authentication protocol.

7. Security Considerations

Maintaining security guarantees of Confidential Computing within the WIMSE reference architecture calls for adding the extensions specified in this draft. * Specifically, Confidential Computing demands strong binding between any Credential of a confidential Workload and the underlying hardware platform, utilizing proof-of-possession of the Workload Credential backed by a key that is safeguarded against disclosure and tampering via data-in-use protections offered by Trusted Execution Environments. * Additionally, Confidential Computing inverts the trust relationship between the Workload and the hosting environment: a confidential Workload cannot trust the hosting environment's claims about its capabilities to obtain a Credential and MUST instead perform its own Remote Attestation. If an Agent is used to perform Remote Attestation and obtain the Credentials on behalf of the Workload, the Agent itself MUST be included in the Remote Attestion of the Workload, as it becomes part of that Workload's TCB. * Finally, without any change to the WIMSE specifications, but to make this

explicit: linkage between a confidential Workload's Credential and that Workload's Provenance can be achieved by treating the pre-existing unique credential ID (e.g., the "jti" claim) as a "hook" by which a Relying Party might discover information about the Workload's Provenance.

8. IANA Considerations

Not applicable for this draft at this time.

9. Acknowledgements

The following persons, in no specific order, contributed to the work directly, participated in design team meetings, or provided valuable comments during the review of this document.

Pieter Kasselman (SPIRL), Arieal Feldman (Google), Mateusz Bronk (Intel), Manu Fontaine (Hushmesh Inc.), Benedict Lau (EQTY Lab), Zvonko Kaiser (NVIDIA), David Quigley (Intel), Sal Kimmich (GadflyAI), Alex Dalton (Shielded Technologies), Eric Wolfe (Mainsail Industries), Nicolae Paladi(Canary Bit), Mark Gentry (JPMorgan Chase), Jag Raman (Oracle), Brian Hugenbruch (IBM), Jens Alberts (Fr0ntierX), Mira Spina (MITRE) and John Suykerbuyk.

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://doi.org/10.17487/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://doi.org/10.17487/RFC8174>.

10.2. Informative References

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

[RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <https://doi.org/10.17487/RFC9334>.

[TWISIGCharter] Confidential Computing Consortium Trustworthy Workload Identity SIG, "Trustworthy Workload Identity (TWI) Special Interest Group -- Charter", n.d., <https://github.com/confidential-computing/governance/blob/main/SIGs/TWI/TWI_Charter.md>.

[TWISIGReq] Confidential Computing Consortium Trustworthy Workload Identity SIG, "Trustworthy Workload Identity (TWI) Special Interest Group -- Requirements", n.d., <https://github.com/confidential-computing/twi/blob/main/TWI_Requirements.md>.

Authors' Addresses

Mark Novak
J.P. Morgan Chase
Email: mark.f.novak@jpmchase.com

Yogesh Deshpande
arm
Email: Yogesh.Deshpande@arm.com

Henk Birkholz
Fraunhofer SIT
Email: henk.birkholz@ietf.contact