Network Working Group
Internet-Draft
Intended status: Informational
Expires: 23 November 2026

L. Iannone
A. Fressancourt
Huawei

22 May 2026

Privacy Considerations for the Discovery of Agents, Workloads, and Named Entities (DAWN)

draft-iannone-dawn-privacy-considerations-00

Abstract

This document describes the privacy issues associated with the Discovery of Agents, Workloads, and Named Entities (DAWN). It provides general observations about typical current privacy practices in similar domains like, DNS, HTTP, and in general privacy in information retrieval.

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 23 November 2026.

Copyright Notice

Table of Contents

1. Introduction

[I-D.akhavain-moussa-dawn-problem-statement] defines the problem space how AI-related entities discover and interact with one another across distributed ecosystems. In particular, it focuses on defining the requirements for a standardized discovery substrate that allows entities to find one another based on attributes like skills, capabilities, policies,communication methods, and collaborate dynamically.

In such a context, privacy is one of the hardest problems in entities discovery because the act of discovery itself can leak sensitive information. When an entity searches for another entity, model, dataset, or compute resource, the query may reveal sensitive intent. For instance, a medical AI agent querying oncology datasets may reveal information about user health, or a company agent searching for GPU clusters before a product launch may reveal business

strategy. This may lead to privacy risks like surveillance of organizational behavior, users/enterprise profiling correlation attacks.

This document focuses on the privacy risks associated to the action of entities discovery, i.e., how to protect the privacy of the entity performing the discovery so there is no information leakage. It provides an threat analysis following the guidelines in [RFC6973].

Published entity properties attributes, such as capabilities, endpoints, availability, geographic location, etc., may be also sensitive information. Entities' published information should be controlled by their operators; privacy considerations about published information will discussed in future revision of this document.

Private communication among entities should be part of the communication protocol itself, hence considered out of the scope of this document.

2. Requirements Notation

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. Definitions of Terms

This document assumes familiarity with the terminology defined in [I-D.farrel-dawn-terminology]:

* Attributes
* Discoverable Object
* Discovery
* Discovery Mechanism
* Entity
* Minimum Discoverable Information

4. Applicability Scope

[RFC6973] clearly states that privacy should be included into protocols from the design phase, not afterwards. This is especially true in open and cross-domain ecosystems. As such any discovery protocol designed in the context of DAWN should natively include mechanisms to provide privacy protection. However, the use of such mechanisms may be relaxed in specific contexts, like for instance closed ecosystems. A closed ecosystem is an environment where discovery, communication, identity, orchestration, and capability advertisement are restricted to a single vendor, platform, enterprise, or tightly controlled federation. Example of such ecosystems, may be internal coding assistants, HR workflow agents, finance analysis agents, procurement bots, security automation agents. In such tightly controlled and closed environments privacy requirements might be less strict and privacy protection can be relaxed. Yet, it is important to understand that privacy protection should be maintained by these closed environments when interfacing with other entities outside the local domain.

5. Privacy Threat Analysis

[RFC6973] provides guidelines about privacy in Internet protocols. It also introduces privacy related terminology and it provides a structured framework for analyzing privacy threats in Internet protocols. Many of the privacy issues emerging in entities discovery, map directly onto the privacy concepts identified in [RFC6973]. As such the following subsections go through the the privacy threats identified in [RFC6973] and describe how they relate to DAWN. The subsections assume that the reader is familiar with [RFC6973].

5.1. Combined Security-Privacy Threats

5.1.1. Surveillance

Surveillance consists in observing communications or interactions over time. In a discovery infrastructure a malicious observer can see which agents query which capabilities, organizational interests, workload patterns, operational behavior. Discovery substrates become surveillance points unless privacy protections are built in.

5.1.2. Stored Data Compromise

Stored protocol data becomes a privacy risk if not adequately protected. Discovery systems may store query logs, interaction histories, and other type of metadata, that may reveal usage patterns, workflows, enterprise operations. Hence, stores data should be minimized and protected.

5.1.3. Intrusion

TBD

5.1.4. Misattribution

TBD

5.2. Privacy-Specific Threats

5.2.1. Correlation

Linking multiple interactions like repeated discovery queries, agent identifiers, capability searches, may lead to identify users or entities behavior.

5.2.2. Identification

Protocols may expose identity unnecessarily. In the context of DAWN, discovery may expose organization names, infrastructure ownership, operator identity. Even capability metadata may uniquely fingerprint an entity. Discovery should avoid identity exposure beyond operational necessity.

5.2.3. Secondary Use

TBD

5.2.4. Disclosure

TBD

5.2.5. Exclusion

TBD

6. Similarities with Domain Name System (DNS) privacy protection

Discussion about DNS privacy and similarities with DAWN from a privacy perspective.

TBD

7. Privacy vs Auditability

TBD

8. [RFC6973] Guidelines Compliance

[RFC6973] provides guidance in the form of a questionnaire about a protocol being designed.

[Replies to RFC6973 questionnaire to be added.]

9. Threats Mitigation

[RFC6973] defines three categories of relevant mitigations, namely (1) data minimization, (2) user participation, and (3) security. They apply also in the context of DAWN.

9.1. Technological building blocks

Some privacy-enhancing technologies can be used at an advantage in improving the privacy of entities discovery. Among those technologies, Private Information Retrieval can be used at a benefit in the development of privacy-preserving discovery protocols.

Private Information Retrieval (or PIR) [PIR95] is a technology initially developed in the database realm. It is used by users to retrieve information stored in a server while hiding the exact retrieved information from the server hosting it. As PIR schemes have gained efficiency and become usable at scale in a distributed setting, their use in DAWN to prevent registrars or entities hosting capacities of interest to retrieve information about requesters' interests and activities. In investigating those developments, specific care need to be taken about the ability of academic PIR schemes to cope with the scale at which DAWN needs to operate.

[Future revision to have DAWN-specific mitigation categories]

10. IANA Considerations

This document does not require any IANA action.

11. Security Considerations

TBD.

12. References

12.1. Normative References

[I-D.akhavain-moussa-dawn-problem-statement]
Akhavain, A., Moussa, H., and D. King, "Problem Statement for the Discovery of Agents, Workloads, and Named Entities (DAWN)", Work in Progress, Internet-Draft, draft-akhavain-moussa-dawn-problem-statement-02, 21 May 2026, <https://datatracker.ietf.org/doc/html/draft-akhavain-moussa-dawn-problem-statement-02>.

[I-D.farrel-dawn-terminology]
Farrel, A., Yao, K., Schott, R., and N. Williams, "Terminology for the Discovery of Agents, Workloads, and Named Entities (DAWN)", Work in Progress, Internet-Draft, draft-farrel-dawn-terminology-01, 21 April 2026, <https://datatracker.ietf.org/doc/html/draft-farrel-dawn-terminology-01>.

[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>.

[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/rfc/rfc6973>.

[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>.

12.2. Informative References

[PIR95] Chor, B., Goldreich, O., Kushilevitz, E., and M. Sudan, "Private information retrieval", Proceedings of IEEE 36th Annual Foundations of Computer Science pp. 41-50, DOI 10.1109/sfcs.1995.492461, November 2002, <https://doi.org/10.1109/sfcs.1995.492461>.

Authors' Addresses

Luigi Iannone
Huawei Technologies France S.A.S.U.
18, Quai du Point du Jour
92100 Boulogne-Billancourt
France

Email: luigi.iannone@huawei.com

Antoine Fressancourt
Huawei Technologies France S.A.S.U.
18, Quai du Point du Jour
92100 Boulogne-Billancourt
France
Email: antoine.fressancourt@huawei.com