

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 6 December 2026

A. Farrel
Old Dog Consulting
K. Yao
China Mobile
R. Schott
Deutsche Telekom
N. Williams
Infoblox
4 June 2026

Terminology for the Discovery of Agents, Workloads, and Named Entities
(DAWN)
draft-farrel-dawn-terminology-02

Abstract

The proliferation of distributed systems, Artificial Intelligence (AI) agents, cloud workloads, and network services has created a need for interoperable mechanisms to discover entities. Entities may include AI agents, software services, compute workloads, and other named resources that need to be found and characterised before interaction can begin.

This document defines terminology for Discovery of Agents, Workloads, and Named Entities (DAWN). The intention is that this common set of terms can be used by other documents related to DAWN and so achieve consistency of meaning across the space.

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 6 December 2026.

Copyright Notice

Copyright (c) 2026 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	2
2. Terminology	3
3. IANA Considerations	6
4. Security Considerations	6
5. Privacy Considerations	6
6. Operational Considerations	7
7. Acknowledgements	7
Authors' Addresses	7

1. Introduction

Distributed systems increasingly rely on the dynamic composition of services, agents, and workloads that may not have pre-configured connectivity relationships. For example, an AI agent may need to find another agent with specific capabilities, a workload orchestrator may need to locate compute resources in a particular jurisdiction, or a service consumer may need to discover providers that support a required protocol or a data schema version. Further use cases and scenarios may be considered, but it is out of scope to enumerate them.

In each case, an entity needs knowledge of remote entities before interaction can proceed: what they are, what they offer, and whether they can be trusted. Such knowledge could be obtained through static configuration, but this approach is impractical at scale and across organisational boundaries. Automated discovery mechanisms are therefore needed.

Today, where automated discovery exists, it is typically handled through proprietary directories or platform-specific mechanisms. These approaches do not scale across organisational boundaries and create fragmented ecosystems where entities cannot find entities managed by other organisations.

An interoperable discovery mechanism is needed that builds on existing protocols and tools, benefits from an established trust model, supports proven delegation and federation architectures, and allows organisations to independently publish discovery information.

This document defines common terminology for use in documents that discuss Discovery of Agents, Workloads, and Named Entities (DAWN).

2. Terminology

The terms presented in this section are in alphabetic order for ease of reference. For those wishing to read this document to gain an understanding of the DAWN scenery, if it recommended to read the terms in the order presented in Figure 1.

Core term	Subsidiary term
Entity	Named Entity Agent Workload Task Discovering Entity Discovered Entity
Discovery	Discovery Information Discoverable Object Minimum Discoverable Information Discovery Mechanism Discovery Scope
Capability	Function Attributes Properties Capability Card Trust Indicator
Registration	Capability Exposure Registrar Discoverable Object Validation
Selection	Capability Exchange

Figure 1: Key DAWN Terms in a Readable Order

Agent: A software entity that acts autonomously or semi-autonomously on behalf of a user, organisation, or system. An agent may initiate interactions with other entities, make decisions, and perform tasks.

AI agents are a specific class of agent that employ artificial intelligence techniques.

Attributes: The properties, features, capabilities, skills, etc., that an entity possess or may have access to such as capabilities, skill type, communication language, capacity, task description, contact information, ID, etc. (See also, properties.)

Capability: A description of the functions, services, or actions that an entity can perform. Capabilities may be described using structured schemas such as capability cards.

Capability Card: A structured, machine-readable description of an entity's capabilities and interface. Variants include agent cards, task cards, resource cards, tool cards, and skill cards depending on the type of entity.

Capability Exposure: The processes by which entities expose their capabilities. Such exposure may be part of the registration or discovery processes, or an achieved through and interaction with an entity. (See also, Capability Exchange.)

Capability Exchange: The processes by which entities exchange details of what they can do, dynamic status information, and which particular features or functions they wish to engage.

Capability exposure, exchange, and negotiation are out of scope for DAWN, but will form an essential part of selection and operation of agents.

Discoverable Object: An information object that is discoverable and includes information that defines what an entity is, what attributes it possess, how to reach to the associated entity, etc. May be represented as a capability card.

Discoverable Object Validation: The process that verifies a discoverable object, ensures its compliance to referenced standards, and makes it available to the discovery substrate.

Discovered Entity: An entity whose properties are returned as the

result of a discovery process. A discovered entity may be a specific instance or a member of a class of entities that can perform a desired function.

Discovering Entity: An entity (or its operator) that initiates the discovery process in order to find other entities to interact with.

Discovery: The process by which an entity or its operator locates other entities that are capable of performing a desired function or providing a desired service, and obtains sufficient information to initiate interaction.

Discovery Information: The information returned by a discovery mechanism that allows the discovering entity to decide whether later interaction is possible and desirable.

Discovery Mechanism: A protocol, system, or method used to perform discovery. Examples include Domain Name System (DNS) based service discovery, directory services, and distributed registries.

Discovery Scope: The explicit domain over which discovery is performed. Discovery scope may be specified in one or more dimensions, including but not limited to administrative identifiers (e.g., DNS domain names, AS numbers), trust domains, topological or distance metrics, geographic or jurisdictional boundaries, and temporal constraints. Discovery scope bounds the search space and supports scalability, relevance, and policy enforcement.

Entity: A system component that communicates with other entities in a peer-to-peer or client-server relationship. Entities include, but are not limited to, AI agents, tools, skills, tasks, compute workloads, software services, task owners, network functions, and application endpoints.

Function: The functional processing capability that an entity offers. Examples include task execution, data transformation, inference, routing, steering, storage, and orchestration.

Minimum Discoverable Information (MDI): The minimum amount of information an entity needs to provide to become discoverable. Think of it as common header of a data structure.

Named Entity: An entity that is identified by a stable name within a naming system. The naming system may be hierarchical (e.g., DNS) or flat.

Properties: The discoverable characteristics of an entity. Properties include, but are not limited to, communication protocols, capability cards, location, trust indicators, and operational status.

Registrar: An entity or system responsible for accepting and maintaining records about entities that wish to be discoverable.

Registration: The steps by which agents can register their existence with a registrar. This should include attestation and other security mechanisms.

Registration is out of scope for DAWN, but the information that can be discovered and the trust with which that information is treated are key to any complete system.

Selection: The mechanisms and policies by which an entity determines which discovered entities it will interact with.

Selection is out of scope for DAWN, but depends on information obtained through discovery.

Task: A legacy term kept for continuity with earlier work. A task may be considered as software service.

Trust Indicator: Information associated with an entity that allows a discovering party to assess the trustworthiness or provenance of the entity and its advertised properties. Examples include digital signatures, certificates, and attestations.

Workload: A unit of compute or processing that is deployed and executed within a hosting environment. Workloads may be transient or long-lived and may move between hosting environments.

3. IANA Considerations

This document does not make any requests of IANA.

4. Security Considerations

This document only defines a set of terms. It does not introduce any issues that require security consideration.

5. Privacy Considerations

This document only defines a set of terms. It does not introduce any issues that require privacy consideration.

6. Operational Considerations

This document only defines a set of terms. It does not introduce any issues that require operational consideration.

7. Acknowledgements

The authors wish to acknowledge the contributions of participants in the DAWN discussions that shaped this document.

Jim Mozley, Med Boucadair, and Chenguang Du, and Daniel King provided useful reviews of this document.

Authors' Addresses

Adrian Farrel
Old Dog Consulting
United Kingdom
Email: adrian@olddog.co.uk

Kehan Yao
China Mobile
China
Email: yaokehan@chinamobile.com

Roland Schott
Deutsche Telekom
Germany
Email: Roland.Schott@telekom.de

Nic Williams
Infoblox
United States of America
Email: nwilliams@infoblox.com