

DNSOP
Internet-Draft
Intended status: Informational
Expires: 30 October 2026

J. Ye
W. Cheng
China Mobile
D. Ma
ZDNS
28 April 2026

Problems Statement and Requirements Analysis of DNS for Internet of
Agents (IoA)
draft-ye-problems-and-requirements-of-dns-for-ioa-02

Abstract

In the AI-driven era, DNS is supposed to evolve with technological advancements to accommodate the complex and diverse requirements of the IoA. This draft analyzes the issues surrounding DNS in supporting agents collaboration and explores corresponding technical requirements.

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 30 October 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
1.1. Terminology	3
1.2. Requirements Language	4
2. Problem Statement	4
2.1. Identifiers Reconstruction	4
2.2. Proliferating AI Agents	4
2.3. High-density and Parallel Interaction	5
2.4. Dynamically Changing Services	5
2.5. Upgrade of Resolution Mode	5
2.6. Security Issues	6
3. Requirements Analysis	6
3.1. Global Unique Identifier	6
3.2. Autonomous Capability Registration and Discovery	6
3.3. Information Exchange	7
3.3.1. Rich-information Metadata	7
3.3.2. Data Freshness Maintenance	7
3.4. High-Performance Resolution System	8
3.5. Multi-Dimensional Dynamic Scheduling Strategies	8
3.6. Authentication and Authorization Mechanism	8
4. Security Considerations	9
5. IANA Considerations	9
6. References	9
6.1. Normative References	9
6.2. Informative References	9
Acknowledgements	10
Contributors	10
Authors' Addresses	10

1. Introduction

In the AI-driven era of intelligence, as intelligent agents with autonomous capabilities in perception, decision-making, execution, and learning enter the network, the network ecosystem is undergoing a profound transformation towards intelligentization and autonomization. Agents, as the core interconnected entities in the Internet of Agents (IoA), autonomously discover, efficiently interact, and harmoniously collaborate with human users, other

agents, and various tools, enabling flexible resource scheduling and promoting the Internet towards intelligentization and autonomization.

Currently, the stable operation of the network heavily relies on IP addresses, with the Domain Name System (DNS) serving as a critical network infrastructure responsible for converting human-readable domain names to machine-readable addresses and acting as a bridge for network resources access. Throughout decades of stable operation and global development, there are a suite of protocols and mechanisms also has been developed to establish a fully matured DNS ecosystem. For examples, DNS-Based Service Discovery (DNS-SD) [RFC6763] is designed to facilitate service discovery, Service Binding (SVCB) and HTTPS [RFC9640] records provide instruction of accessing a service, DNS Security (DNSSEC) [RFC4033] [RFC4034] [RFC4035] ensure data origin authentication and data integrity, and protocols such as DNS over HTTPS (DoH) [RFC8484], DNS over TLS (DoT) [RFC7858] and DNS over QUIC (DoQ) [RFC9250] supply encrypted transmission and enhance privacy and security. Therefore, in the context of the flourishing development of AI agents, the DNS is supposed to evolve with technological advancements to serve the complex and diverse requirements of the IoA. During interactions and collaborations between humans and agents, agents and agents, DNS as an infrastructure will continue to play a key role, providing technical support for efficient agent registration and discovery, real-time data synchronization, and intelligent scheduling and decision-making. Moreover, its capabilities might be further expanded to achieve semantic awareness, and effective orchestration of agents interactions.

This draft aims to conduct an in-depth analysis of the issues surrounding the DNS system in supporting agents collaboration and to explore corresponding technical requirements, thereby providing robust support for the large-scale implementation and efficiency enhancement of the IoA.

1.1. Terminology

DNS: Domain Name System

DNS-SD: DNS-Based Service Discovery

EDNS: Extension Mechanisms for DNS

1.2. Requirements Language

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.

2. Problem Statement

2.1. Identifiers Reconstruction

DNS, centered around domain names and IP addresses, fulfills a crucial role in addressing, mapping domain names to IP addresses. However, the existing IP addresses struggle to identify agents effectively. This because that it is predicted that the number of AI agents will reach the scale of hundreds of billions in the future. Given the limited address space of IPv4, it is inadequate to support large-scale agent deployment and ensure stable connectivity. Even if a full transition to IPv6, which offers abundant address resources, can alleviate the address shortage, issues such as address instability and oversized routing tables will still arise. This is because that the IPv6 interface IDs undergo periodic rotation for security. Additionally, the IPv6 addresses of some physical AI agents, such as embodied intelligent agents, dynamically change with variations of their geographical locations and the network they access. The aforementioned factors make it challenging to uniquely identify an intelligent agent using its IP address.

2.2. Proliferating AI Agents

With the emergence of new agents, corresponding resource records will be generated accordingly. However, manually adding records to authoritative servers are inefficient and cannot keep pace with the rapid generation of agents. Currently, there is no suitable capability-based registration and discovery.

On one hand, the most interactions among agents are based on capabilities, but the current names of each level domains in the hierarchical architecture, which primarily convey information, such as organization and region, fails to intuitively grasp and describe their basic functionalities, making it difficult to directly use domain names to register and discover for agents.

On the other hand, in the Internet, considering the intricate roles and the enormous quantity of devices involved, the current service discovery and registration mechanisms (e.g. DNS-SD) are deficient in essential security authentication capabilities, rendering them

ineffective for application in the IoA. Consequently, these situation underscores an urgent necessity for developing a service discovery and registration mechanism specifically tailored to the IoA.

2.3. High-density and Parallel Interaction

Interactions in IoA exhibit high frequency and complexity. Agent interactions often involve multiple subtask calls within a single task query, triggering multiple DNS queries and significantly increasing the number of queries and densities, accompanied by parallel queries. This high-frequency and concurrent query pattern imposes extremely high demands on the processing performance of the DNS system.

2.4. Dynamically Changing Services

As services undergo continuous evolution and agents engage in frequent interactions, the services of agents are developing and upgrading, and their operational states are also in a state of constant change. New services continually emerge, while existing services may be gradually phased out or subject to updates. Meanwhile, agent states may transition from active to inactive, or continuously fluctuate load conditions. However, the resource records (RR) within the existing DNS system remain static and are updated only infrequently, thereby failing to accurately and promptly reflect the latest situations of agents.

2.5. Upgrade of Resolution Mode

The resolution mechanism of the current DNS system is relatively simplistic and falls short of optimally matching capabilities with resources during agent interactions. When processing multiple resource records associated with the same tree-tuple (name, class, type), existing mechanisms frequently depend on scheduling dimensions such as round-robin, weights, and geographical proximity. Although there are scenarios that take resource load into account, these mechanisms merely offer crude estimations based on simple request counts. Such estimations can significantly deviate from the actual load, thereby leading to reduced scheduling flexibility and accuracy.

2.6. Security Issues

During the interaction of agents, the underlying security issues should not be overlooked. For example, it is crucial whether an agent's identity is forged, whether the capabilities it registers are genuine, and whether these capabilities accurately correspond to its claimed identity. Any lapses in these areas could potentially expose users to attacks, threats, privacy leakage and other risks, or even lead to widespread network breakdown.

3. Requirements Analysis

The proliferation of agents is driving the network towards enhanced efficiency, intelligence, and flexibility. During the processes of autonomous discovery, efficient interaction, and collaborative collaboration among agents as well as between agents and users, new requirements are imposed on the DNS, such as identification, data structure, and resolution mode. The detailed requirement analysis of the key capabilities required for the DNS within the IoA is as below.

3.1. Global Unique Identifier

It is of paramount importance to assign a unique identity identifier to each agent. Firstly, this immutable ID will effectively isolate the impact caused by the frequent changes of IP addresses or URLs. In addition, while facilitating precise identification and differentiation of individual agents, it also provides a solid foundation for the verification of agent identities.

3.2. Autonomous Capability Registration and Discovery

As services keep evolving, new agents are constantly emerging, and there will be new agents registered in the Internet at any given time. Therefore, it is of vital importance to achieve automatic capability registration, discovery, and publication of agents without manual intervention.

One approach is to directly achieve the mapping from agents' capabilities to their identity identifiers by introducing capability-aware domain names. New domain names ought to incorporate additional hierarchical levels that convey capability-related information, beyond the basic information (such as organizations and regions) typically found in conventional domain names. The newly introduced name levels should be capable of intuitively and succinctly representing the capabilities of a specific type of agent or other distinctive attributes. For example, a domain name designated for an image processing agent might include keywords relevant to image processing, such as "...appname.ImageProcess.organizer." This

approach would enable authoritative servers to directly identify and retrieve all image processing agents through domain names when other agents or users are searching for image processing agents, thereby enhancing the efficiency of agent discovery.

The other involves refining the existing DNS-SD mechanism. Crucially, this mechanism should be enhanced by incorporating authentication and rights verification to effectively prevent counterfeiting and impersonation attacks during the registration and discovery of a large number of agents. Furthermore, the structure "< Service >.< Domain >" defined in [RFC6763] also should introduce capability-aware service names to handle the mapping from agents' capabilities to domain names, locating a scope of candidate agents.

3.3. Information Exchange

3.3.1. Rich-information Metadata

In addition to utilizing domain names or PTR [RFC1035] to narrow down the scope of agents that provide specific capabilities, there should also be other data presenting more detailed descriptions of these agents, thereby facilitating further decision-making. Therefore, resource records (e.g. SRV, TXT, SVCB) should be capable of carrying extensive metadata, encompassing detailed capability descriptions, configuration parameters, load conditions, and other pertinentially valuable information about the agents. These metadata serves as an agent's "digital business card," providing other agents, users, and schedulers with a comprehensive insight into the agent's information. For example, an agent's RRset could include its processing performance, supported protocol types, and current workload, thereby assisting schedulers and other agents in making well-informed decisions.

Meanwhile, to gain a better understanding of requester intentions and preferences, it is recommended to incorporate additional information into DNS request messages through Extended DNS (EDNS) [RFC6891] or service-related tags labeled by recursive server or gateways.

3.3.2. Data Freshness Maintenance

Given the frequent changes in agent information, a data update mechanism is imperative to guarantee the freshness of data. Within this mechanism, subscription-push [RFC8490], regular detection and/or periodic reporting should be employed to ensure that the data (e.g. RRset) remains up-to-date. For data with low change frequencies, such as capability descriptions and configurations, real-time updates can be adopted, pushing or reporting relevant resource records when subscribed data undergoes changes. For data with high change

frequencies, such as workload and network performance, periodic updates can be utilized to prevent adverse impacts on network and processing performance. Additionally, to prevent a large number of simultaneous data refreshes across the network, a standby mode can be configured for agents with low usage, thereby reducing network load while maintaining low power consumption.

3.4. High-Performance Resolution System

With the growing number of agents and increasing interaction densities, the entire resolution system must possess high performance, capable of rapidly and accurately processing a large number of concurrent query requests.

3.5. Multi-Dimensional Dynamic Scheduling Strategies

Given the enormous number of agents, selecting the most suitable one from a pool of agents of the same type poses a significant challenge. The DNS for IoA should be equipped with multi-dimensional dynamic scheduling capabilities. It should dynamically select the optimal resolution result based on agent resource records (including capability descriptions, geographical locations, workload, etc.), business demands obtained through EDNS or other labels in packets, and in conjunction with network environments and load-balance strategies to achieve appropriate resource allocation.

3.6. Authentication and Authorization Mechanism

In agent networks, authentication and authorization are crucial for ensuring network security and reliability. This includes verifying agent IDs and capabilities (e.g., resource records). Strict identity authentication guarantees that only legitimate agents whose IDs are corresponding to their feature can access the network for registration. Meanwhile, capability authentication serves to prevent the advertisement of false capability information, thereby ensuring the accuracy of information about agents. These verifications occur during capabilities registration, with authoritative servers validating the relevant information of service providers to ensure their capability and qualification to provide the corresponding services. Additionally, service consumers (e.g. terminals and other agents) also could validate the received data to prevent tampering during transmission by intermediate third parties. By establishing a robust authentication and authorization mechanism, a secure and reliable network ecosystem can be constructed for IoA.

4. Security Considerations

TBD.

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. 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>>.

6.2. Informative References

- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/rfc/rfc6763>>.
- [RFC9640] Watsen, K., "YANG Data Types and Groupings for Cryptography", RFC 9640, DOI 10.17487/RFC9640, October 2024, <<https://www.rfc-editor.org/rfc/rfc9640>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/rfc/rfc4033>>.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <<https://www.rfc-editor.org/rfc/rfc4034>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/rfc/rfc4035>>.

- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/rfc/rfc8484>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/rfc/rfc7858>>.
- [RFC9250] Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", RFC 9250, DOI 10.17487/RFC9250, May 2022, <<https://www.rfc-editor.org/rfc/rfc9250>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/rfc/rfc1035>>.
- [RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <<https://www.rfc-editor.org/rfc/rfc6891>>.
- [RFC8490] Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", RFC 8490, DOI 10.17487/RFC8490, March 2019, <<https://www.rfc-editor.org/rfc/rfc8490>>.

Acknowledgements

Contributors

Authors' Addresses

Jiaming Ye
China Mobile
Email: yejiaming@chinamobile.com

Weiqiang Cheng
China Mobile
Email: chengweiqiang@chinamobile.com

Di Ma
ZDNS
Email: madi@zdns.cn