Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 May 2026

X. Lee
B. Peng
ICT
Y. Fu
Y. Wang
ICT;UCAS

31 October 2025

DNS Resource Records for the Data Interoperability Protocol
draft-analyst-dns-dip-rr-00

Abstract

This document specifies a set of new DNS Resource Record (RR) types to support the Data Interoperability Protocol (DIP). DIP models entities as "Data Objects" (DOs) indexed by "Data Object Identifiers" (DOIs), which are syntactically equivalent to DNS domain names. The RR types specified herein-DLR, DOPK, DOAUTH, DODIGEST, DOCG, and DOALGO-are designed to hold key attributes of a DO. This approach addresses the inherent limitations of repurposing existing RR types (such as TXT) with respect to structured data representation, efficiency in handling binary data, and query granularity. The result is a native, efficient, and extensible mechanism for discovering and managing DO metadata via the DNS.

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 4 May 2026.

Copyright Notice

Table of Contents

1. Introduction

1.1. Background:The Evolution Towards Data Interoperability

The history of the Internet is marked by the progressive refinement of its core interoperability unit. In its early stages, the TCP/IP suite addressed the challenge of reliable data transmission between hosts over unreliable networks, defining the paradigm of "host interconnection." Subsequently, protocols such as HTTP and DNS enabled the "website interconnection" era by providing a structured way to organize and locate information. We are now entering a new

phase, where the central challenge is to enable secure, trusted, and rights-respecting "data interconnection" within a global, untrusted network environment.

However, the full potential of data as a key economic factor is currently hindered by significant obstacles, including data silos, disputes over data sovereignty, data misuse, and privacy breaches. To address these challenges, a new foundational protocol is needed to establish a unified framework that facilitates the secure flow and collaboration of data across diverse entities and jurisdictions, while ensuring that data owners retain control over their assets.

1.2. The Data Interoperability Protocol(DIP)

The Data Interoperability Protocol (DIP) is proposed as a standardized solution to the aforementioned challenges. DIP abstracts any entity in the real or digital world-including individuals, organizations, devices, services, or even an abstract dataset-into a unified construct known as a "Data Object" (DO).

Each DO is identified by a globally unique "Data Object Identifier" (DOI). A core design principle of DIP is that *a DOI is syntactically and semantically equivalent to a DNS domain name*. This design allows DIP to directly leverage the DNS-the world’s largest, most stable, and most widely distributed naming and resolution system-as its trust anchor.

A DO is fully described by a set of "Data Object Attributes" (DOAs). These attributes are structured metadata that define the core characteristics of the DO, such as its network-reachable address (locator), credentials for authentication (public key), access control policies (authorization), a value for content integrity verification (digest), and compliance governance requirements (classification, algorithm governance).

1.3. Rationale for New DNS Resource Record Types

Given that a DOI is equivalent to a DNS domain name, storing DOAs in the DNS and using the DNS protocol for their publication and resolution is a logical and technically advantageous choice. The DNS’s hierarchical structure, high availability, scalability, and mature global operational ecosystem provide a robust infrastructure for DIP’s metadata registration and resolution needs.

A seemingly straightforward implementation approach would be to utilize existing DNS Resource Record (RR) types, such as the TXT record defined in [RFC1035]. However, upon careful analysis, this approach presents significant deficiencies that fail to meet DIP's requirements for structure, efficiency, and security:

* *Lack of Structure and Semantics*: The RDATA of a TXT record is fundamentally a sequence of unstructured characters. Encoding DIP's complex, structured attributes (such as an authorization policy with multiple fields) into a TXT record would necessitate complex, non-standard, application-layer parsing logic. This is highly prone to ambiguity and interoperability issues between different implementations. Moreover, the TXT mnemonic itself provides no semantic information about the purpose of the data it contains.

* *Inefficiency*: DIP attributes include native binary data, such as cryptographic public keys and digital signatures. Storing such data in a TXT record requires text-based encoding (e.g., Base64 or hexadecimal), which significantly increases the size of the RDATA and imposes unnecessary computational overhead for encoding and decoding on both clients and servers.

* *Coarse Query Granularity*: If multiple attributes are bundled into a single TXT record, a client needing only a specific attribute (e.g., only the public key) must still request and receive the entire record's content. The client must then parse and extract the desired piece of information locally. This is wasteful of both network bandwidth and client-side processing resources.

To overcome these limitations and to build an efficient, clear, extensible, and semantically precise foundation for the Data Interoperability Protocol, this document proposes a set of new, dedicated DNS Resource Record (RR) types. Each new RR type corresponds exactly to a core Data Object Attribute (e.g., DLR, DOPK, DOAUTH). This approach provides a solution that is native to structured and binary data, allows for fine-grained queries, and aligns closely with the design philosophy of the DNS.

2. Definitions

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119];[RFC8174].

2.2. Defined Terms

This document uses the following terms:

* *DIP*: Data Interoperability Protocol.
* *DO*: Data Object, a structured digital representation of the objective world. Data objects do not distinguish between people, institutions, or objects, and consist of data objects and data object contents.
* *DOI*: Data Object Identifier, a unique identifier for a data object, which is the index of data object attributes.
* *DOA*: Data Object Attribute, information related to a data object, such as locators, ownership information.
* *DLR*: Data Locator, used to locate the information entry of data object content or related resources, pointing to identity or content resources related to the data object.

3. Data Object Attribute Resource Record Types

3.1. The DLR (Data Locator) Resource Record

The DLR RR specifies a locator for the Data Object’s content or associated resources. This is typically a URI.

* *Mnemonic*: DLR

* *Type Code*: TBD1 (to be assigned by IANA)

* *RDATA Format*: A single, variable-length field containing the URI, encoded as a <character-string> as defined in [RFC1035], Section 3.3.

* *Presentation Format*: <URI>

* *Example*:

```
fuxizhiku.example.com.  DLR
"https://data.example.com/space/fuxizhiku/content"
```

3.2. The DOPK (Data Object Public Key) Resource Record

The DOPK RR stores the public key associated with a Data Object, used for identity authentication and cryptographic operations.

* *Mnemonic*: DOPK

* *Type Code*: TBD2 (to be assigned by IANA)

* *RDATA Format*:

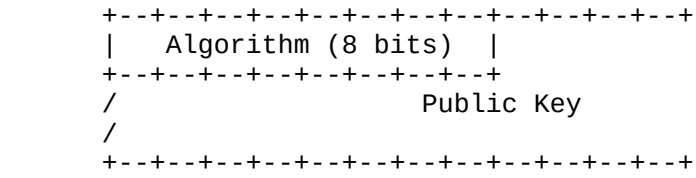

* Algorithm An 8-bit identifier for the public key's cryptographic algorithm. Values are managed by IANA in a new registry. Initial values:

   - 1: RSA
   - 2: ECDSA/P-256
   - 3: Ed25519

* *Public Key*: Variable-length binary data of the public key. The format is algorithm-specific.

* *Presentation Format*: <Algorithm> <Base64-encoded Public Key>

* *Example* (line breaks for clarity):

```
user1.example.com.  DOPK  1 ( AQAB... )
```

3.3. The DOAUTH (Data Object Authorization) Resource Record

The DOAUTH RR defines an authorization (right) granted to a subject (another Data Object) over the owner Data Object.

* *Mnemonic*: DOAUTH

* *Type Code*: TBD3 (to be assigned by IANA)

* *RDATA Format*:

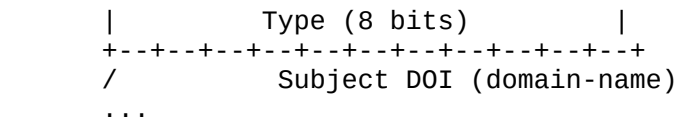

* *Type*: An 8-bit unsigned integer defining the authorization type.

* *Subject DOI*: The domain name of the Data Object being granted the authorization.

* *Permission DOI*: The domain name of a Data Object whose content describes the detailed permissions.

* *Creator DOI*: The domain name of the Data Object that created this authorization record.

* *Parent DOI*: An optional parent Data Object, represented as a domain name. If not present, this field is the root domain ".".

* *Key*: An encrypted symmetric key for accessing the DO content.

* *Confirmation*: The subject's signature over the DO content digest.

* *Presentation Format*: <Type> <Subject DOI> <Permission DOI> <Creator DOI> <Parent DOI> (<Base64-encoded Key>) (<Base64-encoded Confirmation>)

* *Example*:

```
data1.example.com. DOAUTH 192 user1.example.com.
perm1.example.com. owner1.example.com. .
( dGhpcyBpcyBhIGtleQ== ) ( dGhpcyBpcyBhIHNpZw== )
```

3.4. The DODIGEST (Data Object Digest) Resource Record

The DODIGEST RR contains a cryptographic digest (hash) of the Data Object's content.

* *Mnemonic*: DODIGEST

* *Type Code*: TBD4 (to be assigned by IANA)

* *RDATA Format*:

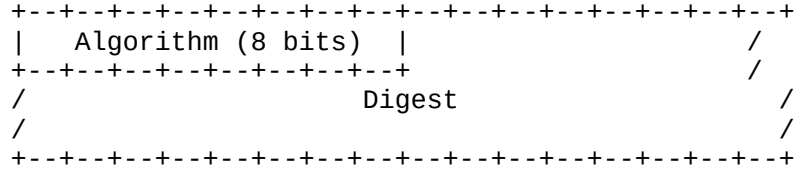

* Algorithm An 8-bit identifier for the digest algorithm. Initial values:

  - 1: SHA-256
  - 2: SHA-384
  - 3: SHA-512

* *Digest*: Variable-length binary data of the content digest.

* *Presentation Format*: <Algorithm> <Hex-encoded Digest>

* *Example*:

```
data1.example.com.  DODIGEST  1
e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855
```

3.5. The DOCG (Data Object Classification and Grading) Resource Record

The DOCG RR specifies the classification and grade of the Data Object.

* *Mnemonic*: DOCG

* *Type Code*: TBD5 (to be assigned by IANA)

* *RDATA Format*:

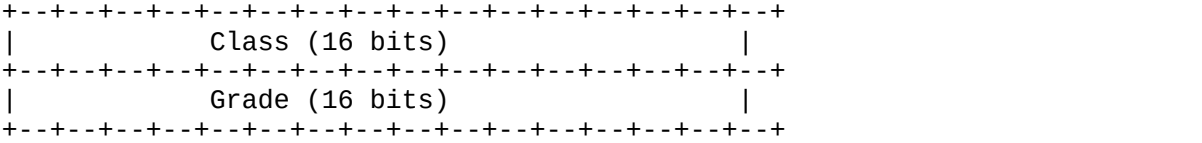

* *Class*: A 16-bit unsigned integer representing the data category.

* *Grade*: A 16-bit unsigned integer representing the data security level.

* *Presentation Format*: <Class> <Grade>

* *Example*:

```
data1.example.com.  DOCG  2048 33792
```

3.6. The DOALGO (Data Object Algorithm Governance) Resource Record

The DOALGO RR specifies the permitted algorithms for processing the Data Object.

* *Mnemonic*: DOALGO

* *Type Code*: TBD6 (to be assigned by IANA)

* *RDATA Format*:

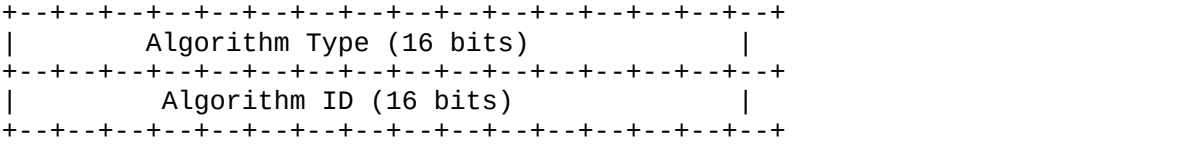

* *Algorithm Type*: A 16-bit unsigned integer representing the algorithm category.

* *Algorithm ID*: A 16-bit unsigned integer representing the specific algorithm.

* *Presentation Format*: <Algorithm Type> <Algorithm ID>

* *Example*:

```
data1.example.com.  DOALGO  1024 4097
```

4. Mechanism

4.1. Declaration Trigger Scenarios

In the entire process of data interoperability, when involving core operations such as data object attribute registration, access authorization, and content query, the declaration of corresponding DNS resource records must be triggered. The specific scenarios are as follows:

Data object initialization phase: When the data owner completes identity registration for the data object, they must simultaneously declare DLR, DOPK records, and DODIGEST, and optionally provide DOAUTH, DOCG, and DOALOG information to clarify the data storage address and identity authentication credentials, ensuring that the data object can be located and identified.

Authorization relationship change phase: When the data owner grants or revokes data usage rights and ownership to users, the DOAUTH record must be declared to record key information such as authorization type and principal DOI, ensuring that the authorization relationship is traceable and verifiable.

Data content update phase: After the data object content is modified, the DODIGEST record must be re-declared to update the content digest information, ensuring that the data integrity can be verified through the digest during subsequent queries and avoiding the use of tampered data.

Security policy configuration phase: When the data classification and grading are adjusted (e.g., from "general data" to "important data") or the allowed algorithms are changed, the DOCG and DOALGO records must be declared respectively to clarify the data security level and algorithm usage scope, supporting compliance detection and access control.

4.2. Declaration Effectiveness Mechanism

The declaration operation takes effect through a four-step process of "submission - verification - storage - synchronization" to ensure that the records are authentic, usable, and trusted throughout the entire process. The specific mechanism is as follows:

Declaration submission: The data owner or authorized principal submits a record declaration request to the data identification system through the data interoperability client in accordance with the format requirements of the corresponding RR type. The request must include the principal’s private key signature for identity verification.

Legality verification: After receiving the request, the data identification system first verifies the validity of the submitter’s signature through the DOPK record to confirm the identity is legal; then checks the record format (e.g., whether DLR conforms to URI specifications, whether the category/level values of DOCG are within predefined ranges) to ensure the record content is compliant.

Trusted storage: After successful verification, the identification system stores the record in the DNS zone and automatically associates it with the data object’s DOI (unique identifier) to form a "DOI - resource record" mapping relationship; at the same time, it triggers the DNSSEC signature process to add a digital signature to the record, ensuring that integrity can be verified during subsequent queries.

*Synchronized across the entire network*The identification system synchronizes the newly declared records to associated DNS server nodes in accordance with the DNS protocol’s zone transfer mechanism, ensuring that data users can obtain the latest and consistent record information when querying in different network environments, avoiding interoperability failures due to unsynchronized records.

4.3. Record Usage Constraints

In the use of record query, invocation, etc., the following constraint rules must be followed to ensure data security and operational compliance:

Query permission constraints: Only principals that have passed identity authentication and obtained authorization (such as data users registered in the DOAUTH record) can query sensitive records such as DOAUTH and DODIGEST; when unauthorized principals query, the data identification system will return a "insufficient permission" error to avoid leakage of sensitive information.

Timeliness constraints: Records such as DODIGEST and DOCG must be consistent with the actual state of the data object. When the data content or security level changes, the principal must update the records within 24 hours; records that are not updated beyond the time limit will be marked as "record to be verified" in the query response, reminding users to verify the current state of the data.

Associated usage constraints: Before querying data content, the storage address must be obtained through the DLR record, and then the content integrity must be verified through the DODIGEST record. The two must be used together; if the address pointed to by DLR is inaccessible or the digest verification fails, the data access operation must be terminated to prevent obtaining invalid or tampered data.

Compliance adaptation constraints: The algorithms declared in the DOALGO record must be adapted to the data classification and grading (DOCG). For example, "core data" only allows the use of encryption algorithms that meet national secret standards; if the algorithm does not match the level, the data exchange service will determine it as "non-compliant" during compliance detection and block the data exchange process.

5. Security Considerations

The DNS resource records defined in this document contain sensitive information, including public keys, authorization policies, and data locators. The integrity and authenticity of these records are critical for the security of the Data Interoperability Protocol.

It is therefore REQUIRED that zones containing these RR types be protected using DNSSEC ([RFC4033];[RFC4034];[RFC4035]). Clients querying for these records MUST perform DNSSEC validation to ensure that the responses have not been tampered with.

The publication of authorization information (DOAUTH) in the public DNS may have privacy implications, as it reveals relationships between data owners and users. Implementers and zone administrators should be aware of these implications and use these records in accordance with applicable privacy policies and regulations.

The DOPK record publishes a public key. The security of the corresponding private key is the responsibility of the Data Object's owner and is outside the scope of this document. Compromise of a private key will allow an attacker to impersonate the Data Object.

6. IANA Considerations

This document requests IANA to take the following actions:

1. Assign RR type codes for DLR, DOPK, DOAUTH, DODIGEST, DOCG, and DOALGO from the "Resource Record (RR) TYPEs" registry.

Record	Value	Meaning
DLR	TBD1	Data Object Locator
DOPK	TBD2	Data Object Public Key
DOAUTH	TBD3	Data Object Authorization
DODIGEST	TBD4	Data Object Digest
DOCG	TBD5	Data Object Class & Grade
DOALGO	TBD6	Data Object Algorithm Governance

2. Create a new registry named "DOPK and DODIGEST Algorithm Identifiers". This registry will manage algorithm identifiers for the DOPK and DODIGEST RRs. Initial registrations are provided in Sections 3 and 5 of this document. The policy for this registry SHOULD be "Specification Required" [RFC8126].

7. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[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/info/rfc2119>.

[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/info/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/info/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/info/rfc4035>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[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/info/rfc8174>.

Acknowledgements

The authors would like to thank the World Internet Conference (WIC) for providing valuable platforms for technical exchange and discussion, which contributed to the refinement of ideas in this document.

Authors' Addresses

Xiaodong Lee
Institute of Computing Technology, Chinese Academy of Sciences
Email: XL@ict.ac.cn

Botao Peng
Institute of Computing Technology, Chinese Academy of Sciences
Email: pengbotao@ict.ac.cn

Yufan Fu
Institute of Computing Technology, Chinese Academy of Sciences
Email: fuyufan20z@ict.ac.cn

Yumeng Wang
Institute of Computing Technology, Chinese Academy of Sciences
Email: wangyumeng24s@ict.ac.cn