

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 August 2026

M. Nottingham
Cloudflare
D. Adrian
Google
19 February 2026

DNS Censorship Transparency
draft-nottingham-dnsop-censorship-transparency-00

Abstract

[I-D.ietf-dnsop-structured-dns-error] introduces structured error data for DNS responses that have been filtered. This specification allows more specific details of filtering incidents to be conveyed.

Discussion Venues

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

Source for this draft and an issue tracker can be found at <https://github.com/mnot/public-resolver-errors>.

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 August 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. Example	3
1.2. Notational Conventions	4
2. Data Types	4
2.1. DNS Filtering Database Entry	4
2.2. DNS Filtering Database Entry List	5
3. Database Entry Resolution Templates	5
4. IANA Considerations	6
4.1. EXTRA-TEXT JSON Names	6
4.2. The DNS Filtering Database Registry	6
5. Security Considerations	6
6. Privacy Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Appendix A. Acknowledgements	8
Authors' Addresses	8

1. Introduction

Internet DNS resolvers are increasingly subject to legal orders that require blocking or filtering of specific names. Because such filtering happens during DNS resolution, there is not an effective way to communicate what is happening to end users, often resulting in misattribution of the issue as a technical problem, rather than a policy intervention.

Some organizations such as Lumen [lumen] monitor legally-mandated filtering as a public service, tracking specific filtering incidents in publicly accessible databases. Other parties may also choose to track filtering requests over time and make them available.

This draft defines a mechanism for DNS resolvers to convey identifiers for entries in such databases, based upon the structured error data for DNS responses introduced by [I-D.ietf-dnsop-structured-dns-error].

A consuming party (e.g., a Web browser) can use this information to construct a link to the specific entry in one or more databases of filtering incidents. This enables user agents to direct users to additional context about the filtering incident they encountered.

The information conveyed is a DNS Filtering Database Entry, specified in Section 2.1. This abstraction is necessary because allowing DNS resolvers to inject links or user-visible messages would bring unique challenges. DNS resolvers are often automatically configured by unknown networks and DNS responses are unauthenticated, so these messages can come from untrusted parties -- including attackers (e.g., the so-called "coffee shop" attack) that leverage many users' lack of a nuanced model of the trust relationships between all of the parties that are involved in the service they are using.

This specification attempts to mitigate these risks by minimising the information carried in the DNS response to abstract, publicly registered identifiers associated with databases of filtering incidents -- the DNS Filtering Database Entry -- rather than arbitrary URLs. An application can choose which database identifiers they support and are willing to direct their users to, without enabling every DNS server to surface arbitrary links and text, and without requiring the consuming party to independently track which URLs are in use.

1.1. Example

In typical use, a DNS query that is filtered might contain an Extended DNS Error Code 17 (see [RFC8914]) and an EXTRA-TEXT field "fdb", which is an array of references to filtering database entries:

```
{
  "fdb": [
    { "db": "example",
      "id": "abc123" },
    { "db": "lumen",
      "id": "def456" }
  ]
}
```

This indicates that the filtering incident can be accessed in two different databases, along with the ID associated with each database. In this example, the data is available in the "example" database at identifier "abc123", and in the "lumen" database at identifier "def456".

An application that decides to present the "example" entry to its users would look up "example" in a local copy of the DNS Filtering Incident Database Registry (see Section 4.2) and obtain the corresponding URI template (see Section 3). For purposes of this example, assume that the registry entry for that value contains:

```
https://example.com/filtering-incidents/{id}
```

That URI template can be expanded using the value of "id" to:

```
https://example.com/filtering-incidents/abc123
```

The application could (but might not) then decide to convey some or all of this information to its user; for example:

The webpage at www.example.net was blocked by your DNS resolver due to a legal request. More information is available here:

```
https://example.com/filtering-incidents/abc123
```

Note that there is no requirement for the resolver to convey any particular information here. The resolver both chooses which database providers it supports, and can evaluate whatever mechanisms it chooses to determine if and when to provide a link to the database entry.

1.2. Notational Conventions

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. Data Types

This section defines the data types used to look up the details of a filtering incident from a DNS error response. Note that these identifiers are not suitable for presentation to end users.

2.1. DNS Filtering Database Entry

A Filtering Database Operator ID is a string identifier for the operator of a database of filtering incidents. It uses the key "db".

A Filtering Incident ID is a string identifier for a particular filtering incident. It might be specific to a particular request, but need not be. It uses the key "id".

An object containing both a Filtering Database Operator ID and a Filtering Incident ID is a Filtering Database Entry.

```
{
  "db": "example",
  "id": "abc123"
}
```

2.2. DNS Filtering Database Entry List

A DNS Filtering Database Entries list is an array of Filtering Database Entry objects. Each entry MUST be related to the same underlying incident.

It is carried in the EXTRA-TEXT field of the Extended DNS Error with the JSON field name "fdb". For example:

```
{
  "fdb": [ { ... }, { ... }, ... ]
}
```

Different applications will implement support for a varying set of database operators. Resolvers provide a list of entries (rather than a single entry) so that they can support as many databases as possible.

3. Database Entry Resolution Templates

An Database Entry Resolution Template is a URI Template [RFC6570] contained in the DNS Filtering Database Registry (Section 4.2) that, upon expansion, provides a URI that can be dereferenced to obtain details about the filtering incident.

It MUST be a Level 1 or Level 2 template (see Section 1.2 of [RFC6570]). It has the following variables from the Filtering Database Entry (see Section 2.1) available to it:

db: the Filtering Database Operator ID

id: the Filtering Incident ID

For example:

<https://resolver.example.com/filtering-incidents/{id}>

Applications MUST store a local copy of the DNS Filtering Database Registry (Section 4.2) for purposes of template lookup; they MUST NOT query the IANA registry upon each use. The registry is keyed by the Filtering Database Operator ID.

4. IANA Considerations

4.1. EXTRA-TEXT JSON Names

IANA will register the following fields in the "EXTRA-TEXT JSON Names" sub-registry established by [I-D.ietf-dnsop-structured-dns-error]:

JSON Name: "fdb"
Short Description: a array of filtering database entries
Mandatory: no
Specification: this document

4.2. The DNS Filtering Database Registry

IANA will establish a new registry, the "DNS Filtering Database Registry." Its registration policy is first-come, first-served (FCFS), although IANA may refuse registrations that it deems to be deceptive or spurious.

It contains the following fields:

Name: The name of the DNS Filtering Database
Contact: an e-mail address or other appropriate contact mechanism
Filtering Database Operator ID: see Section 2.1
Database Entry Resolution Template: see Section 3

The Database Entry Resolution Template can be updated by the contact at any time. However, operators SHOULD accommodate potentially long lag times for applications to update their copies of the registry.

5. Security Considerations

This specification does not provide a way to authenticate that a particular filtering incident as experienced by an application was actually associated with the information presented. This means that an attacker (for example, one controlling a DNS resolver) can claim that a particular filtering incident is occurring when in fact it is not. However, a successful attack would need to reuse an existing DNS Filtering Database Operator ID and Filtering Incident ID that

combine to expand to a URL that can be successfully dereferenced. Doing so is not currently thought to be particularly advantageous to an attacker to do so. Future iterations of this specification may introduce more robust protections.

The details of DNS responses are not available to all applications, depending on how they are architected and the information made available to them by their host. As a result, this mechanism is not reliable; some applications will not be able to display this error information.

Because the registry is first-come first-served, applications (such as Web browsers) will need to exercise judgement regarding which database operators' error messages they display to users. This decision might be influenced by the identity of the resolver producing the error message, the database operator, or local configuration.

6. Privacy Considerations

An application dereferencing an incident URL may reveal the IP address of the user to the filtering database operator, along with the fact that the user attempted to resolve a specific domain whose resolution was filtered or censored. In some jurisdictions this exposure may carry meaningful risk to the user. While allowing users the choice to trust only specific filtering databases may mitigate the risk from a filtering database operator, on-path network observers may still infer that the user encountered a filtering incident for a sensitive domain based on a connection to the filtering database.

To mitigate this risk, applications can choose to route incident URL fetches through a privacy-preserving mechanism like a proxy. In absence of such a proxy, applications **MUST NOT** automatically fetch incident URLs without explicit user action.

Applications and users should be aware that a resolver and filtering database operator acting in collusion, or a single party operating in both roles, could use unique Filtering Incident IDs to correlate user activity across requests. As such, applications should get explicit input from users on which filtering database operators they trust, and consequently disregard any entries from operators the user does not explicitly trust.

7. References

7.1. Normative References

[I-D.ietf-dnsop-structured-dns-error]

Wing, D., Reddy, K., T., Cook, N., and M. Boucadair,
"Structured Error Data for Filtered DNS", Work in
Progress, Internet-Draft, draft-ietf-dnsop-structured-dns-
error-16, 1 February 2026,
<<https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-structured-dns-error-16>>.

[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>>.

[RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M.,
and D. Orchard, "URI Template", RFC 6570,
DOI 10.17487/RFC6570, March 2012,
<<https://www.rfc-editor.org/rfc/rfc6570>>.

[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>>.

[RFC8914] Kumari, W., Hunt, E., Arends, R., Hardaker, W., and D.
Lawrence, "Extended DNS Errors", RFC 8914,
DOI 10.17487/RFC8914, October 2020,
<<https://www.rfc-editor.org/rfc/rfc8914>>.

7.2. Informative References

[lumen] "Lumen", n.d., <<https://lumendatabase.org/>>.

Appendix A. Acknowledgements

Thanks to Lars Eggert, Tommy Pauly, Emily Stark, and Martin Thomson
for their input to this specification.

Authors' Addresses

Mark Nottingham
Cloudflare
Melbourne
Australia
Email: mnot@mnot.net
URI: <https://www.mnot.net/>

David Adrian
Google
United States of America
Email: davidcadrian@gmail.com