Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 26 August 2025

M. Nottingham
Cloudflare

22 February 2025

# DNS Filtering Details for Applications
# draft-nottingham-public-resolver-errors-01

## Abstract

[I-D.ietf-dnsop-structured-dns-error] introduces structured error data for DNS responses that have been filtered. This draft suggests additions to that mechanism that enable applications to convey the details of some filtering incidents to their users.

## 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 26 August 2025.

## Copyright Notice

Table of Contents



## 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.

This draft defines a mechanism to communicate such details when DNS resolver filtering of a name is legally mandated, based upon the structured error data for DNS responses introduced by [I-D.ietf-dnsop-structured-dns-error].

Allowing DNS resolvers to inject user-visible messages brings unique challenges. Because DNS resolvers are often automatically configured by unknown networks and DNS responses are unauthenticated, 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.

Furthermore, lowering the barrier to the presentation of messages explaining why access has been denied by the DNS resolver risks encouraging the wider deployment of DNS-based censorship on the Internet.

This draft attempts to mitigate these risks by minimising the information carried in the DNS response to abstract, publicly registered identifiers -- the DNS Resolver Operator ID and the Filtering Incident ID. A consuming party (e.g., a Web browser) can selectively present messages from those operators that they believe to be using this mechanism for its stated goal -- in particular, those who using it to surface policy-driven filtering, rather than enact discretionary censorship or attack end users.

## 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 containing:

```
{
  "ro": "exampleResolver",
  "inc": "abc123"
}
```

This indicates that the "exampleResolver" resolver has generated the error, and the incident identifier is "abc123".

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

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

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

```
https://resolver.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, with a statement that conveys:

> The webpage at www.example.net was blocked due to a legal request. Your DNS resolver may have more information about the legal request here:

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

Note that there is no requirement for the template to expand to a URL on any particular hostname; for example, it could be hosted by a party other than the resolver's server.

### 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 for presentation to end users.

### 2.1. DNS Resolver Operator ID

A DNS Resolver Operator ID is a short, textual string that uniquely identifies the operator of a DNS resolver. It is carried in the EXTRA-TEXT field of the Extended DNS Error with the JSON field name "ro". For example:

```
{
  "ro": "exampleResolver"
}
```

Generators MUST only use values that are registered in the DNS Resolver Operator registry; see Section 4.2. Consumers MUST ignore unregistered values, and MAY ignore registered values.

### 2.2. Filtering Incident ID

A Filtering Incident ID is an opaque, string identifier for a particular filtering incident. It might be specific to a particular request, but need not be. It is carried in the EXTRA-TEXT field of the Extended DNS Error with the JSON field name "inc". For example:

```
{
  "inc": "abc123"
}
```

## 3. Incident Resolution Templates

An Incident Resolution Template is a URI Template [RFC6570] contained in the DNS Resolver Identifier 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 available to it:

ro: the DNS Resolver Operator ID (see Section 2.1)

inc: the Filtering Incident ID (see Section 2.2)

For example:

https://resolver.example.com/filtering-incidents/{inc}

Applications MUST store a local copy of the DNS Resolver Identifier Registry for purposes of template lookup; they MUST NOT query the IANA registry upon each use.

## 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: "ro"
Short Description: a short, textual string that uniquely identifies the operator of a DNS resolver
Mandatory: no
Specification: this document

JSON Name: "inc"
Short Description: an opaque, string identifier for a particular filtering incident
Mandatory: no
Specification: this document

### 4.2. The DNS Resolver Identifier Registry

IANA will establish a new registry, the "DNS Resolver Identifier 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 resolver operator

Contact: an e-mail address or other appropriate contact mechanism

DNS Resolver Operator ID: see Section 2.1

Incident Resolution Template: see Section 3

The Incident 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 Resolver 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 operators' error messages they display to users. This decision might be influenced by the identity of the resolver (e.g., so-called "public resolvers" are likely to use this mechanism responsibly), its history (e.g., a well-known Internet Service Provider that has been subject to legal filtering orders), or local configuration (e.g., application or operating system settings that indicate that a particular resolver is to be trusted).

## 6. 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-10, 26 November 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-structured-dns-error-10>.

[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>.

## Appendix A. Acknowledgements

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

## Author’s Address

Mark Nottingham
Cloudflare
Prahran
Australia
Email: mnot@mnot.net
URI: https://www.mnot.net/