Network Working Group
Internet-Draft
Intended status: Experimental
Expires: February 1, 2014

E. Hunt
ISC

July 31, 2013

# The DNS Extended Server Diagnostics (ESD) Option
## draft-hunt-dns-server-diagnostics-00

## Abstract

The widespread adoption of DNSSEC implies more frequent DNSSEC failures. Unfortunately, DNSSEC's failure mode is largely opaque to the client: when validation fails, the only signal that the clients of a validating resolver receive is an empty response with a SERVFAIL response code. This note proposes a protocol extension to allow SERVFAIL responses to include additional diagnostic information, giving the client greater insight into what went wrong and a better chance of delivering useful problem reports to DNS operators.

## 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 http://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 February 1, 2014.

## Copyright Notice

Table of Contents

## 1. Introduction

DNSSEC's failure mode is largely opaque to the client: when validation fails, the only signal of this that a resolver's clients receive is a SERVFAIL response code.

With no information provided to explain the exact cause of a SERVFAIL response, there is no straightforward way for an end user to determine whether a failure occurred due to a broken local resolver, a zone misconfiguration such as expired signatures, or a spoofing attack. This makes it difficult to address complaints and problem reports to the right people, slowing the correction of DNSSEC errors while increasing the support burden on validating resolver operators.

This note proposes a protocol extension allowing a name server, upon request from a client, to accompany SERVFAIL responses with detailed diagnostic information indicating what specifically caused the failure. In the typical use case for this mechanism, a validating caching name server would be able to convey specific failure information to a non-validating stub resolver or other client.

### 1.1. Reserved Words

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].

## 2. Protocol

A DNS client such as a non-validating stub resolver may use an EDNS(0) [RFC2671] option, ESD, to solicit extended diagnostic information from a DNS server. The ESD option payload includes a 16-bit "flags" field and a 16-bit diagnostic code; additional payload data may be included in the response.

One bit in the "flags" field is defined as "human-readable": if this bit is set in the solicitation, it indicates a desire for the server to return human-readable text, in addition to machine-readable diagnostic data; this text can be displayed or sent to a logging facility such as syslog [RFC5424]. All other payload data MUST be left empty in the solicitation.

The response payload, in addition to the flags field and the diagnostic code, may also include a supplemental data field whose content and schema are dependent on the diagnostic code being returned. If the "human-readable" flag is set in the response, then the response also includes human-readable text in a counted string,





i.e., a single length octet followed by that number of characters, as in the TXT RDATA format [RFC1035].

## 2.1. Client Behavior

A stub resolver or other DNS client requests extended diagnostic data by sending an ESD option (Section 2.3) in an EDNS(0) OPT pseudo-RR in the query message. The requestor MAY set the "human-readable" bit in the flags field of the request payload. The requestor MUST NOT include any other payload data in the ESD option.

The requestor MUST ignore any ESD option included in a response that does not have response code SERVFAIL.

## 2.2. Server Behavior

A resolver or other name server which encounters a server failure while answering a query that included an ESD option MAY add an ESD option to the SERVFAIL response. If the query did not include an ESD option, then the response MUST NOT include one. The server MUST NOT include an ESD option in any non-SERVFAIL response.

## 2.3. The ESD Option

The OPTION-CODE for the ESD option is [TBD].

The format for the OPTION-DATA portion of an ESD option is shown below:

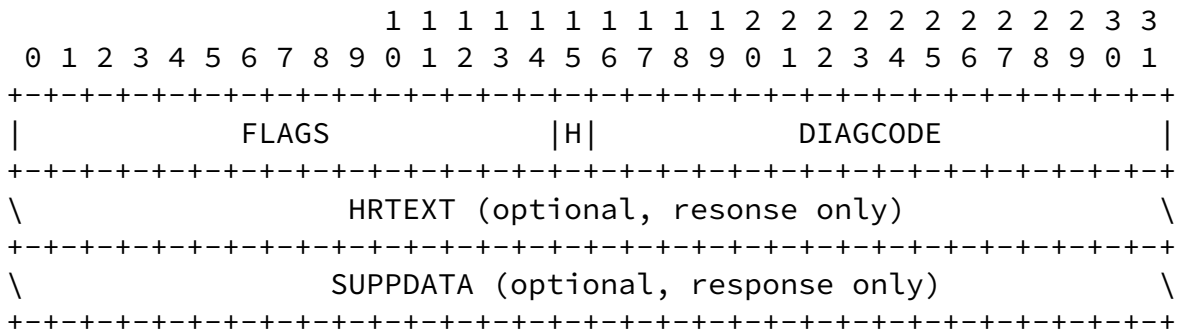


in which the fields are defined as follows:

FLAGS A 16-bit field. Bit 15 (H) is defined to mean "human-readable"; when set in the solicitation, this bit signals a desire for human-readable text to be included in the response; when set in the response, it signals that such text has been included. All other bits in the field MUST be set to zero.





DIAGCODE A 16-bit diagnostic code (Section 2.4) indicating the reason for the SERVFAIL response. In the solicitation payload, this field MUST be set to zero.

HRTEXT An counted string, containing human-readable text suitable for logging. The first octet defines the length of the following text; if the first octet contains 0, the string is emty. This is intended as an opaque string to be passed through to a logging mechanism; content and encoding are outside the scope of the protocol. This field MUST NOT be included in a solicitation payload.

SUPPDATA Optional supplementary data about the cause of the server failure. The presence or absence, content, and schema of this field are entirely dependent on the diagnostic code being returned in the DIAGCODE field (Section 2.4). This field MUST NOT be included in a solicitation payload.

## 2.4. ESD Diagnostic Codes

Diagnostic codes are grouped in five general categories: Internal server error conditions (diagnostic codes 1-255), general DNS

failures (256-511), policy-dependent security errors (512-767), temporary security errors (768-1023), and fatal security errors (1024-1279). Space has been reserved in the namespace for additional categories and codes. All diagnostic codes are optional; there is no requiremnt to impelement all of them.

The DIAGCODE field MUST be set to zero (No Error) in ESD solicitations.

### 2.4.1. Internal Server Errors

These diagnostic codes indicate a system failure in the responding server.

#### 2.4.1.1. Internal Error, Not Otherwise Specified

Diagnostic code 1 indicates an unspecified internal server error unrelated to DNSSEC. A server MAY return this code in place of any other internal server error, at the discretion of the implementor or operator, if information about the internal state of the server is regarded as security sensitive. This code has no supplemental data.

#### 2.4.1.2. Out of Memory

Diagnostic code 2 indicates that the server was not able to dynamically allocate memory. This code has no supplemental data.





#### 2.4.1.3. Resource Unavailable

Diagnostic code 3 indicates that a needed resource was not available. This code has no supplemental data.

#### 2.4.1.4. Zone Not Loaded

Diagnostic code 4: The server has been configured to be authoritative for a zone which is an ancestor of the query name, but was unable to load it. The supplemental data contains the name of the zone the server was unable to load, in DNS wire format.

#### 2.4.1.5. Invalid Zone

Diagnostic code 5: The server has been configured to be authoritative

for a zone which is an ancestor of the query name, but the zone contents are invalid; for example, there is no SOA RR or NS RRset at the zone apex. The supplemental data contains the name of the zone in DNS wire format.

#### 2.4.1.6. Configuration Error

Diagnostic code 6: The server could not be initialized, e.g., as a result of an error in loading or parsing the configuration file. This code has no supplemental data.

#### 2.4.1.7. Timeout

Diagnostic code 7: Query processing timed out. This code has no supplemental data.

#### 2.4.1.8. Shutting Down

Diagnostic code 8: The server is shutting down. This code has no supplemental data.

### 2.4.2. General DNS Errors

These codes describe failure conditions due to bad or inconsistent data in the DNS not specifically related to DNSSEC.

#### 2.4.2.1. Lame Delegation

Diagnostic code 256: The name servers which have been delegated responsibility for a zone cannot be reached or are not performing name service for that zone. The supplemental data contains the zone apex name of the faulty zone.





#### 2.4.2.2. Name Expansion Failure

Diagnostic code 257: A CNAME [RFC1034] or DNAME [RFC6672] record fails sanity checks after name expansion. The supplemental data contains the name and RR type (CNAME or DNAME) of the faulty record, in the following format:

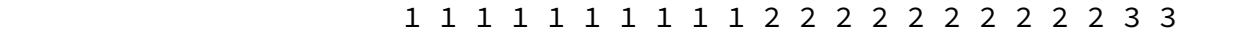

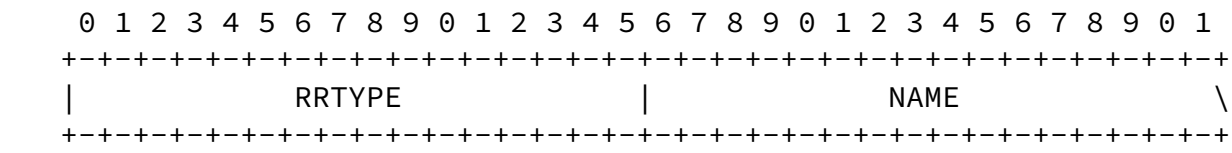


### 2.4.2.3. Protocol Not Supported

Diagnostic code 258: Processing this query requires a protocol extension that is not supported. This code has no supplemental data.

## 2.4.3. Policy-Dependent Security Errors

These are errors returned due to locally-configured policy constraints on DNSSEC processing or other security considerations.

### 2.4.3.1. Key Too Large

Diagnostic code 512: Local policy disallows a DNSKEY being used to validate a record on the grounds that it is too large. The supplemental data specifies the owner name (in DNS wire format) and key tag of the problematic DNSKEY, using the following format:

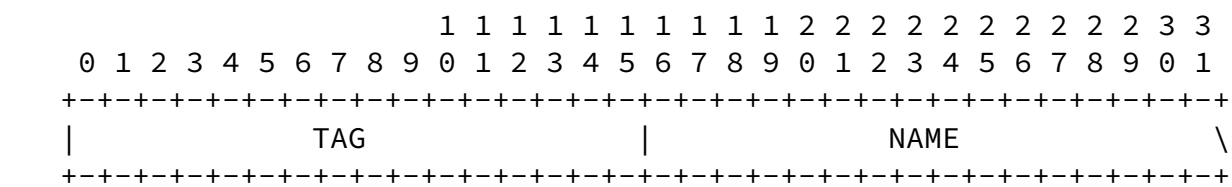


### 2.4.3.2. Key Too Small

Diagnostic code 513: Local policy disallows a DNSKEY being used to validate a record on the grounds that it is too small. The supplemental data contains the nam5 and tag of the problematic key, in the format specified in Section 2.4.3.1.

### 2.4.3.3. Key Not Authorized

Diagnostic code 514: Local policy does not authorize a key to be used for validation. The supplemental data contains the name and tag of the problematic key, in the format specified in Section 2.4.3.1.





### 2.4.3.4. Key Algorithm Refused

Diagnostic code 515: Local policy prohibits validation using a given signing algorithm. The supplemental data contains a 16-bit unsigned integer indicating which algorithm has been disallowed.

#### 2.4.3.5. Unauthorized Signer

Diagnostic code 516: Local policy disallows accepting signatures created by this signer. The supplemental data contains the name of the signer that has been disallowed.

#### 2.4.3.6. No Trust Anchor

Diagnostic code 517: There is no trust anchor for the chain of trust needed to validate this data, but local policy requires validation. This code has no supplemental data.

#### 2.4.3.7. Too Many Links

Diagnostic code 518: The chain of trust is longer than local policy permits. This code has no supplemental data.

#### 2.4.3.8. Response Blocked

Diagnostic code 519: The response to this query has been blocked by local policy. This code has no supplemental data.

### 2.4.4. Temporary Security Errors

These are error conditions occurring from DNSSEC processing which are time-dependent: i.e., problems due to signature validity interval or key expiry.

#### 2.4.4.1. Signature Expired

Diagnostic code 768: An RRSIG is past its useful lifetime. The supplemental data contains the name and covering RR type of the failed RRSIG, in the format specified in Section 2.4.2.2.

#### 2.4.4.2. Signature Not Yet Active

Diagnostic code 769: An RRSIG has not yet begun its useful lifetime. The supplemental data contains the name and covering RR type of the invalid RRSIG, in the format specified in Section 2.4.2.2.

#### 2.4.4.3. Trust Anchor Expired

Diagnostic code 770: A trust anchor can no longer be used. The supplemental data contains the name of the expired trust anchor in wire format.

### 2.4.5. Fatal Security Errors

These error conditions due to DNSSEC processing are always fatal, regardless of time or local policy.

#### 2.4.5.1. Bogus Data

Diagnostic code 1024: Cryptographic validation failed. The supplemental data contains the name and RR type of data which failed to validate, in the format specified in Section 2.4.2.2.

#### 2.4.5.2. Signature Missing

Diagnostic code 1025: There was no RRSIG found for an RRset, but one should have been present. The supplemental data contains the name and RR type of the data that should have been signed, in the format specified in Section 2.4.2.2.

#### 2.4.5.3. DNSKEY Missing

Diagnostic code 1026: No DNSKEY was found, but one should have been present. The supplemental data contains the zone apex name at which the DNSKEY should have been found, in wire format.

#### 2.4.5.4. Key Tag Mismatch

Diagnostic code 1027: RRSIG records have been found for an RRset which is to be validated, but none of them has a key tag matching a valid DNSKEY. The supplemental data contains the name and covering RR type for the faulty RRSIG, in the format specified in Section 2.4.2.2.

#### 2.4.5.5. DS Missing

Diagnostic code 1028: No DS record was found, but there should have been one present. The supplemental data contains the name at which the DS record should have been found.

#### 2.4.5.6. Next-Secure Record Missing

Diagnostic code 1029: No NSEC [RFC4034] or NSEC3 [RFC5155] record was found, but there should have been one present. The supplemental data





contains the name and RR type (NSEC or NSEC3) of the record that was expected, in the format specified in Section 2.4.2.2.

#### 2.4.5.7. Overreaching Next-Secure Record

Diagnostic code 1030: The "next owner" name in an NSEC or NSEC3 record goes beyond another record which is known to exist. The supplemental data contains the name and RR type (NSEC or NSEC3) of the invalid record, in the format specified in Section 2.4.2.2.

#### 2.4.5.8. Next-Secure Record Pointing Backward

Diagnostic code 1031: The ordering of records in the NSEC or NSEC3 chain does not follow canonical ordering rules. The supplemental data contains the name and RR type (NSEC or NSEC3) of the invalid record, in the format specified in Section 2.4.2.2.

#### 2.4.5.9. Irrelevant Proof

Diagnostic code 1032: A proof of non-existence was provided for a record whose non-existence did not need to be proven. This code has no supplemental data.

#### 2.4.5.10. Incomplete Proof

Diagnostic code 1033: Proof of non-existence is incomplete. The supplemental data contains the name and RR type of the data whose non-existence needed to be proven, in the format specified in Section 2.4.2.2.

#### 2.4.5.11. Wrong RRSIG Owner

Diagnostic code 1034: The RRSIG being used for verification is incorrect for the RR in question. The supplemental data contains the name and covering RR type of the invalid RRSIG, in the format specified in Section 2.4.2.2.

#### 2.4.5.12. Unknown DNSKEY Protocol

Diagnostic code 1035: The DNSKEY protocol value is not set to 3. The supplemental data contains the name and tag of the faulty key, in the format specified in Section 2.4.3.1.

#### 2.4.5.13. DS/DNSKEY Mismatch

Diagnostic code 1036: A mismatch was found between the DNSKEY in a zone and the DS record in the parent. The supplemental data contains the name and tag of the DNSKEY that should have been found, in the





format specified in Section 2.4.3.1.

#### 2.4.5.14. Unknown Algorithm

Diagnostic code 1037: An algorithm specified in a DNSKEY, DS, RRSIG, NSEC3 or NSEC3PARAM record is not recognized by the server. The supplemental data contains the name and RR type of the record that referenced the invalid algorithm.

#### 2.4.5.15. Algorithm Not Supported

Diagnostic code 1038: An algorithm specified in a DNSKEY, DS, RRSIG, NSEC3 or NSEC3PARAM record is recognized by the server but is not implemented. The supplemental data contains the name and RR type of the record that referenced the unsupported algorithm.

#### 2.4.5.16. Not a Zone Key

Diagnostic code 1039: The key that is used to verify signatures on zone data does not have the "Zone Key" flag [RFC4034] set. The supplemental data contains the name and tag of the faulty key, in the format specified in Section 2.4.3.1.

#### 2.4.5.17. Key Revoked

Diagnostic code 1040: A key that is required to complete a chain of trust has its REVOKED bit [RFC5011] set. The supplemental data contains the name and tag of the revoked key, in the format specified in Section 2.4.3.1.

## 3. Security Considerations

An ESD option response contains channel diagnostic data, to be used for logging, troubleshooting, and debugging; it falls outside the scope of DNSSEC per se. Ensuring integrity of the response requires the use of a channel security mechanism such as TSIG [RFC2845] or SIG(0) [RFC2931]. In the absence of any such mechanism, ESD responses MUST NOT be used by clients to make policy decisions with respect to DNSSEC validation, as this could allow an attacker to force a security downgrade or denial of service by forging SERVFAIL messages containing particular ESD responses.

Some of the data in an ESD option response may be security sensitive, particularly those responses which increase the transparency of the current state of a running resolver. In the case of SERVFAIL responses due to authoritative server misconfiguration or other non-local conditions, this is generally not a concern, as the data which





caused the failure are readily available in the DNS and can be obtained by other means. However, information about server failures due to local problems such as out-of-memory conditions may be of tactical use to an attacker. Implementors may wish to provide a mechanism for operators to disable certain types of diagnostic response while allowing others.

## 4. IANA Considerations

IANA is requested to make the assignments in this section.

### 4.1. ESD Option Code

This document requests the allocation of an EDNS(0) option code for the ESD option, whose value is [TBD].

### 4.2. Diagnostic Codes

This document requests the creation of a new registry of ESD diagnostic codes. The registry policy is "Specification Required" [RFC5226]. The initial entries in the registry are:

| Value | Description | Reference |
|---|---|---|

| | | |
|---|---|---|
| 0 | No Error | |
| 1 | Internal Error | [This] |
| 2 | Out of Memory | [This] |
| 3 | Resource Unavailable | [This] |
| 4 | Zone Not Loaded | [This] |
| 5 | Invalid Zone | [This] |
| 6 | Configuration Error | [This] |
| 7 | Timeout | [This] |
| 8 | Shutting Down | [This] |
| 9-255 | Unassigned | |
| 256 | Lame Delegation | [This] |
| 257 | Name Expansion Failure | [This] |
| 258 | Protocol Not Supported | [This] |
| 259-511 | Unassigned | |
| 512 | Key Too Large | [This] |
| 513 | Key Too Small | [This] |
| 514 | Key Not Authorized | [This] |
| 515 | Algorithm Refused | [This] |
| 516 | Unauthorized Signer | [This] |
| 517 | No Trust Anchor | [This] |
| 518 | Too Many Links | [This] |
| 519 | Response Blocked | [This] |





| | | |
|---|---|---|
| 520-767 | Unassigned | |
| 768 | Signature Expired | [This] |
| 769 | Signature Not Yet Active | [This] |
| 770 | Trust Anchor Expired | [This] |
| 771-1023 | Unassigned | |
| 1024 | Bogus Data | [This] |
| 1025 | Signature Missing | [This] |
| 1026 | DNSKEY Missing | [This] |
| 1027 | Key Tag Mismatch | [This] |
| 1028 | DS Missing | [This] |
| 1029 | Next-Secure Record Missing | [This] |
| 1030 | Overreaching Next-Secure Record | [This] |
| 1031 | Next-Secure Record Pointing Backward | [This] |
| 1032 | Irrelevant Proof | [This] |
| 1033 | Incomplete Proof | [This] |
| 1034 | Wrong RRSIG Owner | [This] |
| 1035 | Unknown DNSKEY Protocol | [This] |
| 1036 | DS/DNSKEY Mismatch | [This] |

| 1037 | Unknown Algorithm | [This] |
| --- | --- | --- |
| 1038 | Algorithm Not Supported | [This] |
| 1039 | Not a Zone Key | [This] |
| 1040 | Key Revoked | [This] |
| 1041-65535 | Unassigned | |

## 5. Acknowledgments

Thanks to Wes Hardaker, Jakob Schlyter, Sam Weiler and Francis Dupont for assistance in categorizing SERVFAIL error types, and Paul Vixie for design input.

## 6. References

### 6.1. Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.





[RFC2845] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.

[RFC2931] Eastlake, D., "DNS Request and Transaction Signatures ( SIG(0)s)", RFC 2931, September 2000.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", RFC 5011, September 2007.

[RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, March 2008.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", RFC 6672, June 2012.

## 6.2. Informative References

[RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.

## Author's Address

Evan Hunt
ISC
950 Charter St
Redwood City, CA 94063
USA

Email: each@isc.org