

INTERNET-DRAFT

Donald E. Eastlake, 3rd

IBM

Expires February 2000

August 1999

[draft-ietf-dnsind-kitchen-sink-01.txt](#)

The Kitchen Sink DNS Resource Record

--- -----

Donald E. Eastlake 3rd

Status of This Document

This draft, file name [draft-ietf-dnsind-kitchen-sink-01.txt](#), is intended to become a Proposed Standard RFC. Distribution of this document is unlimited. Comments should be sent to <namedroppers@internic.net> or to the author.

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories as listed at <<http://www.ietf.org/shadow.html>>.

Abstract

Periodically people desire to put proprietary, complex, and/or obscure data into the Domain Name System (DNS). This draft defines a

kitchen sink Resource Record that will satisfy this desire for the storage of miscellaneous structured information.

Acknowledgements

The suggestions or information provided by the following persons have improved this document and they are gratefully acknowledged:

Rob Austein
Johnny Eriksson
Phillip H. Griffin
Michael A. Patton
David Singer

Table of Contents

Status of This Document.....	1
Abstract.....	1
Acknowledgements.....	2
Table of Contents.....	2
1 . Introduction.....	3
2 . Kitchen Sink Resource Record.....	3
2.1 The Meaning Octet.....	4
2.2 The Coding and Subcoding Octets.....	5
2.2.1 ASN.1 Subcodings.....	7
2.2.2 MIME Subcodings.....	7
2.2.3 Text Subcodings.....	8
3 . Master File Representation.....	8
4 . Performance Considerations.....	8
5 . Security Considerations.....	9
6 . IANA Considerations.....	9
References.....	10
Author's Address.....	11
Expiration and File Name.....	11

1. Introduction

The Domain Name System (DNS) provides a replicated distributed secure hierarchical database which stores "resource records" (RRs) under hierarchical domain names. This data is structured into zones which are independently maintained. [RFC 1034, 1035, 2535]

Numerous types of RRs have been defined. These support such critical functions as host name to address translation (A, AAAA, etc. RRs), automatic mail routing (MX etc. RRs), and other functions. In addition, there are RRs defined related to the zone structure and administration of the DNS (SOA, NS, and RP RRs), security (SIG, KEY, and NXT RRs), etc. There is a TXT RR for the inclusion of general human readable text.

New RRs that are reasonably simple and designed via the open IETF standards process are periodically added as new needs become apparent. But there are periodically people who want to put some proprietary, complex and/or non-standard structured data in the DNS. In the past they have frequently come up with some way of reinterpreting the TXT RR, since that is one of the least constrained RRs. This is likely a bad idea since all previous ways to reinterpreting the TXT RR have sunk without a trace. (Well, if they actually got an RFC out, it's still there, but, practically speaking, almost nobody actually uses it.)

If a new type of data is needed for a global interoperable use in the DNS, the best course is to design a new RR that meets the need through the IETF standards process. This draft defines an extremely general and flexible RR which can be used for other data, such as proprietary data, where global interoperability is not a consideration. It includes representations of OSI ASN.1, MIME, XML, and, recursively, DNS RRs.

2. Kitchen Sink Resource Record

The symbol for the kitchen sink resource record is SINK. Its type number is 40. This type is defined across all DNS classes.

The RDATA portion of the SINK RR is structured as follows:


```

          1 1 1 1 1 1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           meaning           |           coding           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           subcoding         |                               /
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
/                               data                          /
/                               /                              /
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The "meaning", "coding", and "subcoding" octets are always present. The "data" portion is variable length and could be null in some cases. The size of the "data" portion can always be determined by subtracting 3 from the SINK resource record RDLENGTH. The coding octet gives the general structure of the data. The subcoding octet provides additional information depending on the value of the coding nibble.

All references to "domain name" in this document mean domain name in the IP DNS class.

Although it is unlikely, it is noted that multiple popular uses of SINK might develop that are not distinguished by using different parts of the DNS name space or different DNS classes. If this occurs, retrievals may fetch large sets of SINK RRs which will likely have to be sorted through at the application level. Should this occur, such popular uses of SINK should obtain and migrate to their own RR number using normal RR number allocation procedures or it would be possible to define an extended query operation that used a more fine grained selection method than just the RR type.

[2.1](#) The Meaning Octet

The meaning octet indicates whether any semantic labeling appears at the beginning of the data field and the format of such semantic labeling. This contrasts with the coding and subcoding octets which merely indicate format.

The types of labels available are chosen to be globally unique and under the control of some "owner". The owner designates the meaning associated with the labels they control. Where the label is a URI, it is recommended that a retrieval from the URI fetch material that would be helpful in determining this meaning. No a priori method is defined for determining the meaning of other labels beside an out of band question to the owner.

INITIAL ASSIGNED MEANING VALUES

0 - reserved.

1 - none.

2 - OID.

3 - domain name.

4 - URI.

5-254 - available for assignment, see [section 6](#).

255 - reserved.

A meaning octet value of 1 indicates that there is no semantic labeling at the beginning of the data area. The information, whatever it is, starts at the beginning of the data field and is coded according to the coding and subcoding octets.

Meaning octet values of 2, 3, or 4, indicate, on the other hand, that a semantic label is present. A value of two indicates that a BER [\[X.690\]](#) encoded OID appears prefixed by a single unsigned octet of OID length count. A value of three indicates that a DNS domain name appears in wire format with name compression prohibited. And a value of four indicates that a null (zero) octet terminated URI appears.

[2.2](#) The Coding and Subcoding Octets

The coding octet gives the major method by which the data in the data field is encoded. It should always have a meaningful value. The subcoding octet is intended to give additional coding details. Although the subcoding octet is always present, it must be interpreted in the context of the coding octet. For any coding octet value which does not specify subcoding octet value meanings, the subcoding octet MUST be ignored and SHOULD be zero.

While not explicitly mentioned below, the data field will actually start with a semantic label if indicated by the meaning octet. If such a semantic label is present, any data prefix required by the coding or subcoding octet is placed after the semantic label and before the data.

CODING OCTET VALUES

0 - reserved.

1 - DNS RRs. The data portion consists of DNS resource records as they would be transmitted in a DNS response section. The

subcoding octet is the number of RRs in the data area as an

unsigned integer. Domain names may be compressed via pointers as in DNS replies. The origin for the pointers is the beginning of the RDATA section of the SINK RR. Thus the SINK RR is safe to cache since only code that knows how to parse the data portion of a SINK RR need know of and can expand these compressions.

2 - MIME structured data [RFC 2045, 2046]. The data portion is a MIME structured message. The "MIME-Version:" header line may be omitted unless the version is other than "1.0". The top level Content-Transfer-Encoding may be encoded into the subcoding octet (see [section 2.2.2](#)). Note that, to some extent, the size limitations of DNS RRs may be overcome in the MIME case by using the "Content-Type: message/external-body" mechanism.

3 - Text tagged data. The data portion consists of text formatted as specified in the TXT RR except that the first and every subsequent odd numbered text item is considered to be a tag labeling the immediately following text item. If there are an odd number of text items overall, then the last is considered to label a null text item. Syntax of the tags is as specified in [RFC 2396](#) for the "Authority Component" without the two leading slashes ("//") or trailing slash using the DNS for authority. Thus any organization with a domain name can assign tags without fear of conflict. The subcodings octet specifies the encoding of the labeled text items as specified in [section 2.2.3](#).

4 - HTML. The subcoding octet indicates the version of HTML with the major version number in the upper nibble and the minor version number in the lower nibble. Thus, for example, HTML 3.2 would be indicated by a 0x32 octet.

5 - XML. The subcoding octet is the version of XML, currently 1.

6 - ASN.1 [[X.680](#), etc.]. See [section 2.2.1](#).

7-251 - Available for assignment, see [section 6](#).

252 - Private coding format indicated by an OID. The format of the data portion is indicated by an initial BER encoded OID which is prefixed by a one octet unsigned length count for the OID. The subcoding octet is available for whatever use the private formatting wishes to make of it.

253 - Private coding format indicated by a domain name. The format of the data portion is indicated by an initial wire format domain name with compression prohibited. (Such names are self delimiting.) The subcoding octet is available for whatever

use the private formatting wishes to make of it.

D. Eastlake 3rd

[Page 6]

254 - Private coding format indicated by a URI. The format of the data portion is indicated by an initial URI [[RFC 2396](#)] which is terminated by a zero (null) valued octet followed by the data with that format. The subcoding octet is available for whatever use the private formatting wishes to make of it. The manner in which the URI specifies the format is not defined but presumably the retriever will recognize the URI by some pattern match.

255 - reserved.

NOTE: the existence of a DNS RR coding and the infinite possibilities of ASN.1, XML, and MIME permit one to SINK to even greater depths by nesting.

[2.2.1](#) ASN.1 Subcodings

For ASN.1 [[X.680](#), etc.] data, a specific concrete encoding must be chosen as indicated by the subcoding octet.

ASN.* SUBCODINGS

- 0 - reserved.
- 1 - BER (Basic Encoding Rules [[X.690](#)]).
- 2 - DER (Distinguished Encoding Rules [[X.690](#)]).
- 3 - PER (Packed Encoding Rules) Aligned [[X.691](#)].
- 4 - PER Unaligned [[X.691](#)].
- 5 - CER (Canonical Encoding Rules [[X.690](#)]).
- 6-253 - available for assignment, see [section 6](#).
- 254 - private. This subcoding will never be assigned to a standard set of encoding rules. An OID preceded by a one octet unsigned length of OID appears at the beginning of the data area after the ASN coding OID.
- 255 - reserved.

[2.2.2](#) MIME Subcodings

If the coding octet indicates the data is MIME structured, the precise encoding is given by the subcoding octets as listed below.

MIME SUBCODINGS

- 0 - reserved, see [section 6](#).
- 1 - 7bit.
- 2 - 8bit.
- 3 - binary.

4 - quoted-printable.

D. Eastlake 3rd

[Page 7]

- 5 - base64.
- 6 - 253 - available for assignment, see [section 6](#).
- 254 - private. The data portion must start with an "x-" token denoting the private content-transfer-encoding immediately followed by one null (zero) octet followed by the remainder of the MIME object.
- 255 - reserved, see [section 6](#).

2.2.3 Text Subcodings

If the coding octet indicates the data is text, the exact encoding of the text items is indicated by the subcoding octet as follows:

TEXT SUBCODINGS

- 0 - reserved, see [section 6](#).
- 1 - ASCII.
- 2 - UTF-7 [[RFC 1642](#)].
- 3 - UTF-8 [[RFC 2044](#)].
- 4 - ASCII with MIME header escapes [[RFC 2047](#)].
- 5 - 253 - available for assignment, see [section 6](#).
- 254 - private. Each text item must start with a domain name [RFC 1034] denoting the private text encoding immediately followed by one null (zero) octet followed by the remainder of the text item.
- 255 - reserved, see [section 6](#).

3. Master File Representation

SINK resource records may appear as lines in zone master files. The meaning, coding, and subcoding appear as unsigned decimal integers. The data portion can be quite long. It is represented in base 64 [[RFC 2045](#)] and may be divided up into any number of white space separated substrings, down to single base 64 digits, which are concatenated to obtain the full data. These substrings can span lines using the standard parenthesis notation. (This type of base64 master file data is also required to support the DNS KEY and SIG security RRs [[RFC 2535](#)].)

4. Performance Considerations

Currently DNS is optimized for small data transfers, generally not exceeding 512 octets including overhead. Larger transfers are less

efficient but do work correctly and efforts are underway to make them

D. Eastlake 3rd

[Page 8]

more efficient.

It is easy to create very large RRs or RR sets using SINK. DNS administrators should think about this and may wish to discourage large RRs or RR sets. Consideration should also be given to putting zones from which large RRs or RR sets will be commonly retrieved on separate hosts which can be tuned for the load this will represent.

5. Security Considerations

Since the SINK resource record can be used to store arbitrary data in the DNS, this data could have security consequences, particularly if it is control, executable, macro, or interpretable information or very large and might cause buffer overflow. Due care should be taken.

[RFC 2535] covers data original authentication of the data in the domain name system including SINK RRs.

6. IANA Considerations

Assignment of specific meaning to the values listed herein as "reserved" requires an IETF standards action.

All other assignments of available meaning, coding, or subcoding octet values are by IETF consensus.

The many provisions for private indicita specified by separately allocated OIDs, domain names, or URIs should cover most requirements for private or proprietary values.

References

- [RFC 1034] - P. Mockapetris, "Domain names - concepts and facilities", 11/01/1987.
- [RFC 1035] - P. Mockapetris, "Domain names - implementation and specification", 11/01/1987.
- [RFC 1642] - D. Goldsmith, M. Davis, "UTF-7 - A Mail-Safe Transformation Format of Unicode", 07/13/1994.
- [RFC 2044] - F. Yergeau, "UTF-8, a transformation format of Unicode and ISO 10646", 10/30/1996.
- [RFC 2045] - N. Freed, N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", 12/02/1996.
- [RFC 2046] - N. Freed, N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", 12/02/1996.
- [RFC 2047] - K. Moore, "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text", 12/02/1996.
- [RFC 2396] - T. Berners-Lee, R. Fielding, L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", August 1998.
- [RFC 2535] - D. Eastlake, "Domain Name System Security Extensions", March 1999.
- [X.680] - ITU-T Recommendation X.680 (1997) | ISO/IEC 8824-1:1998, Information Technology - Abstract Syntax Notation One (ASN.1): Specification of Basic Notation
- [X.681] - ITU-T Recommendation X.681 (1997) | ISO/IEC 8824-2:1998, Information Technology - Abstract Syntax Notation One (ASN.1): Information Object Specification
- [X.682] - ITU-T Recommendation X.682 (1997) | ISO/IEC 8824-3:1998, Information Technology - Abstract Syntax Notation One (ASN.1): Constraint Specification
- [X.683] - ITU-T Recommendation X.683 (1997) | ISO/IEC 8824-4:1998, Information Technology - Abstract Syntax Notation One (ASN.1): Parameterization of ASN.1 Specifications
- [X.690] - ITU-T Recommendation X.690 (1997) | ISO/IEC 8825-1:1998, Information Technology - ASN.1 Encoding Rules: Specification of Basic

Encoding Rules (BER), Canonical Encoding Rules (CER) and

D. Eastlake 3rd

[Page 10]

Distinguished Encoding Rules (DER)

[X.691] - ITU-T Recommendation X.691 (1997) | ISO/IEC 8825-2:1998,
Information Technology - ASN.1 Encoding Rules: Specification of
Packed Encoding Rules (PER)

Author's Address

Donald E. Eastlake 3rd
IBM
65 Shindegan Hill Road
Carmel, 10512 USA

Telephone: +1 914-276-2668 (h)
 +1 914-784-7913 (w)
FAX: +1 914-784-3833 (w)
EMail: dee3@us.ibm.com

Expiration and File Name

This draft expires February 2000.

Its file name is [draft-ietf-dnsind-kitchen-sink-01.txt](#).

