

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 3, 2011

N. Borenstein
Mimecast
M. Kucherawy
Cloudmark
June 1, 2011

Reputation Data Interchange using the DNS
draft-kucherawy-reputation-query-dns-00

Abstract

This document defines a mechanism to conduct queries for reputation information using the Domain Name System.

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 December 3, 2011.

Copyright Notice

Copyright (c) 2011 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Document Series	3
3.	Terminology and Definitions	3
3.1.	Keywords	3
3.2.	Other Definitions	3
4.	Description	4
4.1.	Query Format	4
4.2.	Reply Format	5
5.	IANA Considerations	5
6.	Security Considerations	6
6.1.	General	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	6
Appendix A.	Public Discussion	7
Authors' Addresses	7

1. Introduction

This memo defines a method to query a reputation data service for information about an entity, using the Domain Name System (DNS). It is part of a series defining the overall reputation query/response structure as well as the concept of reputation "vocabularies" for particular applications.

2. Document Series

This memo represents the media type registration, part of a series of documents that define the overall service and introduce the initial exemplary applications. The series is as follows:

1. RFCxxxx: A Model for Reputation Interchange
2. RFCxxxx+1: A Media Type for Reputation Information
3. RFCxxxx+2: Using UDP for Reputation Interchange
4. RFCxxxx+3: Using the DNS for Reputation Interchange (this memo)
5. RFCxxxx+4: Using HTTP/XML for Reputation Interchange
6. RFCxxxx+5: A Reputation Vocabulary for Email Identity Reputation
7. RFCxxxx+6: A Reputation Vocabulary for Email Property Reputation

3. Terminology and Definitions

This section defines terms used in the rest of the document.

3.1. Keywords

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 [[KEYWORDS](#)].

3.2. Other Definitions

Other terms of importance in this memo are defined in RFCxxxx, the base memo in this document series.

4. Description

The [DNS] provides a distributed, fault-tolerant, extensible database generally used for retrieving information about services and hosts on the Internet. In the recent past its ability to store arbitrary text data to support various applications has been exploited to store such information as [DKIM] keys, expressions of policy such as [ADSP] and [SPF], or indications of group membership such as [VBR]. This memo defines another such application.

In line with [DNS-EXPAND], the TXT resource record type is used for this application.

4.1. Query Format

When constructing the name to be queried, the following steps are followed:

1. Present the subject of the reputation query, formed per the particular reputation application's rules, to the [SHA1] algorithm, producing a 20-byte blob of binary output.
2. Convert the binary output to a printable ASCII string by expressing each byte, in order, as a two-digit hexadecimal string. Output this string.
3. Append an ASCII period (0x2E).
4. Append either the name of the assertion of interest, defined by the particular reputation application's rules, or the string "_any" (ASCII 0x5F, 0x61, 0x6E, 0x79) if all available assertions are being requested.
5. Append an ASCII period (0x2E).
6. Append the name of the reputation application within which a query is being made. This name **MUST** be one registered with IANA.
7. Append an ASCII period (0x2E).
8. Append the string "_rep" (ASCII 0x5F, 0x72, 0x65, 0x70).
9. Append an ASCII period (0x2E).
10. Append the domain name that constitutes the root of the DNS sub-tree at which the reputation data are available. This is the "base" of the reputation service.

For example, suppose a client wishes to ask for any information the reputation service at "example.com" has about "example.net" within the context of the "email-id" application. A hex-converted SHA1 hash of "example.net" is the string

"c15fd3911e2d2a6ed98d884447782ad67fdb939". The query would be:

c15fd3911e2d2a6ed98d884447782ad67fdb939._any.email._rep.example.com

The hash is done to allow arbitrarily long subjects to be encoded into the name of a DNS query.

4.2. Reply Format

The reply is formatted as one or more TXT resource records. Replies not of type TXT MUST be ignored.

The client MUST decode the TXT reply by concatenating all character-string (see Section 3.3 of [\[DNS\]](#) payloads (i.e., drop all length bytes) into a single composite string. The resultant string is expected to be of the following form, expressed in [\[ABNF\]](#):

rep-result := rep-assertion SP rep-value SP rep-data *rep-extension

rep-assertion := token

rep-extension := SP token ":" token

rep-value := ("0" / "1") ["." 1*4DIGIT]
; MUST be between 0 and 1 inclusive

rep-data := 1*20DIGIT

"token" is imported from [\[MIME\]](#).

When the query was not about a specific assertion within the context of the reputation application, and thus "_any" was used, multiple TXT records MAY be returned, each indicating its own assertion.

Assertions and vocabulary extensions not registered as part of the reputation application in use MUST be ignored.

5. IANA Considerations

This memo presents no actions for IANA.

6. Security Considerations

This memo describes security considerations introduced by the media type defined here.

6.1. General

This memo is part of a series introducing a reputation query and response system (see [Section 2](#)). The Security Considerations sections of the other memos should also be consulted.

7. References

7.1. Normative References

- [ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [DNS] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [SHA1] U.S. Department of Commerce, "Secure Hash Standard", FIPS PUB 180-2, August 2002.

7.2. Informative References

- [ADSP] Allman, E., Fenton, J., Delany, M., and J. Levine, "DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP)", [RFC 5617](#), August 2009.
- [DKIM] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", [RFC 4871](#), May 2007.
- [DNS-EXPAND] Falstrom, P., Ed., Austein, R., Ed., and P. Koch, Ed., "Design Choices When Expanding the DNS", [RFC 5507](#), April 2009.
- [MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", [RFC 2045](#), November 1996.

- [SPF] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", [RFC 4408](#), April 2006.
- [VBR] Hoffman, P., Levine, J., and A. Hathcock, "Vouch By Reference", [RFC 5518](#), April 2009.

[Appendix A.](#) Public Discussion

Public discussion of this suite of memos takes place on the domainrep@ietf.org mailing list. See <https://www.ietf.org/mailman/listinfo/domainrep>.

Authors' Addresses

Nathaniel Borenstein
Mimecast
203 Crescent St., Suite 303
Waltham, MA 02453
USA

Phone: +1 781 996 5340
Email: nsb@guppylake.com

Murray S. Kucherawy
Cloudmark
128 King St., 2nd Floor
San Francisco, CA 94107
USA

Phone: +1 415 946 3800
Email: msk@cloudmark.com

