

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 3, 2012

N. Borenstein
Mimecast
M. Kucherawy
Cloudmark
October 31, 2011

A Media Type for Reputation Interchange
draft-kucherawy-reputation-media-type-02

Abstract

This document defines a media type for exchanging reputation information about an arbitrary class of object.

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 May 3, 2012.

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. Formal Definition	6
5. Scores	6
6. IANA Considerations	7
6.1. application/reputon Media Type Registration	7
6.2. Reputation Applications Registry	8
7. Security Considerations	9
7.1. General	9
8. References	10
8.1. Normative References	10
8.2. Informative References	10
Appendix A. Acknowledgments	10
Appendix B. Public Discussion	10
Authors' Addresses	10

1. Introduction

This memo defines a media type for use when answering a reputation query using the "long form" query defined in RFCxxxx+4, which uses [HTTP]. It is part of a series defining the overall reputation query/response structure as well as the concept of reputation "vocabularies" for particular applications.

Also included is the specification for an IANA registry to contain definitions and symbolic names for known reputation vocabularies.

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 (this memo)
3. RFCxxxx+2: Using UDP for Reputation Interchange
4. RFCxxxx+3: Using the DNS for Reputation Interchange
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

A new media type, "application/reputon", is defined for the representation of reputational data. This media type has two optional parameters: "app", which conveys the specific application of reputation data in use, and usually extends the set of data values that MAY be included in the media object itself; and "format", which specifies the format with which the content are relayed.

The default for "format" is "text", which is defined here. Reputons bearing unrecognized format values MUST be ignored.

The body of the media type consists of [MAIL]-style attribute/value pairs, six of which are standard for all apps:

RATER: The identity of the entity providing the reputation information, generally expressed as a DNS domain name.

RATER-AUTHENTICITY: The level of confidence in that identity being genuine, expressed as a floating-point number between 0 and 1 inclusive.

ASSERTION: A keyword indicating the specific assertion or claim being rated. In the absence of an "app" parameter, the reputon can only indicate generic goodness, with the default assertion "IS-GOOD," but each application is expected to define additional types of ASSERTION.

RATED: The identity of the entity being rated.

RATING: The overall rating score for that entity, expressed as a floating-point number between 0.0 and 1.0 inclusive. See Section 5 for discussion.

SAMPLE-SIZE: The number of data points used to compute that score, possibly an approximation. Expressed as an unsigned 64-bit integer.

A particular application that registers itself with IANA MAY also define extension attribute/value pairs beyond the six mandatory ones.

Thus, the following example:

Content-type: application/reputon

RATER: RatingsRUs.example.com
RATER-AUTHENTICITY: 1.0
ASSERTION: IS-GOOD
RATED: Alex Rodriguez
RATING: 0.99
SAMPLE-SIZE: 50000

...indicates that we are absolutely sure (1.0) that the entity "RatingsRUs.example.com" consolidated 50000 data points (perhaps from everyone in Yankee Stadium) and concluded that Alex Rodriguez is very good (0.99) at something. It doesn't tell us what he's good at, and while it might be playing baseball, it could just as well be paying his taxes on time.

A more sophisticated usage would define a baseball application with a vocabulary of specific assertions, so that this example:

Content-type: application/reputon; app="baseball"

RATER: baseball-reference.example.com
RATER-AUTHENTICITY: 1.0
ASSERTION: HITS-FOR-POWER
RATED: Alex Rodriguez
RATING: 0.99
SAMPLE-SIZE: 50000

...would indicate that 50000 fans polled by the entity baseball-reference.example.com rate A-Rod very highly in hitting for power, whereas this example:

Content-type: application/reputon; app="baseball"

RATER: baseball-reference.example.com
RATER-AUTHENTICITY: 1.0
ASSERTION: CLUTCH-HITTER
RATED: Alex Rodriguez
RATING: 0.4
SAMPLE-SIZE: 50000

...would indicate that a similar poll indicated a somewhat weaker consensus that A-Rod tends to choke in critical baseball situations.

In practice, most usage of reputons is expected to make use of the "app" parameter to target an application-specific set of assertions.

4.1. Formal Definition

More formally, using [ABNF], the content of the application/reputon MIME object MUST conform to the following syntax:

```
reputon := rater rater-auth assertion *extension
          rated rating sampsize

rater := "RATER:"
        *WSP (atom / quoted-string) [CFWS] CRLF

rater-auth := "RATER-AUTHENTICITY:"
              *WSP 1*DIGIT "." 1*DIGIT [CFWS] CRLF
              ; must be a number between -1 and 1 inclusive

assertion := "ASSERTION:"
             *WSP dot-atom-text [CFWS] CRLF

extension := dot-atom-text %x3A *WSP dot-atom-text [CFWS] CRLF
            ; must be registered with IANA within a reputation
            ; vocabulary registration

rated := "RATED:"
        *WSP (atom / quoted-string) [CFWS] CRLF

rating := "RATING:"
         *WSP 1*DIGIT "." 1*DIGIT [CFWS] CRLF
         ; must be a number between 0 and 1 inclusive

sampsize := "SAMPLE-SIZE:"
           *WSP 1*DIGIT [CFWS] CRLF
           ; must be an unsigned 64-bit integer

"atom", "quoted-string" and "dot-atom-text" are imported from [MAIL].
```

5. Scores

The score presented as the value in the RATING parameter appears as a floating point value between 0.0 and 1.0 inclusive. The intent is that the definition of an assertion within an application will declare what the anchor values 0.0 and 1.0 specifically mean. Generally speaking, 1.0 implies full agreement with the assertion, while 0.0 indicates no support for the assertion.

The definition will also specify the type of scale in use when generating scores, to which all reputation service providers for that application space must adhere. This will allow a client to change

which reputation service provider is being queried for a given without having to learn through some out-of-band method what the new provider's values mean. For example, a registration might state that ratings are linear, which would mean a score of "x" is twice as strong as a value of "x/2".

6. IANA Considerations

This memo presents two actions for IANA, namely the creation of the new media type "application/reputon" and the creation of a registry for reputation application types. Another memo in this series creates an initial registry entry for the latter.

6.1. application/reputon Media Type Registration

This section provides the media type registration application from [MIME-REG] for processing by IANA:

To: ietf-types@iana.org

Subject: Registration of media type application/reputon

Type name: application

Subtype name: reputon

Required parameters: none

Optional parameters:

app: Names the reputation application in use within the reputon, which defines the valid assertions and any extensions that may also be valid (i.e., the vocabulary) for that application. These MUST be registered with IANA.

format: Describes the format of the content of the MIME object. The default is "text" which is defined in this memo.

Encoding considerations: "7bit" encoding is sufficient and MUST be used to maintain readability when viewed by non-MIME mail readers.

Security considerations: See Section 7 of [this document].

Interoperability considerations: Implementers MUST ignore any "app" values, attribute/value pairs, or vocabulary items they do not support.

Published specification: [this document]

Applications that use this media type: Any application that wishes to query a service that provides reputation data using the "long form" defined in RFCxxxx. The example application is one that provides reputation expressions about DNS domain names found in email messages.

Additional information: The value of the "app" parameter MUST also be registered with IANA.

Person and email address to contact for further information:

Nathaniel Borenstein <nps@guppylake.com>

Murray S. Kucherawy <msk@cloudmark.com>

Intended usage: COMMON

Author:

Nathaniel Borenstein

Murray S. Kucherawy

Change controller: IESG

6.2. Reputation Applications Registry

IANA is requested to create the "Reputation Applications" registry. This registry will contain names of applications used with the application/reputon media type, as defined by this memo.

New registrations or updates MUST be published in accordance with the "Specification Required" guidelines as described in [IANA-CONSIDERATIONS].

New registrations and updates MUST contain the following information:

1. Name of the application being registered or updated
2. Short description of the application (i.e., the class of entity about which it reports reputation data)
3. The document in which the application is defined
4. New or updated status, which MUST be one of:

current: The application is in current use

deprecated: The application is in current use but its use is discouraged

historic: The application is no longer in current use

5. An optional table of query parameters that are specific to this application; each table entry must include:

Name: Name of the query parameter

Status: (as above)

Description: A short description of the purpose of this parameter

Syntax: A reference to a description of valid syntax for the parameter's value

Required: "yes" if the parameter is mandatory, "no" otherwise

A document creating a reputation application MUST include:

1. A list of one or more assertions registered within this application; each table entry must include:

Name: Name of the assertion

Description: A short description of the assertion, with specific meanings for values of 0.0 and 1.0

Scale: A short description of the scale used in computing the value (see Section 5 of this memo)

7. Security Considerations

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

7.1. 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.

As this message structure is designed for use with the "long form" of

the reputation query, the Security Considerations of RFCxxxx+3 would be of particular interest to implementers.

8. References

8.1. Normative References

[ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 5234, January 2008.

8.2. Informative References

[HTTP] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

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

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

[MAIL] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.

[MIME-REG] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", RFC 4288, December 2005.

Appendix A. Acknowledgments

The authors wish to acknowledge the contributions of the following to this specification: Frank Ellermann.

Appendix B. 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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 23, 2012

N. Borenstein
Mimecast
M. Kucherawy
Cloudmark
October 21, 2011

Reputation Data Interchange using HTTP and XML
draft-kucherawy-reputation-query-http-03

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 April 23, 2012.

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	4
4.2. Response	5
4.2.1. XML Schema	5
4.2.2. Example Reply	7
5. IANA Considerations	8
6. Security Considerations	8
6.1. General	8
7. References	8
7.1. Normative References	8
7.2. Informative References	9
Appendix A. Acknowledgements	9
Appendix B. Public Discussion	9
Authors' Addresses	9

1. Introduction

This memo defines a method to query a reputation data service for information about an entity, using the HyperText Transfer Protocol (HTTP) as the transport mechanism and XML as the payload format. 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
5. RFCxxxx+4: Using HTTP/XML for Reputation Interchange (this memo)
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

4.1. Query

A reputation query made via [HTTP] encodes the question being asked partly in the [URI] and partly within the GET instruction of the protocol.

The components to the question being asked comprise the following:

- o The subject of the query;
- o The name of the host, or the IP address, at which the reputation service is available;
- o The name of the reputation application, i.e., the context within which the query is being made;
- o Optionally, name(s) of the specific reputation assertions or attributes that are being requested.

The name of the application **MUST** be one registered with IANA. A server receiving a query about an unregistered application or one it does not explicitly support **MUST** return a 404 error code.

The syntax for the URI portion of the query is constructed using a template as per [URI-TEMPLATE]. The following variables **MUST** be available during template expansion:

application: The name of the application reputation in whose context the request is being made.

scheme: The transport scheme the client will be using for the query.

service: The hostname or IP address being queried.

Which scheme(s) can be used depends on how the reputation service provider offers its services. Thus, the template could include a specific schema as a fixed string in the template, or it might offer it as a variable in the template. If it is a variable, it is up to the client and server to negotiate out-of-band which schemes are supported for client queries. Implementers should be aware that the template could include a fixed scheme not supported by the client.

The following variables are **OPTIONAL**, but might be required by the template presented for a specific service:

assertion: A list of one or more specific assertions of interest to the client. If absent, the server MUST infer that all available assertion information is being requested.

passwd: The "password" portion of a client credential.

user: The "user" portion of a client credential.

Other required or optional query parameters might be defined by documents that register new vocabularies with IANA.

The template is retrieved by requesting the [WELL-KNOWN-URI] "repute_template" from the host providing reputation service using HTTP. If the template cannot be retrieved, the query should be aborted and/or retried at a later time.

For example, given the following template:

```
{scheme}://{service}/{application}/{subject}/{assertion}
```

A query about the use of the domain "example.org" in the "email-id" application context to a service run at "example.com", where that application declares a required "subject" parameter, requesting the "SENDS-SPAM" reputation assertion using HTTP to conduct the query with no specific client authentication information would be formed as follows:

```
http://example.com/email-id/example.org/sends-spam
```

Matching of the attribute name(s) MUST be case-insensitive.

4.2. Response

The response is expected to be an XML document. The "format" parameter of the "application/reputon" media type MUST be "xml" when used in this mode.

The XML schema definition describing the format of that response is included below.

4.2.1. XML Schema

The following XML schema describes the format of the reply:

```
<?xml version="1.0" encoding="ISO-8859-1" ?%gt;
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">

  <!-- definition of local types -->
  <xs:simpleType name="exttype">
    <xs:restriction base="xs:token">
      <xs:pattern value="\w+(-\w*)*:\s?[\w\p{P}]+"/>
    </xs:restriction>
  </xs:simpleType>

  <!-- definition of simple elements -->
  <xs:element name="rater" type="xs:token"/>
  <xs:element name="rater-authenticity" type="xs:decimal"/>
  <xs:element name="assertion" type="xs:token"/>
  <xs:element name="extension" type="exttype"/>
  <xs:element name="rated" type="xs:token"/>
  <xs:element name="rating" type="xs:decimal"/>
  <xs:element name="sample-size" type="xs:positiveInteger"/>
  <xs:element name="updated" type="xs:positiveInteger"/>

  <!-- definition of complex elements -->
  <xs:complexType name="assertiontype">
    <xs:sequence>
      <xs:element ref="rater" minOccurs="1"/>
      <xs:element ref="rater-authenticity" minOccurs="1"/>
      <xs:element ref="assertion" minOccurs="1"/>
      <xs:element ref="extension"/>
      <xs:element ref="rated" minOccurs="1"/>
      <xs:element ref="rating" minOccurs="1"/>
      <xs:element ref="sample-size" minOccurs="1"/>
      <xs:element ref="updated" minOccurs="1"/>
    </xs:sequence>
  </xs:complexType>

  <xs:complexType name="reporttype">
    <xs:sequence>
      <xs:element name="reputon" type="assertiontype"
        maxOccurs="unbounded" minOccurs="1"/>
    </xs:sequence>
  </xs:complexType>

  <xs:element name="reputation" type="reporttype"/>
</xs:schema>
```

The elements that comprise an "assertion" are used as follows:

rater: The identity of the agent making the assertion.

rater-authenticity: An expression by the rater of its confidence in the report it is giving. Expressed as a decimal value between 0 and 1 inclusive.

assertion: The assertion being made. This MUST be an assertion registered within the specified application by IANA.

extension: (OPTIONAL) One or more application-specific vocabulary extensions and their corresponding values. If present, each of these MUST be a vocabulary extension registered with IANA.

rated: The identity about which an assertion is being made.

rating: The value of the assertion. This is a decimal number from 0 to 1, with 0 meaning the assertion is completely false (according to the agent making the assertion) and 1 meaning the assertion is completely true.

sample-size: The count of data points the asserting agent used to produce the value provided in the previous element.

updated: The time at which the current rating was computed.
Expressed in number of seconds since 00:00:00 UTC, January 1, 1970.

4.2.2. Example Reply

The following is an example reputon generated using the above schema, including the media type definition line:

Content-Type: application/reputon; app="email"; format="xml"

```
<?xml version="1.0" encoding="US-ASCII"?>
```

```
<reputation>
  <reputon>
    <rater>rep.example.net</rater>
    <rater-authenticity>0.95</rater-authenticity>
    <assertion>SENDS-SPAM</assertion>
    <extension>IDENTITY: DKIM</extension>
    <rated>example.com</rated>
    <rating>0.0012</rating>
    <sample-size>16938213</sample-size>
    <updated>1317795852</updated>
  </reputon>
</reputation>
```

Here, reputation agent "rep.example.net" is asserting within the context of email that "example.com" appears to send spam 1.2% of the time, based on just short of 17 million messages analyzed or reported to date. The identity "example.com", the subject of the query, is extracted from the analyzed messages using the [DKIM] "d=" parameter for messages where signatures validate. The reputation agent is 95% confident of this result. (See [RFCxxxx+5] for details about the registered email vocabulary.)

5. IANA Considerations

This memo presents no actions for IANA. Registration of the well-known URI "repute_template" will be done as defined in [WELL-KNOWN-URI] which is not a function of 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

- [HTTP] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", RFC 3986, January 2005.
- [URI-TEMPLATE] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template",

I-D draft-gregorio-uritemplate, September 2011.

[WELL-KNOWN-URI]

Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, April 2010.

7.2. Informative References

[DKIM] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", RFC 6376, September 2011.

Appendix A. Acknowledgements

The authors would like to thank the following for their contributions to this work: Mark Nottingham.

Appendix B. 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

