

REPUTE Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 2, 2014

N. Borenstein
Mimecast
M. Kucherawy
August 29, 2013

A Reputation Query Protocol
draft-ietf-repute-query-http-10

Abstract

This document defines a mechanism to conduct queries for reputation information over the Hypertext Transfer Protocol using JSON as the payload meta-format.

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 March 2, 2014.

Copyright Notice

Copyright (c) 2013 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.	Terminology and Definitions	3
2.1.	Key Words	3
2.2.	Other Definitions	3
3.	Description	3
3.1.	Overview	3
3.2.	URI Template	4
3.3.	Syntax	5
3.4.	Response	6
3.5.	Protocol Support	6
4.	IANA Considerations	6
5.	Security Considerations	6
6.	References	7
6.1.	Normative References	7
6.2.	Informative References	8
Appendix A.	Acknowledgements	8
Appendix B.	Public Discussion	8
	Authors' Addresses	8

1. Introduction

This document 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 JSON as the payload meta-format.

The mechanism is a two-stage query:

1. A client retrieves a template from a server that describes the construction of a Universal Resource Identifier (URI) which will be the actual query;
2. The client then uses the constructed URI to request the reputation data from the server.

2. Terminology and Definitions

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

2.1. Key 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 [[KEYWORDS](#)].

2.2. Other Definitions

Other terms of importance in this document are defined in [[I-D.REPUTE-MODEL](#)] and [[I-D.REPUTE-MEDIA-TYPE](#)].

3. Description

3.1. Overview

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 subject is being evaluated;

- o Optionally, name(s) of the specific reputation assertions or attributes that are being requested.

Assertions are discussed in [[I-D.REPUTE-MEDIA-TYPE](#)].

The name of the application, if given, is expected to be one registered with IANA in the Reputation Applications Registry, which is defined in [[I-D.REPUTE-MEDIA-TYPE](#)]. A server receiving a query about an application it does not recognize or explicitly support (e.g., by virtue of private agreements or experimental extensions) MUST return a 404 error code.

A reputation query made via [[HTTP](#)] encodes the question being asked in an HTTP GET method. The specific syntax of the query itself is specified by retrieving a URI template from the reputation service, completing the template, and then issuing the query.

[3.2.](#) URI Template

The template file is retrieved by requesting the [[WELL-KNOWN-URI](#)] "repute-template" from the host providing reputation service, using HTTP. (The registration for this well-known URI is in [Section 4](#).) The server returns the template file in a reply that MUST use the text/plain media type (see [[MIME](#)]), and SHOULD include an Expires field (see Section 14.21 of [[HTTP](#)]) indicating a duration for which the template is to be considered valid by clients and not re-queried.

Clients SHOULD NOT repeat the query prior to the timestamp in the Expires field, or wait no less than one day if the Expires field is not present.

The template file might contain more than one template. Such a file MUST have each template separated by a carriage return (ASCII 0x0D) and newline (ASCII 0x0A) character, as is typical for most text-based Internet protocols.

Each template in the file is expanded using the variables that are the parameters to the query. These parameters are either the subject about which reputation information is sought (or details associated with it), or other parameters that are established out-of-band with the reputation service; they are not established by any automated discovery described here. The client then attempts to query each expanded template that uses a scheme it is capable of querying, in the order presented in the file, until the client finds one to which it can establish a usable connection and issue the query.

For example, given the following template:


```
http://{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 "SPAM" reputation assertion, would be formed as follows:

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

3.3. Syntax

The syntax for the [\[URI\]](#) of the query is constructed using a template as per [\[URI-TEMPLATE\]](#). (See [Section 3.2.](#)) Clients MUST provide the following values in the expansion of the template:

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

service: The hostname or IP address to which the query is being sent. This MUST be the same as the host to which the template query was issued.

subject: The subject of the query, extracted from some content to be evaluated.

The following variable can also be provided. It is not mandatory in this model, but a specific application (defined in its own extension document) might declare it mandatory in a specific context:

assertion: The name of the specific assertion of interest to the client. If absent, the client is indicating that it requests all available assertion information.

Every application space has a set of assertions applicable to its own context. [\[I-D.REPUTE-MEDIA-TYPE\]](#) defines a single assertion assumed to exist in any application that does not define its own assertion set.

Reputation applications can extend the set of optional or required query parameters as part of their IANA registration actions. The set enumerated above establishes the base set common to all of them. Further, additional required or optional extension query parameters might be defined by specific reputation service providers, though these are private arrangements between client and server and will not be registered with IANA.

Authentication between reputation client and server is outside the scope of this specification. It could be provided through a variety

of available transport-based or object-based mechanisms, including a later extension of this specification.

3.4. Response

The response is expected to be contained in a media type designed to deliver reputons. An media type designed for this purpose, "application/reputon+json", is defined in [[I-D.REPUTE-MEDIA-TYPE](#)].

3.5. Protocol Support

A client has to implement HTTP in order to retrieve the query template as described in [Section 3.2](#). Accordingly, a server can assume the client will be able to handle a URI template that produces a URI for the query using the "http" scheme. The template could yield a query string that uses some other URI scheme, in which case the client could try that URI as well if it supports issuing queries with that scheme.

A server SHOULD include support for providing service over HTTP, and publish templates indicating support for this, as a baseline for interoperability with arbitrary clients.

4. IANA Considerations

This document registers the "repute-template" well-known URI in the Well-Known URI registry as defined by [[WELL-KNOWN-URI](#)], as follows:

URI suffix: `repute-template`

Change controller: `IETF`

Specification document(s): `[this document]`

Related information: `none`

5. Security Considerations

This document defines particular uses of existing protocols for a specific application. In particular, the basic protocol used for this service to retrieve a URI template from a well-known location is basic HTTP, which is not secure without certain extensions. Security issues relevant to use of URI templates are discussed in [[URI-TEMPLATE](#)], and those relevant to well-known URI definitions and retrieval are discussed in [[WELL-KNOWN-URI](#)].

The reputation service itself will use HTTP or other transport methods to issue queries and receive replies. Those protocols have registered URI schemes and, as such, presumably have documented security considerations. The protocol described here operates atop those schemes, and does not itself present new security considerations.

Reputation mechanisms represent an obvious security concern, in terms of the validity and use of the reputation information. These issues are beyond the scope of this specification. General information pertaining to using or providing reputation services can be found in [\[I-D.REPUTE-CONSIDERATIONS\]](#).

6. References

6.1. 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.
- [I-D.REPUTE-MEDIA-TYPE] Borenstein, N. and M. Kucherawy, "A Media Type for Reputation Interchange", [draft-ietf-repute-media-type](#) (work in progress), November 2012.
- [I-D.REPUTE-MODEL] Borenstein, N. and M. Kucherawy, "A Model for Reputation Interchange", [draft-ietf-repute-model](#) (work in progress), November 2012.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", [RFC 2045](#), November 1996.
- [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", [RFC 6570](#), March 2012.

[WELL-KNOWN-URI]

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

[6.2.](#) Informative References**[I-D.REPUTE-CONSIDERATIONS]**

Kucherawy, M., "Operational Considerations Regarding Reputation Services", [draft-ietf-repute-considerations](#) (work in progress), November 2012.

[Appendix A.](#) Acknowledgements

The authors would like to thank the following for their contributions to this work: Simon Hunt, Mark Nottingham, David F. Skoll, and Mykyta Yevstifeyev.

[Appendix B.](#) Public Discussion

[RFC Editor: Please delete before publication]

Public discussion of this set of documents 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
270 Upland Drive
San Francisco, CA 94127
USA

Email: superuser@gmail.com

