REPUTE Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 6, 2013

N. Borenstein
Mimecast
M. Kucherawy

May 5, 2013

A Reputation Query Protocol
draft-ietf-repute-query-http-05

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 November 6, 2013.

Copyright Notice

Table of Contents

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.

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

A reputation query made via [HTTP] encodes the question being asked in an HTTP GET method.

The components to the question being asked comprise the following:

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

The name of the application, if given, MUST be one registered with IANA in the Reputation Applications Registry. A server receiving a query about an unregistered application or one it does not explicitly support MUST return a 404 error code.

3.2. Syntax

The syntax for the [URI] portion of the query is constructed using a template as per [URI-TEMPLATE]. (See Section 3.3.) 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.

subject: The subject of the query.

Which scheme(s) can be used depends on how the reputation service provider offers its services. Thus, the template could include a specific scheme 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 need to be aware that the template could include a fixed scheme not supported by the client.

For example, the following query template includes a fixed scheme, forcing clients to use the "http" URI scheme only:

```
http://{service}/repute.php{?subject,application,assertion}
```

However, this template allows the client to select the scheme to be used if, for example, the service is also available over the "https" URI scheme:

```
{scheme}://{service}/repute.php{?subject,application,assertion}
```

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.

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.

Other required or optional query parameters might be defined by documents that register new response sets with IANA. Further, other

required or optional 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 specificatin. It could be provided through a variety of available transport-based or object-based mechanisms, including a later extension of this specification.

3.3. URI Template

The template 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 MUST return the template in a reply using 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.

If the template cannot be retrieved (i.e., any HTTP error is returned), the reputation query SHOULD be aborted and/or retried at a later time. Clients SHOULD adhere to the expiration time presented in an Expires field, if present, or otherwise assume that the template is valid for no less than one day and SHOULD NOT repeat the query.

The template is expanded, using the variables that are the parameters to the query, and then used as the target for the query itself. 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 "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/spam

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

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 3.3. 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. If the template can yield a query string that uses some other URI scheme, there will need to be some out-of-band negotiation of which scheme(s) are supported by the service, and appropriate protocol support in the client.

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 is basic HTTP which is not secure without certain extensions. As such, the protocol described here does not itself present new security considerations.

Security considerations relevant to email and email authentication can be found in most of the documents listed in the References sections below. Information specific to use of reputation services can be found in [I-D.REPUTE-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.

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-iet-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", draft-gregorio-uritemplate (work in progress), September 2011.

[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: Mark Nottingham, David F. Skoll, and Mykyta Yevstifeyev.

Appendix B. Public Discussion

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
2063 42nd Avenue
San Francisco, CA 94116
USA

Email: superuser@gmail.com