Reputation Data Interchange using HTTP and XML
draft-ietf-repute-query-http-01

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 July 16, 2012.

Table of Contents

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. Terminology and Definitions

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

2.1. 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].

2.2. Other Definitions

Other terms of importance in this memo are defined in
[I-D.REPUTE-MODEL].

3. Description

3.1. 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
   attributies 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. The
server SHOULD return the template in a text/plain reply. If the
template cannot be retrieved, the reputation query SHOULD be aborted
and/or retried at a later time. The server responding to the
template request SHOULD include an Expires field indicating a
duration for which the template should be considered valid by clients
and not re-queried. Clients SHOULD adhere to the expiration time
thus provided or, if none is provided, assume that the template is
valid for no less than one day and not repeat the query.

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.

3.2. Response

The response is expected to be an XML document, contained within a "application/reputon" media type object as defined in [I-D.REPUTE-MEDIA-TYPE]. The XML schema for the document is also specified in that memo.

4. IANA Considerations

This memo 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 memo]

Related information: none

5. Security Considerations

This memo describes security considerations introduced by the query mechanism defined here.

[TBD]

6. 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", I-D draft-ietf-repute-media-type,
          June 2011.

[I-D.REPUTE-MODEL]
          Borenstein, N. and M. Kucherawy, "A Model for Reputation
          Interchange", I-D draft-iet-repute-model, June 2011.

[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.

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 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