               Requirements For Internet Registry Services
                  draft-kucherawy-weirds-requirements-03

Abstract

   This document enumerates a base set of requirements to be included in
   any system that provides registration information for Internet
   registration entities, i.e., network and/or domain name assignments.
   Some of these, in turn, will define requirements for registrars;
   this, however, is an issue outside of the scope of this document.

   It is hoped that this work will also influence the development of
   requirements and specifications for domain name registries at some
   point in the future.

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 September 27, 2012.

Table of Contents

[1](#). Introduction

   The ubiquitous [WHOIS] service can be used today to query for domain
   name registration or network or subnetwork assignment information by
   the general public. It is however a very simple protocol, whose
   output is free-form and thus not amenable to machine parsing. It
   also includes no support for internationalization, and it enables
   only rudimentary (if any) differential service capabilities.

   The CRISP working group created a workable and extensible standard
   for replacing WHOIS, called [IRIS], which attempted to address these
   problems. Unfortunately, IRIS has seen little to no deployment for
   various reasons, mostly its complexity compared to WHOIS and some
   political and technical inertia.

   Thus, this effort confronts anew the need for a better service than
   WHOIS provides, by first laying down a framework of requirements that
   such a service needs to accommodate to become a viable alternative to
   WHOIS.

[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]. In
   particular, since this is not a standards track document, these key
   words are meant to describe requirements for those proposals for a
   WHOIS replacement that seek standards track status.

[2.2](#). Incorporated Requirements

   Many of the requirements distilled from the input provided by various
   communities in [CRISP] will apply to this effort as well. It is
   certainly the case that the research presented there should be
   considered prerequisite reading for this new work.

[3](#). Requirements

   This section enumerates the basic requirements of any WHOIS
   replacement system.

3.1. **Protocol Requirements**

   The protocol requirements are as follows:

   1. To support internationalized values, the protocol SHOULD be able
       to deliver replies that contain data that are not exclusively
       7-bit clean.

   2. The protocol SHOULD be able to deliver a reply that is
       effectively a referral or redirect to another server. The DNS
       and some existing WHOIS extensions have had this capability for
       some time, and this effort would do well to consider those
       methods when developing this capability.

   3. The protocol MUST consist of a reply syntax that uses a well-
       established generic encoding mechanism such as XML or JSON. [We
       need to choose one and standardize on it.]

   4. The protocol MUST define a minimum set of fields and their
       respective syntaxes that are to be included in every reply.
       Context-specific extensions to this set MAY also be defined. The
       set of fields MAY be different for names versus numbers, but a
       common set of fields between the two is expected.

   5. Either the protocol or its underlying transport mechansim MUST be
       capable of authentication of some kind sufficiently robust to
       provide different quality-of-service to different clients once
       they identify themselves in a reliable way.

   6. The protocol SHOULD support the notion of including in the reply
       a suggested time-to-live period during which the client is
       expected to cache the reply and not query for it again.

3.2. **Classes of Service**

   The design of servers supporting the above protocol needs to consider
   support for the following likely classes of service for clients:

   anonymous: Users with no prior arrangement for access to the data;
      typically all available data will be provided in response to a
      query, but the query rate may be severely limited. No
      authentication is typically required. Some data considered to be
      personally identifiable information MAY be elided.

   security: Users that have an interest in a specific subset of a
      registration's data for the purpose of analysis and correlation
      while evaluating the trustworthiness of the source. Examples
      include email client evaluation, email content evaluation, web

site security, etc. The subset will typically include creation/
registration dates, assigned nameserver names and IP addresses,
registrar ID and registrant ID. Users in this class would be
required to authenticate in some way, but such clients would not
typically be subjected to rate limiting given the prior
arrangement.

law enforcement: Users with a bona fide interest in as much
registration data, including change history, as is available.
Typically, queries would be rare but have extremely high priority.
These clients would definitely require authentication and probably
also require encryption.

3.3. Reply Syntax

The reply format needs to conform to the requirements enumerated
below.

NOTE: The standard format is expected to be a significant portion of
the work on the way to describing a new overall WHOIS specification.
In any case, machine-parsability of replies is crucial to the success
of this work.

o All date and/or time fields MUST be formatted as per [DATETIME].

o A server MUST provide a minimum set of data about a given query.
   It is expected that this minimum set will be different for a
   network allocation registry than a domain name registry, however
   the following MUST be provided in either case:

   * The creation date/time of the record

   * The date/time on which the record most recently changed owners/
      registrants

   * The date/time on which any other part of the record was
      modified

   * The identifier of the registrar that created the record (e.g.,
      the registrar ID)

   * The identifier of the registrant that created the record (e.g.,
      the registrant ID)

   * For network registration records, the range of addresses
      comprising the assigned subnet

   o A server MAY provide different output based on the nature of the
      client, where such can be definitively determined.

4. IANA Considerations

   This memo presents no actions for IANA. [RFC Editor: Please remove
   this section prior to publication.]

5. Security Considerations

   This memo introduces an overall protocol model, but no implementation
   details. Specific security considerations of the implementation(s)
   that meet these requirements will be provided in their defining
   documents.

6. Informative References

   [CRISP]   Newton, A., "Cross Registry Internet Service Protocol
              (CRISP) Requirements", RFC 3707, February 2004.

   [DATETIME]
              Klyne, G. and C. Newman, "Date and Time on the Internet:
              Timestamps", RFC 3339, July 2002.

   [IRIS]   Newton, A. and M. Sanz, "IRIS: The Internet Registry
              Information Service (IRIS) Core Protocol", RFC 3981,
              January 2005.

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

   [WHOIS]   Daigle, L., "WHOIS Protocol Specification", RFC 3912,
              September 2004.

Appendix A. Public Discussion

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

Author's Address

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

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