

Network Working Group	A.L. Newton
Internet-Draft	American Registry for Internet Numbers
Intended status: Informational	June 03, 2011
Expires: December 05, 2011	

ARIN's RESTful Web Service for Whois Data
draft-newton-weirds-arin-whoisrws-00

Abstract

This document describes the RESTful Web Service for Whois data as implemented and fielded by the American Registry for Internet Numbers (ARIN). ARIN calls this service Whois-RWS.

The purpose of this document is to facilitate discussion and serve as input into a standards process in this area, currently being discussed on the WHOIS-based Extensible Internet Registration Data Service (WEIRDS) mailing list (<https://www.ietf.org/mailman/listinfo/weirds>). Please excuse this very rough initial draft. It is roughly based on information currently published on ARIN's website. However, future revisions of this document are planned, including discussions on lessons learned by ARIN from the deployment of Whois-RWS and thoughts on a future, unified standard.

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 December 05, 2011.

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. [Requirements notation](#)
- *2. [Background](#)
- *3. [Data Model](#)
 - *3.1. [Networks and Autonomous System Numbers](#)
 - *3.2. [Delegations](#)
 - *3.3. [Organizations and Customers](#)
 - *3.4. [Points of Contact](#)
 - *3.5. [General Query Behavior](#)
- *4. [Use of REST over HTTP](#)
 - *4.1. [Data Formats and Versions](#)
 - *4.2. [Schemas](#)
 - *4.3. [Resource Oriented URLs](#)
 - *4.3.1. [Resources Related to Resources](#)
 - *4.3.2. [Unrelated Lists of Resources](#)
 - *4.3.3. [IP Addresses](#)
 - *4.3.4. [Aggregate Resources](#)
 - *4.3.5. [Query Parameters](#)
 - *4.3.5.1. [Showing More of Detail](#)
 - *4.3.5.2. [Showing Only Points of Contact](#)
 - *4.3.5.3. [Showing ARIN Resources](#)
- *5. [HTTP Caching](#)
- *6. [Security Considerations](#)
- *7. [Lessons Learned](#)

*8. [References](#)

*[Author's Address](#)

[1. Requirements notation](#)

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[\[RFC2119\]](#).

[2. Background](#)

ARIN offers public access to ARIN registration data via a number of services. Traditionally, these services are known in the industry as "Whois" in reference to the public data service of the ARPANET, precursor of today's modern Internet. Whois services are offered by all the Regional Internet Registries (RIRs), most Internet Routing Registries (IRRs) and most domain name registries and registrars. Traditionally, Whois services have been offered using the NICNAME/WHOIS protocol as described by RFC 954 and RFC 3912. This protocol is a simple, text based TCP protocol registered with in the Internet Assigned Numbers Authority (IANA) for well-known port 43. RFC 3912, the most recent specification for the protocol, does not define either data types or data formats. As a consequence, Whois data varies from service provider to service provider and is far from ideal for programmatic consumption.

ARIN provides to the general public services for read-only access to ARIN's registration data. These services include a user-friendly web site (<http://whois.arin.net>), a RESTful Web Service, and a NICNAME/WHOIS port 43 service. For the purposes of programmatic consumption, ARIN recommends the use of the RESTful Web Service and strongly discourages the use of the NICNAME/WHOIS port 43 service. This document describes the RESTful Web Service for ARIN's Whois data, which is known as Whois-RWS.

[3. Data Model](#)

ARIN's Whois-RWS data model is composed of six first order objects: networks, autonomous system numbers, delegations, organizations, points of contact, and customers. Each of these types, except delegations, is directly addressable in a Whois service via a handle (i.e. identifier). Within the Whois RESTful Web Service, these handles are part of URLs that may be used to identify objects in ARIN's Whois registration database by external, non-ARIN systems.

[3.1. Networks and Autonomous System Numbers](#)

Networks and Autonomous System Numbers (ASNs) are collectively referred to as "resources" in ARIN parlance (this should not be confused with the term "resources" in the context of RESTful Web Services and

Resource Oriented Architectures). They are the pieces of information assigned or allocated to organizations for the coordinated administration and operation of the Internet.

Networks signify the allocation or assignment of IP address space and the contiguous IP CIDR blocks that make up that IP address space.

Handles for IPv4 networks start with "NET-", and handles for IPv6 networks start with "NET6-".

Autonomous System Numbers (ASNs) are used for the proper routing of Internet packets. ARIN assigns ASNs in ranges, therefore a single ASN is part of an ASN range allocation. The handles for these registrations start with "AS" and are usually appended with the first number of the ASN range.

3.2. Delegations

Delegations contain the information necessary for Reverse DNS, including the associated nameservers, and DNS DS record information. Unlike the other first order objects, delegations do not have a Handle. Rather, they are directly addressable in a Whois service via a delegation name (i.e. 0.192.in-addr.arpa.).

3.3. Organizations and Customers

Organizations are the registrants of resources and have a direct relationship with ARIN. Organizations may be associated with many resources. Customers do not have a direct relationship with ARIN and, at present, may only be associated with one network registration. Customer handles start with "C", while organization handles have no prefix.

3.4. Points of Contact

A Point of Contact (POC) is the registration of names, mailboxes, and/or phone numbers which fulfill technical or administrative functions on behalf of either an organization or a resource. POC handles are usually appended with "-ARIN" though there are a few exceptions.

3.5. General Query Behavior

For the first order objects addressable via handle, ARIN's Whois services also allow for searching against names and other appropriate fields contained within those objects. Unless otherwise specified, these searches are case insensitive. Because ARIN's networks are composed of multiple CIDR based network blocks, searches by IP address or CIDR notation may sometimes yield what may, at first blush, appear incorrect. For instance, a search for a particular CIDR block may yield a network that covers the given CIDR block and additional CIDR blocks composing the network, while a search by an IP address will always yield the network or networks in which the IP address falls.

[4. Use of REST over HTTP](#)

[4.1. Data Formats and Versions](#)

RESTful web services are not strictly tied to XML. ARIN's Whois RESTful Web Service offers data in XML, JSON, plain text, and HTML (actually, XHTML). However, ARIN's first order data format is XML as there are many format and validation tools readily available for it, and the other formats are offered on a best-effort basis.

If no desired format is specified, XML is the default format.

Specifying a desired data format may be accomplished in one of two ways: either using the HTTP Accept header or using "file extensions". A "file extension" appends a DOS like file extension to the path.

Though there only exists one version of this RESTful interface, it is possible that the "data model" of the structured data such as XML and JSON that is output by this service will need revision. Should it be possible for ARIN to provide a backward compatible version, the HTTP Accept header is the mechanism for specifying the desired version as well as the data format.

The MIME type used in the Accept header will follow the format of "application/arin.whoisrws-{version}+{format}". To use the latest version of this service you would simply use the default MIME types of "application/xml" or "application/json".

The following table lists the data types and their associated MIME types and file extensions.

Data Type	Current Version MIME Type	Version 1 MIME Type	File Extension
XML	application/xml	application/arin.whoisrws-v1+xml	xml
JSON	application/json	application/arin.whoisrws-v1+json	json
plain text	text/plain		txt
HTML	text/html		html

[4.2. Schemas](#)

Relax NG stuff goes here.

Explanation of JSON and Badgerfish goes here.

[4.3. Resource Oriented URLs](#)

URLs point to resources. A typical URL might be `http://whois.arin.net/rest/poc/ALN-ARIN`, which is a URL pointing to the author of this document. Conceptually, this can be broken into a base URL and the resource identifier:

base URL:
http://whois.arin.net/rest

resource: /poc/ALN-ARIN

The base URL is just where ARIN is hosting the service. The real interesting parts are the bits that identify the resources:

/poc this indicates the resource is a Point of Contact.

/ALN-ARIN this is the unique identifier for the Point of Contact, or the POC Handle.

The Whois-RWS data model has six types of addressable resources. These are reflected in Whois-RWS in the following way:

/poc point of contact

/org organization

/net a network

/asn autonomous system number(s)

/customer a customer of an organization

/rdns delegations in the reverse DNS

[4.3.1. Resources Related to Resources](#)

In the ARIN Whois data model, resources have relationships to other resources. Getting references to these resources can be accomplished by addressing the resource in question and applying a resource type qualifier.

Here is the list of relationships (where "XXXX" signifies a handle):

/poc/XXXX

/orgs (i.e. /poc/XXXX/orgs) lists the organizations associated with a given POC

/asns (i.e. /poc/XXXX/asns) lists the autonomous system numbers associated with a given POC

/nets (i.e. /poc/XXXX/nets) lists the networks associated with a given POC

/org/XXXX

/pocs (i.e. /org/XXXX/pocs) lists the points of contact associated with a given organization

/asns (i.e. /org/XXXX/asns)

lists the autonomous system numbers associated with a given organization

/nets (i.e. /org/XXXX/nets) lists the networks associated with a given organization

/asn/XXXX

/pocs (i.e. /asn/XXXX/pocs) lists the points of contact associated with a given autonomous system number

/net/XXXX

/pocs (i.e. /net/XXXX/pocs) lists the points of contact associated with a given network

/parent (i.e. /net/XXXX/parent) lists the parent network of a given network

/children (i.e. /net/XXXX/children) lists the child networks associated with a given network

/rdns (i.e. /net/XXXX/rdns) lists the reverse DNS delegations associated with a given network

/rdns/XXXX

/nets (i.e. /rdns/XXXX/nets) lists the networks associated with a given reverse DNS delegation

4.3.2. Unrelated Lists of Resources

Unrelated lists of resources may be addressed using URL matrix parameters. As an example, to find organizations with the name "ARIN", the /orgs list is appended with the matrix parameter name to form /orgs;name=ARIN (this is without the base URL).

Here is the list of resources and their supported matrix parameters:

/orgs

handle the handle of the organization

name the name of the organization

dba the "doing-business-as" name of the organization

/customers

handle the handle of the customer

name the name of the customer

/pocs

handle the handle of the POC

domain

the domain name of an email address of the POC

first the first name of the POC

middle the middle name of the POC

last the last name of the POC

company the company name registered by the POC

city the city registered by the POC

/asns

handle the handle of the autonomous system number

name the name of the autonomous system number

/nets

handle the handle of the network

name the name of the network

4.3.3. IP Addresses

In the ARIN data model, an IP "network" is a set of associated IP address blocks and the information related to these IP address blocks and the set as a whole. A network can be composed of one IP address block or of multiple IP address blocks. Networks are also hierarchical, meaning that a network can have a parent and can have children. An IP address or range of IP addresses can therefore fall within the IP address block of multiple networks.

To get the networks that a particular IP address may fall within, one can use the /ip/XX.XX.XX.XX resource, where XX.XX.XX.XX signifies the IP address.

Networks may be looked up as well by supplying the full CIDR notation of a range. To use the full CIDR notation, the resource looks like /cidr/XX.XX.XX.XX/YY where XX.XX.XX.XX is the IP address prefix and YY is the CIDR length.

Resources relative to the networks may be fetched using the /less and /more path prefixes. /less will find the networks that are less specific in scope (or wider), and /more will find the networks that are more specific in scope (or narrower).

4.3.4. Aggregate Resources

Though Whois-RWS breaks down the ARIN data model into distinct resources, end users have been accustomed to seeing many related resources with one query over the NICNAME/WHOIS port 43 service. While this can be accomplished by any Whois-RWS clients with multiple

queries, an aggregate resources can reduce the number of queries and make this use case easier on client implementers. For resources under /org, /net, /asn, and /ip, the URL may be appended with "/pft" to produce an aggregate resource which includes the resource being queried and related resources. If you can guess what /pft means, I'll buy you a beer.

4.3.5. Query Parameters

4.3.5.1. Showing More of Detail

By default, lists of resources only show references to the resources. However, lists can be expanded to show full detail by tacking on a showDetails=true URL query parameter (e.g.http://whois.arin.net/rest/org/ARIN/pocs?showDetails=true). This query parameter also forces a resource to list its related resources in line.

4.3.5.2. Showing Only Points of Contact

Some organizations have many, many associated resources. Having all those resources returned when gathering information about an organization maybe very undesirable. To allow getting all the POCs of an organization without taking the penalty of retrieving all the resources of the organization, the showPocs=true parameter may be used. This parameter is only recognized when an organization's information is referenced by handle.

4.3.5.3. Showing ARIN Resources

In the past, ARIN's NICNAME/WHOIS service yielded a "No Match" result for searches of IP address space unallocated by ARIN but within ARIN's IP address pool. This was somewhat confusing and could possibly mislead a person into thinking the address space is simply unallocated to any available address pool (in other words, still held by the IANA). This has been changed in ARIN's NICNAME/WHOIS service so it will return the appropriate ARIN network allocation for IP addresses unallocated by ARIN. By default, the RESTful service will also return the appropriate ARIN network allocation for IP addresses unallocated by ARIN to ARIN constituents.

If the behavior is desired, the showARIN=false query parameter may be used on IP and CIDR queries.

5. HTTP Caching

One of the advantages of a RESTful Web Service is that HTTP infrastructure may be employed to make it more reliable and/or robust. HTTP caching is one such tool to help with these goals. ARIN employs caching at multiple levels and customized to address the query pattern seen against ARIN's services.

If HTTP caching is to be employed on the client side, there are a number of issues to consider to craft a custom caching solution.

First caching should be restricted to URLs that have the highest likelihood of remaining in the cache using the cache's retention strategy. In almost any scenario, caching RESTful resources under /rest/ip will lead to a large dataset and a low cache hit ratio. However, it is likely acceptable to cache resources under /rest/org, /rest/poc, and /rest/net. It is recommended that cache statistics be available for proper cache tuning.

Second, many HTTP caches base the cache objects on URLs. However, the Accept header dictates the content of the object and may not be part of the cache key of cached objects. Therefore it is possible to get a cached object of one type when another type was requested. For instance, /rest/poc/KOSTE-ARIN might return XML when JSON was requested. To overcome this, it is advisable to use file extensions to specify data format as file extensions are part of the URL.

6. Security Considerations

Of course there are some... but I'm not gonna tell you what they are. Seriously.... to be discussed in future revisions.

7. Lessons Learned

to be discussed in a future revision, the lessons we have learned from this

8. References

[RFC2119]	Bradner, S. , " Key words for use in RFCs to Indicate Requirement Levels ", BCP 14, RFC 2119, March 1997.
-----------	---

Author's Address

Andrew Lee Newton
American Registry for Internet Numbers