

Registration Protocols Extensions  
Internet-Draft  
Intended status: Standards Track  
Expires: April 6, 2018

M. Loffredo  
M. Martinelli  
IIT-CNR/Registro.it  
October 3, 2017

**Registration Data Access Protocol (RDAP) Partial Response  
draft-loffredo-regext-rdap-partial-response-00**

**Abstract**

The Registration Data Access Protocol (RDAP) does not include capabilities to request partial responses. In fact, according to the user authorization, the server can only return full responses. Partial responses capability, especially in the case of search queries, could bring benefits to both clients and servers. This document describes a RDAP query extension that allows clients to specify their preference for obtaining a partial response.

**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 <https://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 April 6, 2018.

**Copyright Notice**

Copyright (c) 2017 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 (<https://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

|                      |   |                   |
|----------------------|---|-------------------|
| <a href="#">1.</a>   | Introduction . . . . .                                  | <a href="#">2</a> |
| <a href="#">1.1.</a> | Conventions Used in This Document . . . . .             | <a href="#">3</a> |
| <a href="#">2.</a>   | Approaches to partial response implementation . . . . . | <a href="#">3</a> |
| <a href="#">3.</a>   | RDAP Path Segment Specification . . . . .               | <a href="#">4</a> |
| <a href="#">4.</a>   | Implementation Status . . . . .                         | <a href="#">5</a> |
| <a href="#">4.1.</a> | IIT-CNR/Registro.it . . . . .                           | <a href="#">6</a> |
| <a href="#">5.</a>   | Security Considerations . . . . .                       | <a href="#">6</a> |
| <a href="#">6.</a>   | IANA Considerations . . . . .                           | <a href="#">7</a> |
| <a href="#">7.</a>   | Acknowledgements . . . . .                              | <a href="#">7</a> |
| <a href="#">8.</a>   | References . . . . .                                    | <a href="#">7</a> |
| <a href="#">8.1.</a> | Normative References . . . . .                          | <a href="#">7</a> |
| <a href="#">8.2.</a> | Informative References . . . . .                        | <a href="#">8</a> |
|                      | Authors' Addresses . . . . .                            | <a href="#">9</a> |

## [1.](#) Introduction

The use of partial response in RESTful API [[REST](#)] design is very common. The rationale is quite simple: instead of returning objects in API responses with all data fields, only a subset is returned. The benefit is obvious: less data transferred over the network mean less bandwidth usage, faster server response, less CPU time spent both on the server and the client, as well as less memory usage on the client.

Several leading APIs providers (e.g. LinkedIn [[LINKEDIN](#)], Facebook [[FACEBOOK](#)], Google [[GOOGLE](#)]) implement the partial response feature by providing an optional query parameter by which users require the fields they wish to receive. Partial response is also considered a leading principle by many best practices guidelines in REST APIs implementation ([[REST-API1](#)], [[REST-API2](#)]) in order to improve performance, save on bandwidth and possibly accelerate the overall interaction. In other contexts, for example in digital libraries and bibliographic catalogues, servers can provide responses according to different element sets (i.e. "brief" to get back a short response and "full" to get back the complete response)

Currently, RDAP does not provide a client with any way to request a partial response: the server can only provide the client with the full response ([[RFC7483](#)]). Furthermore, servers cannot define the limits of the results according to partial responses and this causes strong inefficiencies.



The protocol described in this specification extends RDAP search capabilities to enable partial responses, by adding a new query parameter and using a RESTful web service. The service is implemented using the Hypertext Transfer Protocol (HTTP) [[RFC7230](#)] and the conventions described in [RFC 7480](#) [[RFC7480](#)].

Impact on the current state of RDAP implementation is low.

### **1.1. Conventions Used in This Document**

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. Approaches to partial response implementation**

Looking at the implementation experiences described above, two approaches to the implementation of partial response can be detected:

- o the client declares explicitly the data fields to get back;
- o the client declares a name identifying a server pre-defined set of data fields.

The former is more flexible than the latter, because clients can specify all the data fields they need. Anyway, it has some drawbacks:

- o Fields have to be declared according to a given syntax. This is a simple task when the data structure of the object is flat, but it is much more difficult when the object has a tree structure like the one of a JSON object. The presence of arrays and deep nested objects contribute to complicate both the syntax definition of the query and, consequently, the processing phase on the server side.
- o Clients should perfectly know the returned object to avoid cases when the required fields are not compliant with the object data structure.
- o The request of some fields cannot match the user access levels. Clients could put unauthorized fields in their requests and servers should define a strategy for providing a response: to return always an error response or to return a response ignoring the unauthorized fields.

The latter approach seems to facilitate RDAP interoperability. In fact, servers can define some basic field sets which, if known to the clients, can increase the probability to get a valid response. In



addition, the definition of pre-defined sets of fields makes easier to establish the results limits.

Considering that there is not a real need for RDAP users to have the maximum flexibility in defining all the possible sets of logically connected fields (for example, users interested in domains usually need to know the status, the creation date, the expire date of each domain), the latter approach is preferred.

### 3. RDAP Path Segment Specification

The new query parameter is an OPTIONAL extension of search path segments defined in [RFC 7482](#) [RFC7482]. The query parameter is "fieldSet" whose value is a string identifying a server pre-defined set of fields (Figure 1). Values REQUIRED to be implemented are:

- o id: the server provides only the "objectClassName" field and the field identifying the object ("handle" for entities, "ldhName" for domains and nameservers). This field set can be used when the client wants to obtain a collection of object identifiers (Figure 2);
- o brief: it contains the fields that can be included in a "short" response. This field set can be used when the client is asking for a subset of the full response which gives a basic knowledge of each object;
- o full: it contains all the information the server can provide for a particular object.

Fields belonging to brief and full field sets should be provided according to users access levels. Servers MAY implement additional field sets not included in the list above. Servers SHOULD also define a "default" field set.

`https://example.com/rdap/domains?name=example*.com&fieldSet=id`

Figure 1: Example of RDAP search query reporting the fieldSet parameter

```
{
  "rdapConformance": [
    "rdap_level_0",
  ],
  ...
  "domainSearchResults": [
    {
      "objectClassName": "domain",
      "ldhName": "example1.com"
    },
    {
      "objectClassName": "domain",
      "ldhName": "example2.com"
    },
    ...
  ]
}
```

Figure 2: Example of RDAP response according to the "id" field set

#### 4. Implementation Status

NOTE: Please remove this section and the reference to [RFC 7942](#) prior to publication as an RFC.

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC 7942](#) [[RFC7942](#)]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC 7942](#), "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".





#### **4.1. IIT-CNR/Registro.it**

Responsible Organization: Institute of Informatics and Telematics of National Research Council (IIT-CNR)/Registro.it

Location: <https://rdap.pubtest.nic.it/>

Description: This implementation includes support for RDAP queries using data from the public test environment of .it ccTLD. The RDAP server does not implement any security policy because data returned by this server are only for experimental testing purposes.

Level of Maturity: This is a "proof of concept" research implementation.

Coverage: This implementation includes all of the features described in this specification.

Contact Information: Mario Loffredo, [mario.loffredo@iit.cnr.it](mailto:mario.loffredo@iit.cnr.it)

### **5. Security Considerations**

Search query typically requires more server resources (such as memory, CPU cycles, and network bandwidth) when compared to lookup query. This increases the risk of server resource exhaustion and subsequent denial of service due to abuse. Partial response can contribute together with other strategies (e.g. restricting search functionality, limiting the rate of search requests, truncating and paging results) to mitigate this risk.

Furthermore, partial response can help RDAP operators to regulate access control based on client identification, implemented by HTTP basic or digest authentication as described in [RFC 7481](#) [[RFC7481](#)] or by a federated authentication system ([[I-D.hollenbeck-regext-rdap-openid](#)]). In fact, RDAP operators can follow different, not alternative, approaches to the building of responses according to the user access levels:

- o the list of fields for each set (except "id") can be different according to the user access levels. At present, this is already implemented for the full response, but it could be done also for the other defined field sets. In some cases, it might happen that brief and full field sets are exactly the same;
- o some field sets could be available only to some users. In this case, servers could define additional field sets to those



indicated above ("id", "brief", "full"), making them available only to users with specific access levels.

Servers can also define different results limits according to the available field sets, so a more flexible truncation strategy can be realized and users can take advantage of a more efficient results paging implementation ([[I-D.loffredo-regext-rdap-sorting-and-paging](#)]).

Therefore, the new parameter presented in this document provide the RDAP operators with a way to implement a secure server without penalizing its efficiency.

## **6. IANA Considerations**

This document has no actions for IANA.

## **7. Acknowledgements**

The authors would like to acknowledge Scott Hollenbeck for his contribution to this document.

## **8. References**

### **8.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [RFC7480] Newton, A., Ellacott, B., and N. Kong, "HTTP Usage in the Registration Data Access Protocol (RDAP)", [RFC 7480](#), DOI 10.17487/RFC7480, March 2015, <<https://www.rfc-editor.org/info/rfc7480>>.
- [RFC7481] Hollenbeck, S. and N. Kong, "Security Services for the Registration Data Access Protocol (RDAP)", [RFC 7481](#), DOI 10.17487/RFC7481, March 2015, <<https://www.rfc-editor.org/info/rfc7481>>.

- [RFC7482] Newton, A. and S. Hollenbeck, "Registration Data Access Protocol (RDAP) Query Format", [RFC 7482](#), DOI 10.17487/RFC7482, March 2015, <<https://www.rfc-editor.org/info/rfc7482>>.
- [RFC7483] Newton, A. and S. Hollenbeck, "JSON Responses for the Registration Data Access Protocol (RDAP)", [RFC 7483](#), DOI 10.17487/RFC7483, March 2015, <<https://www.rfc-editor.org/info/rfc7483>>.

## 8.2. Informative References

- [FACEBOOK] facebook.com, "facebook for developers - Using the Graph API", July 2017.
- [GOOGLE] google.com, "Making APIs Faster: Introducing Partial Response and Partial Update", March 2010.
- [I-D.hollenbeck-regext-rdap-openid] Hollenbeck, S., "Federated Authentication for the Registration Data Access Protocol (RDAP) using OpenID Connect", [draft-hollenbeck-regext-rdap-openid-04](#) (work in progress), August 2017.
- [I-D.loffredo-regext-rdap-sorting-and-paging] Loffredo, M., Martinelli, M., and S. Hollenbeck, "Registration Data Access Protocol (RDAP) Query Parameters for Result Sorting and Paging", [draft-loffredo-regext-rdap-sorting-and-paging-00](#) (work in progress), May 2017.
- [LINKEDIN] linkedin.com, "Java One 2009: Building Consistent RESTful APIs in a High Performance Environment", July 2009.
- [REST] Fielding, R., "Architectural Styles and the Design of Network-based Software Architectures", 2000.
- [REST-API1] Jobinesh, P., "RESTful Java Web Services - Second Edition", September 2015.
- [REST-API2] Masse, M., "REST API Design Rulebook", October 2011.



[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [BCP 205](#), [RFC 7942](#), DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.

#### Authors' Addresses

Mario Loffredo  
IIT-CNR/Registro.it  
Via Moruzzi,1  
Pisa 56124  
IT

Email: [mario.loffredo@iit.cnr.it](mailto:mario.loffredo@iit.cnr.it)  
URI: <http://www.iit.cnr.it>

Maurizio Martinelli  
IIT-CNR/Registro.it  
Via Moruzzi,1  
Pisa 56124  
IT

Email: [maurizio.martinelli@iit.cnr.it](mailto:maurizio.martinelli@iit.cnr.it)  
URI: <http://www.iit.cnr.it>