
Workgroup: Domain Name System Operations
Internet-Draft: draft-dulaunoy-dnsop-passive-dns-cof-11
Published: 27 April 2024
Intended Status: Informational
Expires: 29 October 2024
Authors: A. Dulaunoy A. Kaplan P. Vixie H. Stern
 CIRCL Farsight Security, Inc. Farsight Security, Inc.

W. Kumari
Google

Passive DNS - Common Output Format

Abstract

This document describes a common output format of Passive DNS Servers that clients can query. The output format description also includes a common semantic for each Passive DNS system. By having multiple Passive DNS Systems adhere to the same output format for queries, users of multiple Passive DNS servers will be able to combine result sets easily.

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 29 October 2024.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Limitation	4
3. Common Output Format	4
3.1. Overview	4
3.2. ABNF grammar	5
3.3. Mandatory Fields	5
3.3.1. rname	5
3.3.2. rrtype	6
3.3.3. rdata	6
3.3.4. time_first	6
3.3.5. time_last	6
3.4. Optional Fields	6
3.4.1. count	6
3.4.2. bailiwick	7
3.5. Additional Fields	7
3.5.1. sensor_id	7
3.5.2. zone_time_first	7
3.5.3. zone_time_last	7
3.5.4. origin	7
3.5.5. time_first_ms	7
3.5.6. time_last_ms	7
3.6. Additional Fields Registry	7
3.7. Additional notes	8
3.8. Suggested MIME Types	8
4. Acknowledgements	8
5. IANA Considerations	8

6. Privacy Considerations	8
7. Security Considerations	9
8. References	9
8.1. Normative References	9
8.2. References	10
8.3. Informative References	10
Appendix A. Examples	10
Authors' Addresses	11

1. Introduction

Passive DNS is a technique described by Florian Weimer in 2005 in [Passive DNS replication, F Weimer - 17th Annual FIRST Conference on Computer Security \[WEIMERPDNS\]](#). Since then, multiple Passive DNS implementations were created and have evolved over time. Users of these Passive DNS servers may query a server (often via [WHOIS \[RFC3912\]](#) or [HTTP REST \[REST\]](#)), parse the results, and process them in other applications.

There are multiple implementations of Passive DNS software. Users of Passive DNS query each implementation and aggregate the results for their search. This document describes the output format of four Passive DNS Systems ([\[DNSDB\]](#), [\[DNSDBQ\]](#), [\[PDNSCERTAT\]](#), [\[PDNSCIRCL\]](#) and [\[PDNSCOF\]](#)) that are in use today and that already share a nearly identical output format. As the format and the meaning of output fields from each Passive DNS need to be consistent, this document proposes a solution to commonly name each field along with its corresponding interpretation. The format follows a simple key-value structure in [JSON \[RFC4627\]](#) format. The benefit of having a consistent Passive DNS output format is that multiple client implementations can query different servers without having to have a separate parser for each individual server. [passivedns-client \[PDNSCLIENT\]](#) currently implements multiple parsers due to a lack of standardization. The document does not describe the protocol (e.g. [WHOIS \[RFC3912\]](#), [HTTP REST \[REST\]](#)) nor the query format used to query the Passive DNS. Neither does this document describe "pre-recursor" Passive DNS Systems. Each of these are separate topics and deserve their own RFC documents. This document describes the current best practices implemented in various Passive DNS server implementations.

1.1. Requirements Language

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 [RFC 2119 \[RFC2119\]](#).

2. Limitation

As Passive DNS servers can include protection mechanisms for their operation, results might be different due to those protection measures. These mechanisms filter out DNS answers if they fail some criteria. The [bailiwick algorithm \[BAILIWICK\]](#) protects the Passive DNS Database from [cache poisoning attacks \[CACHEPOISONING\]](#). Another limitation that clients querying the database need to be aware of is that each query simply gets a snapshot-in-time answer at the time of querying. Clients **MUST NOT** rely on existing answers from different Passive DNS database. Nor should they assume that answers will be identical across multiple Passive DNS Servers.

3. Common Output Format

3.1. Overview

The formatting of the answer follows the [JSON \[RFC4627\]](#) format. In fact, it is a subset of the full JSON language. Notable differences are the modified definition of whitespace ("ws"). The order of the fields is not significant for the same resource type.

The intent of this output format is to be easily parsable by scripts. Each JSON object is expressed on a single line to be processed by the client line-by-line. Every implementation **MUST** support the JSON output format.

[Examples of JSON \(Appendix A\)](#) output are in the appendix.

3.2. ABNF grammar

Formal grammar as defined in [ABNF \[RFC2234\]](#)

```

answer      = entries
entries     = * ( entry newline )
entry       = ws "{" ws keyvallist ws "}" ws
keyvallist  = [ member * ( value-separator member ) ]
member      = field name-separator value
name-separator = ws %x 3 A ws      ; : colon
value-separator = ws %x 2 C ws    ; , comma
field       = field-name | futureField
field-name  = "rrname" | "rrtype" | "rdata" | "time_first" |
             "time_last" | "count" | "bailiwick" | "sensor_id" |
             "zone_time_first" | "zone_time_last" | "origin" |
             "time_first_ms" | "time_last_ms"
futureField = string
newline     = [ CR ] LF
CR         = %x 0 D                ; Carrige return
LF         = %x 0 A                ; Line feed or New line
qm         = %x 2 2                ; " Quotation mark
ws         = * (
             %x 2 0 |              ; Space
             %x 0 9                ; Horizontal tab
             )

```

Note that value is defined in [JSON \[RFC4627\]](#) and has the same specification as there. The same goes for the definition of string. Note the changed definition of ws does not include CR or LF as those are NOT allowed in NDJSON, and thus the definition here MUST be used for other ABNF defitions in [JSON \[RFC4627\]](#).

3.3. Mandatory Fields

Implementation MUST support all the mandatory fields.

Uniqueness property: the tuple (rrname,rrtype,rdata) will always be unique within one answer per server. While rrname and rrtype are always individual JSON primitive types (strings, numbers, booleans or null), rdata MAY return multiple resource records or a single record. When multiple resource records are returned, rdata MUST be a JSON array. In the case of a single resource record is returned, rdata MUST be a JSON string or a JSON array containing one JSON string. Senders SHOULD send an array for rdata, but receivers MUST be able to accept a single-string result for rdata.

3.3.1. rrname

This field returns the name of the queried resource. Represented as a [JSON \[RFC4627\]](#) string.

3.3.2. rrtype

This field returns the resource record type as seen by the passive DNS. The key is `rrtype` and the value is in the interpreted record type represented as a [JSON \[RFC4627\]](#) string. If the value cannot be interpreted, the decimal value is returned following the principle of transparency as described in [RFC 3597 \[RFC3597\]](#). Then the decimal value is represented as a [JSON \[RFC4627\]](#) number. The resource record type can be any values as described by IANA in the DNS parameters document in the section 'Resource Record (RR) TYPES' (<http://www.iana.org/assignments/dns-parameters>). Supported textual descriptions of rrtypes include: A, AAAA, CNAME, etc. A client MUST be able to understand these textual rrtype values represented as a [JSON \[RFC4627\]](#) string. In addition, a client MUST be able to handle a decimal value (as mentioned above) answer represented as a [JSON \[RFC4627\]](#) number.

3.3.3. rdata

This field returns the resource records of the queried resource. When multiple resource records are returned, `rdata` MUST be a JSON array containing JSON strings. In the case of a single resource record is returned, `rdata` MUST be a JSON string or a JSON array containing one JSON string. Each resource record is represented as a [JSON \[RFC4627\]](#) string. Each resource record MUST be escaped as defined in section 2.6 of [RFC4627 \[RFC4627\]](#). Depending on the rrtype, this can be an IPv4 or IPv6 address, a domain name (as in the case of CNAMEs), an SPF record, etc. A client MUST be able to interpret any value which is legal as the right hand side in a DNS master file [RFC 1035 \[RFC1035\]](#) and [RFC 1034 \[RFC1034\]](#). If the `rdata` came from an unknown DNS resource records, the server must follow the transparency principle as described in [RFC 3597 \[RFC3597\]](#).

3.3.4. time_first

This field returns the first time that the record / unique tuple (`rrname`, `rrtype`, `rdata`) has been seen by the passive DNS. The date is expressed in seconds (decimal) since 1st of January 1970 (Unix timestamp). The time zone MUST be UTC. This field is represented as a [JSON \[RFC4627\]](#) number.

3.3.5. time_last

This field returns the last time that the unique tuple (`rrname`, `rrtype`, `rdata`) record has been seen by the passive DNS. The date is expressed in seconds (decimal) since 1st of January 1970 (Unix timestamp). The time zone MUST be UTC. This field is represented as a [JSON \[RFC4627\]](#) number.

3.4. Optional Fields

Implementations SHOULD support one or more fields.

3.4.1. count

Specifies how many authoritative DNS answers were received at the Passive DNS Server's collectors with exactly the given set of values as answers (i.e. same data in the answer set - compare with the uniqueness property in "Mandatory Fields"). The number of requests is expressed as a decimal value. This field is represented as a [JSON \[RFC4627\]](#) number.

3.4.2. **bailiwick**

The bailiwick is the best estimate of the apex of the zone where this data is authoritative. This field is represented as a [JSON \[RFC4627\]](#) string.

3.5. **Additional Fields**

Implementations MAY support the following fields:

3.5.1. **sensor_id**

This field returns the sensor information where the record was seen. It is represented as a [JSON \[RFC4627\]](#) string.

If the data originate from sensors or probes which are part of a publicly-known gathering or measurement system (e.g. RIPE Atlas), a [JSON \[RFC4627\]](#) string SHOULD be prefixed.

3.5.2. **zone_time_first**

This field returns the first time that the unique tuple (rrname, rrtype, rdata) record has been seen via master file import. The date is expressed in seconds (decimal) since 1st of January 1970 (Unix timestamp). The time zone MUST be UTC. This field is represented as a [JSON \[RFC4627\]](#) number.

3.5.3. **zone_time_last**

This field returns the last time that the unique tuple (rrname, rrtype, rdata) record has been seen via master file import. The date is expressed in seconds (decimal) since 1st of January 1970 (Unix timestamp). The time zone MUST be UTC. This field is represented as a [JSON \[RFC4627\]](#) number.

3.5.4. **origin**

Specifies the resource origin of the Passive DNS response. This field is represented as a [Uniform Resource Identifier \[RFC3986\]](#) (URI) in the form of a [JSON \[RFC4627\]](#) string.

3.5.5. **time_first_ms**

Same meaning as the field "time_first", with the only difference, that the resolution is in milliseconds since 1st of January 1970 (UTC).

3.5.6. **time_last_ms**

Same meaning as the field "time_last", with the only difference, that the resolution is in milliseconds since 1st of January 1970 (UTC).

3.6. **Additional Fields Registry**

In accordance with [\[RFC6648\]](#), designers of new passive DNS applications that would need additional fields can request and register new field name at <https://github.com/adulau/pdns-qof/wiki/Additional-Fields>.

3.7. Additional notes

An implementer of a passive DNS Server MAY chose to either return `time_first` and `time_last` OR return `zone_time_first` and `zone_time_last`. In pseudocode: `(time_first AND time_last)` OR `(zone_time_first AND zone_time_last)`. In this case, `zone_time_{first,last}` replace the `time_{first,last}` fields. However, this is not encouraged since it might be confusing for parsers who will expect the mandatory fields `time_{first,last}`. See: [\[github_issue_17\]](#)

3.8. Suggested MIME Types

An implementer of a passive DNS Server SHOULD serve a document in this Common Output Format with a MIME header of "application/x-ndjson".

4. Acknowledgements

Thanks to the Passive DNS developers who contributed to the document.

5. IANA Considerations

This memo includes no request to IANA.

6. Privacy Considerations

Passive DNS Servers capture DNS answers from multiple collection points ("sensors") which are located on the Internet-facing side of DNS recursors ("post-recursor passive DNS"). In this process, they intentionally omit the source IP, source port, destination IP and destination port from the captured packets. Since the data is captured "post-recursor", the timing information (who queries what) is lost, since the recursor will cache the results. Furthermore, since multiple sensors feed into a passive DNS server, the resulting data gets mixed together, reducing the likelihood that Passive DNS Servers are able to find out much about the actual person querying the DNS records. In this sense, passive DNS Servers are similar to keeping an archive of all previous phone books - if public DNS records can be compared to phone numbers - as they often are. Nevertheless, the authors strongly encourage Passive DNS implementors to take special care of privacy issues. [bortzmeyer-dnsop-dns-privacy](#) is an excellent starting point for this. Finally, the overall recommendations in [RFC6973](#) [[RFC6973](#)] should be taken into consideration when designing any application which uses Passive DNS data.

In the scope of the General Data Protection Regulation (GDPR - Directive 95/46/EC), operators of Passive DNS Server needs to ensure the legal ground and lawfulness of its operation.

7. Security Considerations

In some cases, Passive DNS output might contain confidential information and its access might be restricted. When a user is querying multiple Passive DNS and aggregating the data, the sensitivity of the data must be considered.

8. References

8.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [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>>.
- [RFC2234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, DOI 10.17487/RFC2234, November 1997, <<https://www.rfc-editor.org/info/rfc2234>>.
- [RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <<https://www.rfc-editor.org/info/rfc3597>>.
- [RFC3912] Daigle, L., "WHOIS Protocol Specification", RFC 3912, DOI 10.17487/RFC3912, September 2004, <<https://www.rfc-editor.org/info/rfc3912>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", RFC 4627, DOI 10.17487/RFC4627, July 2006, <<https://www.rfc-editor.org/info/rfc4627>>.
- [RFC6648] Saint-Andre, P., Crocker, D., and M. Nottingham, "Deprecating the "X-" Prefix and Similar Constructs in Application Protocols", BCP 178, RFC 6648, DOI 10.17487/RFC6648, June 2012, <<https://www.rfc-editor.org/info/rfc6648>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

8.2. References

- [BAILIWICK]** Edmonds, R., "Passive DNS Hardening", 2010, <https://archive.farsightsecurity.com/Passive_DNS/passive_dns_hardening_handout.pdf>.
- [CACHEPOISONING]** Kaminsky, D., "Black ops 2008: It's the end of the cache as we know it.", 2008, <http://kurser.lobner.dk/dDist/DMK_BO2K8.pdf>.
- [DNSDB]** Security, F., "DNSDB API", 2013, <<https://api.dnsdb.info/>>.
- [DNSDBQ]** Vixie, P., "DNSDB API Client, C Version", 2018, <<https://github.com/dnsdb/dnsdbq>>.
- [github_issue_17]** et.al, P. V. W. A. K., "Discussion on the existing implementations of returning either zone_time{first,last} OR time_{first,last}", 2020, <<https://github.com/adulau/pdns-qof/issues/17>>.
- [PDNSCERTAT]** CERT.at, "pDNS presentation at 4th Centr R&D workshop Frankfurt Jun 5th 2012", 2012, <http://www.centri.org/system/files/agenda/attachment/d4-papst-passive_dns.pdf>.
- [PDNSCIRCL]** Luxembourg, C. -. I. R. C., "CIRCL Passive DNS", 2012, <<https://www.circl.lu/services/passive-dns/>>.
- [PDNSCLIENT]** Lee, C., "Queries 5 major Passive DNS databases: BFK, CERTEE, DNSParse, ISC, and VirusTotal.", 2013, <<https://github.com/chrislee35/passivedns-client>>.
- [PDNSCOF]** Dulaunoy, D. P. A., "Passive DNS server interface using the common output format", 2019, <<https://github.com/D4-project/analyzer-d4-passivedns/>>.
- [REST]** Fielding, R. T., "Representational State Transfer (REST)", 2000, <http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm>.
- [WEIMERPDNS]** Weimer, F., "Passive DNS Replication", 2005, <<http://www.enyo.de/fw/software/dnslogger/first2005-paper.pdf>>.

8.3. Informative References

Appendix A. Examples

The JSON output are represented on multiple lines for readability but each JSON object should be on a single line.

If you query a passive DNS for the rname `www.ietf.org`, the passive dns common output format can be:

```
{ "count": 1 0 2, "time_first": 1 2 9 8 4 1 2 3 9 1, "rrtype": "AAAA",  
  "rrname": "www.ietf.org", "rdata": " 2 0 0 1:1 8 9 0:1 1 1 2:1::2 0",  
  "time_last": 1 3 0 2 5 0 6 8 5 1 }  
{ "count": 5 9, "time_first": 1 3 8 4 8 6 5 8 3 3, "rrtype": "A",  
  "rrname": "www.ietf.org", "rdata": " 4 . 3 1 . 1 9 8 . 4 4",  
  "time_last": 1 3 8 9 0 2 2 2 1 9 }
```

If you query a passive DNS for the rrname ietf.org, the passive dns common output format can be:

```
{ "count": 1 0 9 8 7 7, "time_first": 1 2 9 8 3 9 8 0 0 2, "rrtype": "NS",  
  "rrname": "ietf.org", "rdata": "ns 1 .yyz 1 .afilias-nst.info",  
  "time_last": 1 3 8 9 0 9 5 3 7 5 }  
{ "count": 4, "time_first": 1 2 9 8 4 9 5 0 3 5, "rrtype": "A",  
  "rrname": "ietf.org", "rdata": " 6 4 . 1 7 0 . 9 8 . 3 2",  
  "time_last": 1 2 9 8 4 9 5 0 3 5 }  
{ "count": 9, "time_first": 1 3 1 7 0 3 7 5 5 0, "rrtype": "AAAA",  
  "rrname": "ietf.org", "rdata": " 2 0 0 1:1 8 9 0:1 2 3 a::1:1 e",  
  "time_last": 1 3 3 0 2 0 9 7 5 2 }
```

Please note that the examples imply that a single query returns a single set of JSON objects. For example, two queries were made; one query returned a set of two JSON objects and the other query returned a set of three JSON objects. This specification requires each JSON object individually **MUST** conform to the common output format, but this specification does not require that a query will return a set of JSON objects.

Please note that in the examples above, any backslashes "\" can be ignored and are an artifact of the tools which produced this document.

Authors' Addresses

Alexandre Dulaunoy

CIRCL

122, rue Adolphe Fischer

L-1521 Luxembourg

Luxembourg

Phone: (+352) 247 88444

Email: alexandre.dulaunoy@circl.lu

URI: <http://www.circl.lu/>

L. Aaron Kaplan

A-1170 Vienna

Austria

Email: aaron@lo-res.org

Paul Vixie

Farsight Security, Inc.
11400 La Honda Road
Woodside, California 94062
United States of America
Email: paul@redbarn.org
URI: <https://www.farsightsecurity.com/>

Henry Stern

Farsight Security, Inc.
11400 La Honda Road
Woodside, California 94062
United States of America
Phone: +1 650 542-7836
Email: henry@stern.ca
URI: <https://www.farsightsecurity.com/>

Warren Kumari

Google
Email: warren@kumari.net