Passive DNS Simplified Integration Guide

This guide is suitable for external users wanting to integrate with the public REST API for the mnemonic PassiveDNS database.
This API is available without any authentication, although with a resource limitation.

> If you are a customer of mnemonic with private PDNS data, please see the PassiveDNS Integration Guide in the Argus API integration documentation for how to use PDNS with private data.

> For users/organizations with an agreement, see the authentication section on how to perform authenticated queries.

> The Swagger API documentation is always up-to-date and lets you try out any query with your user session or an API-key.

Introduction

Passive Domain Name System (PDNS) is a database consisting of domain names as used by different IPs. The PDNS service exposes three different endpoints that can be used to search on these records. These are described below. Each record is marked with a TLP, indicating the sensitivity of the record. Records with TLP white are publicly accessible.

Finally, there is an endpoint that can be used to submit PDNS records in bulk. This is described towards the bottom of this document.

Concept	Description
DNS	Domain Name System
Passive DNS	DNS records that the PDNS service collects
TLP	Traffic Light Protocol
Submission	Refers to a DNS record submitted to the PDNS REST API

Simple query

To perform a simple PDNS query, use

```
https://api.mnemonic.no/pdns/v3/<query>?parameters
```

where "<query>" is a DNS query or answer string to lookup, for example

```
https://api.mnemonic.no/pdns/v3/cnn.com?rrType=A
```

Query parameters

There are multiple query parameters that can be passed along:

Parameter	Possible values	Default value	Description

limit	[0-maxint]	25	Limit the number of returned values. The default value of 25 is set to avoid that a queried record with a huge result set is streamed back to the client unnecessarily. A value of "0" means that the client explicitly requests an unlimited result set. The result will return a 412 error code if the limit is set to be more than the server allows. This is 10 000 for authenticated users, while public use has its limit set to 1000.
offset	[0-maxint]	0	Skip the initial <offset> number of values in the result set. Together with the parameter "limit", this allows a client to perform pagination of the results.
rrClass	Any DNS record class	<any>	The most probable record class is IN. By default, this field is not filtered.
rrType	Any DNS record type	<any>	The most used record classes are A, AAAA, PTR, CNAME and MX, but any supported DNS record name can be used here. By default, this field is not filtered.

Result format

The result format is JSON, and consists of a result container, and a number of results.

The result container has the following format

```
{
    "responseCode": 200, # the response code. Normal responses should
return HTTP code 200
    "count": 20,         # the total number of matching records. If this
value is lower than the imposed limit, the result set is truncated!
    "limit": 25,         # the limit imposed on the query results (default
is 25). Use the limit parameter in the request to set a higher limit.
    "offset": 0,         # the offset applied on the query results
(default i 0)
    "currentPage": 1,    # the current "page" (calculated from limit
/offset)
    "size": 20,          # the size of the current result set (should be
same as count or limit)
    "data": [],          # the list of query result objects
    "messages": [],      # any server messages
    "metaData": {},      # any server metadata
}
```

Each query result has the following format:

```
{
    "rrclass": "in",                      # the DNS record class
    "rrtype": "a",                        # the DNS record type
    "query": "cnn.com.",                  # the DNS record query part
    "answer": "157.166.255.19",           # the DNS record answer part
    "firstSeenTimestamp": 1340308340000,  # the first registered timestamp
for this record
    "lastSeenTimestamp": 1377520248000,   # the last registered timestamp
for this record
    "maxTtl": 300,                        # the maximum TTL observed for
this record
    "minTtl": 300,                        # the minimum TTL observed for
this record
    "times": 675,                         # the number of times this
record has been observed
    "tlp": "white",                       # the TLP of this record. Public
records have TLP "white"
    "customer": null,                     # the customer owning this
record. Public records have value null.
    "createdTimestamp": 0,                # always returns 0
    "lastUpdatedTimestamp": 0,            # always returns 0
}
```

Examples

Query	Description
`https://api.mnemonic.no/pdns/v3/cnn.com`	Query for any record concerning the domain "cnn.com" Includes both DNS records for queries for cnn.com and DNS records returning the value "cnn.com" Does not contain subdomains of cnn.com
`https://api.mnemonic.no/pdns/v3/cnn.com?limit=0`	Query for any record concerning the domain "cnn.com". Return all results, no matter how big the resultset is.
`https://api.mnemonic.no/pdns/v3/cnn.com?rrType=A`	Query for A-records for the domain "cnn.com"
`https://api.mnemonic.no/pdns/v3/cnn.com?rrType=A&rrType=MX`	Query for A-records and MX-records for the domain "cnn.com"
`https://api.mnemonic.no/pdns/v3/157.166.255.19`	Query for any record for the IPv4 address "157.166.255.19"
`https://api.mnemonic.no/pdns/v3/2a04:4e42:400::323`	Query for any record for the IPv6 address "2a04:4e42:400::323"

Output in COFF

The PDNS API also supports the PassiveDNS Common Output Format, as specified in https://datatracker.ietf.org/doc/draft-dulaunoy-dnsop-passive-dns-cof/

To retrieve results in the Common Output Format, please use the endpoint `https://api.mnemonic.no/pdns/v3/cof/<query>`

The query input format is the same as for the simple query endpoint, as described above.

JSON query format

An alternative to a simple GET-string is a POST query with a JSON query format to the URL

```
https://api.mnemonic.no/pdns/v3/search
```

Example

```
curl -X POST https://api.mnemonic.no/pdns/v3/search -d '
{
  "query": "cnn.com",
  "rrClass": [
    "IN"
  ],
  "rrType": [
    "a"
  ],
  "limit": 0,
  "offset": 0
}
```

The POST-query is currently permitting the same parameters as the GET query, but new future parameters or advanced/nested parameters may be added only to the JSON query format as an "advanced format".

Querying if a PDNS record has been seen

If you only need to know whether a record matches the search criteria, you can use the seen endpoint

```
https://api.mnemonic.no/pdns/v3/<query>/seen?parameters
```

where "<query>" is a DNS query or answer string to lookup, for example.

```
curl https://api.mnemonic.no/pdns/v3/<query>/seen?ignoreOwnRecords=false
```

This endpoint will return a boolean, true or false, indicating whether a matching record exists.

Query parameters

There are multiple query parameters that can be passed along:

Parameter	Possible values	Default value	Description
tlp	white, green, amber, red	white	Only search in records matching any of the specified TLPs. Multiple values can be submitted.
ignoreOwnRecords	true, false	true	Whether to ignore user's customer's own records, defaults to true

Result format

The result format is JSON, and consists of a result container, and a number of results.

The result container has the following format

```
{
    "responseCode": 200, # the response code. Normal responses should
return HTTP code 200
    "count": 0,          # will always be 0 for this endpoint
    "limit": 0,          # will always be 0 for this endpoint
    "offset": 0,         # the offset applied on the query results
(default i 0)
    "size": 20,          # the size of the current result set (should be
same as count or limit)
    "data": {            # the query result
        "exists": false  # a boolean indicating whether any matching record
(s) was found
    },
    "messages": [],      # any server messages
    "metaData": {},      # any server metadata
}
```

Authenticated queries

The PDNS API is publicly available, and does not require authentication.

However, unauthenticated queries are limited to see only public data (TLP white), and are limited to 1000 requests per day (currently).
To acquire higher resource limits, you need to perform authenticated queries.

To request an API key from mnemonic, contact mss@mnemonic.no.

To access private data (granted that mnemonic is collecting PDNS data from your customer), see the PassiveDNS Integration Guide in the Argus API integration documentation.

If you have an API-key, you can a HTTP header to your request:

```
Argus-API-Key: 1234/1/abcd1234ef012
```

Example:

```
curl -H "Argus-API-Key: 1234/1/abcd1234ef012" -X GET https://api.mnemonic.
no/pdns/v3/cnn.com
```

Resource limits

All users are subject to resource limitations.

- Unauthenticated users are limited to 100 requests per minute, and 1000 requests per day.
- Authenticated users are limited according to their agreement with mnemonic.

If you hit the resource limit, Argus will return a 402 error, with the following JSON response:

```
{
    "responseCode": 402,
    "data": null,
    "messages": [
        {
            "message": "Resource limit exceeded",
            "messageTemplate": "resource.limit.exceeded",
            "type": "ACTION_ERROR"
        }
    ],
    "metaData": {
        "millisUntilResourcesAvailable": 558
    }
}
```

The resource limit is calculated both per minute and per day.

- If you reach the per-minute (short-term) resource limit, you will typically be rejected for a short period, to save resources on our end.
- If you reach the per-day (long-term) resource limit, you will typically be rejected for a period up to 24 hours. This means you have exhausted the data quota granted by mnemonic.

If you find yourself reaching the per-day resource limit a lot, you may want to request a higher resource quota from mnemonic

Please use the metadata key "millisUntilResourcesAvailable" to let your client back off gracefully.