

DNS Response Policy Zones

History, Overview, Usage and Research

Hugo M. Connery
Technical University of Denmark

©
License: CC-BY-SA 3.0 [2]

June 15, 2013

Contents

1	Introduction	1
1.1	DNS RPZ Portal	1
2	Background	1
2.1	DNS: A Very Brief Introduction	1
2.2	rDNS: Normal Operation	1
2.3	Reputation Data	2
2.4	Malware	2
2.5	Conficker	2
2.6	Generalisation and Domain Name Reputation	3
3	RPZ	3
3.1	Zones and Zone Files	4
3.2	Policy	4
3.3	Operation	4
3.4	Example Config	5
3.4.1	named.conf	5
3.4.2	local.rpz	6
4	Implementing RPZ: A Trial	7
4.1	Goals	7
4.2	Architecture	8
4.2.1	Resolvers	8
4.2.2	Policy	8
4.3	Politics	8
4.4	Results	8
4.4.1	Summary	9
5	Phishing Attack Defence	9
6	Response Policy Zone Log Analysis: RPZLA	9
6.1	Overview	10
6.2	Notes on Observed Data	10
6.3	Timing Analysis	10
7	Future Directions	10
7.1	Wider Case Study	10
7.2	Phishing Attack Defence	10
7.3	RPZLA	11
	Appendices	13
A	RPZ RRset Matching	13
B	RPZ Policies	14

C Timing Analysis	14
C.1 Raw Data	14
C.2 Difference Analysis	18

List of Tables

1	Single Client RPZ Timing Analysis	14
2	Single Client RPZ Bucket Analysis	18

1 Introduction

DNS Response Policy Zones (RPZ [5]) is an extension for recursive Domain Name System (DNS) [3] resolvers. RPZ has been implemented in recent BIND [7] name server software.

A brief overview of RPZ, and its origins are presented, followed by a description of a case study of its use in a production environment. A toolset (RPZLA) for displaying and recording log data produced when using RPZ is described.

Observations about using RPZLA for identifying infected systems are presented, along with future research directions.

1.1 DNS RPZ Portal

A portal [1] has been established to provide access to various resources about RPZ.

2 Background

2.1 DNS: A Very Brief Introduction

The Domain Name System is essentially a hierarchical database distributed across the internet consisting of name/value pairs.

One of its primary roles is to translate between human readable domain names (e.g `www.env.dtu.dk`) and IP addresses.

This service, of finding an address for a domain name, is facilitated by two services: authoritative DNS (aDNS) and recursive DNS (rDNS) [4].

aDNS's task is to provide answers to questions about domains for which it is authoritative. For example, say an aDNS is authoritative for the `env.dtu.dk` domain, then one can expect an authoritative answer about `www.env.dtu.dk`, but no useful information about `www.foo.org`.

rDNS's job is to use the DNS hierarchy to find answers. Should a client ask about `www.env.dtu.dk` an rDNS will either generate a response from cache or will perform a fully recursive query asking about the name, starting with the root (.) zone's aDNS, then the `dk.` domain's aDNS, then `dtu.dk.` domain's aDNS, then `env.dtu.dk.` domain's aDNS, and finally respond to the client with the answer obtained.

DNS records (the name/value pairs) come with information describing how long the record's validity persists. rDNS uses this to ensure that records in cache are valid. Being able to re-use 'answers' obviously saves much effort.

2.2 rDNS: Normal Operation

Under normal operation (without RPZ) an rDNS is tasked with accepting queries and responding with either the authoritative answer, or that an error occurred. An authoritative answer may either be an answer with data, or that the domain does not exist.

Within normal operation, there is no flexibility given to an rDNS. It must just do its best to find an authoritative answer and provide it, or report some error.

2.3 Reputation Data

In the fight against Unsolicited Bulk Email (UBE) [8], also known as spam, an opportunity for the internet community to assist other members of the community was developed by Paul Vixie [12]. The project was MAPS [6] (SPAM spelt backwards: Mail Abuse Prevention System), allowing community members to share information about known internet systems (email servers) that delivered only spam. By sharing and combining their reputation knowledge, the community created a improved defence.

There were legal repercussions. Spam was not illegal (certainly not in all jurisdictions) and some businesses based their livelihood on delivering spam.

Despite legal challenges, the mechanism proved valuable. An industry was born in which organisations could pay other organisations to receive reputation data about email servers to protect their personnel's mailboxes receiving unwanted and/or malicious email.

2.4 Malware

Technically skilled and creative people have been devising ways to induce computers to perform unrequested actions since the emergence of computers themselves. With the growth of the internet, its number of connected systems, and the basis of most early (and still many) internet protocols not taking system security as a fundamental concern, the ability for these people to cause unexpected behaviour on computers increased dramatically.

The purpose of these pieces of software turned from advertising the skill of the author, to practical jokes, and continued to the enslavement of the computer which had performed the 'unrequested action'.

Mechanisms for the delivery of these unrequested softwares have also evolved, often using internet protocols (HTTP, FTP, SMTP, etc.), and also using transportable digital storage devices (Floppy disks, USB storage devices etc.).

The control of thousands, or millions, of enslaved computers creates a powerful resource for use in modern societies that use the internet as an augmentation of, or replacement for, many previously non-computer assisted actions (communications, banking etc.). Thus, the collective enslavement of large numbers of computers becomes a potentially profitable tool for use in spamming, fraud, blackmail and other activities.

Current terminology for an enslaved computer is a bot [9] (contraction of robot) or zombie. A collection of bots is known as a botnet [10] (as the collective control is enabled by a network, very commonly the Internet).

Various groups, often with malicious purpose, compete amongst themselves for the identification and infection of computers that are vulnerable to become enslaved, so as to increase the size of their botnet(s).

2.5 Conficker

A botnet, named Conficker [11], was identified in late 2008. It was a well crafted and virulent, self propagating piece of malware (a worm) to enslave computers. The actions of infected systems could be controlled by a central authority (thus a botnet). It grew to several million infected computers.

During remedial efforts to control the spread of Conficker, Paul Vixie was requested to assist neutralising the threat of Conficker (to sinkhole the botnet). At that time, Conficker was known to exist, but not *known* to have performed any malicious action. It was a ticking bomb of several million internet connected computers.

Analysis of Conficker showed that it was programmed to use dates and times to form domain names which a bot would then attempt to contact to receive instruction. These were what are known as *command and control* domains.

To sinkhole the botnet, these domains were registered ahead of time to be in the control of the 'defenders' rather than others who may have been behind Conficker. Initial work required the registration of 500 domain names per day, across some generic top level domains (TLD).

Thereafter, the authors/controllers of Conficker, unable to use DNS to control the bots, re-used the initial attack vector (the owners of the compromised systems has still not protected their systems from the threat which first saw them infected) to re-infect the systems with an update of the Conficker worm software. This used 50 000 domain names per day, that were spread across not just generic top level domains, but also many different country code TLD's (ccTLD).

Countering this 'misuse' of the DNS by the Conficker malware called for a new strategy, as forward registering 50 000 domains per day across many TLD's and ccTLD's is expensive, and often unfeasible due to language barriers and variations in registration mechanisms.

2.6 Generalisation and Domain Name Reputation

As DNS Registrars are a business, and that those businesses exist globally, and that the number of domains required to be 'forward registered' or to be 'not allowed to be registered' was essentially limitless, the only available solution to the problem was within the DNS system itself, rather than with the registrars. Paul Vixie, a founder of Internet Systems Consortium [15], the maintainers of BIND (the most common name server software used on the internet), chose to describe a proposal [14] for the use of reputation data for domain names themselves, and to have that implemented in BIND.

Due to the publication of the specification for RPZ, the mechanism can be implemented by any provider of DNS software, should they wish.

Again, an industry of providing reputation data, this time of domains, rather than of email servers, has been created.

3 RPZ

RPZ has been available in BIND at least from version 9.8.1 (available in late 2010). Numerous improvements to the initial offering have been made in the interim, usually improving speed of operation, but recently offering new features. The initial mechanism's function remains unchanged, which is the focus of this section. New features are beyond the scope of this article.

3.1 Zones and Zone Files

BIND uses zone files to store information (resource records) which describe what it knows about DNS data. aDNS servers describe the domains about which they are authoritative in zone files. rDNS's usually only have zone files for fixed standardized domains (e.g localhost), such that they may directly respond to requests for this standardised information, rather than having to use network bandwidth and systems resources of the aDNS's to acquire the information on behalf of their clients.

RPZ uses zone files to define resource records which are to be *approved and/or filtered*. Four types of record types are supported to allow different matching mechanisms.

The most basic matching mechanism is for the resource record to specify a domain which matches the domain being queried.

See Appendix [B] for a fuller description, or the various matching mechanisms.

An rDNS configured to use RPZ uses locally available zones. However, these zones can be obtained from variable sources using zone file replication mechanisms. This allows for the use of zones that are defined in house, or are provided by external providers.

When external zones are used, incremental zone transfer (IXFR) is commonly used to reduce network traffic. Zones of reputation data are often updated frequently (i.e every few minutes) and often define hundreds of thousands of domains, resulting in zone files that can contain millions of lines of text. IXFR enables updates to transfer only differences in zone file data, greatly improving performance.

3.2 Policy

An rDNS has some defined zones. RPZ allows the description of policy on these zones, if they are to be used for RPZ.

Each zone file defines a collection of resource records. Individual records can specify any of various policy choices.

The zone's response policy specification can provide a policy for *all records* in the zone.

Commonly used zone policies are:

- NXDOMAIN (the default): indicate to the client that the domain does not exist
- CNAME xxx.yyy.zzz: tell the client to visit the xxx.yyy.zzz domain instead (redirect)
- PASSTHRU: explicitly exempt a domain from being affected by other policies. i.e perform normal operation; give the 'true' answer

See Appendix [B] for a fuller description.

Please consult the RPZ specification [14] for precise definitions.

If one uses a redirect (CNAME), this is known in RPZ parlance as a *walled garden*.

3.3 Operation

The normal operation of an rDNS upon receiving a query from a client is:

1. if the answer is in cache and still valid, provide that

2. search local zones for an answer
3. perform a full recursive operation querying aDNS' and provide that answer

RPZ inserts its data and mechanism into step 2, by providing 'local zones', either obtained from external sources, or local files.

If a match is found in these RPZ zones, the normal 'full recursive' behaviour is prevented (unless PASSTHRU), and the policy specified, either in the record (GIVEN), or at the zone file level (all others) takes force, replacing the normal operation of full recursion, and substituting it with another answer based on the configured policy choice.

3.4 Example Config

The following configs are redacted from real configs that have been in use since 2012-09-16.

The `named.conf` uses both a locally defined RPZ zone (provided) and Spamhaus' RPZ data (redacted).

This config is provided as an example, not a recommendation.

3.4.1 `named.conf`

```
// the people that we will serve
acl clients
{
    // IPv4
    127.0.0.0/8;
    192.168.0.0/16;
    // routable IPv4 networks redacted
    // IPv6
    ::1;
    // routable IPv6 networks redacted
};

// redact other acls and logging definitions

options{
    allow-recursion { clients; };
    allow-query-cache { clients; };
    allow-query     { clients; };

    listen-on port 53 { any; };
    listen-on-v6 port 53 { any; };

    directory          "/var/named";
    dump-file           "/var/named/data/cache_dump.db";
    statistics-file    "/var/named/data/named_stats.txt";
    memstatistics-file "/var/named/data/named_mem_stats.txt";
```



```
version "None of your business";

response-policy
{
    zone "local.rpz"          policy CNAME warning.domain.org;
    zone "rpz.spamhaus.org"  policy CNAME warning.domain.org;
};

};

// The locally defined rpz zone
zone "local.rpz" in
{
    type master;
    file "local.rpz.zone";
    allow-query { clients; };
};

// Spamhaus' data
zone "rpz.spamhaus.org" in
{
    type slave;
    masters {
        // redacted: spamhaus rpz DNS'
    };
    file "rpz.spamhaus.org.zone";
    allow-query { clients; };
};
```

3.4.2 local.rpz

This is real data of phishing attack defence. No redaction.

```
$TTL      86400
@         IN SOA  @           root (
                                20130050701    ; serial
                                3H              ; refresh
                                15M            ; retry
                                1W             ; expiry
                                1D )           ; minimum
                                IN NS         LOCALHOST.
; Our locally configured nasty domains
; The first record is a local test case. A non-existent domain to
; be used for testing RPZ.
nastynasty.com      IN CNAME      .
*.nastynasty.com    IN CNAME      .
;
; Phishing domains
```

```
;
; phishing scam (2012-10-29)
webmail-danskebank.com      IN CNAME      .
*.webmail-danskebank.com    IN CNAME      .
; phishing scam (2012-10-31)
ws21.tijdelijke-url.nl     IN CNAME      .
*.ws21.tijdelijke-url.nl   IN CNAME      .
; phishing scam (2012-12-06)
edryecc.com.au             IN CNAME      .
*.edryecc.com.au           IN CNAME      .
; phishing scam (2012-12-11)
shilohstreet.com          IN CNAME      .
*.shilohstreet.com         IN CNAME      .
; more phishing (2012-12-19)
cobern-gaming.nl           IN CNAME      .
*.cobern-gaming.nl         IN CNAME      .
; more phishing (2013-01-16)
naddeoporte.it             IN CNAME      .
*.naddeoporte.it           IN CNAME      .
; more phishing (2013-01-24)
alfakitap.com              IN CNAME      .
*.alfakitap.com            IN CNAME      .
; more phishing (2013-04-16)
lancapris.com              IN CNAME      .
*.lancapris.com            IN CNAME      .
onlinewebshop.net          IN CNAME      .
*.onlinewebshop.net        IN CNAME      .
; more phishing (2013-05-07)
ecdomonline.com            IN CNAME      .
*.ecdomonline.com          IN CNAME      .
```

4 Implementing RPZ: A Trial

The Department of Environmental Engineering (ENV) at the Technical University of Denmark (DTU) decided to trail the use of RPZ in late 2012. Spamhaus graciously provided their RPZ data feed gratis.

A report of the trial has been published by Spamhaus [16].

4.1 Goals

The goal of trialing RPZ was to:

- increase client system security by preventing access to dangerous domains
- increase client system security by informing the community of the risk of visiting dangerous domains

- increase general security by raising the awareness amongst the department's IT staff of dangerous domains

4.2 Architecture

4.2.1 Resolvers

ENV had two in house rDNS' which were used by all internet connected systems. A third was established to subscribe to the Spamhaus data feed. Additionally, a local rpz zone (as shown above) was created, to describe a single fictitious domain to be used in testing.

Once the new third resolver was operational, the two primary rDNS' were set to subscribe to the first as the master for both the Spamhaus data feed (zone) and the locally defined zone.

This architecture was chosen as it reduces the load on Spamhaus, and increases the speed of zone update transfers (1 Gbps local area network).

In accordance with recommendations from ISC, the two existing local resolvers were given an additional 1 GB of RAM (to 2 GB), and the new third resolver was commissioned with 2 GB.

The ability to rapidly access their in memory cache is the highest issue for a resolver, rather than rapidity of CPU operations.

4.2.2 Policy

The default policy for RPZ is `NXDOMAIN`. This provides the defence that one wishes for (clients cannot visit dangerous domains), but does not provide the end user with any manner to differentiate filtering from non-existence.

ENV has an educated community (university) and the author wished to engage the community in understanding why this was being done, and more importantly to let them complain if they felt that they were being inappropriately filtered, or that this technology was impacting their productivity.

Thus, a constant `CNAME` redirection was used. A local website was established, with the departmental logo and information about why people had ended up there, and encouragement to complain if they felt that this was inappropriate.

Thus, a person attempting to visit a filtered domain would end up at the 'warning' web site instead of just being unable to connect to the dangerous domain.

4.3 Politics

A meeting was held to engage the community in the filtering and the goals of the trial. Specifically, the warning web site page was displayed so that should a community member end up visiting it unintentionally, they would already be aware of it and its purpose.

Discussions allowed community concerns to be aired and addressed.

4.4 Results

During the 4 week trial approximately 5 000 attempts to contact dangerous domains were prevented. 75 client systems were defended.

One report of 'inappropriate filtering' was received. It was a false false positive: it was an email harvesting website.

No reports of 'loss of productivity' were received.

Additionally, analysis of the log data (from both the resolvers and the warning site) enabled the identification of malware installed on two systems that had professional anti-virus software installed, with latest updates.

4.4.1 Summary

During the four week trial:

- Approximately 5 000 attempted contacts to dangerous domains were prevented
- 75 systems were protected from contacting dangerous domains
- no loss of productivity reported
- no 'inappropriate filtering' reported
- two infected systems identified
- increased awareness in both the general and IT community was achieved

Due to this entirely positive result, ENV decided to approve the continued use of RPZ and Spamhaus' data. It is still in use 8 months after the trial.

5 Phishing Attack Defence

The university (DTU) came under sustained phishing attacks during and after the RPZ trial. Indeed, the Danish Computer Emergency Response Team (DK*CERT [17]) published on 2013-05-13 that the month of April 2013 had the all time highest number of reported phishing attacks against danish organisations.

The `local.rpz.zone` file was used to provide defence against phishing attacks. The community was asked to forward to their local IT group any email which contained high quality phishing attacks. Given the level of security awareness in the community, there was no need to filter poor quality attacks, just the good ones.

Since the use of RPZ to protect against phishing attacks 68 attempts (as of 2013-05-14 19:12 CEST) to contact phishing sites have been prevented.

6 Response Policy Zone Log Analysis: RPZLA

Once using RPZ, one has the opportunity to use the log data generated by BIND, and a web server, if one is using a 'walled garden' (CNAME redirect).

The availability of this data allows various heuristics to be used on the data to assist in the identification of infected systems, or hazardous human behaviour.

To assist in analysis of this data, the author built the RPZLA toolset [18].

RPZLA is a FOSS project using GPLv3. It is publically published, with no restriction placed upon those who use it, apart from maintaining the GPL type licence.

6.1 Overview

RPZLA comprises a collection of log scrapers, which read BIND recursive resolver logs and/or Apache (walled garden) logs and ship them to a PostgreSQL database, and an Analysis Website which displays the shipped data.

Once established RPZLA builds an ever growing collection of data obtained from the recursive resolvers and 'walled garden' which can be used to identify systems that may be compromised, or systems that are being used by persons in a hazardous manner.

The original diagram, with better resolution, is also viewable [19].

6.2 Notes on Observed Data

The architectural decision to use a walled garden (CNAME) approach was largely based on community awareness, but also on the idea that one could then differentiate between human behaviour (person clicking on a link in a phishing email) and general malware behaviour.

This assumption was poor. Malware sometimes behaves exactly like a browser (HTTP protocol and following DNS redirects).

Indeed, much of the data indicates that community members are visiting news (or similar) web sites that are including malicious advertisements. Thus, DNS records an RPZ hit as the page is loading, but the person is then not clicking on the advertisement and thus no log entry appears in the walled garden. Often one has numerous RPZ hits from numerous dangerous advertisements, at the same time on the same client.

Thus, differentiating between malware and user behaviour is a non-trivial challenge.

6.3 Timing Analysis

The natural distinction between a human and malware is that the malware is automated, and will thus often display regular behaviour in time. This is somewhat obfuscated by browser add-ons which regularly contact dangerous domains, or by Java Script which cause regular page reloads. Browser add-ons which regularly contact dangerous domains are malware embedded in a browser. Some may consider Java Script which regularly reloads a page as malware.

Humans do not visit domains with chronologically tight regularity.

The data of a timing analysis is presented in Appendix [C].

7 Future Directions

7.1 Wider Case Study

There are currently at least 5 commercial producers of RPZ zones. This study was based on the use of just Spamhaus' data. It may be possible to engage other providers to participate in a study which would include more than one commercial zone list.

7.2 Phishing Attack Defence

The current process of the community forwarding phishing attack emails to the IT group to create a defence is poor. Its better than nothing, but it relies on the awareness of the

IT group and their manual effort.

A better strategy may be to build a more automated mechanism.

A delivery point for the community can be created, to which they forward phishing attacks. The delivery point scans the delivered email and removes a known whitelist of domains and then forwards the redacted email, or a list of extracted (non-whitelist) domains to Spamhaus' engine.

Thus, attacks are rapidly redacted and forwarded to Spamhaus so that we can receive their updated domain list after their engine analyses the emails.

7.3 RPZLA

RPZLA is currently a data display engine targetted at engineers.

To use it to identify potentially compromised systems involves human effort (quantity of queries, timing differences etc.)

RPZLA could become a more 'where should I place effort' focussed user interface that identifies potentially infected systems via a collection of heuristics, and highlights suspicious systems.

References

- [1] RPZ Community,
DNS RPZ Portal
<https://dnsrcpz.info>
- [2] Creative Commons,
Creative Commons Attribution-ShareAlike 3.0 Unported
<http://creativecommons.org/licenses/by-sa/3.0/legalcode>
- [3] Wikipedia,
Domain Name System
http://en.wikipedia.org/wiki/Domain_Name_System
- [4] Wikipedia,
Domain Name System, Recursive and Caching Name Server
http://en.wikipedia.org/wiki/Domain_Name_System#Recursive_and_caching_name_server
- [5] Wikipedia, including this author,
Response Policy Zone,
http://en.wikipedia.org/wiki/Domain_Name_System
- [6] Wikipedia
Mail Abuse Prevention System,
http://en.wikipedia.org/wiki/Mail_Abuse_Prevention_System
- [7] Internet Systems Consortium,
BIND Documentation,
<https://www.isc.org/software/bind/documentation>
- [8] Spamhaus
The Definition of SPAM,
<http://www.spamhaus.org/consumer/definition/>
- [9] Wikipedia
Internet Bot,
http://en.wikipedia.org/wiki/Internet_bot
- [10] Wikipedia
Botnet,
<http://en.wikipedia.org/wiki/Botnet>
- [11] Conficker Working Group
Introduction (to Conficker),
<http://www.confickerworkinggroup.org/wiki/pmwiki.php/ANY/Introduction>
- [12] Wikipedia,
Paul Vixie,
http://en.wikipedia.org/wiki/Paul_Vixie

- [13] Paul Vixie, Conference Video,
Taking Back the DNS,
<http://www.youtube.com/watch?v=k5DbNgEXDHo>
- [14] Internet Systems Consortium,
DNS Response Policy Zones, Specification, Format 3,
<https://kb.isc.org/getAttach/21/AA-00512/rpz.txt>
- [15] Internet Systems Consortium,
Internet Systems Consortium,
<https://www.isc.org/>
- [16] Hugo Connery,
A Case Study at DTU Environment
- [17] Danish Computer Emergency Response Team,
Danish Computer Emergency Response Team,
<https://www.cert.dk>
- [18] Hugo Connery,
Response Policy Zone Log Analysis,
<https://github.com/yesxorno/rpzla>
- [19] Hugo Connery,
Response Policy Zone Log Analysis, Pictorial Overview,
<https://github.com/yesxorno/rpzla/blob/master/doc/Pictorial-Overview.odg?raw=true>

Appendices

A RPZ RRset Matching

RPZ provides 4 types of matching a query:

- **domain (qname)**: match a domain, and/or all subdomains of a domain, based on the name that was queried
- **ip (rpz-ip)**: match any domain whose lookup address matches an address, or lies within a collection of related addresses (CIDR)
- **ns (rpz-nsdname)**: match all domains for which an authoritative name server, identified by its domain (host) name, is authoritative
- **ns (rpz-nsip)**: match all domains for which an authoritative name server, identified by its ip address, is authoritative

Please consult the RPZ specification [14] for precise definitions.

B RPZ Policies

Zone level policy specifications are:

- NXDOMAIN (the default): indicate to the client that the domain does not exist
- NODATA: indicate to the client that there is no answer to their question
- GIVEN: honour whatever is specified at the resource record level
- PASSTHRU: explicitly exempt a domain from being affected by other policies. i.e perform normal operation; give the 'true' answer
- CNAME xxx.yyy.zzz: tell the client to visit the xxx.yyy.zzz domain instead (redirect)

Please consult the RPZ specification [14] for precise definitions.

C Timing Analysis

C.1 Raw Data

Here are the RPZ 'hits' of a single client to a single domain with date/time stamps, and the differences in time between successive events, rounded down to 10 second intervals.

Table 1: Single Client RPZ Timing Analysis

Date/Time	Difference	Date/Time	Difference
2013-05-03 12:45:05.985	N/A	2013-05-05 23:20:53.356	2160
2013-05-03 13:15:14.198	1800	2013-05-05 23:30:20.996	560
2013-05-03 13:30:16.426	900	2013-05-05 23:36:48.808	380
2013-05-03 14:15:24.620	2700	2013-05-05 23:45:25.439	510
2013-05-03 14:45:29.251	1800	2013-05-05 23:52:47.140	440
2013-05-03 15:00:31.816	900	2013-05-06 00:08:42.770	950
2013-05-03 15:30:36.445	1800	2013-05-06 00:15:33.463	410
2013-05-03 16:15:46.243	2700	2013-05-06 00:24:38.412	540
2013-05-03 16:30:49.315	900	2013-05-06 00:30:37.384	350
2013-05-03 16:45:52.233	900	2013-05-06 00:40:35.358	590
2013-05-03 17:31:04.298	2710	2013-05-06 01:30:59.335	3020
2013-05-03 18:01:11.023	1800	2013-05-06 01:44:25.020	800
2013-05-03 19:01:20.553	7200	2013-05-06 01:46:04.813	90
2013-05-03 19:31:24.864	1800	2013-05-06 02:00:23.456	850
2013-05-03 19:46:27.930	900	2013-05-06 02:31:23.757	1860
2013-05-03 20:01:31.960	900	2013-05-06 02:46:29.124	900
2013-05-03 20:16:35.373	900	2013-05-06 02:48:11.430	100
2013-05-03 21:16:55.957	7220	2013-05-06 03:04:14.941	960
2013-05-03 21:47:04.504	1800	2013-05-06 03:16:37.205	740
2013-05-03 22:02:07.547	900	2013-05-06 03:31:42.334	900
2013-05-03 22:17:09.849	900	2013-05-06 03:36:11.749	260

2013-05-03 22:32:13.524	900	2013-05-06 03:46:47.660	630
2013-05-03 23:02:22.132	1800	2013-05-06 03:52:07.922	320
2013-05-03 23:17:23.821	900	2013-05-06 04:40:00.618	2870
2013-05-04 00:32:37.110	8110	2013-05-06 04:55:57.582	950
2013-05-04 01:02:43.625	1800	2013-05-06 05:27:50.035	1910
2013-05-04 01:17:46.324	900	2013-05-06 05:32:14.421	260
2013-05-04 02:02:54.258	2700	2013-05-06 05:43:45.724	690
2013-05-04 03:03:21.274	7220	2013-05-06 05:47:19.453	210
2013-05-04 03:33:28.050	1800	2013-05-06 05:59:41.772	740
2013-05-04 04:03:34.274	1800	2013-05-06 06:02:23.934	160
2013-05-04 04:48:46.566	2710	2013-05-06 06:15:43.196	790
2013-05-04 07:04:18.158	15330	2013-05-06 06:17:29.442	100
2013-05-04 07:19:22.972	900	2013-05-06 06:31:40.034	850
2013-05-04 07:49:31.369	1800	2013-05-06 06:32:34.257	50
2013-05-04 08:04:32.954	900	2013-05-06 06:47:36.259	900
2013-05-04 08:19:35.347	900	2013-05-06 06:47:41.267	0
2013-05-04 08:49:43.502	1800	2013-05-06 07:17:49.441	1800
2013-05-04 09:50:01.032	7210	2013-05-06 07:32:55.654	900
2013-05-04 11:20:17.926	9010	2013-05-06 07:48:03.716	900
2013-05-04 11:35:21.551	900	2013-05-06 07:51:21.117	190
2013-05-04 11:50:24.905	900	2013-05-06 08:23:14.454	1910
2013-05-04 12:20:30.271	1800	2013-05-06 08:39:09.720	950
2013-05-04 12:50:35.383	1800	2013-05-06 08:55:06.669	950
2013-05-04 13:20:43.409	1800	2013-05-06 09:11:01.420	950
2013-05-04 13:35:46.305	900	2013-05-06 09:18:32.401	450
2013-05-04 13:50:49.770	900	2013-05-06 09:33:36.719	900
2013-05-04 14:05:52.938	900	2013-05-06 09:42:53.876	550
2013-05-04 14:20:56.024	900	2013-05-06 09:48:41.838	340
2013-05-04 14:51:05.609	1800	2013-05-06 10:14:47.064	1560
2013-05-04 15:51:17.994	7210	2013-05-06 10:18:52.147	240
2013-05-04 16:06:20.969	900	2013-05-06 10:33:56.865	900
2013-05-04 17:51:47.379	9920	2013-05-06 10:46:39.510	760
2013-05-04 18:06:51.101	900	2013-05-06 10:49:01.723	140
2013-05-04 18:21:54.466	900	2013-05-06 11:02:36.197	810
2013-05-04 18:23:11.358	70	2013-05-06 11:18:31.968	950
2013-05-04 19:07:05.963	2630	2013-05-06 11:29:52.997	680
2013-05-04 19:10:57.682	230	2013-05-06 11:34:38.831	280
2013-05-04 19:22:07.829	670	2013-05-06 11:50:34.279	950
2013-05-04 19:26:54.161	280	2013-05-06 12:06:29.977	950
2013-05-04 19:58:44.749	1910	2013-05-06 12:22:29.583	950
2013-05-04 20:07:25.918	520	2013-05-06 12:30:10.535	460
2013-05-04 20:46:31.489	2340	2013-05-06 12:38:26.436	490
2013-05-04 20:52:42.941	370	2013-05-06 12:54:26.591	960
2013-05-04 21:07:45.648	900	2013-05-06 13:10:22.627	950
2013-05-04 21:18:21.741	630	2013-05-06 13:30:25.550	1200
2013-05-04 21:34:17.527	950	2013-05-06 13:58:13.170	1660

2013-05-04 21:52:55.919	1110	2013-05-06 14:46:03.007	2860
2013-05-04 22:38:05.679	2700	2013-05-06 15:33:51.640	2860
2013-05-04 22:54:03.839	950	2013-05-06 15:49:48.903	950
2013-05-04 23:25:56.542	1910	2013-05-06 16:21:41.017	1910
2013-05-04 23:38:21.245	740	2013-05-06 16:53:34.201	1910
2013-05-04 23:53:25.632	900	2013-05-06 17:25:26.761	1910
2013-05-04 23:57:47.987	260	2013-05-06 17:41:24.065	950
2013-05-05 00:08:28.438	640	2013-05-06 17:57:22.080	950
2013-05-05 00:13:42.992	310	2013-05-06 18:13:17.846	950
2013-05-05 00:38:34.858	1490	2013-05-06 18:29:14.637	950
2013-05-05 01:01:29.465	1370	2013-05-06 18:45:11.565	950
2013-05-05 01:53:49.977	3140	2013-05-06 19:33:04.780	2870
2013-05-05 02:21:12.806	1640	2013-05-06 19:49:01.666	950
2013-05-05 02:23:55.663	160	2013-05-06 20:04:57.935	950
2013-05-05 02:37:10.482	790	2013-05-06 20:20:57.160	950
2013-05-05 02:53:06.447	950	2013-05-06 20:36:53.784	950
2013-05-05 02:54:02.879	50	2013-05-06 20:52:50.503	950
2013-05-05 03:25:05.633	1860	2013-05-06 21:08:46.523	950
2013-05-05 04:44:49.829	8380	2013-05-06 21:24:43.009	950
2013-05-05 05:32:36.506	2860	2013-05-06 21:40:39.281	950
2013-05-05 05:39:54.235	430	2013-05-06 21:56:35.531	950
2013-05-05 05:48:33.094	510	2013-05-06 22:28:30.585	1910
2013-05-05 06:04:27.664	950	2013-05-06 22:44:27.415	950
2013-05-05 06:20:23.611	950	2013-05-06 23:00:23.790	950
2013-05-05 06:25:05.064	280	2013-05-06 23:16:21.430	950
2013-05-05 06:40:07.827	900	2013-05-06 23:32:17.541	950
2013-05-05 06:52:16.250	720	2013-05-06 23:48:13.716	950
2013-05-05 07:10:14.738	1070	2013-05-07 00:04:10.694	950
2013-05-05 07:25:17.773	900	2013-05-07 00:20:09.123	950
2013-05-05 07:40:02.417	880	2013-05-07 00:36:04.922	950
2013-05-05 07:55:59.487	950	2013-05-07 01:24:22.769	2890
2013-05-05 08:10:27.754	860	2013-05-07 02:12:16.618	2870
2013-05-05 08:27:56.393	1040	2013-05-07 02:28:16.943	960
2013-05-05 08:40:34.137	750	2013-05-07 02:44:12.109	950
2013-05-05 08:43:54.036	190	2013-05-07 03:00:08.656	950
2013-05-05 08:59:52.538	950	2013-05-07 03:16:05.216	950
2013-05-05 09:10:40.973	640	2013-05-07 03:32:02.672	950
2013-05-05 09:31:50.682	1260	2013-05-07 03:47:58.560	950
2013-05-05 10:03:44.298	1910	2013-05-07 04:03:54.540	950
2013-05-05 10:35:35.857	1910	2013-05-07 04:35:52.950	1910
2013-05-05 10:51:32.025	950	2013-05-07 04:51:50.074	950
2013-05-05 10:56:08.861	270	2013-05-07 05:07:46.704	950
2013-05-05 11:11:13.382	900	2013-05-07 05:23:42.967	950
2013-05-05 11:23:29.280	730	2013-05-07 05:39:40.136	950
2013-05-05 11:39:28.124	950	2013-05-07 05:55:36.524	950
2013-05-05 11:55:23.928	950	2013-05-07 06:43:41.307	2880

2013-05-05 12:11:19.358	950	2013-05-07 06:59:38.808	950
2013-05-05 12:27:18.182	950	2013-05-07 07:47:42.734	2880
2013-05-05 12:43:13.714	950	2013-05-07 08:03:46.261	960
2013-05-05 12:59:08.968	950	2013-05-07 08:51:37.649	2870
2013-05-05 13:15:05.122	950	2013-05-07 09:39:28.437	2870
2013-05-05 14:02:56.029	2870	2013-05-07 09:55:25.297	950
2013-05-05 15:06:41.238	7420	2013-05-07 10:11:22.878	950
2013-05-05 15:38:34.270	1910	2013-05-07 10:59:22.894	2880
2013-05-05 15:54:36.982	960	2013-05-07 11:15:21.881	950
2013-05-05 16:10:32.753	950	2013-05-07 12:03:20.748	2870
2013-05-05 16:27:55.683	1040	2013-05-07 13:07:12.330	7430
2013-05-05 16:42:25.990	870	2013-05-07 13:23:08.337	950
2013-05-05 16:43:00.104	30	2013-05-07 13:39:05.366	950
2013-05-05 16:58:21.404	920	2013-05-07 13:55:01.834	950
2013-05-05 17:13:09.213	880	2013-05-07 14:10:59.281	950
2013-05-05 17:14:17.733	60	2013-05-07 14:26:55.124	950
2013-05-05 17:28:15.028	830	2013-05-07 16:18:43.667	10300
2013-05-05 17:30:14.373	110	2013-05-07 17:22:41.093	7430
2013-05-05 17:46:09.545	950	2013-05-07 17:38:40.054	950
2013-05-05 18:49:53.507	7420	2013-05-07 17:54:36.645	950
2013-05-05 19:21:45.007	1910	2013-05-07 18:10:33.761	950
2013-05-05 19:43:51.973	1320	2013-05-07 18:26:31.167	950
2013-05-05 19:53:39.163	580	2013-05-07 19:30:23.929	7430
2013-05-05 21:29:16.005	9330	2013-05-07 20:45:31.700	8100
2013-05-05 21:29:27.425	10	2013-05-07 21:17:29.607	1910
2013-05-05 21:44:32.964	900	2013-05-07 21:33:26.449	950
2013-05-05 22:01:11.392	990	2013-05-07 21:49:26.046	950
2013-05-05 22:17:07.715	950	2013-05-07 22:21:21.459	1910
2013-05-05 22:29:47.896	760	2013-05-07 22:37:17.664	950
2013-05-05 22:33:04.450	190	2013-05-07 22:53:14.640	950
2013-05-05 22:44:52.861	700	2013-05-07 23:25:07.392	1910

C.2 Difference Analysis

If one summarises the above data by grouping differences into 10 second 'slots', and disregarding slots with only one or two entries in them, one gets the following picture:

Table 2: Single Client RPZ Bucket Analysis

Rounded Difference	Frequency
190	3
260	3
280	3
740	3
900	37
950	79
960	5
1800	17
1910	16
2700	4
2860	3
2870	7
2880	3
7430	3

General statistics are:

- total readings: 287
- number of (disregarded) single or double value time slots: 85

The peaks at 900 and 950 seconds are instructive. They are mirrored by peaks at double their values (1800 and 1910 seconds).

This is almost certainly not human behaviour. It seems likely that, the system is infected with malware, or a browser is visiting a page with dangerous advertising and is regularly reloading the page, amongst other possibilities.