A Case Study at DTU Environment

DNS: Response Policy Zone

Table of Contents

1 Executive Summary

The Response Policy Zone (RPZ) mechanism, designed, implemented and offered as a standard by Internet Systems Consortium (ISC), offers a new mechanism for defence in the challenging world of internet security.

RPZ goes beyond signature based identification of specific pieces of digital 'malware' and uses 'reputational data' to identify internet domain names that are hosting aggressive content.

This concept of using reputational data is not new. Reputational data was first used by P. Vixie (now of ISC) for spam defence in the MAPS RBL project in 1996. It has since become an integral part of current spam defence strategies. RPZ employs reputational data for internet domains, rather than email servers.

Using a domain reputational (RPZ) data feed from Spamhaus, DTU Environment implemented the use of RPZ as a trial. Within the 4 week trial, approximately 5000 attempted contacts to internet sites within domains with a poor reputation were prevented.

The use of RPZ had no impact on productivity, but increased security by:

1. preventing contact with these dangerous sites (attacks on local systems or exfiltration of data)
2 providing identification of locally infected systems (despite their being installed with up-to-date professional anti-virus software and definitions)
3 increasing the breadth of awareness of the local IT organisation to threats

2 Introduction

2.1 ISC and RPZ

RPZ is a facility available with recent versions of the BIND software, maintained and offered by Internet Systems Consortium. It is a technical response from P. Vixie, of ISC, based on his extensive experience with the Domain Name System, and particularly with how DNS is being used by those with criminal intent.

RPZ is essentially a new form of internet filtering, blocking access to all systems within entire internet domain names, where those domain names are believed to exist solely for hosting malicious content.

2.2 Spamhaus

While ISC, through BIND, offer access to the RPZ facility, they leave the provision of the reputational data, upon which automated use of the facility depends, to organisations that specialise in the provision of this reputational data.

DTU Environment had been using Spamhaus reputational data for spam protection when it ran its own email infrastructure. Spamhaus' reputational data feed for spam protection was excellent, and DTU Environment contacted Spamhaus to ask about the availability of RPZ reputational data. Spamhaus graciously provided DTU Environment with the data feed gratis.

2.3 DTU Environment

DTU Environment is the Department of Environmental Engineering at the Technical University of Denmark, hereafter known as ENV.

ENV is an international research organisation. It has significant communities from all 6 of the populous continents. As such, we generate internet traffic of a wide variety in language, culture, subject and geographic location. This is significant, in that any automated internet filtering must be agnostic to these variables, but only focus on the actual malicious content.

ENV has a group of IT professionals tasked with providing IT services and defending IT systems. The author leads this IT group.

3 The Trial

Having been excited by the emergence of RPZ, and encouraged by the availability of data from Spamhaus, ENV's management approved a trial of the use of the RPZ facility with Spamhaus' data feed.

3.1 Objectives

The trial's objectives were to determine whether the use of RPZ would:

1. affect organisational productivity
2. improve computer security
3. improve the ability of the IT group to respond to security threats.

3.2 Consultation

As the use of RPZ was internet filtering, the ENV community were consulted so that they would be aware of what was being done, and why.
The community was encouraged to inform the IT group if they felt that they were being inappropriately filtered, or that their productivity was being impacted.

3.3 TechnicalDescription

ENV has two local recursive resolvers that are used by all internet connected computing systems within its networks. A third resolver was established, which would received the RPZ reputational data from Spamhaus. The two primary resolvers would subscribe to the third resolver to receive the RPZ data.

RPZ places no fixed policy choice on the organisation as to what to do with queries to domains with a poor reputation. ENV chose to use a constant CNAME (redirection) to refer all requests to poor domains to a local web site.

This has two advantages:

1. A person who is redirected to the site (away from a potentially dangerous domain) sees a web site with the departmental logo, and information about why they have arrived there, and encouragement to complain if they think that the filtering is inappropriate.
2. All visits to the site are logged, enabling IT personnel to analyse which systems are visiting, where they were trying to get to, and how often this occurs.

The web server was set to perform hostname lookups on each visit to attempt to identify the hostname of the client visiting.

3.4 Results

Results are presented from two perspectivies:

- grouping by site
- grouping by client hostname

Visit records that include a hostname (successful hostname lookup by the web server) come from what are termed 'named clients' in this document. Visit records in which the hostname lookup was not successful came from what are termed 'unnamed clients' and were excluded from the client analysis.

Item	Data
Trial start	2012-09-16
Trial data collection end	2012-10-13
Total Days	28
Total Site visits [1]	**4967**
Total Site visits by named clients	4383
Number of domains blocked	75
Number of named clients visiting [2]	60
Average site visits per day	**173.8**

Notes:

1. These are visits to the local site based upon RPZ redirection. It does not include valid visits to the site (e.g. to deliberately look at the site to learn about RPZ and the trial)
2. This represents the number of clients that were identified via hostname lookup. 584 (4967 – 4383) requests by unnamed clients were made to the site.

3.4.1 Most Filtered Domains

The full list of filtered domain frequencies is listed in Annex 1.

Frequency for individual domains: Top 20 (of 75)

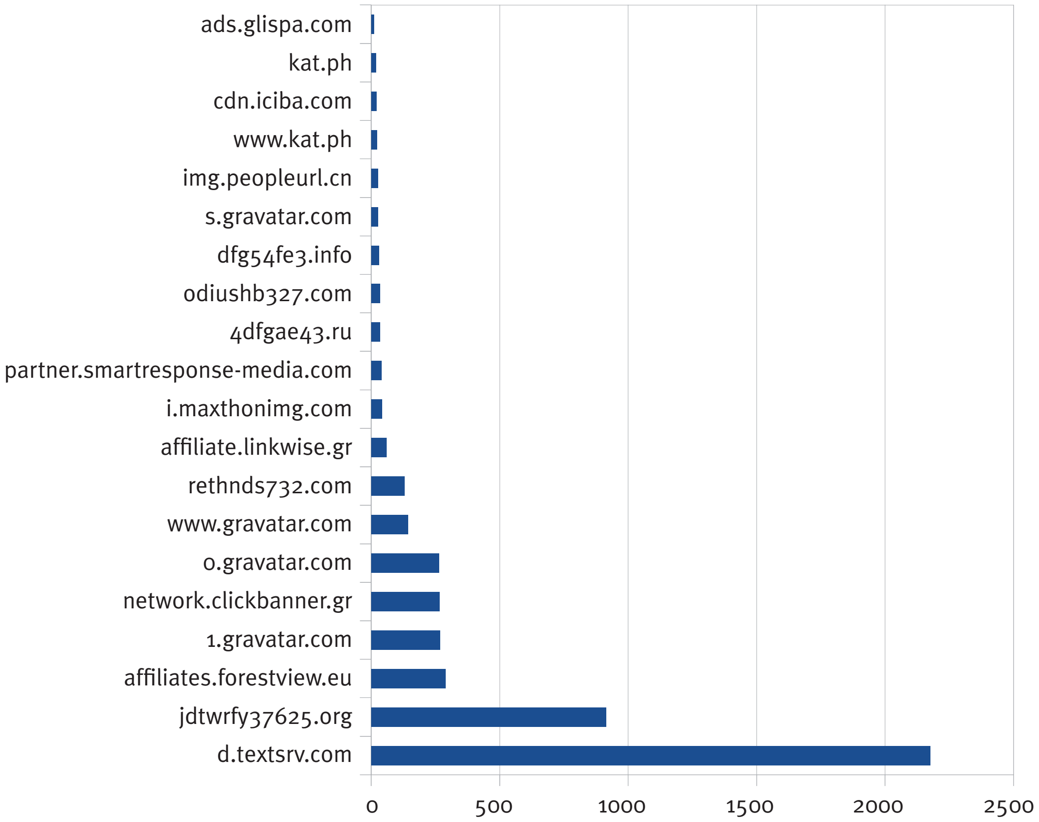

From the full list in Annex 1, it can be seen that:

- The first two domains amounted to 60% of all redirects
- The first eight domains amounted to 80% of all redirects

3.4.2 Most Filtered Clients

The following list shows the frequency of redirected visits to the internal site by the top 20 individual, identifiable client systems (names of the client systems are omitted throughout this report).

A full list of frequencies for nameable clients is listed in Annex 2.

Frequency for individual clients: Top 20 (of 60)

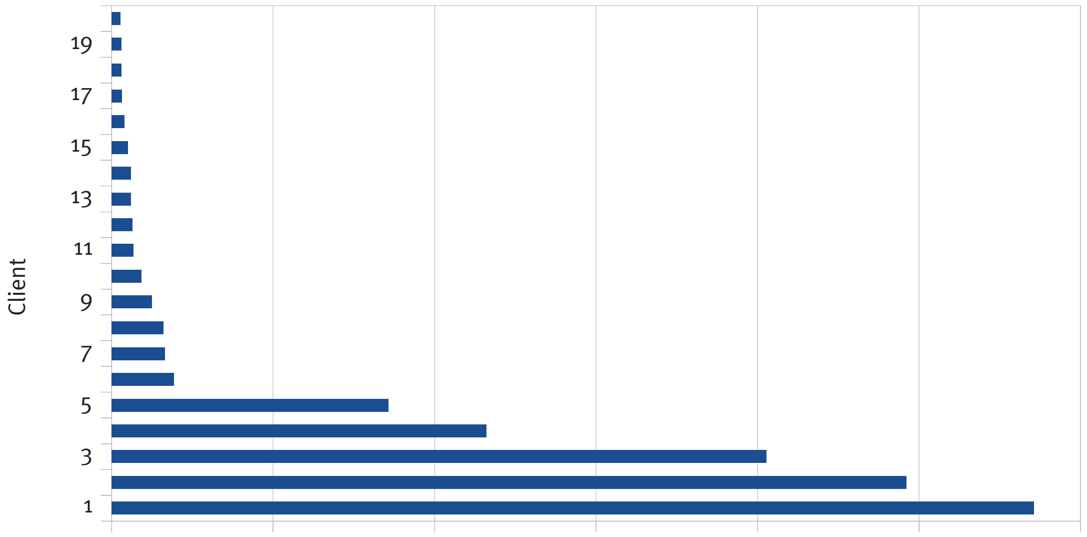

From the full list in Annex 2, one can see that the top 5 clients amount to 70% of all redirections.

4 Conclusions

The trial was a success, with no impact on productivity, or inappropriate filtering.

The trial also identified systems that were infected with malware, despite being installed with professional anti-virus software, and their anti-virus definitions being up to date!

4.1 Complaints/FalsePositives/Productivity

One notification of potentially incorrect filtering was received during the trial. The filtering was valid: the site was an email harvesting web-site.

Thus, there were no inappropriate filtering events reported. There were no reports of loss of productivity.

4.2 Security

The trial gave the ENV IT group a greater insight into the sorts of potentially dangerous domains that people were visiting, and specifically identified the weakness of signature based anti-virus systems.

Several infected client systems were identified during the trial by examining the web site log data. Despite up-to-date professional anti-virus software these systems were infected with malware.

The use of RPZ, and the data received during the trial, has improved the level of security of our systems, and increased the breadth of the local IT organisation's risk awareness.

5 Acknowledgements

The author would like to acknowledge both Internet Systems Consortium, and Spamhaus:

- The provision of the RPZ mechanism continues to highlight ISC's dedication to providing visionary technical solutions to the challenges that face the Internet, and the Domain Name System.
- Spamhaus delivered data and direct, personal support to enable this trial. Spamhaus' data was excellent, and their support personnel were friendly, proactive and professional.

6 References

Document	URL
Taking Back the DNS, Paul Vixie, Blog post (original announcement)	http://www.circleid.com/posts/20100728_taking_back _the_dns/
Response Policy Zones: Taking back the DNS, Paul Vixie, slides	http://www.isc.org/files/TakingBackTheDNSrpz2.pdf
RPZ draft specification,from ISC	ftp://ftp.isc.org/isc/dnsrpz/isc-tn-2010-1.txt
Spamhaus News bulletin on RPZ	http://www.spamhaus.org/news/article/669/
About RPZ, written by the author	https://en.wikipedia.org/wiki/Response_policy_zone
DNS RPZ, Webinar slides from ISC	https://www.isc.org/files/imce/DNSRPZ-2011-03-01- Webinar.pdf

7 Annexes

7.1 Annex 1

A full list of all domains recorded in the access log, along with the frequency of their request.

Domain	Frequency
d.textsrv.com	2176
jdtwrfy37625.org	914
affiliates.forestview.eu	289
1.gravatar.com	268
network.clickbanner.gr	266
0.gravatar.com	264
www.gravatar.com	142
rethnds732.com	130
affiliate.linkwise.gr	60
i.maxthonimg.com	42
partner.smartresponse-media.com	40
odiushb327.com	34
4dfgae43.ru	34
dfg54fe3.info	30
s.gravatar.com	27
img.peopleurl.cn	26
www.kat.ph	22
cdn.iciba.com	20
kat.ph	19
ads.glispa.com	10
img2.imagehyper.com	9
www.inbox.lv	7
www.huotop.com	7
www.hoqotop.com	7
js.miaozhen.com	7
server.netsales.pl	6
medicinepharmacydrug.com	6
partners.shoogloonetwork.com	5
reyy.ucoz.com	4
ipic.staticsdo.com	4
cdn.weather.hao.360.cn	4
2.gravatar.com	4
www.mytwitcam.com	3
www.goo2web.info	3
www.creativediagnostics.net	3
www.converseonmycam.com	3
www.alb-observer.com	3
watchinglivefootball.com	3

v1.jiathis.com	3
unsub-mailing.com	3
partners.linkemann.com	3
medicarecaretab.com	3
mail.inbox.lv	3
clipshardcore.com	3
cdn.globalsurveyfreebies.com	3
c1.youmaker.com	3
ads.yesadvertising.com	3
a4.att.hudong.com	3
2.gravatar.com	2
xoaeryqisj.ru	2
uqmyhiotda.ru	2
serw.myroitracking.com	2
mail.vipoky.com	2
mail.vipmuch.com	2
livekut.ucoz.com	2
khadim.ucoz.com	2
get4cdn.com	2
atmst.net	1
www.myroitracking.com	1
www.ademails.com	1
webmindsmail.com	1
sta.waveca.net	1
static.irs09.com	1
static.csbew.com	1
soikot.net	1
sd.p.360.cn	1
mybooksplace.com	1
list.clk-galaxynewshelp.com	1
gravatar.com	1
digedags.bplaced.net	1
c15.youmaker.com	1
c1.56img.com	1
affiliate.repaymedia.com	1
ads.lzjl.com	1
ads.guava-affiliate.com	1

7.2 Annex 2

A list of the frequencies for all 'named clients' during the trial. There were 584 (of 4967, 11.76 %) requests by unnamed clients.

Rank	Individual Client Frequency	Number of individual clients with this frequency	Cumulative total visits
1	1142	1	1142
2	984	1	2126
3	811	1	2931
4	464	1	3401
5	343	1	3744
6	77	1	3821
7	66	1	3887
8	64	1	3951
9	50	1	4001
10	37	1	4038
11	27	1	4065
12	26	1	4091
13	24	2	4139
14	20	1	4159
15	16	1	4175
16	13	1	4188
17	12	2	4212
18	11	2	4234
19	10	1	4244
20	9	1	4253
21	8	1	4261
22	7	4	4289
23	6	5	4319
24	5	1	4324
25	4	3	4336
26	3	8	4360
27	2	7	4374
28	1	9	4383