DNS Response Policy Zones

Roadmap to Accellerate Adoption

Hugo M. Connery
Based on input from the DNS RPZ Community

June 28, 2013

Contents

1 Introduction

This document presents a brief summary of the state of affairs of the deployment of DNS RPZ and then describes a collection of suggestions to improve uptake of RPZ based on input from the DNS RPZ community.

1.1 Audience

The intended audience are persons what are already involved with RPZ in some way, and thus understand common terminology, abbreviations and related concepts.

2 Background

Despite efforts from senior members of the DNS and security communities (e.g ISC and reputation data providers) the level of awareness in the general community of the strengths of RPZ remains rather low.

Barry Green [1] wrote about the strengths of RPZ, some history of DNS Firewalls, and encouraged people to begin trialling RPZ on 2012-08-20.

The reputation industry's leading organisations are all offering RPZ feeds.

ISC have made numerous improvements since Mr Greene's publication, and some case studies acknowledging the efficacy of RPZ coupled with professional reputation datasets have been published.

3 Roadmap

The following sections list suggestions from the DNS community designed to accelerate RPZ uptake.

It goes without saying that members of the community could take occassions to describe RPZ as opportunity presents.

A community member who has the influence to create opportunities for RPZ to be discussed may consider utilising their influence.

The order of the suggestions listed below is not based on significance.

3.1 Naming

During RPZ'a 2.5 year history it has been known as

- Response Policy Zones
- DNS Firewall

The first leaves anyone who is not intrinsic to the DNS world baffled. The second gives a clearer picture. Perhaps an even better name for the technology can be developed, with a purpose to inspire interest.

3.2 Testimonials

Multiple testimonials (results of trials) are needed.

3.3 Targetted Communications

RPZ is about security via DNS. The target audience are the security people, not the network/DNS people.

3.4 BIND: ESV

An extended support version (ESV) of BIND, which includes response rate limiting (RRL) support and full RPZ support (including all matching and policy rules) could be published by ISC.

As of June 2013, major linux distributions are including versions of BIND which include support for qname type matching. However, support for the other three matching mechanisms (ip, nsdname and nsip) are not entirely supported.

3.4.1 Alternative Repositories

In the interim between basic and full support for RPZ features in BIND, it may be valuable to offer alternative publically open repositories which provide full functionality in relevant distro formats (e.g RPM for RHEL and derivatives, and Deb for Debian and derivatives).

3.5 Data

3.5.1 Availability: Public (free) data

The availability of continuously free data may inspire some to trial the technology. Should professional reputation providers provide continuous access to 'slightly out of date' data, this may encourage the poor, and even the not so plentiful, to trial the technology.

Once convinced of the value, many may decide to invest in obtaining the most recent data.

3.5.2 Data Providers Advertise

Data providers who have an RPZ product could clearly describe it as an RPZ product. This positions the product in the market, and allows potential clients to find more easily discover it.

3.6 Filter Control and Collateral Damage

Organisations that may wish to deploy, need to feel a confidence in their ability to respond to inappropriate filtering.

The community need to provide clear instructions on how an organisation can easily and locally deal with false positives.

3.7 Increased Resolver Support

It is hoped that RPZ will not become a 'BIND' technology, but one that is supported across a wide range of resolver softwares.

3.7.1 FOSS Resolvers: Linux / BSD

Numerous resolvers, both network and local, exist in the FOSS space.

Minimally, support for RPZ is wished for in non-BIND FOSS network resolvers. A key example is Unbound.

Regular outreach, perhaps timed to major Linux distribution release cycles, may be of value, both with non-BIND resolvers and pushing for the inclusion of a feature rich BIND ESV.

3.7.2 Commercial Resolvers

The RPZ community could maintain contact with the commercial resolver vendors to keep RPZ on their radar.

Significant organisations worth maintaining contact with are Nominum and PowerDNS.

3.8 Feedback / PassiveDNS

One challenge for reputation data providers is their lack of access to a local resolver's data. This data may be potentially very useful.

The Security Information Exchange (SIE) could be used for this purpose, with some combination of anonymised local resolver logs and submissions of locally acquired threat intelligence.

Encouraging the use of SIE as a data collection point may prove to be a major, future, industry wide benefit.

3.9 Educational Resources

3.9.1 Upgrading BIND

Small to medium sized organisations that use BIND may require some assistance in upgrading to a recent version of BIND to enable use of RPZ.

A community portal providing educational material and discussion fora could be created. Some basic course material may be:

- Introduction to BIND (basic concepts)
- Best Practices (architecture, spearation of recursion and authority),
- Upgrading (to at least 9.8.2)

This space could then be extended to cover things like RPZ, RRL and DNSSEC.

3.9.2 Deployment Guide

There are a collection of different paths to get from not using RPZ to using RPZ. These could be described both in a general sense, and with technical detail.

One such course of deployment might be like this:

1. Get a resolver that supports RPZ

2. Subscribe to 30 day trials from a few reputation data providers
3. Add the reputation data to the resolver and test that its working
4. Configure to NOT actually block, but to just log the hits against the reputation data
5. Deploy the resolver and watch the log(s). Are you getting hits?

Hopefully, as this point one has seen the both efficiency and efficacy of RPZ and the reputation data, and can then plan a formal roll out.

This may included the following steps:

- Convert 30 day trials to longer agreements
- If you are using multiple resolvers, consider a manner in which the log data of the hits can be centralised and available for review (e.g RPZLA), and setting up sync point(s) to stand between the reputation data provider(s) and the production resolvers.
- Configure production resolvers to use the reputation data, in concert with central logging and/or sync point as desired.
- Watch the data and respond

3.9.3 Recommended Config

For organisations deploying RPZ with BIND for its defensive strengths, it may be valuable to provide a somewhat general recommended configuration. This would include general best practices, (named.conf.local not named.conf etc.) and recommendations for RPZ zone (white/black) ordering.

3.10 Utilities

Utilities to add value to RPZ may not necessarily have direct impact on DNS uptake, but may play a part.

3.10.1 Log Collection and Display

Gathering of the various resolver logs (and possibly walled garden logs) such that their data is centralised and available for viewing, to assist in local client compromise identification.

A proof of concept for this has been created for use with BIND.

3.10.2 Automated Local Phishing Defense

A tool that can be used by an organisation to automate the use of local intelligence by local users forwarding attack mail to a mailbox. Mail items are processed and evaluated against the current entire blacklist, possibly resulting in local blacklist zone update and feedback.

References

[1] Barry Greene
Using DNS to Protect Your Network and Your Customers
http://www.senki.org/archives/966