ISC-TN-2010-1
October-2011

Paul Vixie, ISC
Vernon Schryver, Rhyolite

DNS Response Policy Zones (DNS RPZ, Format 3)

ISC Technical Note Series

This memo describes a method in use inside ISC which may be of use to other members of the Internet technical community.

Use of the methods and formats depicted herein is free of any and all license or other encumbrance by the authors or their employers.

Copyright Notice

Abstract

This memo describes a method for expressing DNS response policy inside a specially constructed DNS zone, and for processing the contents of such zones inside recursive name servers. These response policies are intended for use in fighting Internet crime and abuse. Almost all Internet crime relies on DNS, and many new and existing domains at the time of this writing are malicious.

1 - Overview

This memo specifies a method of expressing DNS policy information inside specially constructed DNS zones, allowing DNS reputation data producers and consumers to cooperate in the application of policies to real time DNS responses. Using this policy information, DNS resolution for low-reputation domain names can be made to deliberately fail or to return local data such as an alias to a "walled garden". A full description of the expressible policies is given below in Section 2.

Configuration examples using ISC BIND Version 9 (BIND9) are given, since this mechanism has already been implemented on that platform. We expect other recursive DNS implementations to also implement these features since the RPZ technology is unencumbered.

The goal of DNS RPZ is to enable a global market in DNS reputation data, since almost all Internet crime and abuse relies on DNS and since many new and existing domain names at the time of this writing are purely malicious.

2 - Zone Format

A DNS Response Policy Zone (RPZ) is a DNS zone, and as such its contents can be transferred between servers (DNS AXFR/IXFR), protected by transaction signatures (DNS TSIG), and expedited by real time change notifications (DNS NOTIFY), all subject to familiar DNS access controls. An RPZ usually does not support query access since it is never required for correct operation. Rather it is the zone transfer of RPZ content from producers to subscribers which effectively publishes the policy data, and it is the transferee’s server configuration which promotes RPZ payload data into DNS control plane data.

To be a valid DNS zone, an RPZ is required to have an SOA record and at least one NS record at its apex. The SOA record is real, with a serial number used for NOTIFY and IXFR and timers used for AXFR. Because query access is never required, an RPZ’s apex NS record will never be used and no parent delegation is required. The zone name itself need not be a unique fully qualified domain name. By convention, a single NS record having the deliberately bogus value "LOCALHOST" is used as a placeholder. The zone’s name can be fully qualified to show the identity of its producer or maintainer.

The remainder of an RPZ consists of expressions of DNS policy. There four kinds of policy triggers: QNAME, IP, NSIP, and NSDNAME. These are expressed as resource record sets (RRsets). Each type of policy trigger can express any one of four policy actions.

All elements described here are from RPZ Format 1 unless otherwise specified. Elements from a higher format number than a name server's implementation level are expected to be invisible to that implementation.

2.1 - Policy Actions

Four policy actions can be encoded by RRsets in an RPZ:

NXDOMAIN Action

A single resource record (RR) consisting of a CNAME whose target is the root domain (.) will cause a response of NXDOMAIN to be returned.

NODATA Action

A single RR consisting of a CNAME whose target is the wildcard top-level domain (*.) will cause a response of NODATA (ANCOUNT=0) to be returned regardless of query type.

PASSTHRU Action

A single RR consisting of a CNAME whose target is the same as the owner name of the policy RRset expresses an exception, preventing any policy action for this name from being triggered by other subscribed response policy zones.

PASSTHRU records are used with other wildcard records that would otherwise match a trusted query name, for trusted IP addresses in CIDR blocks matching other policies, or to exempt a locally trusted name from being affected by externally supplied policy.

PASSTHRU records are incompatible with wildcard owner names of QNAME policy triggers (see below), since the query name and the PASSTHRU record's CNAME target name must match exactly, without wildcarding.

Important note: for response address triggers (whose owner names are subdomains of "rpz-ip"), and for name server address triggers (whose owner names are subdomains of "rpz-nsip"), as well as for name server name triggers (whose owner names are subdomains of "rpz-nsdname"), the PASSTHRU record must not contain the special parent domain ("rpz-ip", "rpz-nsip", or "rpz-nsdname").

Local Data Action

Any other RRset (that is, neither NXDOMAIN, NODATA, nor PASSTHRU) specifies local overriding data which will be used to generate synthetic DNS responses. If any local data policy actions are present then any questions for RR types that are not present will be answered as NODATA (ANCOUNT=0). The common local data is a CNAME RR pointing to a local walled garden. Such CNAME RRs are distinguishable from NXDOMAIN, NODATA, and PASSTHRU actions because the CNAME target name will be neither (.), (*.), nor the question name itself.

[Format 3] A special form of local data involves a CNAME RR with a wildcarded target name. This form causes the QNAME to be prepended to the wildcarded target name, thus allowing the walled garden to learn the triggering QNAME value. For example a policy action of "CNAME *.EXAMPLE.COM." and a query name of "EVIL.ORG." will result in a synthetic "CNAME EVIL.ORG.EXAMPLE.COM." The main purpose for this special form is query logging in the walled garden's authority DNS server.

2.2 - Policy Triggers

Three of the types of policy triggers are based on target data (RDATA). Those policies are applied after recursion, so that the cache contains "truth" even if this truth is hidden by current policy. If the policy changes, the original data is available for processing under the changed policy. The fourth type of policy trigger is based only on the query name (QNAME) and is therefore independent of cache contents or recursion results.

Four policy triggers can be encoded by RRset owner names in an RPZ:

QNAME Trigger

The QNAME policy trigger applies to requested domain names (QNAME). The owner name of an RPZ QNAME policy RRset is the relativized name of the domain name about which policy is being expressed. For example, if the RPZ apex name is RPZ.EXAMPLE.ORG, an RRset at DOMAIN.COM.RPZ.EXAMPLE.ORG would affect responses to requests about DOMAIN.COM. Wildcards also work, so, *.DOMAIN.COM.RPZ.EXAMPLE.ORG would trigger on queries to any subdomain of DOMAIN.COM. To control the policy for both a name and its subdomains, two policy RRsets must be used, one for the domain itself and another for a wildcard sub-domain. An example of this is:

```
$ORIGIN RPZ.EXAMPLE.ORG.
DOMAIN.COM            CNAME .
*.DOMAIN.COM          CNAME .
```

In this example, queries for both DOMAIN.COM and all subdomains of DOMAIN.COM will result in synthetic NXDOMAIN responses.

[Format 2] IP Trigger

The IP policy trigger is based on target data (RDATA). It matches IP addresses that would otherwise appear in A and AAAA records in the "answer" section of a DNS response. IP policy RRsets have owner names that are sub-domains of "rpz-ip" relativized to the RPZ apex name, and an encoded IP address or block of addresses.

IPv4 addresses are encoded as "prefixlength.B4.B3.B2.B1.rpz-ip". The prefix length must be between 1 and 32. All four bytes, B4, B3, B2, and B1, must be present, and are encoded in decimal ASCII. B4 is the low order octet of the IP address and B1 is the high order octet, just as in the IN-ADDR.ARPA naming convention.

IPv6 addresses are encoded in a format similar to the standard IPv6 text representation, "prefixlength.W8.W7.W6.W5.W4.W3.W2.W1.rpz-ip" Each of W8,...,W1 is a one to four digit hexadecimal ASCII number representing 16 bits of the IPv6 address with no leading zeroes and reversed as in IP6.ARPA. All 8 words must be present unless a "zz" label is present which is analagous to the double-colon (::) in standard IPv6 address representation. The "zz" label is expanded so as to zero-fill the middle portion of the IPv6 address. The prefix length must be between 1 and 128.

For example, to force an NXDOMAIN response whenever a truthful response would contain an A RRset having an address in 192.168.1.0/24 unless address 192.168.1.2 is present, the RPZ would contain the following:

```
$ORIGIN RPZ.EXAMPLE.ORG.
24.0.1.168.192.rpz-ip     CNAME .
32.2.1.168.192.rpz-ip     CNAME 32.2.1.168.192.
```

In another example, to answer NODATA (ANCOUNT=0) whenever a truthful response would contain an AAAA RRset having an address 2001:0002::/48 unless address 2001:0002::3 was present, the RPZ would contain these records:

```
$ORIGIN RPZ.EXAMPLE.ORG.
48.zz.2.2001.rpz-ip       CNAME *.
128.3.zz.2.2001.rpz-ip    CNAME 128.3.zz.2.2001.
```

[Format 2] NSDNAME Trigger

The NSDNAME policy trigger matches name server names (NS RR) of any name server which is in the data path for an RRset present in the answer section of a DNS response. The data path shall be construed to mean all delegation points from (and including) the root zone to the closest enclosing NS RRset for the owner name of the answering RRset.

NSDNAME policies are expressed in RRsets in sub-domains of "rpz-nsdname" but otherwise much like QNAME policies. For example, to force an NXDOMAIN answer whenever a name server for the requested domain or one of its parents is NS.DOMAIN.COM, the RPZ would contain the following:

```
$ORIGIN RPZ.EXAMPLE.ORG.
NS.DOMAIN.COM.rpz-nsdname    CNAME .
```

Note: Some implementations of DNS RPZ will exhaustively discover all ancestral zone cuts above the query name and will learn the NS RRset at the apex of each delegated zone. Other implementations of DNS RPZ will know only the zone cut information which has naturally come into the cache, which will often include only parent delegation name server names. Since apex ("below the cut") NS RRsets and delegation NS RRsets need not exactly match, this can lead to instability in DNS RPZ behavior in the presence of zone cuts having differences between the NS RRsets above and below a zone cut. This instability must be taken into account when designing a security policy or testing DNS RPZ.

[Format 2] NSIP Trigger

The NSIP policy trigger matches name server addresses (an A or AAAA RR that's been referenced by an NS RRset). NSIP is much like NSDNAME (described above) except that the matching is by name server address rather than name server name. NSIP policies are expressed as sub-domains of "rpz-nsip" and have the same sub-domain naming convention as described for "IP" policy triggers above.

Note: Some implementations of DNS RPZ will exhaustively discover all IP addresses (V4 and V6) associated with each name server name. Other DNS RPZ implementations will only know the subset of IP addresses which have entered the cache naturally. This can lead to instability in the DNS RPZ behavior since the natural entry of IP addresses into the cache is itself unstable. This instability must be taken into account when designing a security policy or testing DNS RPZ.

3 - Subscriber Behavior

RPZs must be primary or secondary zones. They can only be searched in a recursive server's own storage. By default, policies are applied only on DNS requests that ask for recursion (RD=1) and which either do not request DNSSEC metadata (DO=0) or for which no DNSSEC metadata exists.

Policies are checked at each stage of resolving a domain name defined by a CNAME or DNAME record, stopping at the first CNAME in the chain with an applicable policy.

If a policy trigger results in a modified answer, then that modified answer will include in its "authority" section the SOA RR of the DNS RPZ whose policy was used to generate the modified answer. This SOA RR will tell both the fully qualified name of the DNS RPZ and the serial number of the policy data which was connected to the DNS control plane at the time the answer was modified.

RPZ zones are loaded in the usual way. For primary zones this may mean loading the contents of a local file into memory, or connecting to a database. For secondary zones this means transfering the zone from the primary server using zone transfer (such as IXFR or AXFR). It is strongly recommended that all secondary zone transfer relationships be protected with transaction signatures (DNS TSIG) and that real time change notification be enabled (DNS NOTIFY).

To connect the name server's control plane to the DNS RPZ data plane, an ordered list of RPZ's should be supplied. For each DNS RPZ in this list, it should be possible to specify an overriding policy action to be used for any policy triggers found in that RPZ except in-zone PASSTHRU actions which are always honoured. These override policies should include NXDOMAIN, NODATA, PASSTHRU, GIVEN, and CNAME. GIVEN just explicitly reaffirms the default which is to respect all policy actions found in this DNS RPZ. CNAME is an instance of "local data" which probably points to a walled garden service.

It is possible for more than one policy trigger among the various DNS RPZs connected to the name server's control plane to match a given DNS query or DNS response. The precedence rules for multiple matches are as follows:

RPZ Ordering

Policies from RPZs defined earlier ordered set of DNS RPZs are applied instead of those defined later.

Within An RPZ

Among policies from a single RPZ, QNAME policies are preferred over IP, IP policies are preferred over NSDNAME, and NSDNAME policies are preferred over NSIP.

Name Length

Among applicable QNAME or NSDNAME policies within a DNS RPZ, choose the policy that matches the smallest name.

Prefix Length

Among applicable IP or NSIP policies, use the policy with the longest prefix length.

Tie Breaker

Given equal prefix lengths, use the policy that matches the smallest IP address.

4 - Producer Behavior

A DNS RPZ producer should make every effort to ensure that incremental zone transfer (IXFR) rather than full zone transfer (AXFR) is used to move new policy data toward subscribers. Also, real time zone change notifications (DNS NOTIFY) should be enabled and tested. DNS RPZ consumers are "stealth slaves" as described in RFC 1996, and as such each such server must be explicitly denoted in the master server's configuration. And because DNS NOTIFY is a lazy protocol, it may be necessary to explicitly trigger the master server's "notify" logic after each update to the DNS RPZ. These operational guidelines are to limit policy data latency, since this latency is critical to both prevention of crime and abuse, and to withdrawal of erroneous or outdated policy.

In the data feed for disreputable domains, each addition or deletion or expiration can be handled using DNS UPDATE (see RFC 2136) to trigger normal DNS NOTIFY and subsequent DNS IXFR activity which can keep the subscribing servers well synchronized to the master RPZ. Alternatively, on some primary name servers (such as ISC BIND) it is possible to generate an entirely new primary RPZ file and have the server compute the differences between each new version and its predecessor. In ISC BIND this option is called "ixfr-from-differences" and is known to be performant even for million-rule DNS RPZ's with significant churn on a minute by minute basis.

It's good operational practice to include test records in each DNS RPZ to help that DNS RPZ's subscribers verify that response policy rewriting is working. For example, a DNS RPZ might include a QNAME policy record for BAD.EXAMPLE.COM and an IP policy record for 127.0.0.2. A subscriber can verify the correctness of their installation by querying for BAD.EXAMPLE.COM which does not exist in real DNS. If an answer is received it will be from the DNS RPZ. That answer will contain an SOA RR denoting the fully qualified name of the DNS RPZ itself.

5 - Examples

An existing data feed capable of producing an RHSBL can be trivially used to generate a DNS RPZ. If the desired policy is to alias targeted domains to a local walled garden, then for each domain name, generate the following records, one for the name itself and perhaps also one for its sub-domains:

```
bad.domain.com      CNAME   walled-garden.example.net.
*.bad.domain.com    CNAME   walled-garden.example.net.
```

If it is desirable to return NXDOMAIN for each domain (and its sub-domains in this example), try this:

```
bad.domain.com      CNAME   .
*.bad.domain.com    CNAME   .
```

If there are specific walled gardens for mail versus everything else:

```
bad.domain.com      MX      0 wgmail.example.net.
bad.domain.com      A       192.168.6.66
*.bad.domain.com    MX      0 wgmail.example.net.
*.bad.domain.com    A       192.168.6.66
```

A complete example demonstrating every policy trigger and policy action is as follows:

```
$ORIGIN rpz.example.com.
$TTL 1H
@           SOA LOCALHOST. named-mgr.example.com (1 1h 15m 30d 2h)
          NS  LOCALHOST.

; QNAME policy records.
; Note: There are no periods (.) after the (relativised) owner names.
nxdomain.domain.com CNAME   .           ; NXDOMAIN policy
nodata.domain.com   CNAME   *.          ; NODATA policy
bad.domain.com      A       10.0.0.1    ; redirect to walled garden
                  AAAA    2001:2::1

; do not rewrite OK.DOMAIN.COM (so, PASSTHRU)
ok.domain.com       CNAME   ok.domain.com.
bzone.domain.com    CNAME   garden.example.com.

; redirect X.BZONE.DOMAIN.COM to X.BZONE.DOMAIN.COM.GARDEN.EXAMPLE.COM
*.bzone.domain.com  CNAME   *.garden.example.com.

; IP policy records that rewrite all answers for 127/8 except 127.0.0.1
8.0.0.0.127.rpz-ip  CNAME   .
32.1.0.0.127.rpz-ip CNAME   32.1.0.0.127. ; PASSTHRU for 127.0.0.1

; NSDNAME and NSIP policy records that rewrite to NXDOMAIN all responses
; for domains for which NS.DOMAIN.COM is an authoritative DNS server
; (or server for a parent)
; or that have an authoritative server in 2001:2::/48
ns.domain.com.rpz-nsdname   CNAME   .
48.zz.2.2001.rpz-nsip       CNAME   .
```

6 - Bugs

As of RPZ format 3, there is no way to express a PASSTHRU action for a wildcard query name. This will be addressed in format 4.

7 - Authors

Paul Vixie
Internet Systems Consortium
<vixie@isc.org>

Vernon Schryver
Rhyolite Software
<vjs@rhyolite.com>