

Internet Draft
Expiration: October 26, 2004
Anti-Spam Research Group

J. Levine
Taughannock Networks
April 26, 2004

DNS Based Blacklists and Whitelists for E-Mail
draft-irtf-asrg-dnsbl-00.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 26, 2004.

This document is intended to evolve, based on comments from the Anti-Spam Research Group (ASRG). Comments and corrections are welcome, and may be sent to the ASRG mailing list at <asrg@ietf.org>.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

The rise of spam and other anti-social behavior on the Internet has led to the creation of shared blacklists and whitelists of IP addresses or domains. The DNS has become a de-facto standard method of distributing these blacklists and whitelists. This memo documents the structure and usage of DNS based blacklists and whitelists, and the protocol used to query them.

Table of Contents

[1.](#) Introduction [2](#)

2.	Structure of an IP address DNSBL or DNSWL	2
2.1.	IP address DNSxL	3
2.2.	IP address DNSWL	3
2.3.	Combined IP address DNSxLs	3
2.4.	Test and contact addresses	4
3.	Domain name DNSxLs	4
4.	Typical usage of DNSBLs and DNSWLs	5
5.	Security Considerations	6
6.	Informative References	6
7.	Authors' Address	6

[1.](#) Introduction

In 1997, Paul Vixie, a well known Internet software engineer, started keeping a list of IP addresses that had sent him spam or engaged in other behavior that he found objectionable. Word of the list quickly spread, and he started distributing it as a BGP feed for people who wanted to block all traffic from listed IP's at their routers. The list became known as the Real-time Blackhole List (RBL).[\[3\]](#)

Many network managers wanted to use the RBL to block unwanted e-mail, but weren't prepared to block all traffic from lists in the RBL. Vixie created a DNS-based distribution scheme that quickly became more popular than the original BGP distribution. Other people created other DNS-based blacklists either to compete with the RBL or to complement it by listing different categories of IP addresses. Although some people refer to all DNS-based blacklists as ``RBLs'', that term properly is used for the MAPS RBL, the descendant of Vixie's original list, and the standard term is now DNS Blacklist or Blocklist, or DNSBL. Some people also publish DNS-based whitelists or DNSWLs.

This document describes the structure, operation, and use of DNSBLs and DNSWLs but does not describe or recommend policies for adding or removing addresses to DNSBLs and DNSWLs, nor does it recommend policies for using them, nor does it take a position whether the DNS is the best way to distribute such data.

[2.](#) Structure of an IP address DNSBL or DNSWL

Originally, DNSBLs only listed IP addresses, and most DNSBLs and DNSWLs still list IP addresses, A few DNSBLs now list domain names instead. The structure of a DNSBL and DNSWL are the same, so in the subsequent discussion we use the

abbreviation DNSxL to mean either.

2.1. IP address DNSxL

An IP address DNSxL has a structure adapted from that of the rDNS. Each IP address listed in the DNSxL has a corresponding DNS entry created by reversing the order of the octets of the text representation of the IP address, and appending the domain name of the DNSxL. If, for example, the DNSxL is called bad.example.com, and the IP address to be listed is 192.0.2.99, the name of the DNS entry would be 99.2.0.192.bad.example.com. Each entry in the DNSxL has an A record and often a TXT record. The A record conventionally has the value 127.0.0.2, but may have other values as described below. The TXT record describes the reason that the IP is listed in the DNSxL, and is often used as the text of an SMTP error response when an SMTP client attempts to send mail to a server using the list as a DNSBL. Some DNSxLs use the same TXT record for all entries, while others provide a different TXT record for each entry or range of entries that describes the reason that entry or range is listed. The reason often includes the URL of a web page where more information is available.

If an IP address is not listed in the DNSxL, there is no record for the address. If a /24 or larger range of addresses is listed, the DNSxL may use wildcards to limit the size of the zone file. If for example, the entire range of 192.0.2.0/24 were listed, the DNSBL's zone could contain a single wildcard for *.2.0.192.bad.example.com.

2.2. IP address DNSWL

Since SMTP has no standard way for a server to advise a client why a request was accepted, TXT records in DNSWLs are not very useful. Some DNSWLs contain TXT records anyway to document the reasons that entries are present.

It is possible and occasionally useful for a DNSxL to be used as a DNSBL in one context and a DNSWL in another. For example, a DNSxL that lists all of the IP addresses assigned to dialup or DHCP users on a particular network might be used as a DNSWL on that network's outgoing mail server or intranet web server, and used as a DNSBL for mail servers on other networks.

2.3. Combined IP address DNSxLs

In many cases, a single organization maintains a variety of

DNSxLs for different purposes. There are three common methods of representing multiple sublists, subdomains, multiple A records, and bit encoded entries. Most multiple lists use both subdomains and one of the other methods.

Subdomains are merely subdomains of the main DNSxL domain. If for example, bad.example.com had two sublists ugly and smelly, entries for 192.0.2.99 would be 99.2.0.192.ugly.bad.example.com or 99.2.0.192.smelly.bad.example.com. Sublist names consist of letters, so there is no problem of name collisions with entries in the main domain, where the IP addresses consist of digits.

To minimize the number of DNS lookups, multiple sublists can also be encoded as bit masks or multiple A records. With bit masks, the A record entry for each IP is the logical OR of the bit masks for all of the lists on which the IP appears. For example, the bit masks for the two sublists might be 127.0.0.1 and 127.0.0.2, in which case an entry for an IP on both lists would be 127.0.0.3. With multiple A records, each sublist has a different assigned value such as 127.0.1.1 to 127.0.1.10 for ten sublists, and there is an A record for each sublist on which the IP appears. There is no widely used convention for mapping sublist names to bits or values, beyond the convention that all A values are in the 127/8 range to prevent unwanted network traffic if the value is accidentally used as an IP address.

DNSxLs that return multiple A records generally return multiple TXT records as well; other combined DNSxLs return a single TXT record.

The per-record time-to-live and zone refresh intervals of DNSBLs and DNSWLs vary greatly depending on the management policy of the list. A list of IP addresses assigned to dynamically allocated dialup and DHCP users could be expected to change slowly, so the TTL might be several days and the zone refreshed once a day. On the other hand, a list of IP addresses that had been observed sending spam might change every few minutes, with comparably short TTL and refresh intervals.

2.4. Test and contact addresses

Nearly all IP based DNSxLs contain an entry for 127.0.0.2 for testing purposes. DNSBLs that return multiple values often have multiple test addresses so that, for example, the entry for 127.0.0.5 returns a 127.0.0.5 A record and corresponding TXT record.

Most DNSxLs also contain an A record at the DNSxL's name that points to a web server, so that anyone wishing to learn about

the bad.example.net DNSBL can check http://bad.example.net.

3. Domain name DNSxLs

A few DNSBLs list domain names rather than IP addresses. The names of their entries contain the listed domain name followed

by the name of the DNSBL. If the DNSBL were called `doms.example.net`, and the domain `invalid.edu` were to be listed, the entry would be named `invalid.edu.doms.example.net`. A few named-based DNSBLs encode e-mail addresses using a convention adopted from DNS SOA records, so an entry for `fred@invalid.edu` would have the name `fred.invalid.edu.doms.example.net`.

Name-based DNSBLs are far less common than IP based DNSBLs, There is no agreed convention for a test entry nor for wildcards. Name-based DNSWLs could be created in the same manner as DNSBLs, although to date nobody has done so.

4. Typical usage of DNSBLs and DNSWLs

DNSxLs can be served either from standard DNS servers, or from specialized servers like `rbldns[2]` and `rbldnsd[4]` that accept lists of IP addresses and CIDR ranges and synthesize the appropriate DNS records on the fly. Organizations that make heavy use of a DNSxL usually arrange for a private mirror of the DNSxL, either using the standard AXFR and IXFR or by fetching a file containing addresses and CIDR ranges for the specialized servers.

DNSBL clients are most often mail servers or spam filters called from mail servers. There's no requirement that DNSBLs be used only for mail, and other services such as IRC use them to check clients that are trying to connect.

In practice, mail servers that test combined lists usually handle them the same as single lists and treat any A or TXT record as meaning that an IP is listed without distinguishing among the various reasons it might have been listed.

Most often they check a list of DNSBLs and DNSWLs on every incoming SMTP connection, with the names of the DNSBLs and DNSWLs configured into the server. The server checks each list in turn until it finds one with a DNSBL entry, in which case it rejects the connection, or a DNSWL entry in which case it accepts the connection. If the address appears on no list at all (the usual case for legitimate mail), it accepts the connection. The mail server uses its normal local DNS cache to limit traffic to the DNSxL servers and to speed up retests of IP addresses recently seen. Long-running mail servers may cache DNSxL data internally. When using combined DNSxLs, clients usually only test for the presence or absence of an IP, without regard to the particular value returned.

An alternate approach is to check DNSxLs in a spam filtering package after a message has been received. In that case, the IP(s) to test are usually extracted from Received: headers. The DNSxL results may be used to make a binary accept/reject decision, as when they're tested at SMTP time, or may be used as components in a system that computes an overall score for

each message. Packages that test multiple headers need to be able to distinguish among values in lists with sublists since, for example, an entry indicating that an IP is assigned to dialup users might be treated as a strong indication that a message should be rejected if the IP sends mail directly to the recipient system, but not if the message were relayed through an ISP's mail server.

5. Security Considerations

Any system manager that uses DNSxLs is entrusting part of his or her server management to the parties that run the lists. A DNSBL manager that decided to list 0/0 (which has actually happened) would cause every server that uses the DNSBL to reject all mail. Conversely, if a DNSBL manager removes all of the entries (which has also happened), systems that depend on the DNSBL will find that their filtering doesn't work as they want it to.

As with any other DNS based services, DNSBLs and DNSWLs are subject to various types of DNS attacks which are described in [1].

6. Informative References

- [1] D. Atkins et al, "Threat Analysis of the Domain Name System", [draft-ietf-dnsext-dns-threats-07](#)
- [2] D. J. Bernstein, `rbldns`, in "`djbdns`", <http://cr.yp.to/djbdns.html>.
- [3] Mail Abuse Prevention System, "MAPS RBL", <http://mail-abuse.org/rbl/>
- [4] Michael Tokarev, "`rbldnsd`: Small Daemon for DNSBLs", <http://www.corpit.ru/mjt/rbldnsd.html>.

7. Authors' Address

John R. Levine
Taughannock Networks
PO Box 727
Trumansburg NY 14886 USA
E-mail: johnl@taugh.com
Phone: +1 607 330 5711

Full Copyright Statement

Copyright (C) The Internet Society (2004). All Rights

Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be

prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

\$Id: [draft-irtf-asrg-dnsbl-00](#).n,v 1.5 2004/04/27 04:20:47
johnl Exp \$

