Internet Research Task Force (IRTF)
Request for Comments: 6471
Category: Informational
ISSN: 2070-1721

C. Lewis
Nortel Networks
M. Sergeant
Symantec Corporation

January 2012

Overview of Best Email DNS-Based List (DNSBL) Operational Practices

Abstract

The rise of spam and other anti-social behavior on the Internet has led to the creation of shared DNS-based lists (DNSBLs) of IP addresses or domain names intended to help guide email filtering. This memo summarizes guidelines of accepted best practice for the management of public DNSBLs by their operators as well as for the proper use of such lists by mail server administrators (DNSBL users), and it provides useful background for both parties. It is not intended to advise on the utility or efficacy of particular DNSBLs or the DNSBL concept in general, nor to assist end users with questions about spam.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Anti-Spam Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6471.

Table of Contents

1. Introduction

1.1. DNS-Based Reputation Systems

Due to the rising amount of spam and other forms of network abuse on the Internet, many community members and companies began to create, publish and maintain DNS-based reputation systems (DNS-based lists or DNSBLs) of IP addresses or domain names and make reputation suggestions or assertions about email sourced from these IP addresses or domain names.

The first DNSBLs were almost exclusively intended to be used (by email administrators) as lists of abusive IP addresses to block; however, the DNS publication method has proven to be so robust, popular, and simple to use that it has been extended for use in many different ways, far beyond the imaginings of the designers of DNS or DNS-based blocking IP lists. For example, today, the same basic DNS-based listing technology is commonly used for:

DNSWL: listings of well-behaving email source IP/domain addresses (whitelist).

RHSBL: listings of well/ill-behaving email source domain names (often applied against the domain name part (RHS = Right Hand Side) of the originating email address or DNS PTR (reverse IP) lookups)

URIBL: listings of well/ill-behaving web link domain names or host names used in email

Further, the DNSBL user doesn’t have to use a listing as a pass/fail binary decision -- it can use a listing as one factor in email filters that make decisions based on scoring multiple factors together.

The DNS-based list technology has even been extended to purely informational purposes. For example, there are implementations that return results based on what geographic region an IP/domain is putatively allocated in, implementations that translate an IP/domain address into an Autonomous System Number (ASN) and/or allocation block, implementations that indicate whether the queried domain name is registered through a given domain registrar, implementations that return aggregate numeric reputation for an IP address or domain name from another system’s email system, and so on. The possibilities are virtually endless.

DNS-based listing technology has also been used in areas other than email filtering, such as Internet Relay Chat (IRC), web access control, and transaction verification.

As the terminology in this area has never been well formalized, often overlaps, and lacks precision, this document has been written to use the term "DNSBLs" to refer to DNS-based lists generally, not just DNS-based block (or black) lists. This document is not applicable to some DNSBLs in some areas (mentioned as appropriate), but it is the authors' belief that most of the practices are applicable to almost all DNSBLs.

DNSBLs may be either public or private. A public DNSBL makes its data available to any party seeking information about data on the list, while a private DNSBL is used solely by an organization for its own use, and the data is not made available publicly. There are also commercial DNSBLs, available for a fee. Furthermore, some are free yet require a fee for higher numbers of queries or certain classes of DNSBL users.

The first publicly available DNSBL using the Domain Name System (DNS) for distributing reputation data about email senders emerged in 1997, shortly after spam became a problem for network operators and email administrators. This pioneer DNSBL focused on identifying known spam sources situated at static (unchanging) IP/domain addresses. Due to the broad adoption of this DNSBL, it had a major impact on static spam sources. Consequently, abusers found other methods for distributing their spam, such as relaying messages through unsecured email servers or flawed formmail scripts on web pages. Additional DNSBLs were developed by others in order to address these changing tactics, and today more than 700 public DNSBLs are known to be in operation.

These DNSBLs vary widely in purpose for which the list was intended, the method the list uses to achieve the purpose, the integrity of those overseeing the method, and the stability of the technology used to create and distribute the data. Listing criteria can sometimes be quite controversial; therefore, this document deliberately does not discuss the rightness or wrongness of any criteria. We assert that DNSBL operators are free to choose whatever listing criteria they wish, as long as those criteria are clearly and accurately communicated. It is the responsibility of the DNSBL user to ensure that the listing criteria and other aspects of a DNSBL meets their needs.

This document is intended to provide guidance to DNSBL operators so that they may be able to identify what features users would be interested in seeing as part of a high-quality, well-managed DNSBL --

for example, a clear listing and delisting policy to which the DNSBL operator adheres strictly. This document is intended to be normative rather than prescriptive: it seeks to characterize the features of a well-managed DNSBL rather than setting out rules for how DNSBLs should be operated.

This document is not intended as a protocol specification of DNSBL queries. (See [RFC5782].)

The DNS has been the most popular distribution method for DNSBLs due to its ubiquity and its good scaling and performance characteristics. It is also common to make private arrangements to distribute DNSBL data in bulk to high-volume users, typically by rsync [RSYNC] [RSYNCTHESIS]. The data is the same in either case; the recommendations in this document apply, regardless of distribution method, other than the ones in Sections 3.1 and 3.2 that specifically refer to DNS distribution.

1.2. Guidance for DNSBL Users

When choosing to adopt a DNSBL, a DNSBL user SHOULD keep the following questions in mind:

1. What is the intended use of the list?

2. Does the list have a web site?

3. Are the list's policies stated on the web site?

4. Are the policies stated clearly and understandably?

5. Does the web site function properly, e.g., hyperlinks?

6. Are web pages for removal requirements accessible and working properly?

7. How long has the list been in operation?

8. What are the demographics and quantity of the list's user base? In other words, do other sites like my own use this DNSBL?

9. Are comparative evaluations of the list available? Note: all such evaluations depend on the mail mix used as well as local policy. DNSBL users SHOULD consider trial periods and/or ongoing local monitoring of DNSBL suitability.

10. What do your peers or members of the Internet community say about the list? DNSBLs can sometimes be quite controversial and sometimes considerable misinformation is spread. Ensure that the opinions are knowledgeable and reflect similar goals to yours.

11. Does the DNSBL have a mailing list for announcing changes, outages, etc.?

DNSBLs can, and have, ceased operation without notice. DNSBL users SHOULD periodically check the correct operation of the DNSBL, and cease using DNSBLs that are working incorrectly. See Section 3.3.

The DNSBL user MUST ensure that they understand the intended use of the DNSBL. For example, some IP address-based DNSBLs are appropriate for assessment of only the peer IP address of the machine connecting to the DNSBL user's mail server, and not other IP addresses appearing in an email (such as header Received lines or web links) or IRC connections, etc. While a DNSBL user may choose to ignore the intent of the DNSBL, they SHOULD implement any variance in compliance with the DNSBL usage instructions.

For example, one of the requirements of some DNSBLs is that if the DNSBL is used contrary to the usage instructions, then the DNSBL user should not identify the DNSBL being used. Furthermore, it is the DNSBL user's responsibility to mitigate the effect of the listing locally.

It is the responsibility of the system administrators who adopt one or more DNSBLs to evaluate, understand, and make a determination of which DNSBLs are appropriate for the sites they administer. If you are going to allow a third party's information to guide your filtering decision-making process, you MUST understand the policies and practices of those third parties because responsibility for filter decisions remains ultimately with you, the postmaster.

A DNSBL without DNSBL users does not block (or otherwise impair) email or any other Internet service. A DNSBL user voluntarily uses the DNSBL data to guide their decisions, and the DNSBL user therefore MUST assume responsibility for dealing with the consequences.

DNSBL operators are expressing an opinion through the publication of a DNSBL. However, it is through abiding by the guidelines set forth in this document that the operators of a DNSBL may gain the trust of their users.

These guidelines address only public DNSBLs and do not apply to private-access DNSBLs; however, implementers and users of private-access DNSBLs may wish to use these guidelines as a starting point of things to consider.

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.4. Background

The Anti-Spam Research Group (ASRG) was chartered to address the spam problem. The ASRG charter includes:

"codification of best current practices in spam management"

This note falls within that category by listing guidelines for management of public DNSBLs.

NOTE: This document is a product of the Anti-Spam Research Group (ASRG) of the IRTF.

2. DNSBL Policies

2.1. Transparency

A DNSBL SHOULD carefully describe the criteria for adding and the criteria for removing an entry from the list. Such listing and delisting criteria SHOULD be presented in a clear and readable manner easily accessible to the public on the DNSBL's web site. A DNSBL MUST abide by its stated listing and delisting criteria. Entries that do not meet the published criteria MUST NOT be added to the DNSBL.

In other words, be direct and honest and clear about the listing criteria, and make certain that only entries meeting the published criteria are added to the list. For example, some DNSBL operators have been known to include "spite listings" in the lists they administer -- listings of IP addresses or domain names associated with someone who has insulted them, rather than actually violating technical criteria for inclusion in the list. There is nothing inherently wrong with this practice so long as it is clearly disclosed -- and thus becomes part of the published criteria. For example, a DNSBL described as only listing open relays MUST NOT include IP addresses for any other reason. This transparency

principle does not require DNSBL operators to disclose the precise algorithms and data involved in a listing, but rather the intent behind choosing those algorithms and data.

Furthermore, the DNSBL documentation SHOULD be clear on the intended use of the DNSBL -- whether it be intended for peer addresses of email, IRC, etc.

Availability of documentation concerning a DNSBL SHOULD NOT be dependent on the continued operation of DNS for DNSBL queries.

In other words, if the DNSBL documentation is at "http://dnsbl.example.com", the documentation for the web site should not become unavailable if the DNSBL query name servers are not available (or shut down). See Section 3.1.

2.1.1. Listing/Delisting Criteria SHOULD Be Easily Available

Listing and delisting criteria for DNSBLs SHOULD be easily available and SHOULD be located in a place clearly marked in its own section of the web site affiliated with the DNSBL.

DNSBLs often publish their listing criteria along with additional technical information about using the DNSBL. This additional technical information can confuse end users, so a separate page, section, or query function on its own SHOULD be dedicated to detailing why a specific entry appears in the DNSBL.

2.1.2. Audit Trail SHOULD Be Maintained

A DNSBL SHOULD maintain an audit trail for all listings, and it is RECOMMENDED that it is made publicly available in an easy to find location, preferably on the DNSBL's web site. Please note that making data about an audit trail public does not entail revealing all information in the DNSBL operator's possession relating to the listing. For example, a DNSBL operator MAY make the audit trail data selectively accessible in such a way as to not disclose information that might assist spammers, such as the location or identity of a spam trap.

2.1.3. The Scope and Aggressiveness of Listings MUST Be Disclosed

Some DNSBLs have adopted policies of listing entries that are broader in scope than they have evidence of being involved in abuse. Similarly, some DNSBLs list entries that are "mixed", in that the entry may be behaving in a manner that is both abusive and non-abusive. This is inherent to the techniques that many DNSBLs use.

Examples: Some DNSBLs will list IP address ranges if there is reason to believe that abusive behavior seen from a few IP addresses within the range is (or will be) reflected in the rest of the range. Some DNSBLs utilize scoring to list IP addresses, IP ranges, or domain names that have abusive behavior above some threshold -- often meaning that some of the email corresponding to the listing is not abusive. Even an entry demonstrably infected with email spam or virus-emitting malware may emit non-abusive email.

Inevitably, some of these listings may impact non-abusive email. This has resulted in some labeling of such practices by the emotionally loaded term "collateral damage". No filtering technique is perfect, and an occasional mistake is inevitable no matter what is used, DNSBLs or otherwise.

There is nothing wrong with this practice (of having "collateral damage") because mail server administrators may wish to implement such policies or use them in combination with other techniques (such as scoring). However, a diligent administrator needs information about these policies in order to make an informed decision as to the risk and benefit of using any particularly DNSBL, and to guide them in how to use it for results best reflecting the DNSBL user's requirements.

Therefore, DNSBL listing policies MUST include statements as to the scope and aggressiveness of listings and include, as appropriate, whether the DNSBL operator intends the listings to be used in scoring or other techniques.

2.2. Listings and Removals

2.2.1. Listings SHOULD Be Temporary

In many cases, listings can exist for long periods of time past the conditions leading to the listing's creation, and/or listings can exist after the listed entity has putatively changed ownership.

Generally speaking, listings SHOULD be considered temporary and should expire on their own at some point in the future, unless reasons for listing still exist.

Expiration intervals SHOULD be chosen to be reasonable for the type of listing. For example:

1. It does not make sense to remove entries from DNSBLs where the existence of an entry does not have a direct meaning, that is, DNSBLs that return information in addition to just existence/ non-existence. For example: entries in DNSBLs that return

geographic or assignment information on where the IP address or domain name is located or owned, or DNSBLs that return flow statistics from the DNSBL operator that are intended for the DNSBL user to interpret, need not ever be removed, just kept reasonably current.

2. DNSBLs based on relatively static information, such as block assignment or domain names of demonstrably bad actors, MAY have very long expiration intervals or be removed only upon request after verification that the removal criteria have been met.

3. Automated DNSBLs with highly effective detection and fast listing mechanisms can benefit from very short expiration intervals. Many of the things that these DNSBLs look for are of relatively short duration, and even if they do expire, a resumption of the behavior will be caught quickly by the DNSBL's detection mechanisms and relisted. By utilizing a short expiration interval, after reassignment/problem correction, the listing will automatically expire in short order without manual intervention.

4. Manually created DNSBL entries SHOULD be periodically reviewed in some manner.

It is RECOMMENDED that DNSBL operators publish in general terms their expiration policy, even if it's only "delist on request" or "no expiration is performed". In information-only lists, a method for users requesting corrections to the information (if appropriate) SHOULD be published. Abusers may be able to "game" policy that is too explicit; on the other hand, many DNSBL users wish to have an idea of how "current" the DNSBL is. It is the authors' experience that some automated DNSBLs have increasingly higher error rates as the "last detection date" gets older.

Note that listings being temporary does not mean that all listings will expire after the initial time-out period. If the DNSBL operator determines that the conditions triggering listing still exist, then the timer for determining time outs can be renewed.

2.2.2. A Direct Non-Public Way to Request Removal SHOULD Be Available

Discussions about whether a DNSBL should remove an entry MAY include activity in a public forum. Methods for processing removal requests through private, direct exchanges, such as person-to-person email or a combination of web page requests and email responses, SHOULD be available. As a minimum, the DNSBL SHOULD have a web page that has a removal request function (separate from the page describing listing criteria as per Section 2.1.1). The DNSBL SHOULD also make available an email address to handle issues other than blocking issues.

The DNSBL operator MUST NOT use the list in question in such a way that removal requests would be blocked; and moreover, the operator SHOULD make mailboxes available in order to allow affected users to submit their requests. In some cases, it is impractical not to filter email to accounts due to the amount of spam those mailboxes receive. If filtering should be necessary in such circumstances, filtering methods with as low false positive rate as practical SHOULD be chosen.

DNSBL operators SHOULD be prepared to provide alternate means of contact in case of system failure due to DDoS (distributed denial-of-service) attack or other reasons.

2.2.3. Response SHOULD Be Prompt

A response to removal requests or queries about a listing SHOULD be prompt. A DNSBL operator SHOULD respond within 2 days and MUST respond within 7 days, except in the case that the DNSBL operator has deemed that further discussion of the issue will not result in meeting the conditions for removal and has notified the requestor of that decision.

Consequent removals (if the conditions for removal are met) should be similarly prompt.

A DNSBL MAY impose restrictions on who (e.g., a network operator’s representative or domain name owner) may make valid removal requests. However, in many DNSBLs, this is inadvisable because it requires impractical amounts of effort; hence, it is NOT RECOMMENDED in most cases.

Many DNSBLs (especially those with highly effective detection and fast listing mechanisms) greatly benefit from a "no questions asked" removal policy.

Although this approach allows people to submit a request and have any listed IP address/domain name removed immediately, it does not prevent the DNSBL operator from relisting the IP address/domain name at a later time.

Many DNSBLs can effectively use a "no questions asked" removal policy because by their very nature they will redetect or relist problems almost immediately. They can mitigate more organized attempts to "game" the system by performing elementary checking and rate-limiting procedures, increasing lockout periods, executing re-scans, etc. Furthermore, a adding or removing a few IP addresses usually does not

make a significant difference in the overall effectiveness of a DNSBL. Moreover, a "no questions asked" removal policy provides the huge benefit of a swift reaction to incorrect listings.

As an example, one popular DNSBL uses a "no questions asked" removal policy, but does perform rate-limiting and malicious removal detection and mitigation.

Another important consideration supporting a "no questions asked" self-removal policy is that it forestalls many conflicts between DNSBL operators and organizations whose IP addresses/domain names have been listed. Such a policy may be an effective measure to prevent small issues from becoming big problems.

2.2.4. A Given DNSBL SHOULD Have Similar Criteria for Listing and Delisting

The criteria for being removed from a DNSBL SHOULD bear a reasonable relationship to the factors that were the cause of the addition to the DNSBL. If a listed entity fulfills all published requirements for removal from a DNSBL, then the DNSBL operator SHOULD NOT impose any additional obstacles to remove a given entry from the DNSBL. There SHOULD NOT be any extra rules for delisting other than the ones listed in the published listing criteria.

2.2.5. Conflict of Interest

Some DNSBLs used for blocking/negative reputation have had a practice of requiring fees or donations to charities from the listee for delisting.

It is generally considered entirely appropriate for a DNSBL to charge for access to it by its users -- the definition of a commercial DNSBL.

However, the practice of requiring a listee to pay for delisting from a negative-connotation DNSBL steers perilously close to notions of extortion, blackmail, or a "protection racket". Even when such accusations are entirely unjustified, the practice causes uproar and damage to the DNSBL's reputation, if not the DNSBL mechanism as a whole.

Therefore, negative-connotation DNSBLs MUST not charge fees or require donations for delisting or "faster handling", and it is RECOMMENDED that such DNSBLs that do charge fees or require donations not be used.

3. Operational Issues

3.1. DNSBL Query Root Domain Name SHOULD be a Subdomain

By virtue of using domain names, a DNSBL is a hierarchy with a root anchored in the global Internet. The DNSBL "query root" SHOULD be below the registered domain name, so that the DNSBL information is not conflated with domain name housekeeping information (e.g., name server or MX records) for the domain name. By using this approach, DNSBL queries would take the form of "<query>.dnsbl.example.com" rather than "<query>.example.com". Further, this sub-tree should have its own name servers. Thus, the DNSBL query root has its own zone file containing the DNSBL information, and the registered domain name has its own name servers containing the information (MX records, etc.) for the domain name. This approach facilitates clear delineation of function as well as orderly DNSBL shutdown because the DNSBL name server records can be specified separately from the domain name's principal name servers.

Many DNSBLs support more than one logical zone (DNSBL entries with different meanings) that DNSBL users may wish to treat differently (or even ignore). It is RECOMMENDED that, even if there is a single DNSBL zone with entry type distinguished by return code, separate subdomain names (of the query root) consist only of the corresponding entries. For example, entry types "A" and "B" might return 127.0.0.2 and 127.0.0.3 from the consolidated zone (e.g., dnsbl.example.com), but there should also be zones typeA.dnsbl.example.com and typeB.dnsbl.example.com that contain their respective types only. See also Section 3.3.

3.2. DNSBLs SHOULD Be Adequately Provisioned

The DNSBL SHOULD have sufficient name server capacity to handle the expected loading and have sufficient redundancy to handle normal outages.

Name servers SHOULD provide appropriate glue records, possibly in different Top-Level Domains (TLDs) to protect against single-TLD issues.

If the DNSBL offers zone transfers (in addition to or instead of standard DNSBL query mechanisms), it SHOULD be sufficiently provisioned to handle the expected loading.

Note that some DNSBLs have been subject to DDoS attacks. Provisioning SHOULD take the likelihood of this into account and include plans for dealing with it.

3.3. DNSBLs SHOULD Provide Operational Flags

Most IP address-based DNSBLs follow a convention of query entries for IP addresses in 127.0.0.0/8 (127.0.0.0-127.255.255.255) to provide online indication of whether the DNSBL is operational. Many, if not most, DNSBLs arrange to have a query of 127.0.0.2 return an A record (usually 127.0.0.2) indicating that the IP address is listed. This appears to be a de facto standard indicating that the DNSBL is operating correctly. See [RFC5782] for more details on DNSBL test entries.

If this indicator is missing (query of 127.0.0.2 returns NXDOMAIN), or any query returns an A record outside of 127.0.0.0/8, the DNSBL should be considered non-functional.

There does not appear to be a de facto standard for test entries within domain-name-based DNSBLs. A number of domain-name-based DNSBLs use the same 127.0.0.2 query test mechanism as IP-address-based DNSBLs, and others use a variety of domain-name-based test entries. Due to the way many domain-name-based DNSBLs are used (e.g., hostname parts of URIs in email bodies), using anything likely to appear in a legitimate email message is a bad idea (e.g., http://example.com), especially considering that some email readers will transform bare IP addresses or domain names appearing in the body of an email into links. So, even 127.0.0.2 may be problematic. But a common testing method is desirable.

In the absence of new emerging standards, it is RECOMMENDED that domain-name-based DNSBLs use a test entry of "test". This is chosen because it is a reserved TLD.

Note: In Section 3.4, it is noted that some DNSBLs have shut down in such a way to list all of the Internet. Further, in Section 3.5, DNSBL operators MUST NOT list 127.0.0.1. Therefore, a positive listing for 127.0.0.1 SHOULD indicate that the DNSBL has started listing the world and is non-functional. Similarly, a domain-based DNSBL SHOULD NOT ever list the reserved domain INVALID, and a positive listing for INVALID SHOULD indicate that the DNSBL is non-functional.

Other results, such as 127.0.0.3, may have different meanings. This operational flag usage and meaning SHOULD be published on the DNSBL's web site, and the DNSBL user SHOULD periodically test the DNSBL.

Some mail systems are unable to differentiate between these various results or flags, however, so a public DNSBL SHOULD NOT include opposing or widely different meanings -- such as 127.0.0.23 for "sends good mail" and 127.0.0.99 for "sends bad mail" -- within the same DNS zone.

3.4. Shutdowns MUST Be Done Gracefully

A number of DNSBLs have shut down operations in such a way as to list the entire Internet, sometimes without warning. These were usually done this way to force DNSBL users (mail administrators) to adjust their DNSBL client configurations to omit the now inoperative DNSBL and to shed the DNS query load from the registered domain name servers for the DNSBL. Popular DNSBLs are used by tens of thousands of sites, yet, the correct operation of the DNSBLs are not well monitored by their users. The DNSBL query clients are often not compliant with DNSBL query conventions (e.g., they will treat any A record returned as being "listed", instead of specific 127/8 A record returns), hence shutdowns (or even ordinary domain name expiration) can be quite destructive to all email flow if not done properly.

The DNSBL operator MUST issue impending shutdown warnings (on the DNSBL web site, appropriate mailing lists, newsgroups, vendor newsletters, etc.), and indicate that the DNSBL is inoperative using the signaling given in Section 3.3.

Only after these warnings have been issued for a significant period of time (RECOMMENDED: one or more months), should the DNSBL operator finally shutdown the DNSBL.

The shutdown procedure should have the following properties:

1. MUST NOT list the entire Internet

2. SHOULD shed the DNSBL query load from the DNSBL name servers, permitting the registered domain name to continue being usable.

3. SHOULD, perhaps through increased delays, indicate to the mail administrator that the DNSBL is no longer functional.

4. Name server or query lookups MUST NOT be aimed at third parties unrelated to DNSBL operation. Such behavior is similar to inflicting a DDoS attack.

5. The base domain name SHOULD be registered indefinitely, so as to prevent the domain name from being a "booby trap" for future owners, and/or to prevent a new owner from maliciously listing the entire Internet.

One way of satisfying points 1-4 above is to change the DNS name servers for the DNSBL to point at "TEST-NET" addresses (see [RFC5735]). The below suggested [BIND] declarations will cause a DNSBL query to query non-existent name servers in TEST-NET addresses, which will result in a significant delay (usually more delay as the number of non-existent TEST-NET name servers is increased), but will not return any A records except in very unusual circumstances.

BIND-equivalent DNS declarations for DNSBL shutdown.

```
dnsbl.example.com.  604800  IN  NS  u1.example.com.
u1.example.com.     604800  IN  A   192.0.2.1

dnsbl.example.com.  604800  IN  NS  u2.example.com.
u2.example.com.     604800  IN  A   192.0.2.2

dnsbl.example.com.  604800  IN  NS  u3.example.com.
u3.example.com.     604800  IN  A   192.0.2.3

... [as many NS/A record pairs as you like]
```

This example assumes that the DNSBL is named "dnsbl.example.com". Replace "example.com" and "dnsbl.example.com" as appropriate for the DNSBL.

NOTE: Of course, the above shutdown procedure cannot be implemented if Section 3.1 is not followed.

3.5. Listing of Special and Reserved IP Addresses MUST Be Disclosed

The DNSBL MAY list loopback, [RFC1918], LINK-LOCAL class [RFC3927], class D/E, and any other permanently reserved or special-use IP addresses [RFC5735] (and [RFC5156] for IPv6). Such use MUST be disclosed in the documentation related to the DNSBL.

As additional insurance against listings of space that should not be listed through testing or other unforeseen events, DNSBL operators SHOULD consider implementing facilities to prevent them. At least one popular automated DNSBL has implemented permanent exclusions for such addresses.

A functioning DNSBL MUST NOT list 127.0.0.1. There are a number of mail server implementations that do not cope with this well, and many will use a positive response for 127.0.0.1 as an indication that the DNSBL is shut down and listing the entire Internet.

3.6. Considerations for DNSBLs Listing Insecure Hosts

Some DNSBLs list IP addresses of hosts that are insecure in various ways (e.g., open relays, open proxies). The following recommendations for such DNSBLs may not be relevant to other types of DNSBLs.

The practice of scanning for vulnerabilities can represent a risk in some jurisdictions. The following recommendations for such DNSBLs MAY help alleviate this risk.

3.6.1. DNSBLs MUST NOT Scan without Provocation

DNSBLs MUST NOT automatically probe for insecure hosts without provocation. There is little agreement in the community as to whether or not such activity should be allowed, so this document errs on the side of caution.

Therefore, scanning MUST be targeted, rather than broad-based, where a given scan is motivated by a specific reason to have concern about the address being scanned. Examples of such reasons include delivery of an email, delivery to a spam trap address, receipt of a user complaint, or periodic testing of an address that is already listed.

3.6.2. Re-Scan Periods SHOULD Be Reasonable

If the DNSBL operator re-scans a host in order to determine whether the listing SHOULD time out or not, the re-scan period SHOULD be reasonable. Automated scanning SHOULD NOT occur more often than once every 24 hours.

It is RECOMMENDED that automated re-scanning should cease within a reasonable period of the vulnerability no longer existing and of the targeting conditions no longer being met.

3.6.3. Scans MUST NOT Be Destructive

In the past, some scanning mechanisms have proven to adversely impact the scanned host, sometimes in severe fashion. Scanning methodologies MUST NOT negatively impact the scanned host.

3.7. Removals SHOULD Be Possible in Absence of the DNSBL Operator

If removals cannot be automated (e.g., via robot re-testing or self-removal), then the DNSBL SHOULD have multiple administrators so that a removal request can be processed if the principal list administrator is on vacation or otherwise unavailable.

3.8. Protect against Misconfiguration/Outages

It is not altogether uncommon for DNSBL users to configure their systems improperly for DNSBL queries. The consequences of an error can range from undue (or even damaging) load on the DNSBL servers to accidentally blocking all incoming email.

DNSBL users MUST test their initial DNSBL configurations to ensure that they're working correctly and SHOULD periodically recheck the status of the DNSBLs they use and adjust their configuration as necessary.

Common types of misconfigurations include:

1. Using wrong (sub-)zones for querying (e.g., 4.3.2.1.example.com or 4.3.2.1.dnsbl.exmple.cm instead of 4.3.2.1.dnsbl.example.com).

2. Downloading a local mirror of the data, but failing to set up the local name server infrastructure appropriately, and thus continuing to query the public name servers.

3. Downloading a local mirror of the data, but misconfiguring the local name server infrastructure to query a locally invented zone name (4.3.2.1.dnsbl.local) at the public name servers.

4. Misconfiguring local name servers to not do meaningful caching, thus heavily increasing load on the public name servers.

5. Using the DNSBL query root domain name as the name server for queries.

6. Using the DNSBL incorrectly, e.g., some DNSBLs are suitable only for certain types of filtering. Improper use may result in excessive incorrect filtering.

While in many cases it can be difficult to detect such situations, to protect against such misconfiguration, it is RECOMMENDED that DNSBL operators make design decisions to mitigate the impact of such mistakes and make efforts to contact administrative contacts to remedy the situation where appropriate. But the DNSBL operator SHOULD also prepare to take appropriate steps to protect the operational infrastructure (e.g., have the ability to block abusive users from causing further damage).

Appropriate use of the DNSBL SHOULD be documented on the web site.

3.9. Error Handling

From time to time, DNSBLs have encountered operational data integrity or data collection problems that have resulted in improper listings. For example: data corruption, erroneous restoration of resolved listings, or grossly misfiring detection heuristics. This often results in great consternation over what appear to be nonsensical listings or listings for previously resolved issues.

Many DNSBLs have implemented policies and procedures whereby such situations result in the purging of even slightly doubtful entries, disconnection of untrustworthy components until the entries' validity or correct operation of the component can be verified or corrected, as well as notification of the issue on the DNSBL's web pages.

As an example, one popular DNSBL has a demonstrated track record of disabling faulty data collection mechanisms, purging all listings generated by the faulty mechanism, and publishing a brief description of the problem and course of remediation.

Therefore, DNSBLs SHOULD have policies and procedures in place to treat operational problems conservatively, be prepared to mass purge dubious entries, prevent future erroneous entries, and notify their users by the DNSBL's web page.

4. Security Considerations

Any system manager that uses DNSBLs 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) could 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.

If a registered domain name used for a DNSBL is allowed to lapse, or the DNSBL user spells the DNSBL domain name incorrectly, the system manager's server management is now subject to an entirely different party than was intended. Further, even if there is no malicious intent, some DNSBL query clients will interpret any A record being returned as being listed. DNSBL users SHOULD be prepared to periodically test the DNSBLs they use for correct operation.

Like all DNS-based mechanisms, DNSBLs are subject to various threats outlined in [RFC3833].

5. References

5.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, May 2005.

5.2. Informative References

[BIND] Internet Systems Corporation, "ISC BIND", <http://www.isc.org/software/bind>.

[RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", RFC 3833, August 2004.

[RFC5156] Blanchet, M., "Special-Use IPv6 Addresses", RFC 5156, April 2008.

[RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", BCP 153, RFC 5735, January 2010.

[RFC5782] Levine, J., "DNS Blacklists and Whitelists", RFC 5782, February 2010.

[RSYNC] Tridgell, A., "rsync", <http://rsync.samba.org/>.

[RSYNCTHESIS] Tridgell, A., "Efficient Algorithms for Sorting and Synchronization", <http://samba.org/~tridge/phd_thesis.pdf>.

Appendix A. Acknowledgements

We would like to thank John R. Levine, Alan Murphy, and Dave Crocker for their insightful comments.

We would also like to thank Yakov Shafranovich and Nick Nicholas for editing draft versions of this document.

Authors' Addresses

Chris Lewis
Nortel Networks

EMail: clewisbcp@cauce.org

Matt Sergeant
Symantec Corporation

EMail: matt@sergeant.org