

Network Working Group	J. Damas	
Internet-Draft	ISC	
Intended status: BCP	F. Neves	
Expires: March 5, 2009	Registro.br	
	September 01, 2008	

[TOC](#)

Preventing Use of Recursive Nameservers in Reflector Attacks draft-ietf-dnsop-reflectors-are-evil-06.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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 March 5, 2009.

Abstract

This document describes ways to prevent the use of default configured recursive nameservers as reflectors in Denial of Service (DoS) attacks. Recommended configuration as measures to mitigate the attack are given.

Table of Contents

- [1.](#) Introduction
- [2.](#) Document Terminology
- [3.](#) Problem Description
- [4.](#) Recommended Configuration
- [5.](#) Security Considerations
- [6.](#) Acknowledgments
- [7.](#) IANA Considerations

- [8.](#) References
 - [8.1.](#) Normative References
 - [8.2.](#) Informative References
 - [§](#) Authors' Addresses
 - [§](#) Intellectual Property and Copyright Statements
-

1. Introduction

[TOC](#)

Recently, DNS [\[RFC1034\]](#) (Mockapetris, P., "Domain names - concepts and facilities," November 1987.) has been named as a major factor in the generation of massive amounts of network traffic used in Denial of Service (DoS) attacks. These attacks, called reflector attacks, are not due to any particular flaw in the design of the DNS or its implementations, aside perhaps the fact that DNS relies heavily on UDP, the easy abuse of which is at the source of the problem. They have preferentially used DNS due to common default configurations that allow for easy use of open recursive nameservers that make use of such a default configuration.

In addition, due to the small query-large response potential of the DNS system it is easy to yield great amplification of the source traffic as reflected traffic towards the victims.

DNS authoritative servers which do not provide recursion to clients can also be used as amplifiers; however, the amplification potential is greatly reduced when authoritative servers are used. It is also not practical to restrict access to authoritative servers to a subset of the Internet, since their normal operation relies on them being able to serve a wide audience, and hence the opportunities to mitigate the scale of an attack by modifying authoritative server configurations are limited. This document's recommendations are concerned with recursive nameservers only.

In this document we describe the characteristics of the attack and recommend DNS server configurations that specifically alleviate the problem described, while pointing to the only truly real solution: the wide-scale deployment of ingress filtering to prevent use of spoofed IP addresses [\[BCP38\]](#) (Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," May 2000.).

2. Document Terminology

[TOC](#)

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 [\[RFC2119\]](#) (Bradner, S.,

3. Problem Description

[TOC](#)

Because most DNS traffic is stateless by design, an attacker could start a DoS attack in the following way:

1. The attacker starts by configuring a record on any zone he has access to, normally with large RDATA and TTL.
2. Taking advantage of clients on non-BCP38 networks, the attacker then crafts a query using the source address of their target victim and sends it to an open recursive nameserver.
3. Each open recursive nameserver proceeds with the resolution, caches the record and finally sends it to the target. After this first lookup, access to the authoritative nameservers is normally no longer necessary. The record will remain cached for the duration of the TTL at the open recursive nameserver even if it's deleted from the zone.
4. Cleanup of the zone might, depending on the implementation used in the open recursive nameserver, afford a way to clean the cached record from the open recursive nameserver. This would possibly involve queries luring the open recursive nameserver to lookup information for the same name that is being used in the amplification.

Because the characteristics of the attack normally involve a low volume of packets amongst all the kinds of actors besides the victim, it's unlikely any one of them would notice their involvement based on traffic pattern changes.

Taking advantage of open recursive nameserver that support EDNS0 [\[RFC2671\] \(Vixie, P., "Extension Mechanisms for DNS \(EDNS0\)," August 1999.\)](#), the amplification factor (response packet size / query packet size) could be around 80. With this amplification factor a relatively small army of clients and open recursive nameservers could generate gigabits of traffic towards the victim.

With the increasing length of authoritative DNS responses derived from deployment of DNSSEC and NAPTR as used in ENUM services, authoritative servers will eventually be more useful as actors in this sort of amplification attack.

Even if this attack is only really possible due to non-deployment of BCP38, this amplification attack is easier to leverage because for historical reasons, from times when the Internet was a much closer-knit community, some nameserver implementations have been made available

with default configurations that when used for recursive nameservers made the server accessible to all hosts on the Internet. For years this was a convenient and helpful configuration, enabling wider availability of services. As this document aims to make apparent, it is now much better to be conscious of ones own nameserver services and focus the delivery of services on the intended audience of those services, be they a university campus, an enterprise or an ISP's customers. The target audience also includes operators of small networks and private server managers who decide to operate nameservers with the aim of optimising their DNS service, as these are more likely to use default configurations as shipped by implementors.

4. Recommended Configuration

[TOC](#)

In this section we describe the Best Current Practice for operating recursive nameservers. Following these recommendations would reduce the chances of having a given recursive nameserver be used for the generation of an amplification attack.

The generic recommendation to nameserver operators is to use the means provided by the implementation of choice to provide recursive name lookup service only to the intended clients. Client authorization can be usually done in several ways:

- *IP address based authorization. Use the IP source address of the DNS queries and filter them through an Access Control List (ACL) to service only the intended clients. This is easily applied if the recursive name server's service area is a reasonably fixed IP address range that is protected against external address spoofing, usually the local network.

- *Incoming Interface based selection. Use the incoming interface for the query as a discriminator to select which clients are to be served. This is of particular applicability for SOHO devices, such as broadband routers that include embedded recursive name servers.

- *Use TSIG [\[RFC2845\]](#) (Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)," May 2000.) or SIG(0) [\[RFC2931\]](#) (Eastlake, D., "DNS Request and Transaction Signatures (SIG(0)s)," September 2000.) signed queries to authenticate the clients. This is a less error prone method, which allows server operators to provide service to clients who change IP address frequently (e.g. roaming clients). The current drawback of this method is that very few stub resolver implementations support TSIG or SIG(0) signing of outgoing queries. The effective use of this method implies in most cases running a local instance of a caching nameserver or

forwarder that will be able to TSIG sign the queries and send them on to the recursive nameserver of choice.

*For mobile users use a local caching name server running on the mobile device or use a Virtual Private Network to a trusted server.

In nameservers that do not need to be providing recursive service, for instance servers that are meant to be authoritative only, turn recursion off completely. In general, it is a good idea to keep recursive and authoritative services separate as much as practical. This, of course, depends on local circumstances.

Even with all these recommendations network operators should consider deployment of ingress filtering [[BCP38](#)] ([Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," May 2000.](#)) in routers to prevent use of address spoofing as a viable course of action. In situations where more complex network setups are in place, "Ingress Filtering for Multihomed Network" [[BCP84](#)] ([Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks," March 2004.](#)) maybe a useful additional reference.

By default, nameservers SHOULD NOT offer recursive service to external networks.

5. Security Considerations

[TOC](#)

This document does not create any new security issues for the DNS protocol, it deals with a weakness in implementations. Deployment of SIG(0) transaction security should consider the caveats with SIG(0) computational expense as it uses public key cryptography rather than the symmetric keys used by TSIG. In addition, the identification of the appropriate keys needs similar mechanisms to those for deploying TSIG, or alternatively, the use of DNSSEC signatures (RRSIGs) over the KEY RRs if published in DNS. This will in turn require the appropriate management of DNSSEC trust anchors.

6. Acknowledgments

[TOC](#)

The authors would like to acknowledge the helpful input and comments of Joe Abley, Olafur Gudmundsson, Pekka Savola, Andrew Sullivan and Tim Polk.

7. IANA Considerations

[TOC](#)

This document does not define a registry and does not require any IANA action.

8. References

[TOC](#)

8.1. Normative References

[TOC](#)

[RFC1034]	Mockapetris, P., " Domain names - concepts and facilities ," STD 13, RFC 1034, November 1987 (TXT).
[RFC2119]	Bradner, S., " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[RFC2671]	Vixie, P., " Extension Mechanisms for DNS (EDNS0) ," RFC 2671, August 1999 (TXT).
[RFC2845]	Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, " Secret Key Transaction Authentication for DNS (TSIG) ," RFC 2845, May 2000 (TXT).
[RFC2931]	Eastlake, D., " DNS Request and Transaction Signatures (SIG(0)s) ," RFC 2931, September 2000 (TXT).

8.2. Informative References

[TOC](#)

[BCP38]	Ferguson, P. and D. Senie, " Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing ," BCP 38, RFC 2827, May 2000 (TXT).
[BCP84]	Baker, F. and P. Savola, " Ingress Filtering for Multihomed Networks ," BCP 84, RFC 3704, March 2004 (TXT).

Authors' Addresses

[TOC](#)

	Joao Damas
	Internet Systems Consortium, Inc.
	950 Charter Street
	Redwood City, CA 94063
	US
Phone:	+1 650 423 1300

Email:	Joao_Damas@isc.org
URI:	http://www.isc.org/
	Frederico A. C. Neves
	NIC.br / Registro.br
	Av. das Nacoes Unidas, 11541, 7
	Sao Paulo, SP 04578-000
	BR
Phone:	+55 11 5509 3511
Email:	fneves@registro.br
URI:	http://registro.br/

Full Copyright Statement

[TOC](#)

Copyright © The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.