

Preventing Use of Recursive Nameservers in Reflector Attacks

Status of This Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Abstract

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

Table of Contents

1. Introduction	2
2. Document Terminology	2
3. Problem Description	2
4. Recommended Configuration	4
5. Security Considerations	5
6. Acknowledgments	5
7. References	5
7.1. Normative References	5
7.2. Informative References	6

1. Introduction

Recently, DNS [RFC1034] 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, except that DNS relies heavily on UDP, the easy abuse of which is at the source of the problem. The attacks 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 that 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 impractical 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; 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 real solution: the wide-scale deployment of ingress filtering to prevent use of spoofed IP addresses [BCP38].

2. Document Terminology

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].

3. Problem Description

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 Time to Live (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 at the open recursive nameserver for the duration of the TTL, 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 an open recursive nameserver that supports EDNS0 [RFC2671], 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 [RFC4033] and NAPTR resource records as used in ENUM services, authoritative servers will eventually be more useful as actors in this sort of amplification attack.

Even if this amplification attack is only possible due to non-deployment of BCP38, it is easier to leverage because of historical reasons. When the Internet was a much closer-knit community, some nameserver implementations were 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 one's 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

In this section we describe the Best Current Practice for operating recursive nameservers. Following these recommendations would reduce the chances of any given recursive nameserver being 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 to only the intended clients. Client authorization can usually be done in several ways:

- o 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 nameserver's service area is a reasonably fixed IP address range that is protected against external address spoofing, usually the local network.
- o 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 (Small Office, Home Office) devices, such as broadband routers that include embedded recursive nameservers.
- o TSIG [RFC2845] or SIG(0) [RFC2931] signed queries to authenticate the clients. This is a less error prone method that 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.
- o For mobile users, use a local caching nameserver 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] 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] maybe a useful additional reference.

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

5. Security Considerations

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 [RFC2931] should consider the caveats with SIG(0) computational expense as it uses public key cryptography rather than the symmetric keys used by TSIG [RFC2845]. In addition, the identification of the appropriate keys needs similar mechanisms as those for deploying TSIG or, alternatively, the use of DNSSEC [RFC4033] signatures (RRSIGs) over the KEY RRs if published in DNS. This will in turn require the appropriate management of DNSSEC trust anchors.

6. Acknowledgments

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

7. References

7.1. Normative References

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

- [RFC2931] Eastlake, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, September 2000.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

7.2. Informative References

- [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.
- [BCP84] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

Authors' Addresses

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

Copyright (C) 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.