



TÉCNICO
LISBOA

Dealing with DNS Amplification Attacks using Response Rate Limiting

Pedro Miguel Martins Madeira

Thesis to obtain the Master of Science Degree in

Telecommunications and Informatics Engineering

Supervisors: Prof. Miguel Nuno Dias Alves Pupo Correia
Dr. Assis Guerreiro

Examination Committee

Chairperson: Prof. Ricardo Jorge Fernandes Chaves
Supervisor: Prof. Miguel Nuno Dias Alves Pupo Correia
Member of the Committee: Prof. Fernando Manuel Valente Ramos

October 2018

Resumo

O Sistema de Nome de Domínio (DNS) providencia uma das funcionalidades fundamentais da Internet como a conhecemos, a resolução de nomes para endereços IP e vice-versa. Infelizmente o DNS depende fortemente do protocolo de transporte UDP, tornando-o vulnerável a ataques DDoS de reflexão. Adicionalmente, o DNS também permite aos atacantes amplificar o tráfego malicioso através de servidores DNS. Nestes ataques pequenas queries, feitas aos servidores, podem gerar respostas muitas vezes maiores às do seu tamanho. A combinação destes dois fatores fez do DNS um dos sistemas mais explorados por ataques DDoS, cuja escala também tem vindo a subir em anos recentes. Como entidade que supervisiona a função crítica do domínio de topo (TLD) correspondente a Portugal (.pt), a DNS.pt tem a responsabilidade de evitar causar um grande impacto económico e social nos seus parceiros, não apenas danificando outros negócios como também a sua própria reputação. Um mecanismo introduzido recentemente, chamado Limitação de Ritmo de Resposta (RRL), demonstrou ser eficaz na mitigação de alguns destes tipos de ataques. Este trabalho estuda a aplicação do RRL para o TLD de Portugal, gerido pela DNS.pt.

Palavras-chave: DNS, DDoS, Amplificação, Limitação de Ritmo, Nomes de domínio, Domínio de topo

Abstract

The Domain Name System (DNS) provides one of the most fundamental functionalities of the Internet as we know it, which is resolving names into IP addresses and vice-versa. Unfortunately it relies heavily on the UDP transport protocol making it vulnerable to DDoS reflection attacks. It also allows attackers to amplify their malicious traffic through DNS servers. Small queries can generate responses many times their size. Combining these two factors made DNS one of the most abused systems by DDoS attacks, that have also been increasing in scale in recent years. As the entity that oversees the critical function of the Top-Level-Domain corresponding to Portugal (.pt), DNS.pt has the responsibility to avoid causing a major economical and social impact in its partners, not only damaging other businesses but also their own reputation. A recently introduced mechanism called Response Rate Limiting has shown to be effective in mitigating these kind of attacks. This work studies its application for the Top-Level Domain of Portugal (.pt), managed by DNS.pt.

Keywords: DNS, DDoS, DNS Amplification, Rate Limiting, Domain Names, Top-Level Domain.

Contents

Resumo	iii
Abstract	v
List of Tables	ix
List of Figures	xi
Glossary	xiii
1 Introduction	1
1.1 Motivation	2
1.2 Objectives	2
1.3 Thesis Outline	2
2 Background	3
2.1 Domain Name System	3
2.1.1 Main elements	3
2.1.2 DNS Protocol	5
2.1.3 Query Processing	6
2.1.4 DNS extensions (EDNS)	8
2.1.5 DNS Security Extensions (DNSSEC)	8
2.2 Anycast Routing	10
2.3 Denial of Service Attacks (DoS)	11
2.3.1 Distributed Denial of Service (DDoS)	11
2.3.2 DNS Amplification and Reflection	11
2.3.3 The problem of open recursion	15
2.4 DDoS Prevention, Detection and Mitigation	16
2.4.1 Filtering-based mechanisms	16
2.4.2 Rate-Limiting	19
2.5 Summary	20
3 Response Rate Limiting	21
3.1 Concept and details	21
3.1.1 Response generation process	21
3.1.2 Configuration parameters	22

4 Study Methodology	25
4.1 Testing the effectiveness of RRL	25
4.2 Finding the ideal settings	26
5 Measurements and Results	29
5.1 RRL behaviour in a TLD-like zone	29
5.2 Repeating query attacks	30
5.3 Varying query attacks	31
5.3.1 NXDOMAIN abuse (0% resolvable domain names)	31
5.3.2 Mixed attack (50% resolvable domain names)	32
5.3.3 Fully-indexed zone abuse (100% resolvable domain names)	33
5.4 Results summary	34
6 Conclusions	35
6.1 Challenges and Future Work	36
Bibliography	37

List of Tables

2.1	Potential amplification of different maximum response sizes.	15
4.1	Probability of false positives for different SLIP values	27
5.1	Outgoing traffic volume of repeating query attacks.	31
5.2	Outgoing traffic volume of an attack that results in NXDOMAIN responses.	32
5.3	Outgoing traffic volume of a mixed query attack.	33
5.4	Outgoing traffic volume of attack randomly requests one of all the zone records.	33

List of Figures

2.1	The tree structure of the domain name space.	4
2.2	Structure of a DNS protocol message.	5
2.3	DNS name resolution process.	7
2.4	The chain of trust with DNSSEC.	10
2.5	Network traffic routing schemes.	10
2.6	A Distributed Denial of Service Attack illustration.	12
2.7	DNS protocol messages format [2]: query and response.	13
2.8	Using open DNS resolvers to conduct amplification attacks.	16
2.9	Example of multiple routing domains, or ASes, connected.	17
3.1	Diagram of RRL response generation process.	22
5.1	A repeating query attack against a test server hosting a fictitious zone similar to a TLD. RRL is enabled at 80 seconds using SLIP=1 and then SLIP=2 at 120 seconds.	31
5.2	A varying query attack that results in 100% NXDOMAIN responses. RRL is enabled at 38 seconds using SLIP=1 and then SLIP=2 at 58 seconds.	32
5.3	A varying query attack that returns a mix of NXDOMAIN and NOERROR responses. RRL is enabled at 30 seconds using SLIP=1 and then SLIP=2 at 60 seconds.	33
5.4	A varying query attack that uses all the records of the zone. RRL is enabled at 30 seconds using SLIP=1 and then SLIP=2 at 60 seconds.	34

Glossary

BGP	Border Gateway Protocol
BIND	Berkeley Internet Name Domain
CPU	Central Processing Unit
C&C	Command and Control
DAAD	DNS Amplification Attack Detector
DDoS	Distributed Denial of Service
DNSKEY	DNSKEY record used by DNSSEC
DNSSEC	DNS Security Extensions
DNS	Domain Name System
DS	Delegation of Signing by DNSSEC
DoS	Denial of Service
EDNS	Extensions mechanisms for DNS
FTP	File Transfer Protocol
IP	Internet Protocol
ISP	Internet Service Provider
OSPF	Open Shortest Path First
QR	Query-Response flag in DNS message header
RPF	Router-based Packet Filtering
RRL	Response Rate Limiting
RRSIG	Resource record signature used by DNSSEC
RR	DNS resource record
TCP	Transmission Control Protocol
TLD	Top Level Domain
UDP	User Datagram Protocol

Chapter 1

Introduction

One of the most important systems that make up the Internet as we know it today is the Domain Name System (DNS). It provides one of the most fundamental functions of the Internet which is resolution of domain names into IP addresses and vice versa (whether they are IPv4 or IPv6). In its turn, DNS relies on the UDP transport protocol, making it particularly vulnerable to Distributed Denial of Service (DDoS) attacks.

A DNS server may not only be the target of a DDoS attack, but also the accomplice of one. Since UDP is stateless, there is no validation of the source address of packets in the network, making it impossible for a DNS server to distinguish a forged request from a legitimate one allowing attackers to send forged DNS requests spoofed with their victims address. This results in a reflection of malicious traffic on DNS servers. This could be eliminated by enabling source-based filtering techniques on Internet Service Providers (ISP) networks. Unfortunately, deployment of these is not being widely adopted which makes them ineffective.

Additionally, the DNS service allows attackers to amplify their malicious traffic since simple queries to the DNS service can yield much larger answers. For instance, a 30 bytes long query message can generate an answer over 3000 bytes.

These two factors, reflection and amplification, make DNS an attractive system to carry attacks that may not even target DNS servers themselves but other companies and organizations. When this happens, the victims see the DNS service as their offenders. DNS.PT Association is the national registry for the Internet's Top-Level Domain (TLD) corresponding to Portugal (.pt). They oversee management, assignments, operation and maintenance of the domain registries. Due to DNS.PT's importance in securing essential critical functions at economic and social levels, disturbance of its operation can have a major impact in the wellbeing of the country. It is therefore its responsibility to conduct studies and develop solutions that promote the protection of their services.

1.1 Motivation

Although there are already ways of reducing the impact of DDoS attacks that abuse DNS servers, most of the the solutions suffer from either deployment issues or do not apply in the case of service providers being abused as conveyors of attacks, such as DNS authorities. Network filtering solutions are dependable on being massively adopted or used as defense mechanisms by victims and not by the DNS authorities themselves. On the other hand, the few existing DNS specific approaches to the problem are either ineffective or unable to prevent legitimate clients from being denied service.

A recent mechanism called Response Rate Limiting (RRL), deployed inside DNS servers themselves, was specifically created to deal with the issues of reflection and amplification. There has not been a lot of feedback on how the mechanism behaves and impacts the normal operation of a DNS service. Because of that, deploying it in such an important service as the one provided by a TLD should be preceded by a careful study on how to configure it to avoid harming legitimate users.

1.2 Objectives

The main goal of this work is to study the usage of the RRL feature on the DNS service provided by DNS.pt as a way of mitigating DDoS Attacks that may use DNS servers as reflectors and amplifiers. By doing so, it should be achieved a better understanding of how the feature would work in a TLD server scenario. Additionally, assess its limitations and discuss how the mechanism may be improved.

1.3 Thesis Outline

This document is organized as follows: Section 2 presents background on the DNS as well as DDoS Attacks, particularly the DNS Amplification and Reflection Attack. It also briefly reviews other mechanisms proposed to deal with these types of attacks. Section 3 focuses on RRL, how it functions and what are its configuration options. Section 4 goes over important aspects of each configuration option while giving some ideas on how to fine tune them for specific environments. Additionally, we present a list of attacks scenarios that will be simulated in order to evaluate the effectiveness of different configurations of the RRL feature.

Chapter 2

Background

2.1 Domain Name System

In the early days of the Internet, host name to address mappings were stored in a plain text file called HOSTS.TXT. This file was maintained by the Network Information Center (NIC) which was ran by Stanford Research Institute. Each time a new mapping was created, hosts of the ARPANET would have to transfer the file from an FTP host at NIC. As the number of hosts increased so did the necessary network bandwidth to distribute a new version of the file. It also meant that organizations that managed their host names locally would have to wait for the NIC to update the hosts file so they would be visible to all hosts of the Internet. With an exponentially growing number of hosts and other requirements like name structuring inside organization, a centralized plain text file became impractical.

The Domain Name System (DNS) proposed by Mockapetris in [1, 2] provides a decentralized, scalable and extensible database system that replaced the centralized plain text file scheme. Its main goal was to have a consistent name space using a hierarchical structure that could be accessed and maintained in a distributed manner. DNS provides standard formats for resource data as well as methods for querying and updating the database.

2.1.1 Main elements

DNS can be broadly described as a distributed database system that is composed of three major elements: the domain name space and resource records, the name servers and the resolvers. Each of these elements is described ahead.

Domain name space and resource records

Domain name space is a tree structure that defines the hierarchy of names in the DNS. Names are composed of one or more labels separated by a dot. The root of the tree is the empty name and all names are its children. For example, the domain `example.com` has two labels: `example`, the domain itself and `com`, the name of its parent domain. The right-most label always refers to a Top-Level Domain

(TLD) which is immediately below the root of the name space tree [3]. Figure 2.1 illustrates an portion example of the domain space tree.

Resource records (RR) are the information held by the nodes in the tree. Host name mappings under `example.com` are RRs held by the name server of `example.com`. Different types of information are stored by different RR types [1].

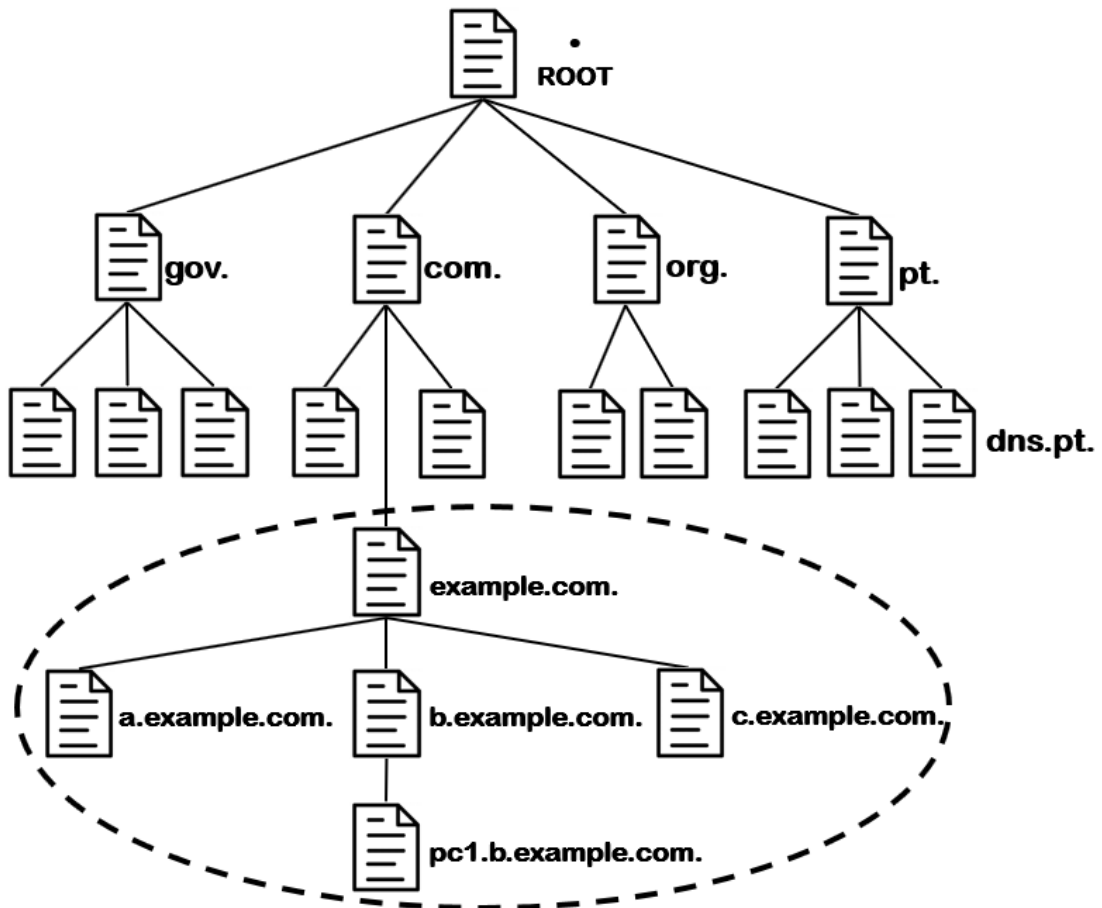


Figure 2.1: The tree structure of the domain name space.

Name Servers

Name servers hold information about a zone of authority. A zone of authority is a group of domains that share a common parent domain and are managed by the same name server. The zone `example.com` has a name server holding the information for the domain `example.com` and its sub-domains. A name server is said to be authoritative for the domains that it manages. A client is able to reach any point or resource of tree by querying at least one name server. If a name server does not have the requested record it either points the requester to a name server that has it or to the ROOT servers¹, allowing any point of the domain name space tree to be reached using any name server.

There are two types of name servers: primary (or master) and secondary (or slave). Primary name

¹<https://www.iana.org/domains/root/servers>

servers are authoritative name servers that contain zone master files that define zones. Updates or additions to the database are made in these zone files stored on the primary servers. Secondary name servers are authoritative name servers that do not have the master file and instead receive it from the primary server. A secondary server does not update RRs, it only holds them and serves them to clients.

Resolvers

Resolvers are the programs that, by client request, extract information from the name space tree by consulting at least one name server. If a client needs the IP address corresponding to the name `web.example.com`, a resolver will craft the query and send it to the necessary name servers in order to obtain the resource record containing the requested address. Resolvers implement a recursion algorithm that from a single request from the client may result multiple queries to different name servers in order to reach an answer. Resolvers are usually DNS servers that offer recursion (see section 2.1.3). Resolver and recursive name server are used interchangeably throughout this document. A simpler resolver, or stub-resolver, is commonly used by operating systems that usually just build the query to be sent but do not go through the recursion algorithm. Instead, the stub-resolver offloads that work to a resolver.

2.1.2 DNS Protocol

Although the term DNS is usually employed as the name of the system, it can also be used as the name of the protocol itself, i.e., how DNS operates. DNS protocol essentially has two operating modes: lookups or zone transfers.

For lookups it is a stateless protocol meaning that all the information that is needed is contained in individual messages, with the structure shown in Figure 2.2. Although both TCP and UDP protocols can be used to transport lookup messages, UDP is the most common. If messages are lost when using UDP, a retransmission scheme should be used by the client.

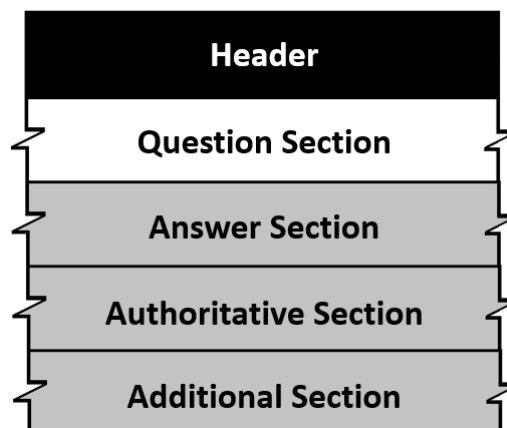


Figure 2.2: Structure of a DNS protocol message.

- **Header:** Compose by an opcode that states what operation is being made, a response code (RCODE)², the id of the request and corresponding reply, a 1 bit field (QR) to distinguish questions from responses, additional flags related to message truncation and recursion as well as the count of RRs for each of the sections following.
- **Question:** These are always Name, Type and Class tuples. Name is a fully-qualified domain name (FQDN), Type is a valid RR type and Class is IN for Internet applications. The question is the only one section besides the header section that is present in query messages. All others are only used for reply messages.
- **Answer:** RRs that match the tuple in the question section. It may have multiple RRs has it is possible to have multiple RRs for the same label as, for example, multiple IPv4 addresses for the same host name.
- **Authority:** Name server (NS) records that point to name servers closer to the target name in the hierarchy.
- **Additional:** RRs that the name server considers useful to the client. Usually IPv4 (A) and IPv6(AAAA) RRs for the NS records of the authority section.

Zone transfers are made between name servers, most commonly at the request of a secondary server. Using a very similar message to a normal query but using TCP and the special Type fields AXFR or IXFR in the question section, RRs are transferred from a name server to another, with the transfer of the Start of Authority (SOA) record marking the end of the zone transfer.

2.1.3 Query Processing

To access and retrieve information from DNS, a client must query at least one name server. There are two types of query that can be made: a recursive or non-recursive query. It is by client request which query type should be performed but because not all name servers provide both methods, it is ultimately decided by them how the query will be processed.

Non-recursive method

The most basic method of requesting information is the non-recursive query and it must be supported by all name servers. It is a simple, single question and response process in which a client either receives the requested information or a referral to another name server that may have the answer or will get the client closer to getting the information. If the name server has the records that point to the authoritative name server of one of the labels, it refers the client to that name server. If it does not have records for any of the labels, it refers the client to the authoritative name servers of the root of the tree which are known by all name servers.

²<https://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml>

Recursive method

Another way of retrieving a resource from the name space tree is to request recursion to the queried name server, i.e., to perform a recursive query. Not all name servers will provide recursion to their clients. A name server that provides recursion to its clients is also considered a resolver or a recursive name server. If the requested resource exists, the client will receive it even if the name server does not have it. A recursive name server will perform all the necessary queries to retrieve the information and send it to the client.

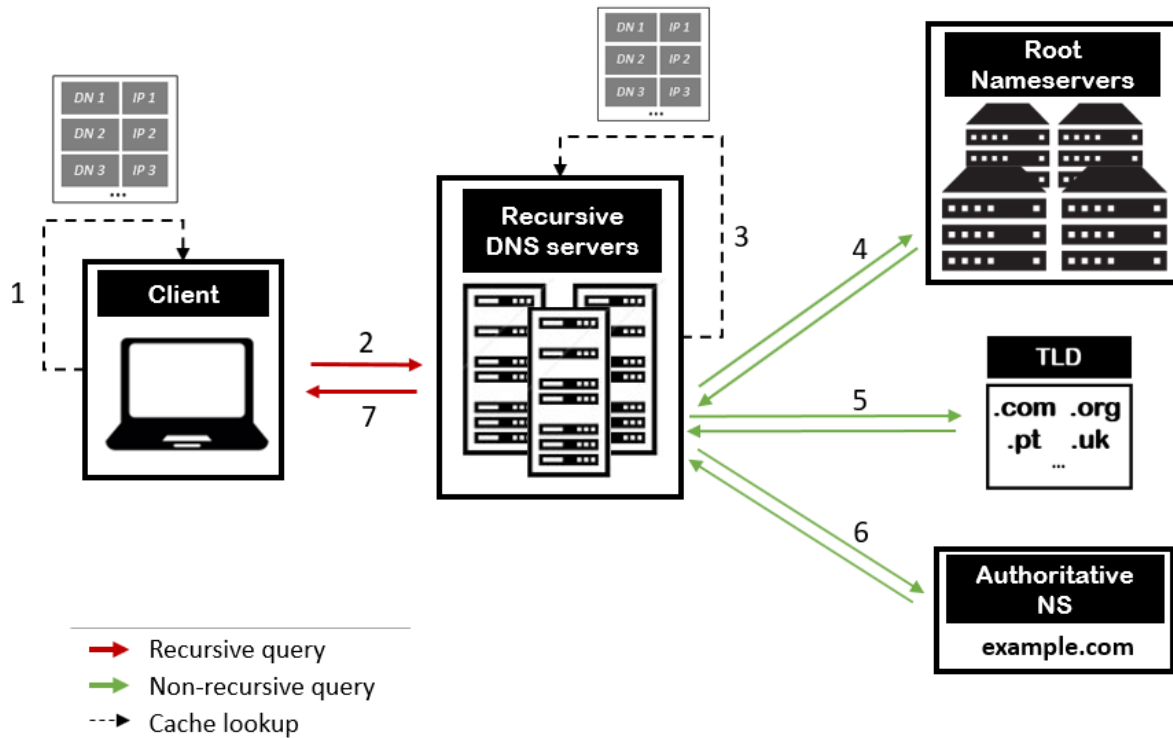


Figure 2.3: DNS name resolution process.

When a user or application needs to resolve, for example, the host name `web.example.com`, the lookup-and-handoff process to obtain the correct IP address in order to reach it is shown in Figure 2.3 and follows these steps:

1. The client searches its stub resolver cache to see if it already has that resource. Caching recently obtained resource records is a way of reducing load on name servers and saving time by not having to repeat the process each time a resource record is requested;
2. After a cache miss, the clients stub resolver will query a name server, usually one provided by its ISP that also offers recursion;
3. The resolver, will also search for the resource record in its cache and if does not find a cached answer it will start the recursive querying process to get it;
4. The resolver queries a ROOT server, which will not answer the question but will send a referral to the authoritative name servers of the rightmost domain, i.e., the TLD `com`;

5. The resolver now queries the TLDs authoritative name server and receives a referral pointing to the authoritative name server for the domain `example.com`;
6. The resolver queries the `example.com` authoritative name server which returns the requested resource records for `web.example.com`;
7. After receiving the answer, the resolver sends it back to the requesting client.

2.1.4 DNS extensions (EDNS)

The initial DNS protocol specification defined a small set of 1-bit flags on the header section of the messages. As time passed, additional features were introduced and for DNS servers and clients to use those features, additional flags had to be introduced. Changing the header section structure meant that older implementations would not be able to communicate with newer ones.

To maintain backwards compatibility, any additional flags would have to be added outside the header section. EDNS [4] extends the DNS protocol messages through the creation of the OPT pseudo-resource-record including it in the already existing additional section of the messages. It is called a pseudo-resource-record because it does not contain any resource data and carries only control information as, for instance, additional flags.

One particular feature introduced by EDNS is the option to specify a UDP payload size larger than the original 512 bytes limit using the 2-byte CLASS field of the OPT resource record. Earlier research [5] as shown that most of the clients using EDNS advertise a maximum UDP payload size over 4000 bytes.

This means the usage of EDNS can potentially result in higher amplification, because it allows for significantly larger responses to be sent. In [6], the authors show that even though open resolvers have a higher amplification potential (see Section 2.3.3), due to higher control of the responses by the attackers, authoritative name servers that supported EDNS could allow much higher amplification ratio than open resolvers that do not support EDNS.

2.1.5 DNS Security Extensions (DNSSEC)

To protect DNS against most of the threats described in [7], DNSSEC [8–10] was introduced as a security extension to DNS providing origin authentication and integrity protection for DNS data while still maintaining backwards compatibility.

DNSSEC protects clients and applications from using forged or manipulated DNS data as it happens, for instance, with DNS cache poisoning attacks. DNSSEC protected zones return their answers with digital signatures. A DNS resolver checks the digital signatures in order to confirm if the information is unmodified and complete as published by the zone owner on its authoritative DNS server.

The following RR types [9] were created in order to facilitate signature validation:

- **RRSIG**: Resource Record Signature containing the digital signature for a resource record set (RRset);

- **DNSKEY:** Contains the public key that should be used to verify the signatures on RRSIG records;
- **DS:** Placed on parent zones alongside with the delegating NS records;
- **NSEC and NSEC3:** Used to verify the non-existence of a record name and type;
- **CDNSKEY and CDS:** Used when updating Delegation Signer records in the parent zone.

Zone-Signing Keys (ZSK)

To secure a zone with DNSSEC a Zone-Signing key pair (ZSK) is created. Resource records are group by label and type, forming resource record sets (RRset) which are digitally signed by the private key of ZSK thus creating RRSIG records for each RRset. The public key of ZSK is published as a DNSKEY record in the same zone file.

Key-Signing Keys (KSK)

Before validating RRSIG records, the DNSKEY record itself needs to be validated. An additional key pair, the Key-Signing pair, is created to validate the ZSK in the same way the former validates RRSIG records. The private key is used to digitally sign the public ZSK creating an RRSIG for the DNSKEY record. The public KSK is then published in another DNSKEY record which together with the DNSKEY record of the public ZSK form the DNSKEY RRset. This RRset is signed with the private KSK.

Delegation Signer (DS) Records

DNSSEC introduces the Delegation Signer (DS) record for the purpose of transferring trust from a parent zone to a child zone. A zone operator hashes the public key of the KSK pair and gives it to the parent zone to be published as a DS record. When a resolver is referred to the child zone by the parent zone it also receives the DS record that it can then use to verify that the child's public KSK has not been tampered.

Chain of Trust

Similarly to X.509 certificates, to verify that a certificate can be trusted it is necessary to transfer the trust from an already known trusted authority [11]. In DNSSEC this is done by establishing a chain of trust through DS records. Resolvers can establish this chain by verifying DNSKEY records beginning in the ROOT zone and going down the tree during the recursive algorithm. Figure 2.4 illustrates an example of a chain. The ROOT zone DNSKEY is the only one which doesn't have a corresponding DS record, since there is no parent to it. Instead, the ROOT DNSKEYs are signed and published in a ceremony³. These keys should be obtained and installed by operators of DNS resolvers or shipped with DNS resolver software packages.

³<https://www.iana.org/dnssec/ceremonies>

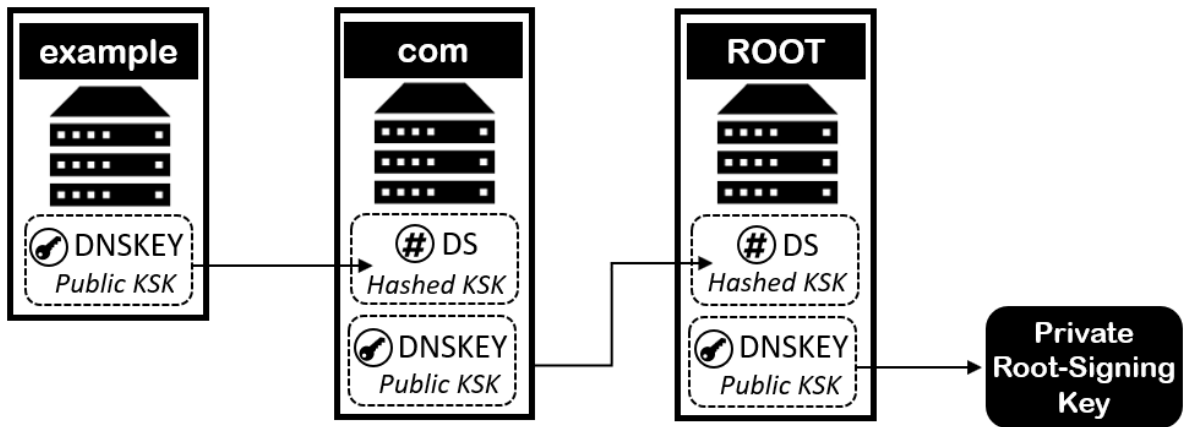


Figure 2.4: The chain of trust with DNSSEC.

2.2 Anycast Routing

Anycast is an addressing and routing technique for networks that allows hosts in different locations to share the same IP address. A sender routes datagrams (using OSPF, static routing, etc.) to the nearest node, or nodes if the destination address is the same, of a group of receivers, this process is called one-to-nearest association. Figure 2.5 shows how anycast works compared to other routing techniques. The used route is based on the routing protocol costs and state of the advertising servers. Anycast exploits the per-flow load-balancing capabilities of common router vendors when the routes between nodes all have the same weights, but also acts as a fail-over mechanism when the routes between nodes have different weights.

Technique	Description	Example
Unicast (or Point-to-Point)	A single sender communicates with a single receiver in the network	
Multicast	A single sender communicates with multiple receivers in the network	
Anycast	A single sender communicates with the nearest receivers	

Figure 2.5: Network traffic routing schemes.

Anycast is mostly used by UDP based services. When a user sends a datagram it sends it to an anycast "cloud" with a single IP address. Depending on the anycast configuration used that datagram will be sent to the nearest server belonging to that cloud, reducing latency for the user and spreading the load on the servers.

In DNS, anycast works by announcing the same route out of multiple datacenters in the world. Behind one IP address there is a highly redundant, high performance, geographically spread cluster of physical servers. This means it's not necessary to configure different DNS servers for computers deployed in different parts of the world. This also creates a certain level of availability, because if, while doing a DNS query for a very common domain the server fails, the route is removed and quickly forwarded to the next closest server without any intervention or reconfiguration. The route calculation and transmission of packets from the DNS client to the DNS server is handled dynamically by layer 3 routing.

Anycast can also help deal with server overload because it allows horizontal scalability. If a server is overloaded with requests additional ones can be seamlessly deployed and ease the load.

2.3 Denial of Service Attacks (DoS)

A DoS is a type of attack where the goal is to make a service, machine, website or network resource unavailable to its users. DoS attacks can be divided into two major categories: attacks that exploit vulnerabilities in the software causing a service to crash or function in unpredictable ways and attacks that cause a system or service to be unusable or unresponsive by exhausting critical system resources like memory, CPU or network bandwidth. Most cases in the first category can be avoided by simply patching those vulnerabilities, but the same does not apply to the second category. Even in the absence of vulnerabilities, resource exhaustion may still occur.

2.3.1 Distributed Denial of Service (DDoS)

One of the most common types of DoS is the DDoS in which multiple attacking agents cooperate to cause excessive load to a victims service or network. This kind of attack usually takes advantage of botnets composed of multiple compromised hosts running one or more bots controlled by one master program called Command and Control (C&C) program. Figure 2.6 illustrates this kind of attack. The attacker, through the C&C program, sends instructions to the compromised hosts running the bots to perform tasks that actually consist of the attack. This not only hides the real attacker but allows it to scale the impact on its victim(s). In fact, recent DDoS attacks have reached traffic volumes close to 1 Tbps [12] with hundreds of thousands of devices participating simultaneously in them [13].

2.3.2 DNS Amplification and Reflection

A DNS Amplification and Reflection attack is a type of DDoS attack that takes advantage of DNS servers and protocol in order to flood a victim's system or network with a large volume of traffic [14].

This type of attack is based on the DNS protocol, which is a request-response protocol relying heavily on UDP [2]. Unlike TCP, there is no handshake between the system that sends a request to a server. This means the server will immediately send a response to the source IP address that comes in the request message. Attackers exploit this by spoofing the source IP address in order to make the DNS server send the response to their targets.

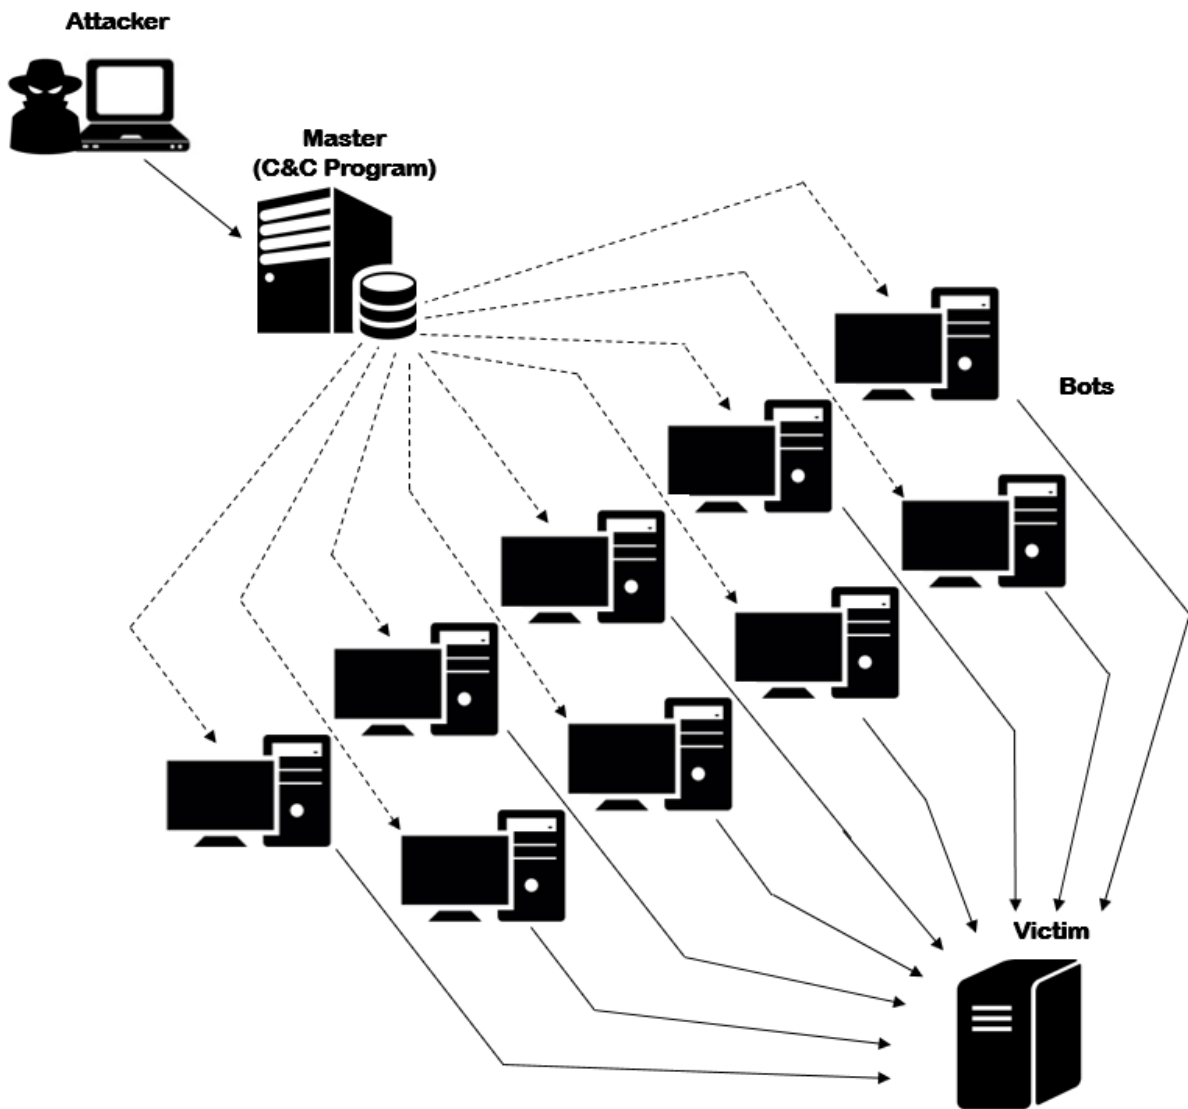


Figure 2.6: A Distributed Denial of Service Attack illustration.

Reflection itself is not the most attractive exploit for attackers. Because DNS response messages are almost always larger than their requests (see Figure 2.7), a DNS server acts as a traffic volume amplifier. Response messages carry the same question section as the request message plus all the records in the answer, authority and additional sections. The amplification factor is given by the ratio between the size of the response message and the request message.

Not including IP and UDP headers, a request message usually has a few dozen bytes depending on the size of the QNAME field. A small response message with only one resource record in each of the three previously mentioned sections can double or triple the size of the request message. However, a study by MacFarland et al., in which resource records of nine TLDs were queried, shows that much greater amplification can be easily obtained [15].

Due to this amplification effect, attackers are able to generate enormous amounts of traffic towards their victims; the spoofing of the source IP address prevents the attackers from being blacklisted because their real IP address is unknown; and the victims may not be able to blacklist the IP address of the DNS

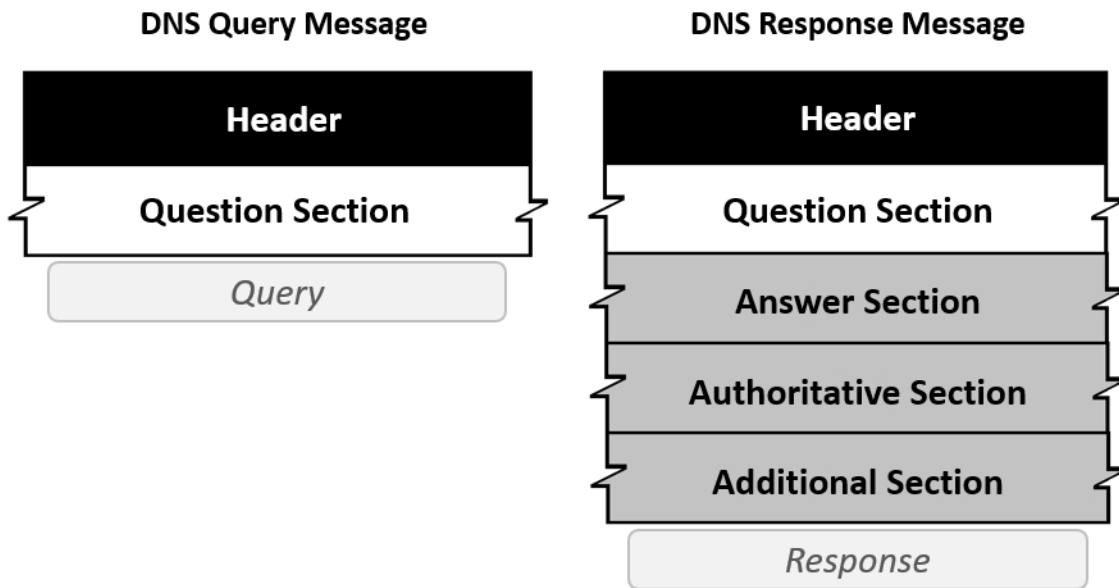


Figure 2.7: DNS protocol messages format [2]: query and response.

servers because it eliminates all legitimate DNS resolutions [6, 15].

Query types and amplification

One aspect attackers consider in order to greatly increase the volume of traffic sent to their victims is the types of query they send to the servers. A particular query type, the *ANY* query, asks the server to return all known resource records associated with a name. For instance, issuing *ANY* queries for base domain names commonly returns *SOA*, *NS* and *MX* records along with *A* records associated with the host names [16]. Because of this it is very common to see attackers using the *ANY* query type. In fact, in most of the observed attacks of this type the spoofed queries sent by attackers were of type *ANY* [14].

However, attacks continuously increase in sophistication and there are other ways to achieve higher amplification ratios. The *TXT* resource record, for example, allows arbitrary text to be stored using the DNS [17]. This had led to a recently observed tendency of attackers crafting *TXT* records to increase the impact of their attacks [18].

Effect of EDNS and DNSSEC

As stated before in Sections 2.1.4 and 2.1.5, DNSSEC introduces additional resource records and most of them can be very large due to carrying digital signatures. Since it implicitly requires the use of EDNS, signaling support for DNSSEC allows the usage of the UDP payload extension. DNSSEC-signed domains may be used to achieve even higher amplification ratios.

Below is a DNS response message, obtained using the `dig` tool, after requesting the *A* record of `tecnico.ulisboa.pt` without DNSSEC support. The received response has a total of 239 bytes (Listing 2.1).

Listing 2.1: DNS Response for tecnico.ulisboa.pt without DNSSEC.

```
dig @a.ul.pt tecnico.ulisboa.pt

; <<>> DiG 9.8.3-P1 <<>> @a.ul.pt tecnico.ulisboa.pt
; (1 server found)
;; global options: +cmd
;; Got answer:
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 41245
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 6
;; WARNING: recursion requested but not available

;; QUESTION SECTION:
;tecnico.ulisboa.pt.      IN  A

;; ANSWER SECTION:
tecnico.ulisboa.pt. 3600 IN  A 193.136.128.169

...

;; Query time: 10 msec
;; SERVER: 194.117.0.150#53(194.117.0.150)
;; WHEN: Wed May 31 04:21:34 2017
;; MSG SIZE rcvd: 239
```

Without additional resource records introduced by DNSSEC, the amplification is already significant. The request message size is around 30-40 bytes. The response message achieves an amplification ratio of roughly 7 times the size of the request.

Performing the same query to the same server with DNSSEC support (do flag set) returns a response message with 1470 bytes as shown in Listing 2.2.

Listing 2.2: DNS Response for tecnico.ulisboa.pt with DNSSEC.

```
dig @a.ul.pt tecnico.ulisboa.pt +dnssec

; <<>> DiG 9.8.3-P1 <<>> @a.ul.pt tecnico.ulisboa.pt +dnssec
; (1 server found)
;; global options: +cmd
;; Got answer:
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 11343
;; flags: qr aa rd; QUERY: 1, ANSWER: 2, AUTHORITY: 4, ADDITIONAL: 12
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 1480
;; QUESTION SECTION:
;tecnico.ulisboa.pt.      IN  A

;; ANSWER SECTION:
tecnico.ulisboa.pt. 3600 IN  A 193.136.128.169
tecnico.ulisboa.pt. 3600 IN  RRSIG A 8 3 3600 20170614013524 20170515000508 35313 tecnico.ulisboa.pt.
    cM7fVmHz4jArON11xGTF8fhnGjj+Jrf71pSM0pZQPpIhqZ/RofCa8cem
    htpkRHoMOVBeecJeRz8a58C1WIRc5j0UnBTZjEFYOR17LOKiM6iqDiUS q1REVCPdcQVLL/T9toJBRMD6Zk1ddUV7/
    DyRre2NAGlaYwDd2WX8Aexz 1iQ=

...


```

```
;; Query time: 12 msec
;; SERVER: 194.117.0.150#53(194.117.0.150)
;; WHEN: Wed May 31 04:32:58 2017
;; MSG SIZE rcvd: 1470
```

In this case, the DNSSEC enabled response was roughly 6 times larger than the non-DNSSEC one. However, DNSSEC enabled requests may yield responses to query messages that are over 50 times larger than responses without DNSSEC records [19]. Table 2.1 shows examples of theoretical volumes of traffic that can be achieved with commonly advertised maximum response sizes.

Table 2.1: Potential amplification of different maximum response sizes.

Query Size (bytes)	Response Size (bytes)	Amplification Ratio	Attacker (bits/s)	Victim (bits/s)
40	512	12,8	100M	1,28G
40	1472	36,8	100M	3,68G
40	4096	102,4	100M	10,24G

2.3.3 The problem of open recursion

Name servers can operate in two modes: recursive or non-recursive. A non-recursive name server will either respond with a resource record it holds or a referral to other name servers. On the other hand, a recursive name server, i.e., a resolver, will try to resolve the clients request by recursively querying other name servers before answering back to the client. An attacker is able to craft resource records or even entire domains and use an open resolver to send them to the victim (Figure 2.8). The resolver caches the crafted records as a way to reduce load on name servers, meaning the attacker's name server does not need to provide an answer to every request sent to the resolver.

Disabling recursion in name servers or at least restricting it is highly recommended [20] when trying to avoid becoming a reflection and amplification agent in a DDoS attack. Open resolvers offer a way for attackers to control the degree of amplification they want to achieve. Google Public DNS⁴ and Open DNS⁵ are popular DNS resolvers open to the public. When a resolver is needed, it should only be available as a non-authoritative server accessible only by their intended clients, preferably ones with IP addresses that cannot be spoofed as, for instance, private networks.

However, authoritative name servers are required to be accessible to the Internet and do not restrict who can reach them, so they may still be used to conduct these attacks. In addition with an increasing number of secured domains with DNSSEC, authoritative DNS responses become more useful in providing great amplification.

⁴<https://developers.google.com/speed/public-dns/>

⁵<https://www.opendns.com/setupguide/>

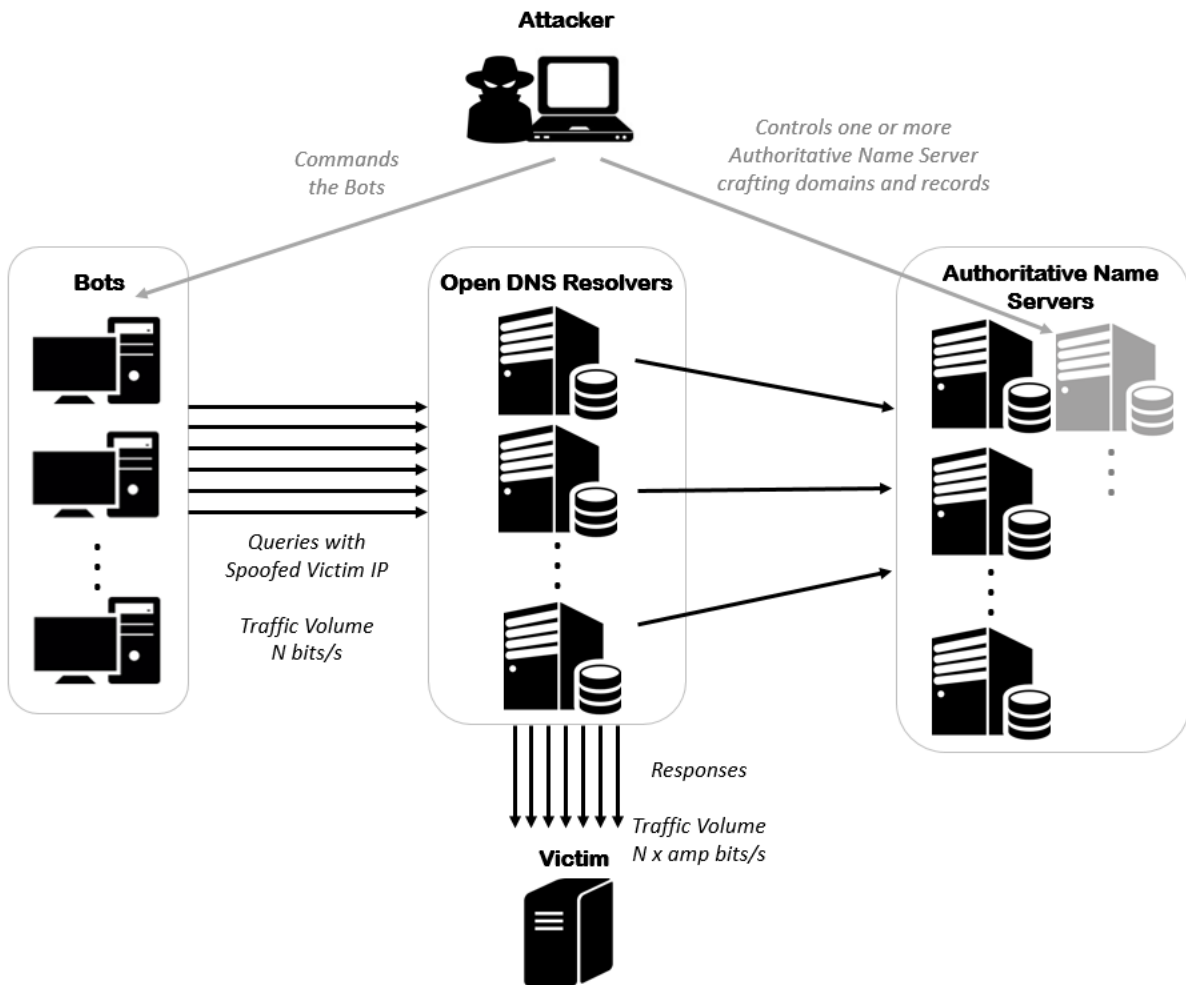


Figure 2.8: Using open DNS resolvers to conduct amplification attacks.

2.4 DDoS Prevention, Detection and Mitigation

We will look into two major types of defense mechanism against DDoS attacks: filtering and rate limiting. Although similar, since both employ some sort of packet/message dropping, these two types of defense usually work at different points or layers.

2.4.1 Filtering-based mechanisms

One way of preventing DDoS attacks is to perform filtering of network traffic at some point before it reaches the intended victim. The following are some of the filtering mechanisms mainly focused on dealing with traffic amplification and/or reflection.

Network Ingress Filtering

Network Ingress Filtering [21] is a very effective way of dealing with the key aspect of reflection attacks which is source IP address spoofing. Also known as BCP38 [21], it allows routers to check if the source address in a packet is a valid address of the network it is coming from. Internet Service Providers

(ISP) customers receive IP addresses from a certain range, e.g., 1.0.0.0/16. If the ISP's router receives packets from the customers with source addresses outside that range, e.g., 2.0.0.1, it should drop the packets. Preventing source address spoofing would not allow reflection attacks to happen. Amplification would still happen due to protocol nature but attackers would not be able to direct the amplified traffic to their victims unless their victims were on the same subnet as them.

Route-Based Packet Filtering (RPF)

Border routers are the Internet topology connectors, connecting the various routing domains (AS) that compose the Internet with each other. Routes are propagated throughout these using BGP. As the name implies, RPF performs filtering for flooding spoof traffic at these Internet border routers. RPF does this by having access to the BGP routing topology.

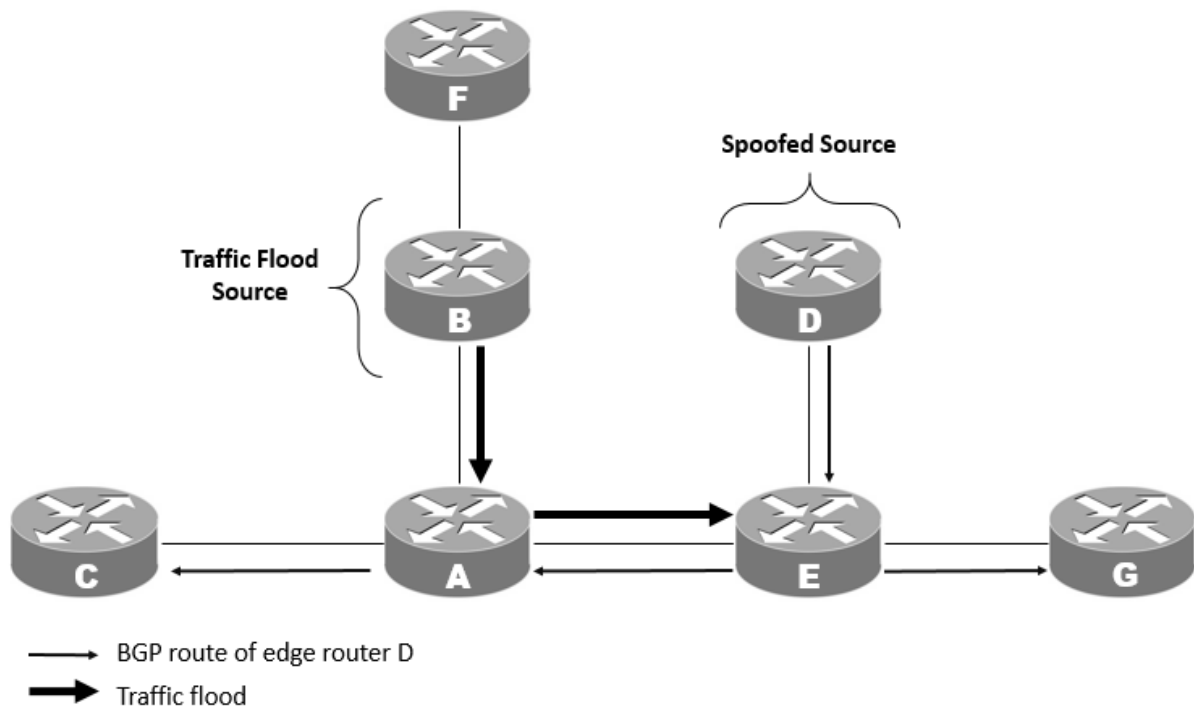


Figure 2.9: Example of multiple routing domains, or ASes, connected.

In the routing topology of Figure 2.9, consider a host inside AS B is sending traffic to a host in AS E with spoofed source address from AS D. A border router in AS A, aware of the topology, upon receiving those packets would recognize them as spoofed traffic because traffic from AS D should not come from the connection with AS B. It is essentially a sparse variant of ingress filtering. Instead of filtering in customer edge routers, filtering only needs to be done in AS border routers which are much fewer than the former.

RPF can be effective, but its major setback is that its effectiveness depends heavily in a deployment strategy, not only in number of AS that do it but also which [22, 23]. The effect of these source-address based mechanisms depends heavily on universal deployment. If only a few AS deploy this mechanism, attackers may simply choose other networks or hosts that do not stop spoofed traffic. It requires a global

combined effort for these mechanism to be successful in stopping DDoS attacks that use source IP address spoofing.

Dropping ANY queries

ANY type queries have the highest amplification potential since answers to it return all resource records associated with the domain name. Attackers frequently use this query type to increase the amplification ratio, making it the most used query type in observed DNS Amplification attacks [14]. Based on the fact that most clients do not rely on these queries, a very simple way of reducing the impact of amplification attacks would be to have a firewall drop all requests of type ANY before they reach the name servers. However, there are at least two issues with this approach. First, it would probably result in false positives, dropping legitimate queries as well. Second, and most importantly, is that with the rise of DNSSEC and other resource records like RRSIG, DNSKEY or TXT that cause large amplifications, attackers could simply switch to DNS queries of those types. An approach that consists in having a firewall drop malicious traffic would require a much higher level of sophistication to deal with equally sophisticated attacks.

D-WARD

Mirkovic et al. [24, 25] presented a scheme which aims to detect DDoS attack traffic by monitoring inbound and outbound traffic of a source network and comparing traffic information with predefined normal flow models. Malicious traffic flows are detected and filtered when they don't match the models. For instance, a TCP packet should always be acknowledged. A normal TCP traffic flow would not exceed a certain ratio of sent and received packets for one peer. Calibration for each kind of traffic reduces the chance of false positives.

D-WARD may be effective at stopping unusual traffic flows but it may not detect traffic that although being malicious and part of a reflection-based DDoS, still follows a normal pattern of a UDP traffic flow. Massive DDoS attacks, carried out by a large number of bots could easily avoid a solution that tries to detect abnormalities in traffic flows close to the source. Large volumes of traffic in DDoS reflection attacks are mostly achieved by a large number of bots sending traffic at reasonable rates instead of a few sending large volumes that could easily be identified as malicious.

DAAD

DNS Amplification Attack Detector (DAAD) [26] is based on the idea that a DNS request should create a corresponding response and any other responses that do not pair with a previous request should be marked as suspicious. It collects DNS requests and responses using the IPtraf⁶ tool. Once an IP address has raised enough suspicion due to unrequested DNS responses, DAAD blocks that IP address. This proposal aims to mitigate the effects of DDoS flooding attacks on the victim side, most likely at the edge of its network. It is an effective solution at stopping unrequested traffic. From a DNS

⁶<http://iptraf.seul.org/>

server operator perspective it becomes useless because it not only the servers receive the requests but the responses usually go out through the same path the request came.

2.4.2 Rate-Limiting

Another approach to deal with malicious or abnormal traffic, specifically traffic sent to DNS servers, is to impose rate limits. Rate-Limiting is a softer defense against DDoS that can be deployed closer to the destination of the traffic. Currently, there are two mechanisms that employ the idea of rate-limiting on DNS servers.

DNS Dampening

DNS dampening obtains query data and parameters from all the requests made and assigns them penalty points. Queries most used to perform attacks or that have a larger sized response receive higher penalty points, and they will increase if the query ID keeps repeating. The server has a limit of penalty points that can be reached. If the penalty points are higher than that limit the server will start dampening, dropping all the queries coming from that IP address. The penalty points will decrease in time, and when they get to a value under the server limit the queries will start being processed again. The problem with DNS dampening is that it's impossible to control the dampening of legitimate queries if they happen to reach the penalty limit. There is no mechanism to reduce the impact of false positive occurrences.

DNS Response Rate Limiting

Response Rate Limiting (RRL) was introduced in 2012 by Paul Vixie and Vernon Schryver [27] after observing attacks that abused DNSSEC-signed domains. The idea behind RRL is that authoritative name servers would rate limit identical outgoing responses to the same address or subnet. When the rate limit is exceeded, the server either drops the response or sets the TC (truncated) flag to signal the clients to retry using TCP instead of UDP thus denying the chance for traffic reflection.

RRL is strongly recommended for authoritative name servers which do not and should not offer recursion (see section 2.3.3). Authoritative name servers are mostly queried by other name servers, usually resolvers, meaning the vast majority of their clients will properly cache the records making reasonably safe to assume that authoritative name servers would not need to repeat the same responses to same clients or subnets. Resolvers, on the other hand, are usually queried by stub-resolvers that perform very little caching of records or applications that make the requests using third-party libraries that may not perform caching at all. Because of that, resolvers may legitimately need to send identical responses to the same client.

RRL can significantly reduce the impact of DNS Amplification Attacks that abuse authoritative name servers but it also has a few limitations or disadvantages. First, attackers are able to avoid triggering rate limiting by crafting their attacks to use a large spread of queries instead of just a few very repetitive ones. Second, it can potentially affect legitimate clients that use resolvers that do not properly cache

RRs. Third and most significant one, it can become itself the weapon of a DoS attack if an authoritative name server using it is flooded with queries using the source address of legitimate resolvers, for instance, the resolver of a large ISP. Clients of a rate limited resolver might see service degradation when trying to resolve names from the domain being flooded by the attackers.

2.5 Summary

The stateless nature of the DNS protocol paired with its amplification from request to response make DNS servers a great way for attackers to scale their attacks. A DNS server can become either the target of a DDoS attack or the accomplice of one. In this work, we study the second scenario in which a DNS server may not actually be affected by malicious, abnormal queries, but may be harming others, especially if the amount of resources behind a DNS authoritative server are far greater than the victim's. Root servers and TLDs are commonly deployed as anycast groups with dozens of servers spread worldwide capable of producing enormous volumes of traffic.

Network Ingress Filtering or Router-Based Packet Filtering could greatly reduce the occurrence of these attacks, depending on the degree of worldwide deploy. Both eliminate or severely reduce the ability to reflect traffic on DNS servers, killing the scenario of a victim seeing DNS servers as the source of an attack which is the primary goal of this work. However, these solutions are only effective if deployed by a lot of providers managing edge routers and assigning IP addresses to their customers.

Since in this work we look specifically at the authoritative name servers of a TLD, rate limiting solutions are the most obvious and attractive to explore. Because it is not fundamental for ANY queries to be answered back - all the records can be obtained through multiple queries for the different record types - dropping them might cut attack traffic. However, it deals only with the worst scenario which is the exploitation of the query with highest amplification ratio. All attacks that are crafted to not use the ANY query will not be affected by this solution. DNS Dampening and RRL provide much better results in cutting DDoS attack traffic but with the latter being able to deal much better with the false positives problem thus making it the more interesting choice of the two.

In the next section we analyze in more detail the RRL feature details, specifically the configuration options provided for it in the most popular and widely used DNS server, the Berkeley Internet Name Domain (BIND) server⁷ [28].

⁷<https://www.isc.org/downloads/bind/>

Chapter 3

Response Rate Limiting

Response Rate Limiting is mechanism designed to be used on authoritative DNS servers. Typical behaviour of resolvers should include caching records received from authoritative servers, thus the same response should not have to be repeated so frequently by authoritative servers. In this chapter it is explained the RRL feature, how it works and how it can be configured through different parameters.

3.1 Concept and details

Conceptually, RRL uses a credit or token bucket scheme. To each client-response combination is assigned an account that earns credit every second that is spent every time an identical response is sent. While an account credit is below zero, responses become rate-limited and can either be dropped, truncated or sent(leaked), depending on configuration parameters. Truncating or leaking responses gives the victims a chance of still getting answers to their legitimate requests.

3.1.1 Response generation process

When generating a response, a server imputes a name to a client/subnet which can either be a wildcard name, the query name (QNAME) or the zone name.

When the server has records that match the query name it imputes both query name and the requested RR type. If the name exists in the server but the requested RR type for it does not then only the query name is imputed and a NODATA message is sent.

When the server does not have matching records to the query name it imputes the zone name. As an example, considering two separate queries asking for the A records of "pc1.example.com" and "pc2.example.com", both received by the authoritative name server of the com zone. Even though the requests are different, the answer to both is a referral message with the authoritative servers of "example.com" (see section 2.1.3). In this case, the imputed name of both responses is the zone name "example.com". The same behaviour applies to NXDOMAIN answers, all of them being relative to the zone name instead of the query name itself.

To keep track of clients and responses, RRL stores information in state-blobs. These are essentially key-value pairs in a table. The key is a tuple of subnet, imputed name and an error status boolean. The value is a counter for how many times the same client-response combination has been considered.

RRL goes through the following steps:

1. The sender address is masked with a subnet prefix. By default, IPv4 addresses on the same /24 subnet are grouped, while IPv6 addresses are grouped from the same /56 subnet;
2. It imputes a name and a boolean error indicator (true if RCODE is set to REFUSED, FORMERR or SERVFAIL);
3. Using the tuple $\langle \text{mask}(\text{ip}), \text{impute}(\text{name}), \text{errorstatus} \rangle$, it looks up for a state blob, creates one if it doesn't exist and increments its counter;
4. Checks if the rate limit for that state blob has been exceeded, sending the response if it has not;
5. If the limit is exceeded, it decides dropping, truncating or leaking the response.

Figure 3.1 illustrates the response generation process with response rate limiting enabled.

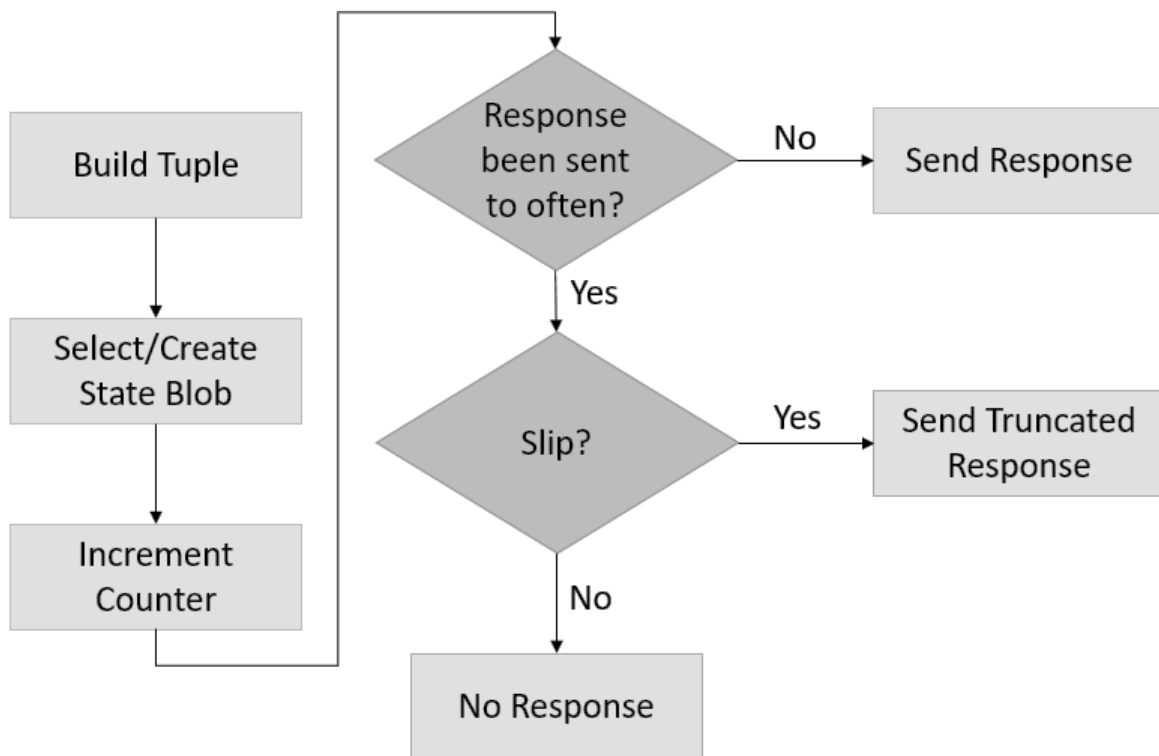


Figure 3.1: Diagram of RRL response generation process.

3.1.2 Configuration parameters

The behaviour of a responder with RRL enabled can be customized through configuration parameters. Different servers may have face different scenarios, whether in the number of hosted zones, types of

records it holds or overall request traffic. The RRL draft introduced the parameters that may be used to control the rate limiting behaviour [27]. They are described as follows:

- **RESPONSES-PER-SECOND**: the maximum allowed identical responses. A value of zero disables rate limiting for a view or restricts throttling only to NXDOMAIN and/or ERROR responses in case they are set.
- **NXDOMAINS-PER-SECOND**: sets the limit for NXDOMAIN responses. Defaults to the same as RESPONSES-PER-SECOND.
- **ERRORS-PER-SECOND**: sets the limit of error (REFUSED, FORMERR or SERVFAIL) responses.
- **WINDOW** - The period in seconds over which rates are measured. Assuming a responses-per-second value of 5 and window value of 5 then a single client is limited to 5 identical responses in any second and no more than 25 (5 x 5) identical responses within any 5 second window period.
- **IPv4-PREFIX-LENGTH**: The prefix length used to mask and group IPv4 addresses in the same bucket. Default is a /24 prefix.
- **IPv6-PREFIX-LENGTH**: The prefix length used to mask and group IPv6 addresses in the same bucket. Default is a /56 prefix.
- **LEAK-RATE**: Rate at which responses are sent normally when the client is rate limited. A value of 3 means a leaked response will be sent for every 3 dropped while the rate limit is exceeded.
- **TC-RATE or SLIP**: The same as LEAK-RATE but for truncated responses.
- **MAX-TABLE-SIZE**: Maximum number of state blobs the server can keep. It is recommended to be set to the product of WINDOW and the maximum queries per second to handle the worst case scenario where every single response is unique and requires a state blob for itself.
- **MIN-TABLE-SIZE**: The initial size to be allocated for an empty state blob table.

Support for this list is not mandatory. Depending on the DNS name server implementation, some parameters may not be available and others may be added. For instance, the most widely used name server, BIND, does not provide LEAK-RATE [29] and therefore only having support for sending truncated responses when applying rate limiting. BIND also adds ALL-PER-SECOND which sets the maximum number of all combined responses in a second, including truncated. PowerDNS and Microsoft's DNS server do not have NXDOMAINS-PER-SECOND with the latter adding MAXIMUM-PER-WINDOW which is similar to BIND's ALL-PER-SECOND but for total responses per WINDOW seconds.

Chapter 4

Study Methodology

The goal of this work is to study the usage of the RRL feature in the DNS service provided by DNS.pt¹, the manager of the .pt TLD, as a mitigation mechanism against DNS Amplification Attacks that abuse but don't necessarily target its servers. The first task is to verify the behaviour of RRL when used in a server that hosts a very delegation centric domain as, for example, a TLD server. Then, it measures how effective RRL can be in reducing the impact of DNS amplification attacks.

4.1 Testing the effectiveness of RRL

To test the RRL feature, an isolated server is setup to host a generated zone similar to a TLD that is then hit with request generated at a client host.

The server is setup in a virtual machine running on top a Macbook Air with a dual-core hyperthreaded Intel processor. The virtual machine has 2vCPUs and 1 GB of RAM. The guest OS is Debian 9.4.0 and the name server software is BIND 9.10.3-P4. Recursion is disabled for all clients and a single zone is created on the server. Inbound and outbound network traffic on the server are measured and sampled using the IO graphing ability of Wireshark, a packet capture and network protocol analyzer.

A TLD name server usually has just one very large zone file containing the collection of all its sub-zones. For instance, the .pt zone has over 1 million² registered domain names in total. In an attempt to increase the sophistication of an attack, the attacker may previously gather information about the zone using webcrawlers or by conducting a zone walk³. To study the effectiveness of RRL on a TLD like .pt, a single zone file is generated containing 1000 domain names with their delegation (NS) and glue (A and AAAA) records. The zone is then signed so that the increasingly present effect of DNSSEC can also be considered.

The client host, running in a separate machine, uses Python and Scapy⁴ scripts to generate and send forged DNS requests at different rates. Only a single client is used since this study looks mostly at the behavior of the server making the reflection part of the attacks irrelevant. Attacks are separated into

¹<https://www.dns.pt/pt/>

²<https://www.dns.pt/pt/estatisticas/?tipo=0&ordem=2&graph=0&subm=Filtrar>

³<https://www.farsightsecurity.com/2017/09/01/stsauver-zone-walking/>

⁴<https://scapy.net/>

two different types: repeating query attacks in which the same request is repeated and varying query attacks in which the set of requests uses a particular percentage of the existing sub-domain records, varying from 0% to a 100% of resolvable domain names.

4.2 Finding the ideal settings

Although there are recommended values for each setting, there is no global set of options that suites every scenario. Some of the options can be fine-tuned while using the LOG-ONLY option. In this mode, RRL does not actually drop any responses but instead logs its activity which enables post-analysis to be performed to see the effects of different settings without the risk of causing an impact on normal operation of the DNS service.

Responses per-second

To determine what PER-SECOND settings to use, the best approach is to analyze DNS traffic captures or logs, under normal conditions, of the specific environment. For a fully delegation-centric server, such as a TLD authoritative server, responses can be controlled with REFERRALS-PER-SECOND and NXDOMAINS-PER-SECONDS separately.

A simple solution is to determine what is the largest response that can be sent, as an example, to /24 IPv4 or /56 IPv6 subnets. With that information decide how much DNS outgoing traffic is considered acceptable for a single client or subnet. For instance, considering a maximum of 1 Mbps of DNS traffic and with an observed maximum response size of 3500 bytes then the maximum number of identical responses for that subnet can be calculated using the formula:

$$\frac{\textit{Maximum acceptable bandwidth}}{\textit{Largest response size}} \approx \textit{Maximum number of responses} \quad (4.1)$$

Window size

The WINDOW setting acts as the punishment for an offending client. A bigger punishment means the client can start receiving normal answers after a longer period. If RESPONSES-PER-SECOND has a low setting and WINDOW is high, the DNS server will be extremely protective, assigning long punishments to rather low offenses. The opposite is a DNS server that is slow to start rate limiting offenders and fast at letting them receive answers. Preferably, it should be kept to smaller values meaning victims of attacks can return to normal behavior quicker.

IPv4 and IPv6 prefix length

RRL groups clients into buckets because it is protecting them at a distance. The default values of /24 for IPv4 and /56 IPv6 represent small enough groups that their DNS traffic can be combined into a single bucket while avoiding false positives. Smaller length prefixes will group more clients together increasing

Table 4.1: Probability of false positives for different SLIP values

SLIP	False positive probability	TCP response probability
SLIP 1	0%	100%
SLIP 2	50%	≈ 97%
SLIP 3	66.6%	≈ 87%
SLIP 5	80%	≈ 67%
SLIP 10	90%	≈ 41%

the chance of legitimate clients being grouped with offenders. Bigger length prefixes do the opposite, decreasing that chance.

SLIP

With the SLIP setting there is a trade-off between how much outgoing traffic to allow and chances of false positives. Higher values for SLIP will result in less outgoing traffic but increase the chance of a legitimate request being dropped, i.e., a false positive occurrence. The chance is related to the SLIP setting in the following expression:

$$\frac{SLIP - 1}{SLIP} \times 100 = \text{Chance of false positive} \quad (4.2)$$

An increased chance of having legitimate requests being dropped also means a lower chance for clients to retry using TCP. Setting SLIP to a value higher than 1 means less responses with the TC bit set are sent, decreasing the chance for the client to receive one. The chance that a rate-limited client reconnects and receives a TCP response can be seen as a binomial probability of having 1 successful attempt out of a maximum of n retries with the probability of success being equal to chance of receiving a response with the TC bit set.

Table 4.1 shows the probability of false positive occurrences for different values of SLIP and the corresponding probability of a client receiving a TCP response assuming the client does 5 UDP retries before giving up.

Due to security concerns regarding DNS cache poisoning at resolvers[30], the recommended value for SLIP is 1. With a SLIP value of 1 there are no false positives. Responses are still rate-limited, meaning amplification is still cut but for every request, a truncated response is sent while an attack is occurring. SLIP 1 completely eliminates the chance of false positive occurrences at the cost of reflecting all the traffic. If a repeated request is being sent at a rate of 100 Mbps, a DNS server configured with a SLIP value of 1 will send approximately 100 Mbps worth of truncated responses to the victim.

Chapter 5

Measurements and Results

5.1 RRL behaviour in a TLD-like zone

A TLD is a special type of zone. Its zone file should only contain the delegation records for the domain names registered under the TLD. This means the server holds only NS records and, in most cases, their respective glue A and AAAA records. This means that all of the requests for resolvable domains names should return referral responses. To verify this behaviour RRL was enabled on the test server with default values and the log-only option enabled. The client host sends requests using random QNAMEs under the same domain-name at a rate of 10 requests per second which is enough to trigger rate-limiting for the demonstration.

Listing 5.1: Grouping of different qnames for the same referral response.

```
root@ns1:~# tail -n 5 /var/log/named/queries.log
172.16.156.1#58982 (hclgtixc.67.tld): query: hclgtixc.67.tld IN A + (172.16.156.10)
172.16.156.1#59336 (piruzcoa.67.tld): query: piruzcoa.67.tld IN A + (172.16.156.10)
172.16.156.1#53353 (ndhpquof.67.tld): query: ndhpquof.67.tld IN A + (172.16.156.10)
172.16.156.1#51309 (klmgudbx.67.tld): query: klmgudbx.67.tld IN A + (172.16.156.10)
172.16.156.1#55025 (stfvcqle.67.tld): query: stfvcqle.67.tld IN A + (172.16.156.10)
root@ns1:~# tail -n 5 /var/log/named/rate-limit.log
172.16.156.1#58982 (hclgtixc.67.tld): would rate limit drop referral response to 172.16.156.0/24 for
67.tld IN
172.16.156.1#59336 (piruzcoa.67.tld): would rate limit slip referral response to 172.16.156.0/24 for
67.tld IN
172.16.156.1#53353 (ndhpquof.67.tld): would rate limit drop referral response to 172.16.156.0/24 for
67.tld IN
172.16.156.1#51309 (klmgudbx.67.tld): would rate limit slip referral response to 172.16.156.0/24 for
67.tld IN
172.16.156.1#55025 (stfvcqle.67.tld): would rate limit drop referral response to 172.16.156.0/24 for
67.tld IN
```

In the log lines on listing 5.1 we can see different, non-repeating requests for existing records being received by the server. The behaviour of a TLD server is observed here as every request, although different, results in the same response which is a referral for "67.tld" with information for its authoritative name servers. This shows that RRL in a TLD zone will group responses by domain name regardless of what record type is being requested since all the server will ever do is send referral responses.

A similar behaviour is observed when the requests are made for non-existing domains on the TLD. The sample on listing 5.2 shows completely different requests being received by the server resulting in NXDOMAIN responses. Because NXDOMAIN responses are not bound to a domain name or set of domain names, they are seen and grouped under the TLD name meaning that all NXDOMAIN responses to a single client or subnet are treated as the same.

Listing 5.2: Grouping of NXDOMAIN responses for the whole TLD.

```

root@ns1:~# tail -n 10 /var/log/named/queries.log
172.16.156.1#51316 (ephpsffzgb.205851.tld): query: ephpsffzgb.205851.tld IN A + (172.16.156.10)
172.16.156.1#64146 (sorzqloepv.256681.tld): query: sorzqloepv.256681.tld IN NS + (172.16.156.10)
172.16.156.1#55326 (niqvmygdo.239178.tld): query: niqvmygdo.239178.tld IN NS + (172.16.156.10)
172.16.156.1#57485 (sguupwjsnr.294212.tld): query: sguupwjsnr.294212.tld IN AAAA + (172.16.156.10)
172.16.156.1#54719 (oxofpeeskj.256061.tld): query: oxofpeeskj.256061.tld IN AAAA + (172.16.156.10)
root@ns1:~# tail -n 10 /var/log/named/rate-limit.log
172.16.156.1#51316 (ephpsffzgb.205851.tld): would rate limit drop NXDOMAIN response to 172.16.156.0/24
      for tld
172.16.156.1#64146 (sorzqloepv.256681.tld): would rate limit slip NXDOMAIN response to 172.16.156.0/24
      for tld
172.16.156.1#55326 (niqvmygdo.239178.tld): would rate limit drop NXDOMAIN response to 172.16.156.0/24
      for tld
172.16.156.1#57485 (sguupwjsnr.294212.tld): would rate limit slip NXDOMAIN response to 172.16.156.0/24
      for tld
172.16.156.1#54719 (oxofpeeskj.256061.tld): would rate limit drop NXDOMAIN response to 172.16.156.0/24
      for tld

```

5.2 Repeating query attacks

The most simple attacks to craft are the ones that repeat the same request over and over. These can be referred as repeating query attacks and it will usually try to maximize amplification by making requests that are known to provide a great deal of amplification. The ANY request is the most commonly abused RTYPE in DNS Amplification attacks [14] since it triggers the server to return all resource records it has for the domain name.

Figure 5.1 shows a measurement of inbound and outbound traffic on the test server. A repeating ANY query attacks begins flooding the server resulting in an average of 600 Kb/s of incoming traffic and 3.5 Mb/s of outgoing traffic. RRL is enabled and starts rate-limiting to 10 responses per second with a SLIP setting of 1 which is then increased to 2. Table 5.1 shows the effect of RRL using different values for SLIP against a repeating query attack. As expected, when RRL is enabled with SLIP=1 it reduces outgoing attack traffic to the same level of incoming traffic as the repeated responses are now truncated and therefore the same size as the corresponding requests. When SLIP is set to 2, outgoing offending traffic is reduced even more due to only 1 out of every 2 responses being sent truncated.

RRL effectively stops the not so sophisticated attacks, that rely almost entirely on a single query to maximize amplification. This is the ideal scenario for RRL effectiveness and the results on table 5.1 show that as outgoing traffic is reduced to almost exactly what is expected for each SLIP value.




Figure 5.1: A repeating query attack against a test server hosting a fictitious zone similar to a TLD. RRL is enabled at 80 seconds using SLIP=1 and then SLIP=2 at 120 seconds.

Table 5.1: Outgoing traffic volume of repeating query attacks.

SLIP	False positive probability	Incoming traffic	Outgoing traffic
SLIP 1	0%	618 Kb/s	625 Kb/s
SLIP 2	50%	625 Kb/s	312 Kb/s
SLIP 3	66.6%	630 Kb/s	210 Kb/s
SLIP 5	80%	635 Kb/s	126 Kb/s
SLIP 10	90%	640 Kb/s	64 Kb/s

5.3 Varying query attacks

Repeating the same request over and over is a simple flooding attack. RRL was designed from the premise that authoritative name server should not have to repeat the same answer frequently. In this section three experiments are conducted to understand how effective RRL can be when the sophistication of attacks increases. By randomly changing the requested record name and type, an attack may avoid triggering rate-limiting for its intended victim. The request records can either be randomly generated names or randomly picked from already known existing records. Prior to an attack, the attacker can gather information about the a zone by abusing NSEC records and performing a so called zone walk.

5.3.1 NXDOMAIN abuse (0% resolvable domain names)

The first experiment consists of send varying requests that result in NXDOMAIN responses. Although the requests and responses are different, RRL groups all NXDOMAIN responses for a single zone on the same "bucket" (see 5.1). Figure 5.2 and table 5.2 shows RRL reducing outbound traffic from an attack by which all requests result in NXDOMAIN. The results resemble the obtained ones with the repeating query attack, since outgoing NXDOMAINs are all treated as the same response, therefore triggering rate-limiting for all NXDOMAIN responses for that client.




Figure 5.2: A varying query attack that results in 100% NXDOMAIN responses. RRL is enabled at 38 seconds using SLIP=1 and then SLIP=2 at 58 seconds.

Table 5.2: Outgoing traffic volume of an attack that results in NXDOMAIN responses.

SLIP	False positive probability	Incoming traffic	Outgoing traffic
SLIP 1	0%	347 Kb/s	347 Kb/s
SLIP 2	50%	347 Kb/s	174 Kb/s
SLIP 3	66.6%	348 Kb/s	115 Kb/s
SLIP 5	80%	347 Kb/s	69 Kb/s
SLIP 10	90%	348 Kb/s	35 Kb/s

5.3.2 Mixed attack (50% resolvable domain names)

In the second experiment, a slightly more sophisticated attack results in 50% NXDOMAIN responses and 50% unique NOERROR responses for existing records. This covers the scenario of an attacker having some knowledge of records it can abuse on the zone. In order to avoid triggering rate-limit for the known records requests, the attacker mixes the attack with generated names that result in NXDOMAIN responses. Figure 5.3 and table 5.3 show that having knowledge and using a larger set of different requests results in higher outgoing traffic. Roughly 50% of the responses avoid getting rate-limited as they are not repeated very often. The other 50% that result in NXDOMAIN are still grouped but because they are now less frequent, triggering rate-limit on NXDOMAIN also does not happen as easily. Taking into account the generated domain is signed, NXDOMAIN responses are much more amplified than requesting, for example, existing A records. NXDOMAIN responses with DNSSEC carry, in most cases, 3 NSEC3 records to prove denial of existence. If a percentage of these responses can escape RRL while being diluted with other requests it also adds to the overall amplification ratio of the attack.




Figure 5.3: A varying query attack that returns a mix of NXDOMAIN and NOERROR responses. RRL is enabled at 30 seconds using SLIP=1 and then SLIP=2 at 60 seconds.

Table 5.3: Outgoing traffic volume of a mixed query attack.

SLIP	False positive probability	Incoming traffic	Outgoing traffic
SLIP 1	0%	347 Kb/s	1 Mb/s
SLIP 2	50%	348 Kb/s	940 Kb/s
SLIP 3	66.6%	347 Kb/s	920 Kb/s
SLIP 5	80%	347 Kb/s	892 Kb/s
SLIP 10	90%	347 Kb/s	881 Kb/s

5.3.3 Fully-indexed zone abuse (100% resolvable domain names)

The worst case scenario may happen when the attacker has managed to index all of the records in the zone. Figure 5.4 and table 5.4 demonstrate that an attack using a very large set of unique requests may completely avoid being rate-limited. This raises a question of how RRL effectiveness may decrease when the number of records in a target zone or name server increases. The more information an attacker acquires about a sufficiently large zone (millions of records) the more easily its attack can fly under the radar of RRL.

Table 5.4: Outgoing traffic volume of attack randomly requests one of all the zone records.

SLIP	False positive probability	Incoming traffic	Outgoing traffic
SLIP 1	0%	630 Kb/s	2.45 Mb/s
SLIP 2	50%	630 Kb/s	2.4 Mb/s
SLIP 3	66.6%	630 Kb/s	2.38 Mb/s
SLIP 5	80%	631 Kb/s	2.36 Mb/s
SLIP 10	90%	630 Kb/s	2.34 Mb/s




Figure 5.4: A varying query attack that uses all the records of the zone. RRL is enabled at 30 seconds using SLIP=1 and then SLIP=2 at 60 seconds.

5.4 Results summary

Through a simple experiment in a virtual lab environment, it is shown how the RRL feature will typically behave when enabled in an authoritative TLD DNS server. Since a TLD almost exclusively sends referral responses, it is expected that a lot of similar, yet different, requests may be grouped together.

Put against simple repetitive query attacks, RRL shows that it is able to cut malicious traffic exactly as it is expected. By increasing the SLIP value, it is possible to reduce even more the outgoing traffic related to the offending requests although at the cost of possibly denying service to legitimate requests.

By crafting increasingly complex attacks, it was possible to understand the limitations of the current RRL approach. Through techniques such as zone-walking, or simply crawling the web, attackers may gather sufficient knowledge about existing records such that they can avoid being rate limited by not repeating the same requests too often. TLDs typically have a very large number of unique records and that sheer size may actually work against the RRL feature. With a generated TLD zone with 1000 sub-domains, by targeting 50% of its records it is already noticeable how the RRL feature begins to struggle.

Chapter 6

Conclusions

DNS amplification attacks rely on two points, reflection and amplification. The DNS protocol is mostly used over UDP meaning that the source IP address can be easily spoofed, making DNS servers traffic reflectors. Additionally, the DNS protocol has an amplification behaviour in its request-response mechanism. Small requests are able to be returned as very large responses. The adoption of DNSSEC greatly contributes to this amplification effect since it adds very large records, necessary to authenticate information on the client side. By controlling a large number of bots, attackers can achieve network traffic volumes in excess of 1 Tbps [12].

To tackle the issue, RRL was proposed and recommended for usage on authoritative DNS servers. RRL limits the number of similar responses within a configurable window. If the same response is seen too many times for the same client within the window period, RRL triggers rate-limiting during a certain time, stopping potential offending traffic from coming out of the DNS server. In this work, in a lab setting, the RRL feature has shown to be effective against common, yet simple attacks that repeat the same request in an attempt for the greatest amplification ratio possible on the server. However, its effectiveness starts decreasing as attacks begin to use larger sets of requests. If an attacker knows a large enough set of existing domain names within a zone, it may be able to completely avoid triggering rate-limiting. Using the SLIP=1 setting, false positives can be avoided by never dropping responses and simply truncate them. Increasing the SLIP value may reduce outbound traffic but at the cost of denying service to legitimate clients as they may not be able to reconnect using TCP.

Even though it does not pose as a great barrier against smart and sophisticated attacks, RRL does have a positive effect on authoritative name servers with minimal side effects for either legitimate clients or the servers. When configured to never drop responses, i.e., with the SLIP=1 setting, a DNS server has the guarantee that the clients are always able to reconnect but at the cost of an increase in TCP traffic. In the best scenario, RRL can completely stop common repeating query attacks and in the worst scenario it acts as if it is disabled.

6.1 Challenges and Future Work

One particular challenge of this work was the fact that it was not carried on real production DNS servers so it could not study how a baseline of traffic from a real TLD server actually affects RRL and vice-versa. In other words, it did not evaluate RRL in a highly heterogeneous environment of clients/resolvers and traffic patterns. In such environment, it is difficult to establish a baseline or traffic pattern per client and response which is what the per-second settings of RRL should portray. An addition that RRL could benefit from is a process in which continuous analysis of traffic patterns causes the mechanism to automatically adjust its per-second limits to either allow abnormal yet legitimate client behaviour or stop large volume attacks that currently avoid triggering rate-limiting.

Bibliography

- [1] Domain names - concepts and facilities. RFC 1034, Nov. 1987. URL <https://rfc-editor.org/rfc/rfc1034.txt>.
- [2] Domain names - implementation and specification. RFC 1035, Nov. 1987. URL <https://rfc-editor.org/rfc/rfc1035.txt>.
- [3] ICANN. List of top-level domains, 03 2017. URL <http://data.iana.org/TLD/tlds-alpha-by-domain.txt>.
- [4] J. Damas, M. Graff, and P. A. Vixie. Extension Mechanisms for DNS (EDNS(0)). RFC 6891, Apr. 2013. URL <https://rfc-editor.org/rfc/rfc6891.txt>.
- [5] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: illuminating the edge network. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 246–259. ACM, 2010.
- [6] C. Rossow. Amplification hell: Revisiting network protocols for ddos abuse. In *NDSS*, 2014.
- [7] D. Atkins and R. Austein. Threat Analysis of the Domain Name System (DNS). RFC 3833 (Informational), August 2004. URL <http://www.ietf.org/rfc/rfc3833.txt>.
- [8] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirements. RFC 4033 (Proposed Standard), Mar. 2005. URL <http://www.ietf.org/rfc/rfc4033.txt>. Updated by RFC 6014.
- [9] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Resource Records for the DNS Security Extensions. RFC 4034 (Proposed Standard), Mar. 2005. URL <http://www.ietf.org/rfc/rfc4034.txt>. Updated by RFCs 4470, 6014.
- [10] S. Rose, M. Larson, D. Massey, R. Austein, and R. Arends. Protocol Modifications for the DNS Security Extensions. RFC 4035, Mar. 2005. URL <https://rfc-editor.org/rfc/rfc4035.txt>.
- [11] S. Boeyen, S. Santesson, T. Polk, R. Housley, S. Farrell, and D. Cooper. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280, May 2008. URL <https://rfc-editor.org/rfc/rfc5280.txt>.

- [12] S. Khandelwal. World's largest 1 tbps ddos attack launched from 152,000 hacked smart devices, September 2016. URL <https://thehackernews.com/2016/09/ddos-attack-iot.html>.
- [13] D. Goodin. Record-breaking ddos reportedly delivered by over 145k hacked cameras, September 2016. URL <https://arstechnica.com/security/2016/09/botnet-of-145k-cameras-reportedly-deliver-internets-biggest-ddos-ever/>.
- [14] US-CERT. Ta13-088a : Dns amplification attacks, 2013. URL <https://www.us-cert.gov/ncas/alerts/TA13-088A>.
- [15] D. C. MacFarland, C. A. Shue, and A. J. Kalafut. Characterizing optimal dns amplification attacks and effective mitigation. In *International Conference on Passive and Active Network Measurement*, pages 15–27. Springer, 2015.
- [16] A. J. Kalafut, C. A. Shue, and M. Gupta. Touring dns open houses for trends and configurations. *IEEE/ACM Trans. Netw.*
- [17] R. Rosenbaum. Using the Domain Name System To Store Arbitrary String Attributes. RFC 1464, May 1993. URL <https://rfc-editor.org/rfc/rfc1464.txt>.
- [18] Security Bulletin: Crafted DNS Text Attack. Technical report, Akamai's Prolexic Security Engineering and Research Team (PLXsert), 11 2014.
- [19] R. van Rijswijk-Deij, A. Sperotto, and A. Pras. Dnssec and its potential for ddos attacks: a comprehensive measurement study. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 449–460. ACM, 2014.
- [20] F. Neves and J. Damas. Preventing Use of Recursive Nameservers in Reflector Attacks. RFC 5358, Oct. 2008. URL <https://rfc-editor.org/rfc/rfc5358.txt>.
- [21] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (Best Current Practice), May 2000. URL <http://www.ietf.org/rfc/rfc2827.txt>. Updated by RFC 3704.
- [22] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed dos attack prevention in power-law internets. In *ACM SIGCOMM computer communication review*, volume 31, pages 15–26. ACM, 2001.
- [23] T. Peng, C. Leckie, and K. Ramamohanarao. Survey of network-based defense mechanisms countering the dos and ddos problems. *ACM Computing Surveys (CSUR)*, 39(1):3, 2007.
- [24] J. Mirkovic, G. Prier, and P. Reiher. Attacking ddos at the source. In *Network Protocols, 2002. Proceedings. 10th IEEE International Conference on*, pages 312–321. IEEE, 2002.
- [25] J. Mirkovic, G. Prier, and P. Reiher. Source-end ddos defense. In *Network Computing and Applications, 2003. NCA 2003. Second IEEE International Symposium on*, pages 171–178. IEEE, 2003.

- [26] G. Kambourakis, T. Moschos, D. Geneiatakis, and S. Gritzalis. Detecting dns amplification attacks. In *International Workshop on Critical Information Infrastructures Security*, pages 185–196. Springer, 2007.
- [27] P. Vixie and V. Schryver. Dns response rate limiting (dns rrl). URL: <http://ss.vix.su/~vixie/isc-tn-2012-1.txt>, 2012.
- [28] D. B. Terry, M. Painter, D. W. Riggle, and S. Zhou. The berkeley internet name domain server. In *Proceedings USENIX Summer Conference, Salt Lake City, Utah*, pages 23–31, 1984.
- [29] ISC. Bind 9 administrator reference manual, 2017. URL <https://ftp.isc.org/isc/bind9/9.12.0a1/doc/arm/Bv9ARM.pdf>.
- [30] Florian Maury, Mathieu Feuillet. Blocking dns messages is dangerous. URL https://www.ssi.gouv.fr/uploads/IMG/pdf/DNS-OARC-2013-Blocking_DNS_Messages_Is_Dangerous.pdf.

