International Journal of Wireless Networks and Broadband Technologies
Volume 9 • Issue 1 • January-June 2020

A Demonstration of Practical DNS Attacks and their Mitigation Using DNSSEC

Israr Khan, Letterkenny Institute of Technology, Letterkenny, Ireland

William Farrelly, Letterkenny Institute of Technology, Letterkenny, Ireland
https://orcid.org/0000-0002-6675-2040

Kevin Curran, Ulster University, London, UK

ABSTRACT

The authors implement common attacks on a DNS server and demonstrate that DNSSEC is an effective solution to counter DNS security flaws. This research demonstrates how to counter the zone transfer attack via the generation of DNSSEC keys on the name servers which prevent attackers from obtaining a full zone transfer as its request for the transfer without the keys was denied by the primary server. This article also provides a detailed scenario of how DNSSEC can be used as a mechanism to protect against the attack if an attacker tried to perform Cache Poisoning. The authors ultimately show that a DNSSEC server will not accept responses from unauthorised entities and would only accept responses which are authenticated throughout the DNSSEC chain of trust.

KEYWORDS

Cybersecurity, Security, DNS, Domain Name System, Internet, Routing, DNSSEC

1. INTRODUCTION

DNS is a critical part of network and internetwork infrastructure. However, it is vulnerable, and attackers have exploited vulnerabilities within the protocol to launch various kinds of attacks against it (Gupta, 2018). The DNS protocol does not provide origin of data for authenticity and it also lacks the mechanism to provide data integrity. Taking advantage of these vulnerabilities, the attackers can forge the DNS records and direct legitimate clients to malicious domains to fulfil their own vested interests. To overcome the problems of origin authentication and data integrity, DNSSEC was proposed. It is the result of focused and continuous efforts of the security communities to secure the DNS protocol (Krishnaswamy et al., 2009). DNSSEC solves these vulnerabilities wherein security parameters are added to the DNS responses from the server which allows the client to verify that the responses originated from the intended server and that the data in the responses is not forged. Over the past decade, attacks on the Internet and private networks are on the rise. Attackers look for vulnerabilities within protocols and software, which in turn assist them in exploiting those vulnerabilities to launch attacks (Stergiou et al., 2016; Tewari & Gupta, 2017;

DOI: 10.4018/IJWNBT.2020010104

Memos et al., 2017). Two of the most fundamental and popular attacks against the DNS protocol are Cache Poisoning and Man-in-the-Middle (MITM) attacks. In a Cache Poisoning attack, the DNS server is manipulated in a way so that it accepts and stores false data in its cache. This data does not come from an authoritative DNS server but instead it comes from a malicious user who tries to corrupt the DNS server cache by providing false information. Best practice for DNS server administrators is to randomize the UDP source port number from which caching DNS servers send out query packets as a mitigation against cache poisoning attacks. In effect, the UDP port used for a query should not be the default port 53, but instead a port randomly chosen from the entire range of UDP ports (less the reserved ports). This UDP source port randomization (SPR) makes it more difficult attackers to guess query parameters (RFC5452, 2009).

Once the DNS server cache has been corrupted, the false information will remain in the cache until the Time to Live (TTL) expires. This attack has adverse effects on the clients wanting to access the domain names from the servers. DNS data that is provided by name servers lacks support for data origin authentication and data integrity. This makes DNS vulnerable to man in the middle (MITM) attacks, as well as a range of other attacks (Ariyapperuma and Mitchell, 2007). In MITM attacks, an attacker can intercept and modify the network traffic between the resolver and the server. This occurs because the DNS protocol does not provide integrity checks and hence it is possible for the attacker to intercept and modify the data within DNS requests or responses. In 2018, a major DNS spoofing attack left the MyEtherWallet (MEW) service compromised (Nation, 2018).

The purpose of DNSSEC is to add security mechanisms such as origin authenticity and data integrity to make sure that the users can verify the origin of the data and be assured that the data has not be tempered with by anyone from its inception to reception. DNSSEC was designed to mitigate the Cache Poisoning attacks against DNS (Yu et al., 2011; Gupta, 2016). It also addresses the problems of Man-in-the-Middle attacks. DNSSEC adds digital signatures to the responses by the server and hence protects clients from deception by the malicious user who provide false data (Zou et al., 2016; Plageras et al., 2018). This is achieved by the introduction of new Resource Record (RR) types such as DNSKEY, DS, NSEC, RRSIG and other new records (Larson et al., 2005). Forensics play a significant role in determining the nature of the attacks on systems. Proper deployment and use of forensic tools can mitigate those attacks and provide a means to specify the type and degree of the attacks. Consequently, the information gained during the forensic investigation can then be used by the security experts to plan countermeasures against the attacks. (Shulman and Waidner, 2014) claim that DNSSEC is not only the most appropriate defence mechanism against Cache Poisoning attacks, but it also provides a means to carry out analysis of the attacks and delivers evidence based on the information gained. DNSSEC has been around for more than a decade, but much research has focused more on the challenges in deploying DNSSEC, the operational impacts and performance of deploying DNSSEC instead of providing practical solutions using DNSSEC as a mitigation mechanism against DNS attacks. We therefore demonstrate the implementation of practical attacks against DNS and their prevention using DNSSEC.

In this paper, we perform attacks on DNS, show mitigations using DNSSEC and determine whether the captured network data can be used as a mechanism to create a fingerprint to identify an attack. The novelty lies specially in how we show how a Zone Transfer attack can be countered with DNSSEC keys preventing attackers from obtaining a full zone transfer and how a Cache Poisoning can be prevented using the DNSSEC chain of trust. The research questions are:

RQ1: What effect does cache poisoning and man-in-the-middle attacks have on the DNS messages interchange

RQ2: How can these attacks be mitigated using DNSSEC?

RQ3: Does the captured network data provide an opportunity to create an attack fingerprint and provide a mechanism to identify the attack?