



Injection Attacks Reloaded: Tunnelling Malicious Payloads over DNS

Philipp Jeitner, *TU Darmstadt*; Haya Shulman, *Fraunhofer SIT*

<https://www.usenix.org/conference/usenixsecurity21/presentation/jeitner>

**This paper is included in the Proceedings of the
30th USENIX Security Symposium.**

August 11-13, 2021

978-1-939133-24-3

**Open access to the Proceedings of the
30th USENIX Security Symposium
is sponsored by USENIX.**

Injection Attacks Reloaded: Tunnelling Malicious Payloads over DNS

Philipp Jeitner
TU Darmstadt

Haya Shulman
Fraunhofer SIT

Abstract

The traditional design principle for Internet protocols indicates: “Be strict when sending and tolerant when receiving” [RFC1958], and DNS is no exception to this. The transparency of DNS in handling the DNS records, also standardised specifically for DNS [RFC3597], is one of the key features that made it such a popular platform facilitating a constantly increasing number of new applications. An application simply creates a new DNS record and can instantly start distributing it over DNS without requiring any changes to the DNS servers and platforms. Our Internet wide study confirms that more than 1.3M (96% of tested) open DNS resolvers are standard compliant and treat DNS records transparently.

In this work we show that this ‘transparency’ introduces a severe vulnerability in the Internet: we demonstrate a new method to launch string injection attacks by encoding malicious payloads into DNS records. We show how to weaponise such DNS records to attack popular applications. For instance, we apply string injection to launch a new type of DNS cache poisoning attack, which we evaluated against a population of open resolvers and found 105K to be vulnerable. Such cache poisoning cannot be prevented with common setups of DNSSEC. Our attacks apply to internal as well as to public services, for instance, we reveal that all eduroam services are vulnerable to our injection attacks, allowing us to launch exploits ranging from unauthorised access to eduroam networks to resource starvation. Depending on the application, our attacks cause system crashes, data corruption and leakage, degradation of security, and can introduce remote code execution and arbitrary errors.

In our evaluation of the attacks in the Internet we find that all the standard compliant open DNS resolvers we tested allow our injection attacks against applications and users on their networks.

1 Introduction

Domain Name System (DNS) is a key component of the Internet. Originally designed to translate domain names to

IP addresses, DNS has developed into a complex infrastructure providing platform to a constantly increasing number of applications. The applications that are built over DNS range from Internet specific services, such as location of hosts using GPOS record [RFC1712] [1] to security mechanisms, such as authentication with certificates using TLSA record [RFC6698] [2]. The core design feature that allows DNS to support new applications without involving any changes to its infrastructure is the requirement that the handling of the DNS records is done transparently [RFC3597, RFC1035] [3, 4]. Namely, DNS should not attempt to interpret nor understand the records that it is serving. Thanks to this feature new DNS records can be easily added to the DNS infrastructure without requiring any modifications, and novel applications can instantly run over DNS using the newly added records.

In this work we show that the transparency-feature of DNS, while critical for fast and smooth deployment of new technologies, introduces a gaping hole in Internet security.

Exploiting transparency to encode injections. We exploit the transparency of the DNS lookups to encode injection strings into the payloads of DNS records. The attacker places the malicious records in the zonefile of its domain. When provided by the attacker’s nameserver the records appear to contain legitimate mappings under the domain controlled by the attacker. However, when the record is processed by the receiving victim application, a misinterpretation occurs - resulting in the injection attack. Our attacks exploit two key factors caused by the transparency of DNS: (1) the DNS resolvers do not alter the received records hence the malicious encoding is preserved intact and (2) the receiving applications do not sanitise the received records. We devise injection payloads to attack popular applications.

Applications do not sanitise DNS records. Classical injection attacks are well known and have been extensively studied: the attacker provides a malicious input through a web application to alter the structure of a command, hence subverting the logic of the application, e.g., [5, 6]. Such injection attacks are easy to mitigate in practice: the input of the user is validated and invalid characters are filtered before

reaching the application. Due to the long history of injection vulnerabilities in web applications and the awareness to the potential risks, most applications validate user input [7].

We show that in contrast to user input, the inputs provided by the DNS resolvers are not validated. For instance, user credentials provided to LDAP via a web interface to authenticate the user and enable it to use services, are validated, while DNS values that are provided to LDAP to route the authentication request to an authentication server are not validated. We show how to construct malicious payloads to launch injection attacks, such as XSS and cache poisoning, against a variety of applications and services, including DNS caches, LDAP, eduroam.

Attacker model. The attacker causes the victim resolvers to issue queries for records that encode malicious payloads, e.g., by deploying an ad-network or by sending an Email from attacker’s domain to the victim. The resolvers cache the records received in DNS responses and provide them to applications and users. We illustrate the attacker model and the setup with eduroam as example victim application, in Figure 1. Using our “weak” attacker we demonstrate a range of attacks against popular applications and services that use DNS lookups, including DNS cache poisoning, applications’ crashes, downgrade of security mechanisms, remote code execution vulnerabilities, XSS.

Contributions

The core issue that we explore in this work is the balance between security and the requirement to enable easy deployment of new applications over DNS. Our contributions include:

- **Analysis of components in resolution chain.** We analyse the interaction between the applications and the components in DNS resolution chain. We find that the processing applied by the DNS resolvers over DNS records is compliant with the requirement in [RFC3597,RFC1035] and preserves the structure of the malicious inputs encoded by the attackers - this property is key to our attacks. We validate this also in the Internet against 3M open DNS resolvers. Our measurement study reveals that more than 96% of the open DNS resolvers do not modify the records that they receive from the nameservers, and serve them intact to the calling applications.

- **Study of DNS input validation.** We find that although DNS delivers untrusted data from potentially malicious Internet servers the applications trust the data returned by the DNS resolvers. Our study shows that the lack of input validation is systematic and prevalent and is not a bug mistakenly introduced by developers in some isolated cases – this includes custom functions in applications as well as standardised function calls of IEEE POSIX, e.g., `gethostbyname()`.

- **Implementation of injection attacks over DNS.** We demonstrate that the attackers can systematically and efficiently construct attack vectors and show how to integrate them into the zonefile of a malicious domain operated by

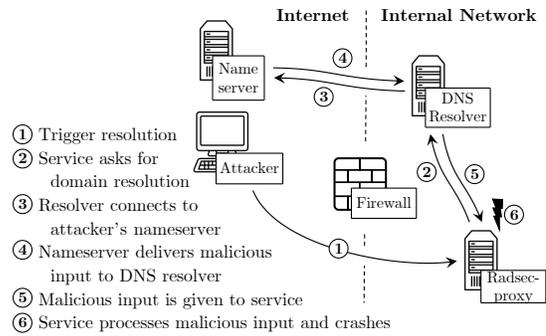


Figure 1: Attack and setup overview, with eduroam radsecproxy as example application.

the attacker. We then demonstrate injection attacks over DNS against popular applications using these malicious DNS records.

- **Injections into DNS caches.** We show how to encode payloads for injecting malicious records into DNS caches. When cached by the victim DNS resolver, a misinterpretation occurs mapping a resource of some victim domain to an attacker’s IP address. In contrast to classical cache poisoning, [8–10], which require a strong (on-path) attacker or assume specific network properties, such as side channels or fragmentation, our cache poisoning attacks do not make any requirements on attacker capabilities. We also do not need to spoof IP addresses in DNS responses - a requirement which is essential in prior cache poisoning attacks. We automated the evaluation of our poisoning attacks, which allowed us to launch them against a large set of 3M target DNS resolvers, performing successful poisoning against 105K resolvers. Implementation of previous cache poisoning attacks had to be manually tailored per each target - automating the attacks would result in a negligible success probability. Hence the previous attacks were carried out against at most a handful of targets, and the rest of the resolvers’ population was merely checked for properties that make them potential targets. Furthermore, in contrast to previous poisoning attacks, ours cannot be prevented with common setups of DNSSEC [11–13].

- **Evaluation of injection attacks over DNS.** We evaluate our injection attacks against popular applications (listed in Table 1). Our analysis of the vulnerabilities in applications, where suitable, combines fuzzing, source code review and dynamic (black box) execution. We evaluate our injection attacks against a population of more than 3M open resolvers in the Internet. We provide additional information on our evaluations at <https://xdi-attack.net>.

Ethics and Disclosure

We have already taken preliminary steps to address these vulnerabilities by contacting the DNS software vendors as well the applications evaluated in this work. We experimentally evaluated the attacks reported in this work against servers that we set up as well as against open DNS resolvers in the Internet using domains that we control. This allowed us to

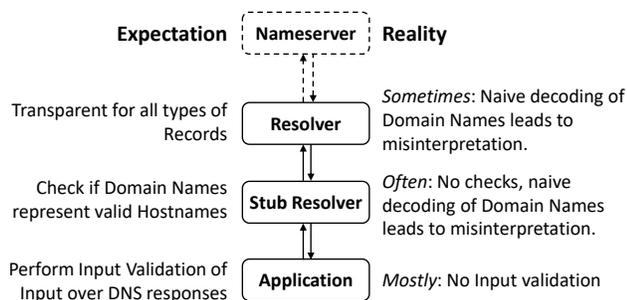


Figure 2: Expected vs. actual behaviour in DNS lookup.

validate the presence of the vulnerabilities without exploiting them against real victims and without causing damage to the networks nor services in the Internet. Our attacks similarly apply also to non-open DNS resolvers.

Prior to performing the validation of the vulnerabilities in the wild we received an approval from our research institution. In the next steps we will be coordinating countermeasures with the DNS and applications vendors, as well as the IETF community.

Organisation

In Section 2 we analyse the interaction between components in DNS resolution chain. In Section 3, we demonstrate injection attacks against popular applications. In Section 4 we evaluate our attacks against open resolvers in the Internet. We propose countermeasures in Section 5, review related work in Section 6 and conclude in Section 7.

2 Analysis of DNS Resolution Chain

In this section we analyse the interaction between the components relevant to processing DNS records in responses. In our analysis we use popular DNS resolvers and stub resolver implementations built into operating systems and programming languages, and experimentally test how they handle control characters in domain names and if they modify any of the maliciously crafted payloads needed to conduct the application-specific exploits that we evaluate in this work.

2.1 Components in DNS Lookup

We consider 3 different types of software components which fulfil different roles during a DNS lookup: (recursive) DNS resolvers, stub DNS resolvers and applications. In our setup the DNS nameservers are controlled by the attacker and provide maliciously-encoded DNS responses. The victim resolvers serve these records to the stub resolvers in applications. We provide an illustration of these software components together with the expectations on their behaviour and their actual behaviour discovered in this work in Figure 2. DNS lookups by system stub resolvers are implemented in various ways in applications. For standard A, AAAA and PTR queries, system C libraries include POSIX-standardised [14, 15] functionality in form of the `gethostbyname()`, `getaddrinfo()`,

`gethostbyaddr()` and `getnameinfo()` functions. When considering full fledged resolution functionality, which also supports other query types, like MX, SRV or TXT, there is no standardised API so applications need to rely on third party libraries for constructing the DNS packets or for parsing the DNS responses they receive from the network (e.g., recursive resolvers). DNS software implementations in recursive resolvers and forwarders typically implement their own packet decoding logic since they do not interface with the applications directly and do not need to decode parts of the DNS records at all.

2.2 System Stub Resolvers

Stub resolvers provide the interface between the applications and the DNS resolvers. Applications typically do not issue DNS requests to the recursive resolvers directly, but instead use a POSIX standardised API [14, 15] to supply the hostname they want to resolve, and the system's standard C library translates this into a DNS request and parses the response from the recursive DNS resolver. The applications can perform hostname-to-address (A, AAAA) and reverse (PTR) lookups via the `gethostbyname()`, `gethostbyaddr()`, `getaddrinfo()` and `getnameinfo()` calls. This API is defined to return CNAME aliases and results of reverse lookups as null-terminated 'host names' [14, p. 320] in its returned `hostent` structure, defined [16, 17] to only contain Latin characters, digits and hyphens ("A-Z", "a-z", "0-9", "-").

According to this definition the expected behaviour of a system stub resolver is to check that any domain name returned by a call to any of the POSIX-standardised resolver functions must be checked before returning it to the application.

DNS Record Processing. We analyse the DNS record processing done by stub-resolvers in Section 3.6.1 in detail. We find that while one implementation (glibc) is fully conforming with our expectations, most stub-resolvers are not. Stub resolvers that do not validate the value of the hostname expose the applications to attacks; we show these in Section 3.6.2.

2.3 Recursive Resolvers and Forwarders

Recursive resolvers provide lookup services to system DNS resolvers: they locate nameservers and look up the requested records.

DNS resolvers and forwarders process DNS records in their line-format, i.e., they store and handle domain names in their encoded form and do not try to parse the records that they do not need to understand, like TXT or MX. This makes them transparent for any binary values inside domains or other record data as is required by the DNS standard [4, 18, 19]. For internal caching purposes, a DNS resolver may choose to decode a domain name for further processing or to store it inside a cache.

DNS Record Processing. DNS resolvers are expected to transparently handle known and unknown DNS record types [RFC3597] to ensure forward-compatibility and compatibility

with mechanisms such as DNSSEC, as any change in the encoded record would invalidate the DNSSEC signatures. Moreover, [RFC1034] states that software like DNS resolvers should not try to decode domain names into a string as stub-resolvers would do. Most resolvers we tested are [RFC1034] compliant: they handle any payload transparently or only change the case of letters to lowercase, which is allowed since domain names are defined to be case insensitive [18]. One tested resolver software, MaraDNS Deadwood 3.2.14, was unable to handle `inject\000`.

We also identified an ‘[RFC1034] non-compliant’ behaviour, which we describe next. We take as an example a systemd-resolved forwarding DNS resolver which when receiving a DNS record decodes and escapes all included domain names into zero-terminated strings like a system stub resolver would do. Then it caches the records as decoded strings. The resolver unescapes and re-encodes the records when sending them to a requesting client or application. This behaviour ensures that any misinterpretation during decoding (step 2 in Figure 7) will cause the domain name to be modified. This modification cannot be detected by downstream resolvers or applications because the misinterpreted record is re-encoded instead of just being passed as binary.

2.4 Applications

Applications are the source of DNS lookups and process the records in the response after it has traversed all the other components in the resolution chain. The records may contain unexpected characters which the application cannot process correctly. The impact of injection attacks on the applications depends on the use case of the DNS lookup. For example, applications which do service discovery or authentication lookups typically need to parse DNS packets by themselves, because system stub resolvers do not support queries for such record types. This can also be abused to cause misinterpretation of domain names or other data in DNS records.

DNS Record Processing. When applications perform DNS lookups of types other than A, AAAA or PTR, they implement DNS lookup functionality by themselves as there is no standardised API for this. While the standard libraries of some programming languages like java, go or nodejs include functions for query types like MX, SRV and TXT, there is no recommended behaviour for such functions and they do not perform any validation of the data which is passed to the application. This also applies to DNS lookup implementations done in applications directly. Domain names are typically not validated nor escaped when decoded into a string.

Theoretically, when applications use the system resolver to do DNS lookups, they could implement validation (step 3 in Figure 7) by themselves, to ensure no malicious input is processed. However, applications would still not be able to detect decoding errors in step 2 in Figure 7. For example, since the application does not see the binary DNS data, it cannot determine if the example domain in Figure 7 is `"a\b`.

`<>.com." or "a.b.<>.com."`.

3 Injection Attacks Against Applications

In this section we demonstrate injection attacks using malicious payloads tunnelled over DNS. We first explain our study methodology (Section 3.1) and then show attacks against selected popular applications (Sections 3.2 – 3.6) taking DNS software as the first example application.

3.1 Study Methodology

3.1.1 Attack overview

The attack is illustrated in Figure 1. The target victim application, e.g., radsecproxy of eduroam, is behind a firewall on the victim network. The attack is initiated by causing the target service to issue a request via its DNS resolver to the attacker’s domain, e.g., by trying to authenticate at eduroam’s wireless access point (steps ①, ② and ③ in Figure 1). In the zonefile of its domain, the attacker encodes malicious payloads into the DNS records. The records are then provided in responses to the queries of the DNS resolvers (step ④), and are subsequently relayed to the requesting services, in the example in Figure 1, to the radsecproxy server (step ⑤) which processes the attacker’s authentication request. The payload then causes the application to divert from a standard behaviour (step ⑥), e.g., causing it to allow unauthenticated network access.

3.1.2 Selecting Target Applications

We evaluate injection attacks against popular services and applications. In this work we present attacks against some selected applications listed in Table 1. We select them based on the following considerations:

DNS Use-Case. We identify 4 different use-cases of DNS (address lookup, service discovery, reverse lookup and authentication). We select a few popular applications and services for each DNS use case.

Triggering query. The attacker must be able to trigger a DNS lookup, e.g., via a script in a browser, via an Email to a target Email Server. We summarise methods for triggering query and setting the query domain in Table 1, column ‘Trigger/Set query’. We also prefer target applications which allow the attacker to trigger queries to attacker-selected domains, e.g., by sending an Email or triggering a query via javascript in browsers.

Attack surface. To find meaningful attacks, we focus on applications where input from DNS is used for some interesting action, e.g., for implementing a cache, creating a URL, etc. We do not analyse applications which only do standard address lookups (without caching), as such a scenario does not create a meaningful attack surface, even if no input validation is performed. We list the applications, along with how the DNS inputs are used by those applications, in Table 1, column ‘Input use’.

Usage of vulnerable resolvers. For applications which use the system’s libc resolver for DNS lookups, we prefer those

DNS Use-Case	Application	Trigger	Set	Uses libc	Validates	Input use	Attack found
		Query					
Address lookups (A, CNAME)	Chrome	js,html		yes	no	cache	no
	Firefox	js,html		yes	no	cache	no
	Opera	js,html		yes	no	cache	no
	Edge	js,html		yes	no	cache	no
	unscd	client app		yes	no	cache	no
	java	client app		both	no	cache	no
	ping(win32)	X	X	yes	no	display	yes
discovery (MX, SRV, NAPTR)	openjdk	login	X	no	no	create URL	yes
	ldapsearch	login	X	no	no	create URL	no
	radsecproxy	login		no	no	configure	yes
Reverse lookups (PTR)	ping(linux)	X	X	yes	no	display	yes
	trace(linux)	X	X	yes	no	display	yes
	OpenWRT	X	ping	yes	no	display	yes
	openssh	login		yes	no	display,log	yes
Authentication (TXT, TLSA)	policyd-spf	SMTP		no	no	text protocol	no
	libspf2	SMTP		no	-	parse	yes
All	Resolvers	client app		no	some	cache	yes

Table 1: Analysed software and tools.

which are often used on systems with vulnerable libc implementations. For example, OpenWRT was chosen because it uses a vulnerable libc implementation with uclibc.

3.1.3 Vulnerabilities Analysis

After identifying a target application, we analyse its DNS usage and whether input from the DNS is validated, as follows: (1) source code review, (2) fuzzing and (3) by executing the application, feeding it with inputs and analysing the resulting behaviour and the outputs. We first test if an application does not validate DNS records received in input. For such applications we then check how the input is used by the application, and construct attack vectors accordingly, e.g., XSS injection. The results of this analysis are listed in Table 1, the found vulnerabilities in Table 2.

3.2 DNS Caches

The attacks exploit the fact that domains and hostnames are not restricted to characters, and implements misinterpretation of domain names due to presence of "\." and of "\000" characters. These characters cause the appearance of "." to be altered hence manipulating the subdomains of a given parent domain.

The attacker can trigger a DNS query directly when launching the attack against open resolver or can initiate the attack via an application which uses the target DNS resolver, e.g., a web browser or an Email server.

3.2.1 DNS Cache Poisoning Attacks

In this section we present two types of cache-injection attacks which are based on domain name misinterpretation and verify them against popular DNS resolvers' software as well as against 3M open DNS resolvers in the Internet. We also show how to extend our poisoning attacks against forwarders and provide an example of the poisoning attack we launched using the public Verisign DNS resolver.

• **Attack #1: Period injection.** To inject a malicious DNS record or to overwrite a cached DNS record with a new

value (controlled by an attacker), we design the following record set `inject\.: www\.\target.com. A 6.6.6.6`. This attack requires the attacker to control a specially-malformed domain `www\.\target.com` under the same parent domain (in this example `com`.) as the domain of its victim, say `www.target.com`. Since most client software does not allow triggering a query for a domain `www\.\target.com` directly, to perform injection of a malicious record into the victim's cache, the attacker can set up a CNAME record with arbitrary subdomain, e.g., `injectdot.attacker.com`, as follows:

```
injectdot.attacker.com. CNAME www\.\target.com.
www\.\target.com. A 6.6.6.6
```

When decoding these records naively without escaping the period ("\.") it appears that `www.target.com` has IP address `6.6.6.6`. Caching this misinterpreted record after decoding leads to DNS cache injection.

• **Attack #2: Zero-byte injection.** We design the following record set `inject\000`, which indicates end of data, for performing DNS cache poisoning.

```
injectzero.attacker.com CNAME
www.target.com\000.attacker.com
www.target.com\000.attacker.com A 6.6.6.6
```

When naively decoded and fed into a victim cache this record enables an attacker to inject records for arbitrary domains into the cache. In this attack we also use a CNAME alias mapped to some secondary domain `injectzero.attacker.com`, since triggering a query to `www.target.com\000.attacker.com` without direct access to the resolver is not possible with most client software. When decoding this record set into a C-string without escaping the zero-byte after `www.target.com`, the `.\attacker.com` is removed since it is after the end of data `\000` value, the DNS software misinterprets the record and caches a record mapping `www.target.com` to IP address `6.6.6.6`.

3.2.2 Evaluation of the Attacks

Every application-level DNS-cache running on a system which misinterprets the `inject\.` or `inject\000` payloads (See Table 5) is vulnerable to these attacks. In our Internet study (see Section 4) we found that 105,854 open DNS resolvers (or 8% of 1,3M) are vulnerable to our attacks. Our attack evaluation was automated hence did not include potentially vulnerable resolvers which could result in a successful attack when the evaluation was manually tailored per resolver. These cases include lost packets (we sent only one response to avoid loading the network), resolvers with multiple caches (the attack was tested once against each client-exposed-IP of the resolver). Adjusting our attack to these cases is straightforward, would however generate much more traffic to the tested systems.

3.2.3 No Countermeasures Against Cache Poisoning

Classic countermeasures against DNS cache poisoning do not mitigate our cache poisoning attacks. The situation is even more risky when the same host is configured as nameserver and resolver, [20], as a lack of validation by the DNS resolver

Section - Category	DNS use-case	Software	Mis-interpretation is in	Record type	Attacker can choose domain	Possible outcome(s)
3.2 - DNS	Address-lookup	Verisign DNS	Resolver	CNAME	yes	Cache injection
3.3 - Eduroam	Service discovery	radsecproxy	Application	NAPTR, SRV	yes	Strip TLS, hijack connection, Crash
3.4 - LDAP	Service discovery	openjdk	Application	SRV	no	Crash
3.5 - Email	Authorization	libspf2	Application	TXT (SPF)	yes	Crash, (potential code execution)
3.6 - Admin tools	Address-, Reverse-lookup	ping,openssh,trace	Stub resolver	CNAME	no	Terminal Escape Code injection
3.6 - Web-interface	Reverse-lookup	OpenWRT luci	Stub resolver	PTR	yes	XSS in Admin web-interface

Table 2: Applications’ categories with vulnerabilities and attacks exploiting them.

can allow the attacker to also manipulate the zonefile which is hosted on the same machine.

Defences against off-path attackers. Defences against off-path attackers, such as [RFC5452] [21], are not effective against our attacks: we do not send the malicious DNS responses from spoofed IP addresses but respond from a nameserver that we control. Hence in our attacks the attacker does not need to guess the randomisation values, such as UDP source port and the TXID. The bailiwick check [22], which prevents the attackers from responding with values not under their domains is also ineffective against our attacks since the bailiwick check is applied over the records *before* the misinterpretation occurs.

Defences against on-path attackers. Cryptographic defences, most notably DNSSEC [RFC4033-RFC4035], can not prevent our cache poisoning attacks in common setups: in situations where upstream resolvers are used the misinterpreted records are not detected by the downstream DNS forwarders, since those typically do not perform DNSSEC validation¹. DNSSEC validation is performed by the recursive resolvers over the DNS records in line-format, *before* the decoding and the misinterpretation occur. After the records successfully pass DNSSEC validation, they are cached in a “misinterpreted” form.

Cross-zone CNAME caching. Additionally to the misinterpretation, these attacks require resolvers to cache and process CNAME records across zone-boundaries, i.e., the resolvers must use the misinterpreted second record `www.target.com\000.attacker.com` from zone `attacker.com` to answer queries for `www.target.com`. While this is not typically the case for recursive resolvers, we validated such behaviour in `dnsmasq`, the most frequently used forwarder on our open resolver dataset: given the records `injectdot.attacker.com CNAME www.victim.com` and `www.victim.com A 6.6.6.6` in response to a query for `injectdot.attacker.com`, `dnsmasq` will answer queries for `www.victim.com` with `6.6.6.6`. We illustrate how this leads to a vulnerable configuration of forwarder and recursive resolver in the case `dnsmasq` is combined with a misinterpreting recursive resolver like Verisign Public DNS in Figure 3; our resolver evaluation is in Section 4.

¹Neither `dnsmasq`, `systemd-resolved` nor `OpenWRT` or `Fritz!Box SOHO` routers perform DNSSEC validation by default.

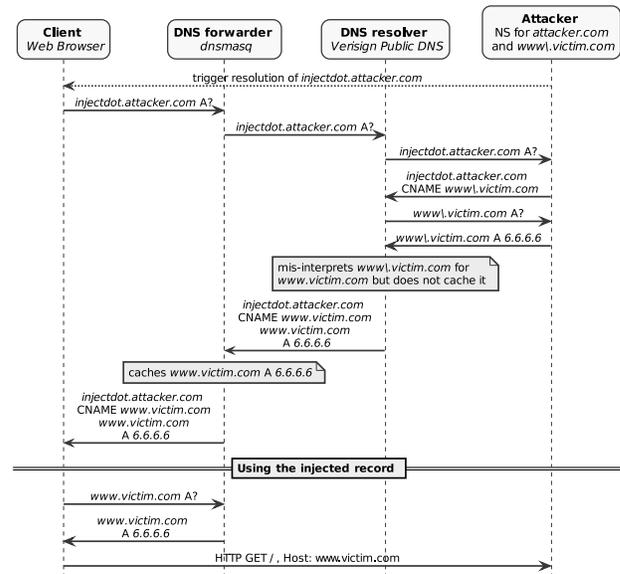


Figure 3: Downstream forwarder attack using `dnsmasq` with Verisign Public DNS misinterpretation of `inject\` payload.

3.2.4 Required attacker capabilities

To launch the `inject000` attack, the adversary only needs to control a nameserver for an arbitrary domain in the internet. In contrast, in order to launch the `inject\` attack the adversary has to control a specially crafted malicious sub-domain under the same parent-domain (e.g., `com.`) as his target. This means that conducting this attack requires registering a sub-domain like `www\`.`target` via a domain registry. Applicability of this attack depends on ability of the attacker to register such sub-domains. For instance, a registry.pw for `.pw` reported that registering domain `www\`.`asd.pw` was possible, while `www`.`asd.pw` or `asd.pw` was not (indicating that they are existing registered domains). Namely the attacker can register `www\`.`asd.pw` and use it to attack the existing victim domain `asd.pw`.

3.3 Eduroam Peer Discovery

Eduroam federation uses Remote Authentication Dial-In User Service (Radius) [23] for authentication of guest access. Radsecproxy is an application that implements Radius transport over TCP and TLS as well as dynamic peer discovery

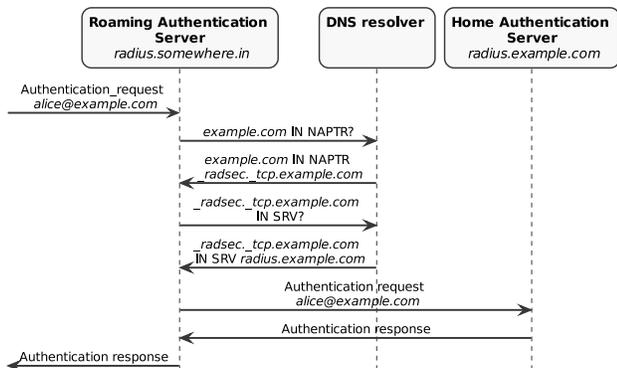


Figure 4: Radius Dynamic Peer discovery.

for servers which do not support these features themselves. Radsecproxy uses a shell-script-based method for dynamically updating the configuration to support the DNS lookups needed for Dynamic Peer discovery. The script is invoked with the domain component of the user’s network access identifier (i.e., `example.com` in Figure 4) as its first argument by the radsecproxy server and outputs a new dynamic radsecproxy configuration for the user’s realm. Example output of this script (called `naptr-eduroam.sh`) when invoked from shell is below:

```

$ ./naptr-eduroam.sh example.com
server dynamic_radsec.example.com {
    host radius1.example.com:2083
    host radius2.example.com:2083
    type TLS
}
  
```

3.3.1 Radius Dynamic Peer Discovery

In this section we provide a detailed explanation of the radius dynamic peer discovery process illustrated in Figure 4, as well as how an adversary can abuse the mechanism to trigger queries. First, a client (alice) connects to a wireless access point at the campus of `somewhere.in` (the domain of that university), providing authentication material including her network access identifier (NAI) `alice@example.com`. The access point then forwards the authentication request to the roaming authentication server at `somewhere.in`. From alice’s NAI this server defers that alice’s authentication request needs to be routed to the home authentication server of `example.com`. To find the home authentication server, it issues DNS queries for `example.com IN NAPTR?` followed by an SRV query of the domains listed in the NAPTR record. Finally, the roaming authentication server forwards the authentication request to the home authentication server at `radius.example.com`, which answers the request. The attacker sets up his own domain `attacker.com` and configures his nameserver to answer with one of the attack payloads from Table 3. The attacker provides a username `user@attacker.com` when connecting. This leads the roaming authentication server to send DNS requests via its resolver to the attacker’s nameserver. Depending on the payload, a corresponding attack is launched against the roaming authentication server.

3.3.2 Attacks Against Radsecproxy

We found multiple security vulnerabilities in the script [24] for Dynamic Peer Discovery in eduroam, see vulnerabilities in Table 3. The vulnerabilities allow an attacker to control various variables inside the script as well as in the generated dynamic configuration. These vulnerabilities are caused by the lack of input validation of the resulting output of `dig` as well as the usage of `printf`, which negates the escaping of special characters done by `dig`.

To initiate an attack the attacker causes the target system to issue a query to its domain. To make radsecproxy query for a domain of attacker’s choice, the attacker just needs to attempt to log-in at an eduroam access point with a username ending with the malicious domain. This triggers NAPTR and SRV queries to locate the correct authentication server, see messages exchange in Figure 4.

No validation of `dig` output: By changing the NAPTR record’s replacement field (`$HOST`), the attacker can control one argument to `dig`, which is not checked for its format. We exploited this to launch two attacks:

- **Attack #1 in Table 3.** Make `dig` query an attacker-chosen DNS resolver, instead of the default resolver from `/etc/resolv.conf`. We used this attack to verify the vulnerability remotely without causing damage to the tested eduroam network.

- **Attack #2 in Table 3.** Make `dig` use *any file* on the radsecproxy-system as a “batch-file” (option `-f`), thereby querying all lines in this file as DNS queries. Attackers which are located on-path to the DNS resolver can apply the second attack to read arbitrary files from the system.

Vulnerable usage of `printf`: The (double) use of `printf` in `naptr-eduroam.sh` allows the attacker to inject arbitrary strings into the dynamically generated configuration file via a format-string attack. The reason is that `printf` removes the escaping done by `dig` over the user input, which is given in format specifier argument. This allows the attacker to access radsecproxy’s configuration parser and subsequent `confserver_cb` function, which can be used to make the process read any file on the filesystem using `include /path/file`.

- **Attack #3 in Table 3.** We evaluated a ‘resource starvation’ attack using `/dev/zero` as an input, and caused an infinite 100% CPU usage loop in the configuration file parser.

- **Attack #4 in Table 3.** In this attack we demonstrate how the attacker can manipulate the generated dynamic configuration file, specifying a TLS certificate CommonName (CN) regular expression. This expression is passed to the `libc’s regcomp()` function, which on many implementations of `libc` (e.g., `glibc`) has known, unfixed vulnerabilities² which can be used, e.g., to crash radsecproxy via stack consumption.

- **Attack #5 in Table 3.** When the attacker provides a functional server configuration it can also override parameters of the dynamically generated server entry, most importantly the

²E.g., [25] can still be exploited on current Ubuntu and used in attack #4.

the resulting URL becomes `ldap://ldap.example.com/uid=admin,gid=users,dc=example,dc=com????./uid=john,gid=users,dc=example,dc=com`, which means the client will search for user `admin` instead of `john`, enabling `john` to execute actions with `admin` privileges. This attack enables to circumvent security mechanisms like LDAP over TLS (`ldaps://`) because it changes the information in the URL before it is transmitted over the TLS secured channel.

- **Attack #2: Denial-of-Service via malformed records.**

The LDAP SRV lookup function calls itself recursively after looking-up an SRV record, see Algorithm 1. We manipulate an SRV record so that the resulting URL does not contain a hostname-component, which then causes an infinite recursion and crashes the ‘LDAP-using’ application with a stack overflow.

3.4.2 Evaluation of the Attacks

We tested attack #1 against two LDAP library implementations (`ldapsearch` and `openjdk`’s `javax.naming`). We find that both applications use a potentially vulnerable LDAP peer discovery algorithm, which just concatenates the SRV record with the rest of the URL and do not check the contents of the SRV record for sanity. However, in both implementations, the DN (e.g., `uid=john,gid=users`) from the LDAP URL is actually ignored and must be given in an additional function call or parameter in order to allow execution of multiple search queries after the connection to the server has been established. We verified attack #2 experimentally using the record we constructed, and evaluated it against `openjdk`’s 11.0.6 `javax.naming` API:

```
_ldap._tcp.attacker.com. IN SRV /dc=attacker,dc=com.
```

Triggering a query. A query for the LDAP SRV record is triggered when a new connection to the LDAP server is created, i.e., when a user triggers an action which requires an LDAP-lookup such as logging into a web application which uses LDAP for user management. However, to execute the attack, the attacker must either be able to (1) control the full LDAP DN or (2) modify the SRV record on the network via a MitM position. In our evaluation we tested the implementation of LDAP middleware/libraries, which do not restrict how the LDAP DN is set. However, typical applications will restrict control over the LDAP DN to the components relevant for the user³, such that control over the necessary `dc=` components is not available to the attacker.

3.5 Domain-Based Anti-Spam Validation

Sender Policy Framework (SPF) [27] is a domain-based mechanism to prevent forgery of SMTP envelope headers. To trigger an SPF DNS query, the attacker needs to send an Email to an SPF-supporting Email server.

We provide a detailed explanation of the Email SPF validation process shown in Figure 5 in the case where an in-

³<https://docs.spring.io/spring-ldap/docs/current/reference/>

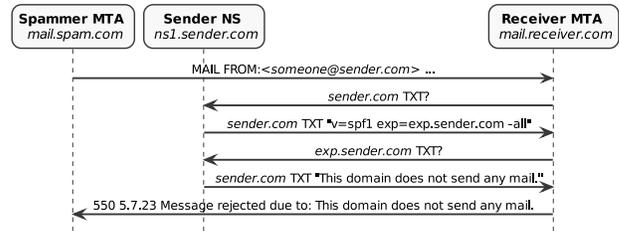


Figure 5: SPF resolution example.

coming Email is rejected: first, an non-authorized Email transfer agent at `mail.spam.com` (Spammer MTA) connects to the mail server at the receiver domain (`mail.receiver.com`) and tries to send a Mail coming from `someone@sender.com` to a mailbox at `receiver.com` using SMTP. To check if the Spammer MTA is authorized to send mail from `sender.com`, the Receiver MTA will query the DNS for the SPF records for `sender.com`. In this case the record indicates that no one is authorized to send mail from that domain (option `-all`) and that a detailed explanation why the mail is rejected is stored at `exp.sender.com`. After the receiver MTA has received this record it will decide to reject the Email and fetch the explanation from `exp.sender.com` via DNS to include it together with the rejection message. Finally, it parses the rejection message, replaces any included macros and sends it back to the spammer MTA notifying it that the mail was rejected and why.

3.5.1 Attacks Against Checks of SPF Records

- **Attack #1: Injection against `policyd-spf`.** The default implementation, e.g., Ubuntu [28] used by postfix, for checking SPF records is based on a separately running daemon called `policyd-spf`. This daemon is listening at a unix socket for responses to determine whether an Email should be rejected or not according to the SPF records [29]. The interface to this daemon is line based, the client (postfix) provides properties of the received Email line-by-line and submits the request with an empty line, as shown below. In this example `policyd-spf` session, the client’s request lines are marked with `>`, daemon response lines with `<`.

```
> request=smtpd_access_policy
> protocol_state=RCPT
> client_address=192.168.234.20
> sender=someone@hardfail.example.com
> recipient=vagrant@postfix
...
< policy_context=
< action=550 5.7.23 Message rejected due to: SPF fail -
not authorized. Please see http://www.openspf.net/Why?s=hello;
id=hardfail.example.com;ip=192.168.234.20;r=<UNKNOWN>
```

The server (`policyd-spf`) will answer with a single line providing information on how to proceed. We construct a malformed SPF record using the SPF `exp=` parameter. This parameter allows to include an explanation why an Email was rejected. We inject control characters into this client-server interface by including them in the rejection message, specified as a separate DNS record. The attacker can further include

```
attacker.com TXT "v=spf1 exp=exp.attacker.com"
exp.attacker.com TXT "AAAAA..." ; (510 times)
```

Figure 6: libspf2 exploit with malicious SPF record payload.

newline characters ("`\n`") to create additional output lines in the line-based connection to `policyd-spf`. This is interpreted as the responses to requests asked in the future thereby changing the SPF result for the next Email.

• **Attack #2: Stack buffer overflow in libspf2.** Libspf2 is a library for checking SPF records for incoming Email messages for Mail Transfer Agents. Libspf2 is used by some versions of the command line utility `spfquery` and also directly from SMTP server source code. Because of the complexity of the library we used fuzz-testing with `afl-fuzz` against a custom-built application calling libspf2 functions. This allowed us to provide the SPF record as a file rather than via the network to test libspf2 against potential vulnerabilities exploitable via malformed SPF records. The evaluations showed that libspf2 is vulnerable to the malformed records attack, see malicious SPF record payload in Figure 6. We performed attacks against the `spfquery` command line utility using the vulnerable record set which resulted in ‘stack-smashing detected’ error and crashes, further allowing remote-code-execution. This attack exploits a stack-buffer overflow while parsing the SPF explanation macro.

3.5.2 Evaluation of the Attacks

We evaluated attack #1 against postfix using `policyd-spf-perl`, and were able to inject additional lines of output to the unix socket, showing that `policyd-spf-perl` does not verify the contents of the SPF explanation record. In contrast to other attacks in this work, attack #2 cannot be prevented by validating the DNS records since the malicious SPF record presents a theoretically valid SPF explanation message.

3.6 Administrative Tools

In the attacks that we presented until now, the applications implemented the DNS lookup themselves, not by using an API like `gethostbyname()`, where the behaviour is standardised and data validation is performed by the systems’ stub resolvers. In this section we present vulnerabilities in applications which do not implement DNS lookups themselves but use the system stub resolver to do so. First, we analyse different stub resolver implementations in Section 3.6.1 and then show vulnerabilities in applications using these stub-resolvers in Section 3.6.2.

3.6.1 DNS Record Processing in stub resolvers.

Libc is the C standard library for C programming language, specified in ANSI C and is a subset of C library POSIX specification. There are different implementations of the standard C system library. We experimentally tested, as well as analysed the source code of, all the major implementations of the C system library and except two found them to be vulnerable, see Table 5. We explain our analysis of the DNS record

processing on two implementations: `glibc`⁴ and on `uClibc`⁵. We selected those implementations as examples because they represent two distinct methodologies that we observed in processing DNS records. The other tested implementations are similar to `uClibc`. We demonstrate the processing on the `gethostbyname()` library function as an example; the same applies to other calls.

After a DNS response has been received by `glibc` (resp `uClibc`) library, it is first checked against the length field, DNS transaction identifier and the return code (e.g., OK, NX-DOMAIN). The libraries then go through the resource record sets in the answer section and process each record. For each domain name in a DNS record the following steps are done, as shown with our example domain name `036123...6d00` in Figure 7: (1) domain name compression is removed; (2) domain name is decoded from DNS line format into a (zero-terminated) string; (3) domain name is validated. We explain these steps next.

DNS decompression and decoding into a string. These two steps are typically done simultaneously in one function, e.g., `dn_expand` for `glibc`, or `__decode_dotted` for `uClibc`. When decoding a domain name into a string, the resolver must ensure that the characters which cannot be represented in an ASCII string, must be escaped appropriately [RFC4343] [18]. This also applies to zero-bytes (which would otherwise be interpreted as string terminators) and period characters (which would otherwise be interpreted as label-separators). Escaping values outside the range of `0x21` ("!") to `0x7E` ("~") is required by [18].

In our example in Figure 7 this means that to avoid confusion with label separators the second byte (`0x3e`) of the first label must be expressed as `"\."` instead of `"."`. The final decoded domain name is then `"a\.b.<>.com."` if decoded correctly applying escaping to non-printable characters, or `"a.b.<>.com."` if decoded incorrectly when escaping is not applied.

Our analysis shows that `glibc` applies escaping to the decoded domain name, while `uClibc` does not. This means that any record which contains zero-bytes or dots inside labels will be misinterpreted during decoding when processed by `uClibc` and other non-printable characters will be included in the returned string unescaped. Such incorrect decoding logic can allow cache-injection attacks when the misinterpreted record for the domain `a.b.<>.com.` is cached and re-used in another context, as we show in Section 3.2.

Domain name validation. After the record has been decoded, it should be validated to check that it represents a valid hostname. The POSIX standard defines `getnameinfo()` to return ‘hostnames’, not ‘domain names’. Notice that domain names are defined in [RFC1034, RFC2181] [22, 30] as a list of binary labels that can contain any value. The only limitation over domain names is on the length of the name

⁴The GNU’s project implementation of the C standard Library

⁵The Linux standard library for mobile and embedded devices.

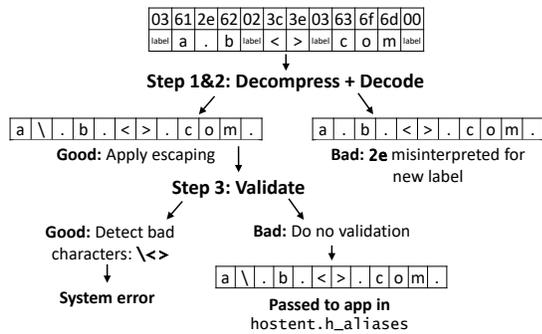


Figure 7: DNS record processing.

[RFC2181] [22]: 63 octets per component, 255 octets for a domain name. Values like, brackets ([,]), colons, NULs (\000), newlines, backslashes, and so on are all legal. In contrast, hostnames are only allowed to contain alphanumeric characters ("A-Z", "a-z", "0-9"), hyphens ("-") and dots (".") to separate labels, as defined in [RFC952] [16], this specification is referenced by newer standards [RFC1123, RFC2181, RFC3492] [17, 22, 31]. As a result the system library should validate that all the returned values represent valid hostnames, not domains. Our evaluation shows that most stub resolvers use a naive domain name decoding logic which misinterprets "\." for "." and "\000" as a string delimiter; we show how to exploit this for cache poisoning attacks in Section 3.2.

In our example in Figure 7 step 2, we assume that the library has decoded the domain name correctly, hence the value `a.b.<>.com.` is passed on to the validation step. Based on the logic that the POSIX standard defines the return value of `gethostbyname()` as a hostname the library should ensure that the hostname does not contain invalid characters, and if it finds any, it should signal an error to the application. If this is not done, the string `a.b.<>.com.` is passed to the application unchanged and exposes to a range of vulnerabilities as shown in Sections 3.6. Again we review the functions in `glibc` and `uClibc` and find that only `glibc` implements this validation correctly, `uClibc` does not validate the decoded domain name at all.

Notice that the steps 1-3 are always needed to transform a domain name from its line format to a (zero-terminated) ASCII string. An implementation might choose to switch the order of steps 2 (decoding) and 3 (validation), or combine these steps into one, but this does not change the fact that both steps are needed to correctly implement parsing of a line-format domain names into a hostname string. Therefore, our analysis of `glibc` and `uClibc` can be extended to other resolver implementations. We validated this assumption experimentally and via code review in popular system resolver implementations, and summarise the results in Table 5. Our results indicate that most stub resolvers do not check that domain names constitute valid hostnames. The other `libc` implementations, such as `dietlibc` and `windows`, all result in the same incorrectly processed output. We show how to exploit

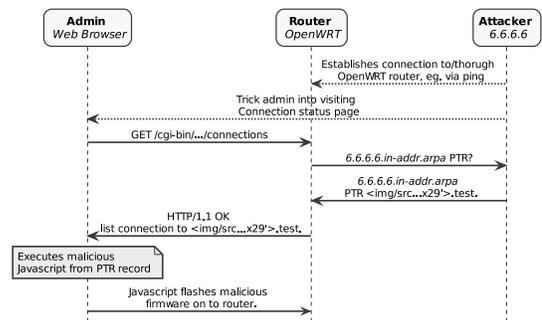


Figure 8: Attack flow against OpenWRT.

lack of validation to attack applications which do not expect special characters inside hostnames, launching different injection attacks, such as XSS or ANSI terminal escape code injection.

3.6.2 Attacks Against Administrative Utilities

We demonstrate attacks against two vulnerable example applications from different contexts: Windows 10's `ping` as an example of a simple command-line utility and `OpenWRT LuCi` as an example of a integrated web-based administration interface.

- **Attack #1: Windows ping.** `Ping` shows the CNAME alias of the ping-ed host without checking the alias for disallowed characters. We use this to inject arbitrary bytes into the output of `ping`, which is then displayed by the terminal. By including ANSI terminal escape codes in the CNAME record, an attacker controlling the DNS response can manipulate the terminal output by moving the cursor, replacing already printed data or changing the Window title⁶.

We tested a similar attack against linux `ping`, `traceroute` and `openssh`. Different than windows `ping`, these applications query the reverse PTR record of the remote address. We found that none of the applications perform input validation of the returned domain name and therefore allow the same kind of attack.

- **Attack #2: Cross-site scripting (XSS) attack in OpenWRT LuCi.** `OpenWRT` is an aftermarket operating system for residential-gateway routers based on linux. It uses `uClibc` as its standard C library. This means that any application included in `OpenWRT` which does not check CNAME or reverse-DNS responses itself is vulnerable to our attacks. We explore the web-based `LuCi` [34] administration interface which can be used to configure all of a router's settings via a web-browser.

Apart from letting the user change the router's settings, `LuCi` also contains a status page listing all currently active connections through the router at `/cgi-bin/luci/admin/status/realtime/connections`. When a user opens this page, `LuCi` performs reverse-DNS lookups for all IP addresses currently connected to or through the router. The attacker abuses this to trigger reverse queries for its own IP addresses in the

⁶Terminal-escape injection vulnerabilities were found in applications like web-servers [32, 33], however they were exploitable via direct input over HTTP.

Use-Case	Apps	validates	attack prevented by
Address Lookup	Chrome, Firefox, Opera, Edge, unscd, java	no	Cache does not evaluate CNAME aliases for indexing
Service discovery	ldapsearch	no	Generated URL not used to specify LDAP query
Authentication	policyd-spf + postfix	no	postfix resets connection if second answer is injected

Table 4: Non-validating applications without meaningful exploits.

reverse-DNS tree and delivers malicious DNS records from there. The attacker also maintains a connection to or through the OpenWRT router when the victim opens the ‘Connections page’, by continuously sending ICMP echo-request (‘ping’) messages to the router. Using the `6.6.6.6.in-addr.arpa` record shown in Figure 9 we were able to successfully launch this attack by placing javascript code inside the PTR record of the attacker’s IP address which is not validated by the LuCi web interface. This record is queried when the user views the ‘Connections page’ and is placed in the page’s HTML code which executes the malicious code. A successful attack allows different exploits, as an example we created a record which, when injected, loads a third party script from an external HTTP server. This script then issues various requests to the OpenWRT configuration interface in name of the user which finally results in replacing the OpenWRT firmware running on the device with a malicious attacker-provided firmware. We show an example of a full attack flow against OpenWRT using this vulnerability in Figure 8.

3.7 Impact of the Attacks

Our analysis of the attacks shows that none of the tested applications perform DNS input validation. Some attacks do not result in meaningful exploits, see examples in Section 3.7.1, the causes are not defences against potentially malicious inputs, but rather not security related implementation decisions. We also showed attacks with limited impact, such as the ping exploit, to demonstrate that the problem (lack of DNS-input validation) is systematic and prevalent, affects different systems, services and tools and is not limited to isolated cases or a specific application. These examples show that developers do not check DNS results, which contain untrusted data. Furthermore, using such tools in certain scenarios could expose to the exploitable attack. For instance, output of utilities like Ping, can also be saved in a format which allows command injection, like an HTML report, or SQL database, this allows for a much more severe attack.

3.7.1 Non-validating applications without exploits

In this section we provide examples (listed in Table 4) where DNS input validation vulnerability does not lead to meaningful attacks. We caution that the factors preventing meaningful attacks are not security related, but are due to different implementation considerations: (1) in browsers and

app-caches, the cached records are only indexed by their query domain, not by additional domain aliases inside the response. (2) in `ldapsearch`, a potentially vulnerable peer discovery algorithm is used to generate an LDAP URL, but depending on the configuration this URL may not be used to specify the LDAP query, which would prevent the attacker from changing the query. (3) in `policyd-spf` if implementations are configured to check for additional data on unix socket, the result will be discarded to prevent desynchronisation with the `policyd-spf` daemon.

4 Internet Evaluation on Open Resolvers

In Section 3 we examined how the popular DNS resolvers and stub resolver implementations built into operating systems and programming languages handle control characters in domain names and if they modify any of the maliciously crafted payloads needed to conduct our application-specific exploits. In this section we extend our evaluation to open DNS resolvers in the Internet, and confirm the results of our in-lab study in the Internet. Specifically, we evaluate behaviour when processing a set of crafted DNS records designed to trigger classic input validation vulnerabilities like SQL injection, as well as DNS-specific special cases like handling of period characters inside DNS labels. We do not evaluate attacks against the applications using the vulnerable resolvers, since this would result in an attack against an application or service in the network which we do not own. We do not risk evaluating even the ‘more benign’ attacks, which when run against the servers that we setup do not cause critical damage. This is since our attacks can trigger unexpected outcome when evaluated against the servers in the Internet, e.g., due to differences in configurations. We observe this phenomena during the DNS cache poisoning evaluation of domains that we control against open resolvers in the Internet – some of the DNS software which was found secure when running the cache poisoning attack against our servers, resulted in cache poisoning vulnerabilities in the Internet.

4.1 Methodology

We run the same set of queries as for the in-lab evaluation, but employ additional logging at the nameserver-level to gain knowledge about the nature of the resolvers we test. As this study targets the same class of resolver as our in-lab evaluation, the expected behaviour is the same as described in Section 2.

Dataset. To conduct the study, we use a dataset from Censys [35] with 3M open resolvers. We include baseline tests for each record type (A, CNAME, SRV, TXT) to ensure that the resolver supports the record type we use in our payloads and only consider resolvers who respond to our queries and return the correct result for all of these baseline tests. This results in 1,328,146 open resolvers from 228 different countries.

Tests are conducted by using custom test applications which send DNS queries to the resolvers over the network

or by calling the respective stub resolvers using the appropriate API, by calling the POSIX `gethostbyname()` and `getnameinfo()` functions. To test handling of commonly used control characters like slash ("`/`"), at ("`@`"), zero-byte ("`\000`"), etc. we use the CNAME and PTR records with the injection payloads. We furthermore test all the resolvers against the applications'-specific records listed in Section 3 whenever applicable. All tested payloads are listed in Figure 9 in the order they appear in Table 5 and 6.

Forward-lookups. We evaluate hostname-to-address records for all 3 groups of resolvers (resolver, stub and open resolver) by triggering a query to the domain name of the payload (e.g., `cnameslash.example.com`) and observing the response from the resolver.

Injection payloads. For the DNS-specific injection payloads, the test takes place in two stages: First we trigger a query to the domain name of the injection payload (i.e., `injectzero.example.com`) twice and observe if the result was misinterpreted, i.e. if the result CNAME is `www.target.com\000.example.com` (marked **X**) or just `www.target.com` (marked **X⁵**). We then trigger a query to the potentially misinterpreted domain name (i.e., `www.target.com` instead of `www.target.com\000.example.com`) and observe if the IP address was successfully injected into the resolvers' cache (marked **✓** or given in **%** for open resolvers). We test each payload (`inject\000` and `inject\.`) in both scenarios: via a CNAME-record which points to the malicious record (i.e., `injectzero.attacker.com`) and by triggering a query to the malicious domain (i.e., `victim.com\000.attacker.com`) directly.

Reverse-lookups. Reverse-DNS lookups (PTR) were tested against system stub resolvers only by setting the upstream DNS server to a custom controlled nameserver directly providing the records under `in-addr.arpa`-tree and triggering a reverse PTR lookup for the respective IP address, e.g., `1.1.1.1.in-addr.arpa` and observing the response.

Additional considerations. To enhance the robustness of our tests, we randomise all the queried domain names by prepending a random subdomain to ensure that the query is not cached before the test is conducted and processed by all the components of the DNS lookup chain. This also allows us to link the open resolver we sent the query to, to the final recursive resolver which connects back to our nameserver by matching these random prefixes. To prevent other users of the resolver to be negatively affected by our tests, we run these tests only against domains we own, which allows us to validate the full injection attack without performing an attack against any other domain.

4.2 Evaluation Results

We present the results for forward-lookups in Table 5 and the results of reverse-lookups in Table 6. For each test, ticks (**✓**) mark that the resolver is vulnerable to this kind of payload and crosses (**X**) mark that the resolver is not vulnerable. Note

```

cnamebase.example.com CNAME works.cnameslash.example.com
cnameslash.example.com CNAME t/t.cnameslash.example.com
cnameat.example.com CNAME t\t.cnameat.example.com
cnamexss.example.com CNAME <img/src='' /onerror='alert
&#x28;&#x22xss&#x22&#x29'>.
cnamexss.example.com
cnamesql.example.com CNAME 'OR''='--.cnamesql.example.com
cnameansi.example.com CNAME \027[31\;1\;4mHello\027[0m.
cnameansi.example.com

injectdot.example.com. CNAME www.target.com.
www.target.com. A 6.6.6.6
injectzero.example.com CNAME www.target.com\000.example.com
www.target.com\000.example.com A 6.6.6.6

_ldap._tcp.example.com. IN SRV /dc=example,dc=com.
_radsec._tcp.example.com IN SRV 6.6.6.6\|n\|ttype\|tTCP\|n\|t
secret\|t\something\|n\|n%$p.

exp.example.com IN TXT "AAAAA..." (510 times)

1.1.1.1.in-addr.arpa PTR works.test
2.2.2.2.in-addr.arpa PTR te/st.test
3.3.3.3.in-addr.arpa PTR te\@st.test
4.4.4.4.in-addr.arpa PTR t\000t.test
5.5.5.5.in-addr.arpa PTR t\..t.test
6.6.6.6.in-addr.arpa PTR <img/src='' /onerror='alert
&#x28;&#x22xss&#x22&#x29'>.test
7.7.7.7.in-addr.arpa PTR 'OR''='--.test
8.8.8.8.in-addr.arpa PTR \027[31\;1\;4mHello\027[0m.test

```

Figure 9: Injection payloads based on CNAME and PTR records.

however that depending on the test, 'vulnerable' means that the resolver performs as expected and conforms to the DNS standard (in case of non-stub resolvers and special character tests like `cnameslash`) or that it misinterprets a domain name which allows for a cache-poisoning attack (in case of injection payloads, such as `inject\000`). For injection payloads, we only call a resolver vulnerable (**✓**) when the malicious IP address was cached, otherwise we use (**X⁵**) to show that the misinterpretation occurs, but can only be exploited in conjunction with a caching downstream resolver like `dnsmasq` (See Section 3.2.3). We list the percentage and absolute number of open resolvers vulnerable to each payload in the bottom of Table 5.

Transparent handling of DNS records. The results from Internet evaluation of open resolvers correspond to our in-lab evaluation: Around 96% of all tested open DNS resolvers are transparent for application-specific payloads in DNS records.

Comparing the CNAME-based `cnameslash` and SRV-based LDAP and Eduroam payloads, we can observe that there is some difference of resolver behaviour in handling these records, even though they are based on the same property (including non-standard characters in a domain name field). This can be seen as a confirmation that many resolvers indeed ignore the contents of record type field that is not directly needed for the DNS resolution (like SRV), but do not do so for records which are important to finish the lookup (like A or CNAME).

Cache injection vulnerabilities. For the injection specific payloads `inject\.` and `inject\000`, the evaluation results also mostly match the expectation from the in-lab evaluation. Like all resolver implementations except Verisign Public DNS, most Open DNS resolvers which do handle the injection pay-

loads transparently and are not vulnerable to our injection payloads.

Nevertheless, we found 1.3% to 4.6% of the open DNS resolvers to be vulnerable to a cache poisoning attack without any further requirements, which was verified by querying for the maliciously injected record. Overall, 8.0% (or 105,854) of the open resolvers are vulnerable to cache poisoning via *any* of the injection payloads. This result is alarming considering that the inject\000 attack does not require any attacker resources other than control over an arbitrary nameserver in the internet.

Misinterpretation analysis. In our Internet measurement we observed the following phenomenon: the resolvers respond with the non-misinterpreted value in return to the first query for a record with injection payload in our domain. However, subsequent queries for the same resource record (with the same domain name) which are responded from the cache (without issuing a query to our nameserver) can result in one of the two outcomes: (1) a misinterpreted value or (2) non-misinterpreted value. Nevertheless a subsequent validation of the cached record confirms that in both cases the injection was successful and the target cache stored a misinterpreted record.

Vulnerable DNS software in Internet. We use `version.bind` special query to infer the implementation and version of the open resolvers. While most servers do not respond to these queries (in our study 65%, see column ‘our study’ in Table 5), we find that, depending on the implementation and attack type, even those exact versions we have found not vulnerable during our lab evaluation were found vulnerable during our evaluation of open resolvers in the Internet, e.g., 19.2% of vulnerable resolvers were Bind resolvers. In these cases we find that the misinterpretation causing the vulnerability does not occur in the ‘visible’ software implementations, but rather in one of the upstream (forwarding) resolvers in the resolution chain, which we cannot identify.

In some cases upstream forwarders is a public DNS service, like Google, OpenDNS and Cloudflare. We infer this by mapping the IP addresses of the DNS queries received by our nameservers to the Autonomous System (AS) numbers of the resolvers sending the queries.

[RFC1034] non-compliant resolvers. 12% of the open resolvers in the Internet exhibit the same flaws as we found in stub resolver implementations. These are misinterpretations when decoding line-format domain names into zero-terminated string. Not only this allows for cache injection attacks, but it also indicates that those resolvers do not follow the recommendations in [RFC1034] [30], requiring that the DNS software should store domain names in their line format and not as decoded strings.

5 Root Causes, Insights and Mitigations

Missing Specifications

The study that we carried out in this work showed us that one of the root factors allowing our attacks is a lack of threat model in the standard RFCs as well as a lack of specifications on DNS and on its interactions with applications:

Threat modelling. There is a lack of threat modelling in the DNS infrastructure and in the interaction of DNS with the applications. The RFCs should provide a threat model discussing potential pitfalls. For instance, most applications typically expect hostnames, but in the RFCs, this is not considered.

DNS record parsing. There is a lack of detailed specification on how to parse DNS records. This critical functionality should be specified in the RFCs.

Validation of DNS records. There is a lack of standardised implementation for validation of DNS records. This is important esp. for non-address record types as these are not supported by the OS resolver (e.g., `libc`). Similarly to, say libraries for generating DNSSEC keys, there should be an implementation for validation of received DNS records.

Domain names vs hostnames. There is a discrepancy between definitions of domain names and hostnames, which leads to confusion in DNS software in how to parse the DNS records. To avoid pitfalls the same rules should apply to both.

Mitigations

Applications. Since DNS resolvers serve data from untrusted Internet sources, the applications should always treat data from DNS the same way they treat user input, hence validation of formatting and escaping should always be performed (ie. validate html special characters, etc.), regardless if the API used to receive the data indicates an already checked result (POSIX) or not.

System stub resolvers. Stub resolvers should be modified to check if domain names returned by POSIX calls like `gethostbyname()` are valid hostnames, [17]. If not, the domain name should not be given to the application, like it is implemented in `glibc` already. Merely escaping special characters as done by `netbsd` can still be vulnerable. For instance, we demonstrate this with the `6.6.6.6.in-addr.arpa` Cross-Site Scripting payload in Figure 9 which does not contain any character typically escaped in domain names and, for example, can be used to execute an attack against OpenWRT. Furthermore non-`libc` DNS libraries should follow the same rules (i.e., only allow hostnames as per [17]) even for non-standard lookup types like SRV to prevent confusion among developers who only used `libc` resolvers before. Guidance on how to implement such checks should ideally be given by standardisation bodies.

DNS resolvers. Filtering DNS responses on the DNS resolver or forwarder level is possible but is against the DNS standard [3, 4, 17]. Changing this requires a discussion in

DNS Payload shown in		Fig. 9						Section 3.2.1		Section 3.2.1		Sec. 3.4.2	Tab. 3	Fig. 6
Test	Base	/	@	XSS	SQL	ANSI	inject\	inject\	inject\	inject\	LDAP	Eduroam	libspf2	
							CNAME Direct	CNAME Direct	CNAME Direct	CNAME Direct				
Recursive resolvers	BIND (9.14.0)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	MaraDNS Deadwood (3.2.14)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✗ ²	
	Unbound (1.9.1)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	PowerDNS Recursor (4.3.0)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Windows Server (2012 R2)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Windows Server (2016)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Windows Server (2019)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
Forwarders	pdnsd (1.2.9a)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	dnsmasq (2.79)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	NxFilter (4.3.3.9)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	systemd resolved (237)	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
Public resolvers	OpenDNS	✓	✓	✓	✓	✓ ¹	✗	✗	✗	✗	✓	✓	✓	
	Cloudflare Public DNS	✓	✓	✓	✓	✓ ¹	✗	✗	✗	✗	✓	✓	✓	
	Comodo Secure DNS	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Google Public DNS	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Hurricane Electric	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Neustar UltraRecursive	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Norton ConnectSafe	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	Oracle Dyn	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	SafeDNS	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	
	VeriSign Public DNS	✓	✓	✓	✓	✓	(✗) ⁵	✗	✗	✗	✓	✓	✓	
Yandex DNS	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓		
Stub resolvers	glibc	✓	✗	✗	✗	✗	✗	-	✗	-	n/a no SRV/TXT support or does not apply			
	musl	✓	✗	✗	✗	✗	(✗) ⁵	-	✗	-				
	dietlibc	✓	✓	✓	✓	✓	(✗) ⁵	-	(✗) ⁵	-				
	uclibc	✓	✓	✓	✓	✓	(✗) ⁵	-	(✗) ⁵	-				
	Windows	✓	✓	✓	✓	✓	(✗) ⁵	-	(✗) ⁵	-				
	NetBSD	✓	✓	(✓) ³	✓	✓	✗ ³	-	✗ ³	-				
	Mac OS X	✓	✓	✓	✓	✓	✗ ³	-	✗ ³	-				
	go*	✓	✓	✓	✓	✗	(✗) ⁵	-	✗	-				
	nodejs	✓	✓	✓	✓	✓	✗ ³	-	(✗) ⁵	-				
openjdk8*	InetAddress.getCanonicalHostName() gives PTR instead of CNAME										✓			
Open resolvers in Internet	100%	96.6%	96.3%	96.4%	95.9%	96.3%	1.3%	2.7%	3.6%	4.6%	99.6%	99.6%	56.1%	
	1,328,146	1,283,447	1,279,573	1,279,835	1,273,710	1,279,100	45,596	51,049	82,589	88,864	1,322,529	1,322,203	745,599	

✓: Vulnerable. ¹: record converted to lower case. ²: NXDOMAIN/no response. ³: output was escaped. ⁵: Record is misinterpreted, injection is not cached.
*: Uses system stub resolver by default but offers a builtin one.

Table 5: Forward-lookup test results for all groups of resolvers.

Test	Base	/	@	\.	\000	XSS	SQL	ANSI
Payload (Fig.9)	1.1.1.1	2.2.2.2	3.3.3.3	5.5.5.5	4.4.4.4	6.6.6.6	7.7.7.7	8.8.8.8
glibc	✓	✗	✗	✗	✗	✗	✗	✗
musl	✓	✓	✓	✓	✓	✓	✓	✓
dietlibc	✓	✓	✓	✓	✓	✓	✓	✓
uclibc	✓	✓	✓	✓	✓	✓	✓	✓
windows	✓	✓	✓	✓	✓	✓	✓	✓
netbsd	✓	✓	(✓) ²	(✓) ²	(✓) ²	✓	✓	(✓) ²
mac os x	✓	✓	✓	(✓) ²	✓	✓	✓	✓
go*	✓	✓	✓	✓	(✓) ³	✓	✓	✓
openjdk8*	✓	✗	✓	(✓) ²	(✓) ³	✓ ⁴	✓	✓
node	✓	✓	✓	(✓) ²	✓	✓	✓	✓

✓: Vulnerable. ²: output was escaped. ³: Zero-byte did not stop output.
⁴: Alternative XSS payload with " " instead of "/".
*: Uses system stub resolver by default but offers a builtin-one.

Table 6: Reverse-lookup results for different stub resolvers.

the corresponding working groups within the IETF, which we are initiated within our disclosure efforts. Nevertheless, performing checks on DNS records is challenging: some applications, like SRV service discovery [38], require domain names with characters that are not allowed in hostnames (e.g., underscore). Defining a list of allowed characters so that legitimate applications would still work but injection attacks would be blocked should be further investigated and is not straightforward. In particular, it is difficult to foresee what characters and formats will be needed by future applications,

hence a ‘too-restrictive’ list of allowed characters would make DNS less transparent, possibly introducing obstacles in deployment of new applications, or when adding new versions or new features to existing applications. On the other hand, leaving this completely transparent may lead to confusion about what values a field can actually have, and can even introduce vulnerabilities – our work can be generalised to other Internet protocols. We show that the decision to enable easy future deployment of new applications by not restricting the domain names to alphanumeric characters exposes to attacks.

Nevertheless, as an immediate protection against our attacks, operators which are unable to implement changes on the application- or stub-resolver-level might use filtering proxies which implement those validation steps on the network-level instead⁷.

Mitigations against cache-poisoning. Since the cache-poisoning attacks cannot be reliably detected by downstream forwarders, these attacks must be mitigated by patching the resolvers causing the misinterpretation. If the resolver is only misinterpreting malicious records (but not caching them, like Verisign Public DNS), switching to a DNS forwarder which does not cache cross-CNAME records can prevent the attack. This however does not fix the root cause of the issue.

⁷We provide a proof-of-concept implementation of such a proxy at <https://xdi-attack.net>.

When the resolver operator cannot not fix the vulnerability, switching to another DNS resolver is the best option.

6 Related Work

DNS cache poisoning. Kaminsky provided the first demonstration of DNS cache poisoning attack [39]. Since then DNS resolvers have been patched to support best practices [RFC5452] [21]: randomising fields in requests, such as source port and DNS TXID, and validating them in responses, and also to apply checks such as bailiwick [22]. This makes DNS resilient to off-path cache poisoning attacks. Nevertheless, recent works developed cache poisoning attacks when DNS responses are served over UDP [40]. The attacks use different side channels to predict the randomisation parameters, as well as other methodologies like fragmentation to bypass guessing the parameters altogether, [8–10,41–45]. Our attacks are not limited by the transport protocol and apply to DNS over TCP as well as DNS over UDP. In contrast to all existing DNS cache poisoning attacks which evaluate the cache poisoning on one victim DNS resolver and then check if some selected population of DNS resolvers have the properties that could potentially make them vulnerable, our attack is the first to have been fully automated and evaluated on a large set of target networks, 3M, and the first to have been successfully launched against 105K resolvers. Prior attacks cannot be automated since they need to be tailored per each target victim resolver [8–10,44,45]. For instance, the servers set the fragmentation offset slightly differently hence making fragmentation difficult to match, the servers randomise the records in responses making the UDP checksum extremely difficult to match, UDP ports need to be measured per each target separately, overwriting cached records with new values depends on already cached records and caching policies, and so on. Our attack is not restricted by these hurdles.

None of the proposed non-cryptographic defences prevent our cache poisoning attacks. Even the cryptographic protection with DNSSEC [RFC4033-RFC4035] [11–13], which blocks all previous DNS cache poisoning attacks, does not prevent our attacks in common settings. Furthermore, DNSSEC deployments were showed to often use weak cryptographic algorithms or vulnerable keys, [46–48]. Cipher-suite negotiation schemes were proposed to allow easy adoption of stronger cryptographic ciphers [49].

Recent proposals for encryption of DNS traffic, such as DNS over HTTPS [50] and DNS over TLS [51], although vulnerable to traffic analysis [52,53], may also enhance resilience to cache poisoning but do not prevent our injection attack.

User input injections in web applications. Injection vulnerabilities [54] are the primary medium for performing remote exploits, including SQL injection attacks [54], Cross Site Scripting (XSS) [55], buffer overflow [56], XPath injections [57], LDAP injections [58], HTTP header injection [59], Email header injection [60], SMTP injection [61]. All these

differ from our injection attacks. User injection attacks via the web interfaces are typically blocked as user input is sanitised prior to being accepted by applications. Our attacks apply even when user input is properly validated and also where the users cannot provide any meaningful input at all, since we deliver malicious payloads by encoding them into DNS records.

DNS rebinding attacks. DNS rebinding attack [62,63] uses a script on the victim network and an external attacker to create a confusion in web browsers bypassing Same Origin Policy (SOP), say by mapping the external attacker to an internal IP address. This allows the attacker to impersonate internal hosts in order to bypass filtering that is applied on external packets, e.g., for spam or for Denial of Service (DoS) attacks. Our attacks target the internal services directly, without impersonation of internal devices. DNS rebinding are prevented, e.g., with filtering private IP addresses and blocking the resolution of external hostnames into internal IP addresses, or via DNS pinning [62] in web browsers - none of which prevent our attacks.

7 Conclusions

Our work shows that central transparency-related principles of development of Internet systems should be reconsidered:

Flexibility. Be strict when sending and permissive when receiving is a good principle in the Internet. Systems and protocols that are too rigid are much more difficult to use and require significant changes to the existing infrastructure for adoption of new technologies or mechanisms. The huge success of DNS in providing platform to new applications is thanks to its transparent handling of DNS records. If DNS is made less transparent, e.g., by requiring that the records are checked for invalid characters, it would make the roll out of new applications in the Internet much more challenging. For instance, if DNS parsed each record, new applications using records containing not yet supported characters, e.g., not just alphanumeric characters, like in SRV record type, would require changes to the DNS servers all over the Internet to enable support for new characters. Unupgraded servers would risk failures or even crashes when processing the new records. On the other hand, making systems too tolerant can expose to vulnerabilities. We showed that leaving the specification completely open exposes DNS and the applications using DNS to attacks. Hence, a balance should be found between the ease of deployment and security.

Layering. Although it is a known networking principle that each layer provides services to the layer above it, and the upper layer does not have to worry about the data provided by lower protocols, we show that when it comes to security this principle may result in vulnerabilities. We recommend that the validation of DNS data is integrated into applications directly not relying on the lower layers to do this for them. For instance, it may not always be possible for DNS to predict all the applications of the data that it provides and scenarios

where it will be used. Hence even if DNS is changed to apply checks over the data in DNS records, the applications should nevertheless do the validation also themselves.

Acknowledgements

We are grateful to Yuval Yarom and to the anonymous referees for their thoughtful feedback on our work.

This work has been co-funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) SFB 1119.

References

- [1] C. Farrell, M. Schulze, S. Pleitner, and D. Baldoni, “DNS Encoding of Geographical Location,” RFC 1712 (Experimental), Internet Engineering Task Force, Nov. 1994. [Online]. Available: <http://www.ietf.org/rfc/rfc1712.txt>
- [2] P. Hoffman and J. Schlyter, “The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA,” RFC 6698 (Proposed Standard), Internet Engineering Task Force, Aug. 2012, updated by RFCs 7218, 7671. [Online]. Available: <http://www.ietf.org/rfc/rfc6698.txt>
- [3] A. Gustafsson, “Handling of Unknown DNS Resource Record (RR) Types,” RFC 3597 (Proposed Standard), Internet Engineering Task Force, Sep. 2003, updated by RFCs 4033, 4034, 4035, 5395, 6195, 6895. [Online]. Available: <http://www.ietf.org/rfc/rfc3597.txt>
- [4] P. Mockapetris, “Domain names - implementation and specification,” RFC 1035 (INTERNET STANDARD), Internet Engineering Task Force, Nov. 1987, updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2673, 2845, 3425, 3658, 4033, 4034, 4035, 4343, 5936, 5966, 6604, 7766. [Online]. Available: <http://www.ietf.org/rfc/rfc1035.txt>
- [5] W. G. Halfond, J. Viegas, A. Orso *et al.*, “A classification of SQL-injection attacks and countermeasures,” in *Proceedings of the IEEE international symposium on secure software engineering*, vol. 1. IEEE, 2006, pp. 13–15.
- [6] J. Grossman, S. Fogie, R. Hansen, A. Rager, and P. D. Petkov, *XSS attacks: cross site scripting exploits and defense*. Syngress, 2007.
- [7] T. Pietraszek and C. V. Berghes, “Defending against injection attacks through context-sensitive string evaluation,” in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2005, pp. 124–145.
- [8] A. Herzberg and H. Shulman, “Fragmentation Considered Poisonous: or one-domain-to-rule-them-all.org,” in *IEEE CNS 2013. The Conference on Communications and Network Security, Washington, D.C., U.S.* IEEE, 2013.
- [9] X. Zheng, C. Lu, J. Peng, Q. Yang, D. Zhou, B. Liu, K. Man, S. Hao, H. Duan, and Z. Qian, “Poison over troubled forwarders: A cache poisoning attack targeting DNS forwarding devices,” in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 577–593.
- [10] K. Man, Z. Qian, Z. Wang, X. Zheng, Y. Huang, and H. Duan, “DNS Cache Poisoning Attack Reloaded: Revolutions with Side Channels,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security CCS*. ACM, 2020.
- [11] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, “DNS Security Introduction and Requirements,” RFC 4033 (Proposed Standard), Internet Engineering Task Force, Mar. 2005, updated by RFCs 6014, 6840. [Online]. Available: <http://www.ietf.org/rfc/rfc4033.txt>
- [12] —, “Resource Records for the DNS Security Extensions,” RFC 4034 (Proposed Standard), Internet Engineering Task Force, Mar. 2005, updated by RFCs 4470, 6014, 6840, 6944. [Online]. Available: <http://www.ietf.org/rfc/rfc4034.txt>
- [13] —, “Protocol Modifications for the DNS Security Extensions,” RFC 4035 (Proposed Standard), Internet Engineering Task Force, Mar. 2005, updated by RFCs 4470, 6014, 6840. [Online]. Available: <http://www.ietf.org/rfc/rfc4035.txt>
- [14] “Standard for Information Technology–Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7,” *IEEE Std 1003.1, 2016 Edition (incorporates IEEE Std 1003.1-2008, IEEE Std 1003.1-2008/Cor 1-2013, and IEEE Std 1003.1-2008/Cor 2-2016)*, pp. 1–3957, Sep. 2016.
- [15] R. Gilligan, S. Thomson, J. Bound, J. McCann, and W. Stevens, “Basic Socket Interface Extensions for IPv6,” RFC 3493 (Informational), Internet Engineering Task Force, Feb. 2003. [Online]. Available: <http://www.ietf.org/rfc/rfc3493.txt>
- [16] K. Harrenstien, M. Stahl, and E. Feinler, “DoD Internet host table specification,” RFC 952, Internet Engineering Task Force, Oct. 1985, updated by RFC 1123. [Online]. Available: <http://www.ietf.org/rfc/rfc952.txt>
- [17] R. Braden, “Requirements for Internet Hosts - Application and Support,” RFC 1123 (INTERNET STANDARD), Internet Engineering Task Force, Oct. 1989, updated by RFCs 1349, 2181, 5321, 5966, 7766. [Online]. Available: <http://www.ietf.org/rfc/rfc1123.txt>
- [18] D. E. 3rd, “Domain Name System (DNS) Case Insensitivity Clarification,” RFC 4343 (Proposed Standard), Internet Engineering Task Force, Jan. 2006. [Online]. Available: <http://www.ietf.org/rfc/rfc4343.txt>
- [19] —, “Domain Name System (DNS) IANA Considerations,” RFC 6895 (Best Current Practice), Internet Engineering Task Force, Apr. 2013. [Online]. Available: <http://www.ietf.org/rfc/rfc6895.txt>
- [20] H. Shulman and M. Waidner, “Towards security of internet naming infrastructure,” in *European Symposium on Research in Computer Security*. Springer, 2015, pp. 3–22.
- [21] A. Hubert and R. van Mook, “Measures for Making DNS More Resilient against Forged Answers,” RFC 5452 (Proposed Standard), Internet Engineering Task Force, Jan. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5452.txt>
- [22] R. Elz and R. Bush, “Clarifications to the DNS Specification,” RFC 2181 (Proposed Standard), Internet Engineering Task Force, Jul. 1997, updated by RFCs 4035, 2535, 4343, 4033, 4034, 5452. [Online]. Available: <http://www.ietf.org/rfc/rfc2181.txt>
- [23] C. Rigney, A. Rubens, W. Simpson, and S. Willens, “Remote Authentication Dial In User Service (RADIUS),” RFC 2058 (Proposed Standard), Internet Engineering Task Force, Jan. 1997, obsolete by RFC 2138. [Online]. Available: <http://www.ietf.org/rfc/rfc2058.txt>
- [24] radsecproxy, “naptr-eduroam.sh,” 2012. [Online]. Available: <https://github.com/radsecproxy/radsecproxy/blob/8d287300f510e0559f01a2e7a4dec90674215f25/tools/naptr-eduroam.sh>
- [25] MITRE, “CVE-2010-4052 in GNU C Library (glibc),” 2010. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2010-4052>
- [26] M. Smith and T. Howes, “Lightweight Directory Access Protocol (LDAP): Uniform Resource Locator,” RFC 4516 (Proposed Standard), Internet Engineering Task Force, Jun. 2006. [Online]. Available: <http://www.ietf.org/rfc/rfc4516.txt>
- [27] S. Kitterman, “Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1,” RFC 7208 (Proposed Standard), Internet Engineering Task Force, Apr. 2014, updated by RFC 7372. [Online]. Available: <http://www.ietf.org/rfc/rfc7208.txt>

- [28] Ubuntu Community Help Wiki. (2013) Postfix/SPF. <https://help.ubuntu.com/community/Postfix/SPF>, accessed 2020-02-04.
- [29] The Postfix Home Page. (2020) Postfix SMTP Access Policy Delegation. http://www.postfix.org/SMTDP_POLICY_README.html, accessed 2020-02-04.
- [30] P. Mockapetris, "Domain names - concepts and facilities," RFC 1034 (INTERNET STANDARD), Internet Engineering Task Force, Nov. 1987, updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936. [Online]. Available: <http://www.ietf.org/rfc/rfc1034.txt>
- [31] A. Costello, "Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)," RFC 3492 (Proposed Standard), Internet Engineering Task Force, Mar. 2003, updated by RFC 5891. [Online]. Available: <http://www.ietf.org/rfc/rfc3492.txt>
- [32] MITRE, "CVE-2009-4487 in nginx 0.7.64," 2009. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2009-4487>
- [33] —, "CVE-2013-1862 in Apache HTTP Server," 2013. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2013-1862>
- [34] The OpenWrt Project, "LuCI - OpenWrt Configuration Interface," 2020. [Online]. Available: <https://github.com/openwrt/luci>
- [35] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman, "A search engine backed by Internet-wide scanning," in *22nd ACM Conference on Computer and Communications Security*, Oct. 2015.
- [36] M. Kührer, T. Hupperich, J. Bushart, C. Rossow, and T. Holz, "Going wild: Large-scale classification of open dns resolvers," in *Proceedings of the 2015 Internet Measurement Conference*, ser. IMC '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 355–368. [Online]. Available: <https://doi.org/10.1145/2815675.2815683>
- [37] D. Tatang, C. Schneider, and T. Holz, "Large-scale analysis of infrastructure-leaking dns servers," in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2019, pp. 353–373.
- [38] A. Gulbrandsen, P. Vixie, and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)," RFC 2782 (Proposed Standard), Internet Engineering Task Force, Feb. 2000, updated by RFC 6335. [Online]. Available: <http://www.ietf.org/rfc/rfc2782.txt>
- [39] D. Kaminsky, "It's the End of the Cache As We Know It," Presentation at Blackhat Briefings, 2008.
- [40] Y. Gilad, A. Herzberg, and H. Shulman, "Off-path hacking: The illusion of challenge-response authentication," *IEEE Security & Privacy*, vol. 12, no. 5, pp. 68–77, 2013.
- [41] A. Herzberg and H. Shulman, "Security of patched DNS," in *European Symposium on Research in Computer Security*. Springer, 2012, pp. 271–288.
- [42] —, "Socket overloading for fun and cache-poisoning," in *Proceedings of the 29th Annual Computer Security Applications Conference*, 2013, pp. 189–198.
- [43] —, "Vulnerable delegation of dns resolution," in *European Symposium on Research in Computer Security*. Springer, 2013, pp. 219–236.
- [44] M. Brandt, T. Dai, A. Klein, H. Shulman, and M. Waidner, "Domain validation++ for mitm-resilient pki," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 2060–2076.
- [45] F. Alharbi, J. Chang, Y. Zhou, F. Qian, Z. Qian, and N. Abu-Ghazaleh, "Collaborative Client-Side DNS Cache Poisoning Attack," in *INFOCOM*. IEEE, 2019.
- [46] T. Dai, H. Shulman, and M. Waidner, "Dnssec misconfigurations in popular domains," in *International Conference on Cryptology and Network Security*. Springer, 2016, pp. 651–660.
- [47] H. Shulman and M. Waidner, "One key to sign them all considered vulnerable: Evaluation of DNSSEC in the internet," in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, 2017, pp. 131–144.
- [48] T. Chung, R. van Rijswijk-Deij, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "Understanding the role of registrars in dnssec deployment," in *Proceedings of the 2017 Internet Measurement Conference*, 2017, pp. 369–383.
- [49] A. Herzberg, H. Shulman, and B. Crispo, "Less is more: cipher-suite negotiation for dnssec," in *Proceedings of the 30th Annual Computer Security Applications Conference*, 2014, pp. 346–355.
- [50] P. Hoffman and P. McManus, "Rfc 8484: Dns queries over https (doh)," 2018.
- [51] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. Hoffman, "Rfc 7858-specification for dns over transport layer security (tls)," 2016.
- [52] H. Shulman, "Pretty bad privacy: Pitfalls of dns encryption," in *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, 2014, pp. 191–200.
- [53] S. Siby, M. Juarez, C. Diaz, N. Vallina-Rodriguez, and C. Troncoso, "Encrypted dns-> privacy? a traffic analysis perspective," *arXiv preprint arXiv:1906.09682*, 2019.
- [54] Z. Su and G. Wassermann, "The essence of command injection attacks in web applications," *Acm Sigplan Notices*, vol. 41, no. 1, pp. 372–382, 2006.
- [55] Y. Nadji, P. Saxena, and D. Song, "Document structure integrity: A robust basis for cross-site scripting defense," in *NDSS*, vol. 20, 2009.
- [56] M. Dalton, H. Kannan, and C. Kozyrakos, "Real-world buffer overflow protection for userspace and kernelspace," in *USENIX Security Symposium*, 2008, pp. 395–410.
- [57] J. Blasco, "Introduction to xpath injection techniques," in *Hakin9, Conference on IT Underground, Czech Republic*, 2007, pp. 23–31.
- [58] J. M. Alonso, R. Bordon, M. Beltran, and A. Guzman, "LDAP injection techniques," in *2008 11th IEEE Singapore International Conference on Communication Systems*, Nov. 2008, pp. 980–986, iSSN: null.
- [59] M. Johns and J. Winter, "Requestrodeo: Client side protection against session riding," in *Proceedings of the OWASP Europe 2006 Conference*, 2006.
- [60] S. P. Chandramouli, P.-M. Bajan, C. Kruegel, G. Vigna, Z. Zhao, A. Doupe, and G.-J. Ahn, "Measuring e-mail header injections on the world wide web," in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, 2018, pp. 1647–1656.
- [61] T. Terada, "Smtp injection via recipient email addresses," *MBSD White Paper (December 2015)*, 2015.
- [62] C. Jackson, A. Barth, A. Bortz, W. Shao, and D. Boneh, "Protecting browsers from dns rebinding attacks," *ACM Transactions on the Web (TWEB)*, vol. 3, no. 1, pp. 1–26, 2009.
- [63] G. Acar, D. Y. Huang, F. Li, A. Narayanan, and N. Feamster, "Web-based attacks to discover and control local iot devices," in *Proceedings of the 2018 Workshop on IoT Security and Privacy*, 2018, pp. 29–35.