PDoT: Private DNS-over-TLS with TEE Support

Yoshimichi Nakatsuka
University of California, Irvine
nakatsuy@uci.edu

Andrew Paverd*
Microsoft Research
andrew.paverd@ieee.org

Gene Tsudik
University of California, Irvine
gene.tsudik@uci.edu

ABSTRACT

Security and privacy of the Internet Domain Name System (DNS) have been longstanding concerns. Recently, there is a trend to protect DNS traffic using Transport Layer Security (TLS). However, at least two major issues remain: (1) how do clients authenticate DNS-over-TLS endpoints in a scalable and extensible manner; and (2) how can clients trust endpoints to behave as expected? In this paper, we propose a novel Private DNS-over-TLS (*PDoT*) architecture. *PDoT* includes a DNS Recursive Resolver (RecRes) that operates within a Trusted Execution Environment (TEE). Using *Remote Attestation*, DNS clients can authenticate, and receive strong assurance of trustworthiness of *PDoT* RecRes. We provide an open-source proof-of-concept implementation of *PDoT* and use it to experimentally demonstrate that its latency and throughput match that of the popular Unbound DNS-over-TLS resolver.

CCS CONCEPTS

• **Security and privacy** → **Web protocol security**; *Hardware-based security protocols*; *Network security*; *Privacy protections*.

KEYWORDS

Domain Name System, Privacy, Trusted Execution Environment

ACM Reference Format:

Yoshimichi Nakatsuka, Andrew Paverd, and Gene Tsudik. 2019. **PDoT**: Private DNS-over-TLS with TEE Support. In *2019 Annual Computer Security Applications Conference (ACSAC '19)*, December 9–13, 2019, San Juan, PR, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3359789.3359793>

1 INTRODUCTION

The Domain Name System (DNS) [29] is a distributed system that translates human-readable domain names into IP addresses. It has been deployed since 1983 and, throughout the years, DNS privacy has been a major concern.

In 2015, Zhu et al. [39] proposed a DNS design that runs over Transport Layer Security (TLS) connections [13]. DNS-over-TLS protects privacy of DNS queries and prevents man-in-the-middle (MiTM) attacks against DNS responses. [39] also demonstrated the

practicality of DNS-over-TLS in real-life applications. Several open-source recursive resolver (RecRes) implementations, including Unbound [23] and Knot Resolver [12], currently support DNS-over-TLS. In addition, commercial support for DNS-over-TLS has been increasing, e.g., Android P devices [18] and Cloudflare’s 1.1.1.1 RecRes [9]. However, despite attracting interest in both academia and industry, some problems remain.

The first challenge is that clients need a way to authenticate the RecRes. Certificate-based authentication is natural for websites, since the user (client) knows the URL of the desired website and the certificate securely binds this URL to a public key. However, this approach cannot be used for a DNS RecRes because the RecRes does not have a URL or any other unique long-term user-recognizable identity that can be included in the certificate. One way to address this issue is to provide clients with a white-list of trusted RecRes’s public keys. However, this is neither scalable nor maintainable, because the white-list would have to include all possible RecRes operators, ranging from large public services (e.g., 1.1.1.1) to small-scale providers, e.g., a local RecRes provided by a coffee-shop.

Even if the RecRes can be authenticated, the second major issue is the lack of means to determine whether a given RecRes is trustworthy. For example, even if communication between client stub (client) and RecRes, and between RecRes and the name server (NS) is authenticated and encrypted using TLS, the RecRes must decrypt the DNS query in order to resolve it and contact the relevant NS-s. This allows the RecRes to learn unencrypted DNS queries, which poses privacy risks of a malicious RecRes misusing the data, e.g., profiling users or selling their DNS data. Some RecRes operators go to great lengths to assure users that their data is private. For example, Cloudflare promises “We will never sell your data or use it to target ads” and goes on to say “We’ve retained KPMG to audit our systems annually to ensure that we’re doing what we say” [9]. Although helpful, this still requires users to trust the auditor and can only be used by operators who can afford an auditor.

In this paper, we use Trusted Execution Environments (TEEs) and *Remote Attestation* (RA) to address these two problems. By using RA, the identity of the RecRes is no longer relevant, since clients can check what software a given RecRes is running and make trust decisions based on how the RecRes behaves. RA is one of the main features of modern hardware-based TEEs, such as Intel Software Guard Extensions (SGX) [27] and ARM TrustZone [4]. Such TEEs are now widely available, with Intel CPUs after the 7th generation supporting SGX, and ARM Cortex-A CPUs supporting TrustZone. TEEs with RA capability are also available in cloud services, such as Microsoft Azure [28]. Specifically, our contributions are:

- We design a Private DNS-over-TLS (*PDoT*) architecture, the main component of which is a privacy-preserving RecRes that operates within a commodity TEE. Running the RecRes

*Work done while visiting University of California, Irvine, as a US-UK Fulbright Cyber Security Scholar.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACSAC '19, December 9–13, 2019, San Juan, PR, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7628-0/19/12.

<https://doi.org/10.1145/3359789.3359793>

inside a TEE prevents even the RecRes operator from learning clients' DNS queries, thus providing query privacy. Our RecRes design addresses the authentication challenge by enabling clients to trust the RecRes based on how it behaves, and not on who it claims to be. (See Section 4).

- We implement a proof-of-concept *PDoT* RecRes using Intel SGX and evaluate its security, deployability, and performance. All source code and evaluation scripts are publicly available [21]. Our results show that *PDoT* handles DNS queries without leaking information while achieving sufficiently low latency and offering acceptable throughput (See Sections 5 and 6).
- In order to quantify privacy leakage via traffic analysis, we performed an Internet measurement study. It shows that 94.7% of the top 1,000,000 domain names can be served from a *privacy-preserving* NS that serves at least two distinct domain names, and 65.7% from a NS that serves 100+ domain names. (See Section 7).

2 BACKGROUND

2.1 Domain Name System (DNS)

DNS is a distributed system that translates host and domain names into IP addresses. DNS includes three types of entities: *Client Stub* (client), *Recursive Resolver* (RecRes), and *Name Server* (NS). Client runs on end-hosts. It receives DNS queries from applications, creates DNS request packets, and sends them to the configured RecRes. Upon receiving a request, RecRes sends DNS queries to NS-s to resolve the query on client's behalf. When NS receives a DNS query, it responds to RecRes with either the DNS *record* that answers client's query, or the IP address of the next NS to contact. RecRes thus recursively queries NS-s until the record is found or a threshold is reached. The NS that holds the queried record is called: *Authoritative Name Server* (ANS). After receiving the record from ANS, RecRes forwards it to client. It is common for RecRes to cache records so that repeated queries can be handled more efficiently.

2.2 Trusted Execution Environment (TEE)

A Trusted Execution Environment (TEE) is a security primitive that isolates code and data from privileged software such as the OS, hypervisor, and BIOS. All software running outside TEE is considered untrusted. Only code running within TEE can access data within TEE, thus protecting confidentiality and integrity of this data against untrusted software. Another typical TEE feature is remote attestation (RA), which allows remote clients to check precisely what software is running inside TEE.

One recent TEE example is Intel SGX, which enables applications to create isolated execution environments called *enclaves*. The CPU enforces that only code running within an enclave can access that enclave's data. SGX also provides RA functionality.

Memory Security. SGX reserves a portion of memory called the Enclave Page Cache (EPC). It holds 4KB pages of code and data associated with specific enclaves. EPC is protected by the CPU to prevent non-enclave access to this memory region. Execution threads enter and exit enclaves using SGX CPU instructions, thus ensuring that in-enclave code execution can only begin from well-defined call gates. From a software perspective, untrusted code can

make *ECALLs* to invoke enclave functions, and enclave code can make *OCALLs* to invoke untrusted functions outside the enclave.

Attestation Service. SGX provides two types of attestation: local and remote. Local attestation enables one enclave to attest another (running on the same machine) to verify that the latter is a genuine enclave actually running on the same CPU. Remote attestation involves more entities. First, an application enclave to be attested creates a *report* that summarizes information about itself, e.g., code it is running. This report is sent to a special enclave, called *quoting enclave* which is provided by Intel and available on all SGX machines. Quoting enclave confirms that requesting application enclave is running on the same machine and returns a *quote*, which is a report with the quoting enclave's signature. The application enclave sends this quote to the Intel Attestation Service (IAS) and obtains an *attestation verification report*. This is signed by the IAS confirming that the application enclave is indeed a genuine SGX enclave running the code it claims. Upon receiving an attestation verification report, the verifier can make an informed trust decision about the behavior of the attested enclave.

Side-Channel Attacks. SGX is vulnerable to side-channel attacks [24, 36], and various mechanisms have been proposed [11, 32, 34] to mitigate them. Since defending against side-channel attacks is orthogonal to our work, we expect that a production implementation would include relevant mitigation mechanisms.

3 ADVERSARY MODEL & REQUIREMENTS

3.1 Adversary Model

The adversary's goal is to learn, or infer, information about DNS queries sent by clients. We consider two types of adversaries, based on their capabilities:

The first type is a malicious RecRes operator who has full control over the physical machine, its OS and all applications, including RecRes. We assume that the adversary cannot break any cryptographic primitives, assuming that they are correctly implemented. We also assume that it cannot physically attack hardware components, e.g., probe CPU physically to learn TEE secrets. This adversary also controls all of RecRes's communication interfaces, allowing it to drop/delay packets, measure the time required for query processing, and observe all cleartext packet headers.

The second type is a network adversary, which is strictly weaker than the malicious RecRes operator. In the passive case, this adversary can observe any packets that flow into and out of RecRes. In the active case, this adversary can modify and forge network packets. DNS-over-TLS alone (without *PDoT*) is sufficient to thwart a passive network adversary. However, since an active adversary could redirect clients to a malicious RecRes, clients need an efficient mechanism to authenticate the RecRes and determine whether it is trustworthy, which is the main contribution of *PDoT*.

We do not consider Denial-of-Service (DoS) attacks on RecRes, since these do not help to achieve either adversary's goal of learning clients' DNS queries. Connection-oriented RecRes-s can defend against DoS attacks using cookie-based mechanisms to prevent SYN flooding [39].

3.2 System Requirements

We define the following requirements for the overall system:

R1: Query Privacy. Contents of client’s query (specifically, domain name to be resolved) should not be learned by the adversary. Ideally, payload of the DNS packets should be encrypted. However, even if packets are encrypted, their headers leak information, such as source and destination IP addresses. In Section 7.1, we quantify the amount of information that can be learned via traffic analysis.

R2: Deployability. Clients using a privacy-preserving RecRes should require no special hardware. Minimal software modifications should be imposed. Also, for the purpose of transition and compatibility, a privacy-preserving RecRes should be able to interact with legacy clients that only support unmodified DNS-over-TLS.

R3: Response Latency. A privacy-preserving RecRes should achieve similar response latency to that of a regular RecRes.

R4: Scalability. A privacy-preserving RecRes should process a realistic volume of queries generated by a realistic number of clients.

Note: query privacy guarantees provided by *PDOT* rely on the forward-looking assumption that communication between RecRes and respective NS-s is also protected by DNS-over-TLS. The DNS Privacy (DPriv) Working Group is working towards a standard for encryption and authentication of DNS resolver-to-ANS communication [6], using essentially the same mechanism as DNS-over-TLS. We expect an increasing number of NS-s to begin supporting this standard in the near future. Once *PDOT* is enabled at the RecRes, it can provide incremental query privacy for queries served from a DNS-over-TLS NS. As discussed in Section 5, with small design modifications, *PDOT* can be adapted for use in NS-s.

4 SYSTEM MODEL & DESIGN CHALLENGES

4.1 *PDOT* System Model

Figure 1 shows an overview of *PDOT*. It includes four types of entities: client, RecRes, TEE, NS-s. We now summarize *PDOT* operation, reflected in the figure: (1) After initial start-up, TEE creates an attestation report. (2) When client initiates a secure TLS connection, the attestation report is sent from RecRes to the client alongside all other information required to setup a secure connection. (3) Client authenticates and attests RecRes by verifying the attestation report. It checks whether RecRes is running inside a genuine TEE and running trusted code. (4) Client proceeds with the rest of the TLS handshake procedure only if verification succeeds. (5) Client sends a DNS query to RecRes through the secure TLS channel it has just set up. (6) RecRes receives a DNS query from client, decrypts it into TEE memory, and learns the domain name that the client wants to resolve. (7) RecRes sets up a secure TLS channel to the appropriate NS in order to resolve the query. (8) RecRes sends a DNS query to NS over that channel. If NS’s reply includes an IP address of the next NS, RecRes sets up another TLS channel to that NS. This is done repeatedly, until RecRes successfully resolves the name to an IP address. (9) Once RecRes obtains the final answer, it sends this to client over the secure channel. Client can reuse the TLS channel for future queries.

Note that we assume RecRes is not under the control of the user. In some cases, users could run their own RecRes-s, which would side-step the concerns about query privacy. For example,

Figure 1: Overview of the proposed system.

modern home routers are sufficiently powerful to run an in-house RecRes. However, this approach cannot be used in public networks (e.g., airports or coffee shop WiFi networks), which are the target scenarios for *PDOT*.

4.2 Design Challenges

The following key challenges were encountered in the process of *PDOT*’s design:

- C1: TEE Limited Functionality.** In order to satisfy their security requirements, TEE-s often limit the functionality available to code that runs within them. One example is the inability to fork within the TEE. Forking a process running inside the TEE forces the child process to run outside the TEE, breaking RecRes security guarantees. Another example is that system calls, such as socket communication, cannot be made from within the TEE.
- C2: TEE Memory Limitations.** A typical TEE has a relatively small amount of memory. Although an SGX enclave can theoretically have a large amount of in-enclave memory, this will require page swapping of EPC pages. The pages to be swapped must be encrypted and integrity protected in order to meet the security requirements of SGX. Therefore, page swapping places a heavy burden on performance. To avoid page swapping, enclave size should be less than the size of the EPC – typically, 128MB. Since RecRes is a performance-critical application, its size should ideally not exceed 128MB. This limit negatively impacts RecRes throughput, as it bounds the number of threads that can be spawned in a TEE.
- C3: TEE Call-in/Call-out Overhead.** Applications requiring functionality that is not available within the TEE must switch to the non-TEE side. This introduces additional overhead, both from the switching itself, and from the need to flush and reload CPU caches. Identifying and minimizing the number of times RecRes switches back and forth (whilst keeping RecRes functionality correct) is a substantial challenge.

Figure 2: Overview of *PDoT* implementation.

5 IMPLEMENTATION

Figure 2 shows an overview of the *PDoT* design. Since our design is architecture-independent, it can be implemented on any TEE architecture that provides the features outlined in Section 2.2. We chose the off-the-shelf Intel SGX as the platform for the proof-of-concept *PDoT* implementation in order to conduct an accurate performance evaluation on real hardware. (See Section 6). Therefore, our implementation is subject to performance and memory constraints in the current version of Intel SGX, and is thus best suited for small-scale networks, e.g., the public WiFi network provided of a typical coffee shop. However, as TEE technology advances, we expect that our design will scale to larger networks.

5.1 *PDoT*

PDoT consists of two parts: (1) trusted part residing in TEE enclaves, and (2) untrusted part that operates in the non-TEE region. The former is responsible for resolving DNS queries, and the latter – for accepting incoming connections, assigning file descriptors to sockets, and sending/receiving data received from the trusted part.

Enclave Startup Process. When the application enclave starts, it generates a new public-private key-pair within the enclave. It then creates a *report* that summarizes enclave and platform state. The report includes a SHA256 hash of the entire code that is supposed to run in the enclave (called *MRENCLAVE* value) and other attributes of the target enclave. *PDoT* also includes a SHA256 hash of the previously generated public key in the report. The report is then passed on to the SGX quoting enclave to receive a *quote*. The quoting enclave signs the report and thus generates a quote, which cryptographically binds the public key to the application enclave. The quoting enclave sends the quote to the application enclave, which forwards it to the Intel Attestation Service (IAS) to obtain an *attestation verification report*. It can be used in the future by clients to verify the link between the public key and the

MRENCLAVE value. After receiving the attestation verification report from IAS, the application enclave prepares a self-signed X.509 certificate required for the TLS handshake. In addition to the public key, the certificate includes: (1) attestation verification report, (2) attestation verification report signature, and (3) attestation report signing certificate, extracted from (1). *MRENCLAVE* value and hash of public key are enclosed in the attestation verification report.

TLS Handshake Process.¹ Once the application enclave is created, *PDoT* can create TLS connections and accept DNS queries from clients. The client initiates a TLS handshake process by sending a message to *PDoT*. This message is captured by untrusted part of *PDoT* and triggers the following events.² First, untrusted part of *PDoT* tells the application enclave to create a new TLS object within the enclave for this incoming connection. This forces the TLS endpoint to reside inside the enclave. The TLS object is then connected to the socket where the client is waiting to be served. RecRes then exchanges several messages with the client, including the self-signed certificate that was created in the previous section. Having received the certificate from RecRes, the client authenticates RecRes and validates the certificate. (For more detail, see Section 5.2). Only if the authentication and validation succeed, the client resumes the handshake process.

DNS Query Resolving Process. The client sends a DNS query over the TLS channel established above. Upon receiving the query, RecRes decrypts it within the application enclave and obtains the target domain name. RecRes begins to resolve the name starting from root NS, by doing the following repeatedly: 1) set up a TLS channel with NS, 2) send DNS queries and receive replies via that channel. Once RecRes receives the answer from NS, RecRes returns it to the client over the original TLS channel.

¹In implementing this process, we heavily relied on SGX RA TLS [20] whitepaper.

²Since we consider a malicious RecRes operator, it has an option not to trigger these events. However, clients will notice that their queries are not being answered and can switch to a different RecRes.

Figure 3: Overview of *PDoT* threading model.

Figure 3 illustrates how *PDoT* divides DNS query resolution process into three threads: (1) receiving DNS query – *ClientReader*, (2) resolving it – *QueryHandler*, and (3) returning the answer – *ClientWriter*.

ClientReader and *ClientWriter* threads are spawned anew upon each query. Dividing receiving and sending processes and giving them a dedicated thread is helpful because many clients send multiple DNS queries within a short timespan without waiting for the answer to the previous query.³ When *ClientReader* thread receives a DNS query from the client, it stores the query and a client ID in a FIFO queue, called *inQueryList*.

QueryHandler threads are spawned when *PDoT* starts up. The number of *QueryHandler* threads is configured by *RecRes* operator. *QueryHandler* threads are shared among all current *ClientReader* and *ClientWriter* threads. When a *QueryHandler* thread detects an entry in the *inQueryList*, it removes this entry and retrieves the query and the client ID. *QueryHandler* first checks whether this client is still accepting answers from *RecRes*. If not, *QueryHandler* simply ignores this query and moves on to the next one. If the client is still accepting answers, *QueryHandler* resolves the query and puts the answer into a FIFO queue (called *outQueryList*) dedicated to that specific client. In some cases, NS response might be too slow. When that happens, *QueryHandler* thread drops that query, since it is very likely that the request will also be dropped. This also prevents resources (such as mutex) from being locked up by this *QueryHandler* thread. In our implementation, this timeout was set to be the same as the client’s timeout.

Once an answer is added to *outQueryList* dedicated to its client, *ClientWriter* uses that answer to compose a DNS reply packet and sends it to the client. The reason we have N *outQueryList*s for N clients is to improve performance. With only one *outQueryList*, *ClientWriter* threads must search through the queue to find the answer for the connected client. This takes $O(M \times N)$ time, where N is the number of clients and M is the number of queries each client sends. Instead, with N *outQueryList*s, we reduce complexity to $O(1)$ because *ClientWriter* thread merely selects the query at the head of the list.

Caching. Some DNS recursive resolvers can cache query results. Caching is beneficial from the client’s perspective, since, in case of a cache hit, the *RecRes* can answer immediately, thus reducing query latency. The *RecRes* also benefits from not having to establish connections to external NS-s. However, irrespective of how

³For example, a client has received a webpage that includes images and advertisements that are served from servers located at different domains. This triggers multiple DNS queries at the same time.

it is implemented, caching at the *RecRes* causes potential privacy leaks, e.g., timings can reveal whether a certain domain record was already in the cache. This is an orthogonal challenge, discussed in Section 7.2.

To explore caching in a privacy-preserving resolver, we implemented a simple in-enclave cache for *PDoT*. It uses a red-black tree data structure and stores all records associated with the clients’ queries, indexed by the queried domain. This results in $O(\log_2(N))$ access times with N entries in the cache. In practice, *PDoT* could also use current techniques to mitigate against side-channel attacks on cache’s memory access patterns, e.g., [10, 31, 34]. During remote attestation, clients can ascertain whether the resolver has enabled caching, and which mitigations it uses.

***PDoT* ANS with TEE support.** With minor design changes, *PDoT* *RecRes* design can be modified for use as an ANS. Similar to the caching mechanism described above, an *PDoT* ANS can look up the answers to queries in an internal database, rather than contact external NS-s. The same way that clients authenticate *PDoT* *RecRes*, the *RecRes* can authenticate the *PDoT* ANS. Clients can thus establish trust in both *RecRes* and ANS using *transitive attestation* [2].

5.2 Client with *PDoT* Support

We took the *Stubby* client stub from the *getdns* project [22] which offers DNS-over-TLS support and modified it to perform remote attestation during the TLS handshake. We now describe how the client verifies its *RecRes*, decides whether the *RecRes* is trusted, and emits the DNS request packet.

***RecRes* Verification.** After receiving a DNS request from an application, the client first checks whether there is an existing TLS connection to its *RecRes*. If so, the client reuses it. If not, it attempts to establish a new connection. During the handshake, the client receives a certificate from *RecRes*, from which it extracts: 1) attestation verification report, 2) attestation verification report signature, and 3) attestation report signing certificate. This certificate is self-signed by IAS and we assume that the client trusts it. From (3), the client first retrieves the IAS public key and, using it, verifies (2). Then, the client extracts the SHA256 hash of *RecRes*’s public key from (1) and verifies it against a copy from (3). This way, the client is assured that *RecRes* is indeed running in a genuine SGX enclave and uses this public key for the TLS connection.

Trust Decision. The client also extracts the MRENCLAVE value from (1), which it compares against the list of acceptable MRENCLAVE values. If the MRENCLAVE value is not listed or one of the verification steps fail, the client stub aborts the handshake, moves on to the next *RecRes*, and re-starts the process. Note that the trust decision process is different from the normal TLS trust decision process. Normally, a TLS server-side certificate binds the public key to one or more URLs and organization names. However, by binding the MRENCLAVE value with the public key, the clients can trust *RecRes* based on its behavior, and not its organization (recall that the MRENCLAVE value is a hash of *RecRes* code). There several options for deciding which MRENCLAVE values are trustworthy. For example, vendors could publish lists of expected MRENCLAVE values for their resolvers. For open-source resolvers like *PDoT*, anyone can re-compute the expected MRENCLAVE value by recompiling the software, assuming a reproducible build process. This would

allow trusted third parties (e.g., auditors) to inspect the source code, ascertain that it upholds required privacy guarantees, and publish their own lists of trusted MRENCLAVE values.

Sending DNS request. Once the TLS connection is established, the client sends the DNS query to RecRes over the TLS tunnel. If it does not receive a response from RecRes within the specified timeout, it assumes that there is a problem with RecRes and sends a DNS reply message to the application with an error code `SERVFAIL`.

5.3 Overcoming Technical Challenges

As discussed in Section 4.2, *PDOT* faced three main challenges, which we addressed as follows:

Limited TEE Functionality. The inability to use sockets within the TEE is a challenge because the RecRes cannot communicate with the outside world. We address this issue by having a process running outside the TEE, as described in Section 5.1. This process forwards packets from the client to TEE through ECALLs and sends packets received from TEE via OCALLs. However, this process might redirect the packet to a malicious process or simply drop it. We discuss this issue in Section 6.1. Another function unavailable within TEE is forking a process. *PDOT* uses pthreads instead of forking to run multiple tasks concurrently in a TEE.

Limited TEE Memory. We use several techniques to address this challenge. First, we ensure no other enclaves (other than the quoting enclave) run on RecRes machine. This allows *PDOT* to use all available EPC memory. Second, we fix the number of QueryHandler threads in order to save space. This is possible because of dis-association of QueryHandler and ClientReader/Writer threads.

OCALL and ECALL Overhead. ECALLs and OCALLs introduce overhead and therefore should be avoided as much as possible. For example, all threads mentioned in the previous section must wait until they receive the following information: for ClientReader thread – DNS query from the client, for QueryProcessor thread – query from `inQueryList`, and for ClientWriter thread – response from `outQueryList`. *PDOT* was implemented so that these threads wait inside the enclave. If we were to wait outside the enclave, we would have to use an ECALL to enter the enclave each time the thread proceeds.

6 EVALUATION

6.1 Security Analysis

This section describes how query privacy (Requirement R1) is achieved, with respect to the two types of adversaries, per Section 3.1.

Malicious RecRes operator. Recall that a malicious RecRes operator controls the machine that runs *PDOT* RecRes. It cannot obtain the query from intercepted packets since they flow over the encrypted TLS channel. Also, because the local TLS endpoint resides inside the RecRes enclave, the malicious operator cannot retrieve the query from the enclave, as it does not have access to the protected memory region.

However, a malicious RecRes operator may attempt to connect the socket to a malicious TLS server that resides in either: 1) an untrusted region, or 2) a separate enclave that the operator itself created. If the operator can trick the client into establishing a TLS connection with the malicious TLS server, the adversary can obtain

the plaintext DNS queries. For case (1), the verification step at the client side fails because the TLS server certificate does not include any attestation information. For case (2), the malicious enclave might receive a legitimate attestation verification report, attestation verification report signature, and attestation report signing certificate from IAS. However, that report would contain a different MRENCLAVE value, which would be rejected by the client. To convince the client to establish a connection with *PDOT* RecRes, the adversary has no choice except to run the code of *PDOT* RecRes. Therefore, in both cases, the adversary cannot trick the client into establishing a TLS connection with a TLS server other than the one running a *PDOT* RecRes.

Network Adversary. Recall that this adversary captures all packets to/from *PDOT*. It cannot obtain the plaintext queries since they flow over the TLS tunnel. The only information it can obtain from packets includes cleartext header fields, such as source and destination IP addresses. This information, coupled with a timing attack, might let the adversary correlate a packet sent from the client with a packet sent to an NS. The consequent amount of privacy leakage is discussed in Section 7.1

6.2 Deployability

Section 5 shows how *PDOT* clients do not need special hardware, and require only minor software modifications (Requirement R2). To aid deployability, *PDOT* also provides several configurable parameters, including: the number of QueryHandle threads (to adjust throughput), the amount of memory dedicated to each thread (to serve clients that send a lot of queries at a given time), and the timeout of QueryHandle threads (to adjust the time for a QueryHandle thread to acquire a resource). Another consideration is incremental deployment, where some clients may request DNS-over-TLS without supporting *PDOT*. *PDOT* can handle this situation by having its TLS certificate **also** signed by a trusted root CA, since legacy clients will ignore *PDOT*-specific attestation information.

On the client side, an ideal deployment scenario would be for browser or OS vendors to update their client stubs to support *PDOT*. The same way that browser vendors currently include and maintain a list of trusted root CA certificates in their browsers, they could include and periodically update a list of trustworthy MRENCLAVE values for *PDOT* resolvers. This could all be done transparently to end users. As with root CA certificates, expert users can manually add/remove trusted MRENCLAVE values for their own systems. In practice, there are only a handful of recursive resolver software implementations. Thus, even allowing for multiple versions of each, the list of trusted MRENCLAVE values would be orders of magnitude smaller than the list of public keys of every trusted resolver, as would be required for standard DNS-over-TLS.

6.3 Performance Evaluation

We ran *PDOT* on a low-cost Intel NUC consisting of an Intel Pentium Silver J5005 CPU with 128 MB of EPC memory and 4 GB of RAM. We used Ubuntu 16.04 and the Intel SGX SDK version 2.2. We configured our RecRes to support up to 50 concurrent clients and process queries using 30 QueryHandle threads. For comparison, we performed the same benchmarks using Unbound [23], a popular open source RecRes.

Figure 4: Latency comparison of *PDoT* and Unbound

6.3.1 Latency Evaluation. The objective of our latency evaluation is to assess overhead introduced by running RecRes inside an enclave. To do so, we measure the time to resolve a DNS query using *PDoT* and compare with Unbound. To meet requirement R3, *PDoT* should not incur a significant increase in latency compared to Unbound.

Experimental Setup. The client and RecRes ran on the same physical machine to remove network delay. We conducted the experiment using *PDoT* and Unbound as the RecRes, and Stubby as the client. We measured latency under two different scenarios: cold start and warm start. In the former, the client sets up a new TLS connection every time it sends a query to the RecRes. In the warm start scenario, the client sets up one TLS connection with the RecRes at the beginning, and reuses it throughout the experiment. In other words, the cold start measurements also include the time required to establish the TLS connection. In this experiment, the caching mechanisms of both *PDoT* and Unbound were disabled.

We created a python program to feed DNS queries to the client. The program sends 100 queries sequentially for ten different domains. That is, the program waits for an answer to the previous query before sending the next query. We used the top ten domains of the Majestic Million domain list [26].

The python program measures the time between sending the query and receiving an answer. For the cold start experiment, we spawned a new Stubby client and established a new TLS connection for each query. In the warm start scenario, we first established the TLS connection by sending a query for another domain (not in the top ten), but did not include this in the timing measurement.

Note that the numeric latency values are specific to our experimental setup because they depend on network bandwidth of our RecRes, and latency between the latter and relevant NS-s. The important aspect of this experiment is the ratio between the latencies of *PDoT* and Unbound. Therefore it is not meaningful to compute average latency over a large set of domains. Instead, we took multiple measurements for each of a small set of domains (e.g., 100

measurements for each of 10 domains) so as to analyse the range of response latencies for each domain.

Results and observations. Results of latency measurements are shown in Figure 4. Red boxes show latency of *PDoT* and the blue boxes – of Unbound. In these plots, boxes span from the lower to upper quartile values of collected data. Whiskers span from the highest datum within the 1.5 interquartile range (IQR) of the upper quartile to the lowest datum within the 1.5 IQR of the lower quartile. Median values are shown as black horizontal lines inside the boxes.

For the cold-start case in Figure 4a, although Unbound is typically faster than our proof-of-concept *PDoT* implementation, the range of latencies is similar. For 7 out of 10 domains, the upper whisker of *PDoT* was lower than that of Unbound. Overall, *PDoT* shows an average 22% overhead compared to Unbound in this setting.

For the warm-start case in Figure 4b, the median latency is lower across the board compared to the cold-start setting because the TLS tunnel has already been established. In this setting, *PDoT* shows an average of 9% overhead compared to Unbound. In practice, once the client has established a connection to RecRes, it will maintain this connection; thus, the vast majority of queries will see only the warm-start latency.

6.3.2 Throughput evaluation. The objective of throughput evaluation is to measure the rate at which the RecRes can sustainably respond to queries. *PDoT*'s throughput should be close to that of Unbound to satisfy requirement R4.

Experiment setup. The client and RecRes were run on different machines, so that the RecRes could use all available resources of a single machine. This is representative of a local RecRes running in a small network (e.g., a coffee shop WiFi network). We conducted this experiment using the same two RecRes-s as in the latency experiment. Stubby was configured to reuse TLS connections. To simulate a small to medium-scale network, we varied the number of

Figure 5: Throughput comparison of *PDoT* (red) and Unbound (blue)

concurrent clients between 1 and 25 and adjusted the query arrival rate from 5 to 100 queries per second. Query rates were uniformly distributed among the clients, e.g., for an overall rate of 100 queries per second with 10 clients, each client sends 10 queries per second. To eliminate any variability in resolving the query, all queries were for the domain `google.com`. We maintained constant query rate for one minute. Caching mechanisms of both *PDoT* and Unbound were disabled.

Results and observations. Results of throughput experiments are shown in Figure 5. Each graph corresponds to a different number of clients. Horizontal axis shows different query rates and vertical axis shows the range of response latencies for each query rate. Measurement are plotted using the same box plot arrangement as in latency evaluation. Blue boxes show results for Unbound and red boxes – for *PDoT*.

If queries arrive faster than RecRes can process them, they start clogging the queue, and the latency of each successive response

increases as the queue grows. In this case, the average latency continues to increase indefinitely until queries begin to timeout or RecRes runs out of memory. We view this rate of query arrival as *unsustainable*. On the other hand, if RecRes can sustain the rate of query arrival, average response latency would remain roughly constant irrespective of how long RecRes runs. For this experiment, we define a *sustainable* rate of query arrival as the one for which the average response latency is constant over time, and below one second – well below a typical DNS client timeout. Figure 5 only shows cases where query arrival rate is sustainable for the respective RecRes. In other words, the presence of a box in Figure 5 shows that the RecRes can achieve that level of throughput.

Surprisingly, we observed that Unbound cannot handle query rates exceeding 10 queries per second per client, i.e., its maximum sustainable rate was $10n$ queries per second distributed among n clients. This is because Unbound’s design only uses one query processing thread per client. In contrast, *PDoT* handled more than

Figure 6: Latency comparison of *PDoT* (red) and Unbound (blue) with different number of domains in cache

100 queries per second in all cases because its design uses a separate pool of QueryHandle threads.

Overall, Figure 5 confirms that our proof-of-concept implementation achieves at least the same throughput as Unbound across the range of clients and query arrival rates, and can achieve higher throughput when the number of clients is low. Although Unbound again achieves slightly lower latency, this is consistent with our latency measurements in Section 6.3.1 and is likely due to the fact that Unbound is an optimized production-grade RecRes.

6.3.3 Caching evaluation. We evaluated performance of both resolvers with caching enabled; Unbound with its default caching behavior, and *PDoT* with our simple proof-of-concept cache.

Experiment setup. The experimental setup is similar to that of the latency evaluation described earlier. We pre-populated resolvers’ caches with varying numbers of domains and measured response latency for a representative set of 10 popular domains.

Results and observations. As shown in Figure 6, Unbound serves responses from cache with a consistent latency irrespective of the number of entries in the cache. Although *PDoT* achieves lower average latencies when the cache is relatively empty, it has higher variability than Unbound. This is probably due to the combination of our unoptimized caching implementation and latency of accessing enclave memory. Nevertheless, Figure 6 shows that – even with the memory limitations of current hardware enclaves – *PDoT* can still benefit from caching a small number of domains.

7 DISCUSSION

7.1 Information Revealed by IP Addresses

Even if the connections between the client, RecRes, and NS-s are encrypted using TLS, some information is still leaked. The most prominent and obvious is source/destination IP addresses. The network adversary described in Section 3.1 can combine these cleartext IP addresses with packet timing information in order to correlate packets sent from client to RecRes with subsequent packets sent from RecRes to NS.

Armed with this information, the adversary can narrow down the client’s domain name query to one of the records that could be served by that specific ANS. Assuming the ANS can serve R

domain names, the adversary has a $1/R$ probability of guessing which domain name the user queried. When $R > 1$, we call this a *privacy-preserving ANS*. This prompts two questions: 1) what percentage of domains can be answered by a privacy-preserving ANS; and 2) what is the typical size of anonymity set (R) provided by a privacy-preserving ANS?

To answer these questions, we designed a scheme to collect records stored in various ANS-s. We sent DNS queries for 1,000,000 domains from the Majestic Million domain list [26], and gathered information about ANS-s that can possibly provide the answer for each. By collecting data on possible ANS-s, we can map domain names to each ANS, and thus estimate the number of records held by each ANS. Following the *Guidelines for Internet Measurement Activities* [8], we limited our querying rate, in order to avoid placing undue load on any servers.

As shown in Figure 7, only 5.7% of domains we queried were served by non-privacy-preserving ANS-s, i.e., those that hold only one record). Examples of domain names served from non-privacy-preserving ANS-s included: `tinyurl.com`⁴, `bing.com`, `nginx.org`, `news.bbc.co.uk`, and `cloudflare.com`. On the other hand, 9 out of 10 queries were served by a privacy-preserving ANS, and 65.7% by ANS-s that hold over 100 records.

These results are still approximations. Since we do not have data for domains outside the Majestic Million list, we cannot make claims about whether these would be served by a privacy-preserving ANS. We hypothesize that the vast majority of ANS-s would be privacy-preserving for the simple reason that it is more economical to amortize the ANS’s running costs over multiple domains. On the other hand, we can be certain that our results for the Majestic Million are a strict lower bound on the level of privacy because the ANS-s from which these are served could also be serving other domains outside of our list. It would be possible to arrive at a more accurate estimate by analyzing zone files of all (or at least most) ANS-s. However, virtually all ANS-s disable the interface to download zone files because this could be used to mount DoS attacks. Therefore, this type of analysis would have to be performed by an organization with privileged access to all ANS-s’ zone files.

⁴Since `tinyurl.com` is a URL shortening service, this is actually still privacy-preserving because the adversary can not learn which short URL was queried.

Figure 7: Percentage of Majestic Million domains answered by an ANS with at least N records

7.2 Caching & Timing attacks

Introducing a cache into an RecRes would allow the adversary to launch timing attacks and help guess the domain name queried by the end-user. We consider two types of timing attacks:

- **Measuring time between query and response.** This is the simplest attack, whereby the adversary monitors the network between client and RecRes, and records the time for the RecRes to respond to the client. If the response time is shorter (compared to other queries), it likely has been served from a cache. This attack can be launched by both adversary types described in Section 3.1. One obvious countermeasure is to artificially delay the response to match the latency of NS-served responses.
- **Correlating client and RecRes requests.** To counter the above countermeasure the adversary may attempt to correlate DNS requests sent from client to RecRes with those sent from RecRes to NS-s e.g., using the times at which the packets were sent. If successful, the adversary can distinguish requests that involve contacting an NS from those that were served from the cache. This attack can be also launched by a malicious RecRes or a network sniffer. One way to counter this is to always send a query to one NS (although not necessarily the correct NS). This diminishes the benefits of caching, but still reduces the number of NS queries since the real answer may have required more than one NS query. Another way is to batch and randomize the order of requests to the NS, creating a type of DNS mix network.

For both of these types of attacks, the information leakage depends on whether the adversary is passive or active. A passive adversary can (at most) guess the domain name. If the caching strategy is Most Recently Used (MRU), the domain name must be one of the popular ones. The active adversary can generate its own DNS queries for a wide range of domain names and keep a list of those that result in cache hits, thus improving the chances of inferring the user’s query target.

8 RELATED WORK

There has been much prior work aiming to protect the privacy of DNS queries [7, 15, 16, 25, 33, 37, 38]. For example, Lu et al. [25] proposed a privacy-preserving DNS that uses distributed hash tables, different naming schemes, and methods from computational private information retrieval. Federrath et al. [16] introduced a dedicated DNS Anonymity Service to protect the DNS queries using an architecture that distributes the top domains by broadcast and uses low-latency mixes for requesting the remaining domains. These schemes all assume that all parties involved do not act maliciously.

There have also been some activities in the Internet standards community that focused on DNS security and privacy. DNS Security Extensions (DNSSEC) [3] provides data origin authentication and integrity via public key cryptography. However, it does not offer privacy. Bortzmeyer [5] proposed a scheme. Also, though not Internet standards, several protocols have been proposed to encrypt and authenticates DNS packets between the client and the RecRes (DNSEncrypt [30]) and RecRes and NS-s (DNSCurve [1]). Moreover, the original DNS-over-TLS paper has been converted into a draft Internet standard [19]. All these methods assume that the RecRes operator is trusted and does not attempt to learn anything from the DNS queries.

Furthermore, there has been some research on establishing trust through TEEs to protect confidentiality and integrity of network functions. Specifically, SGX has been used to protect network functions, especially middle-boxes. For example, Endbox [17] aims to distribute middle-boxes to client edges: clients connect through VPN to ensure confidentiality of their traffic while remaining maintainable. LightBox [14] is another middle-box that runs in an enclave; its goal is to protect the client’s traffic from the third-party middle-box service provider while maintaining adequate performance. Finally, ShieldBox [35] aims to protect confidential network traffic that flows through untrusted commodity servers and provides a generic interface for easy deployability. These efforts focus on protecting confidential data that flows in the network, and do not target DNS queries.

9 CONCLUSION & FUTURE WORK

This paper proposed *PDoT*, a novel DNS RecRes design that operates within a TEE to protect privacy of DNS queries, even from a malicious RecRes operator. In terms of query throughput, our unoptimized proof-of-concept implementation matches the throughput of Unbound, a state-of-the-art DNS-over-TLS recursive resolver, while incurring an acceptable increase in latency (due to the use of a TEE). In order to quantify the potential for privacy leakage through traffic analysis, we performed an Internet measurement study which showed that 94.7% of the top 1,000,000 domain names can be served from a *privacy-preserving* ANS that serves at least two distinct domain names, and 65.7% from an ANS that serves 100+ domain names. As future work, we plan to port the Unbound RecRes to Intel SGX and conduct a performance comparison with *PDoT*, as well as to explore methods for improving *PDoT*’s performance using caching while maintaining client privacy.

ACKNOWLEDGMENTS

We thank Geonhee Cho for the initial data collection for the privacy-preserving ANS analysis in Section 7.1. We are also grateful to the paper's shepherd, Roberto Perdisci, and ACSAC'19 anonymous reviewers for their valuable comments. First and third authors were supported in part by NSF Award Number:1840197, titled: "CICI: SSC: Horizon: Secure Large-Scale Scientific Cloud Computing". The first author was also supported by The Nakajima Foundation. The second author was supported by a US-UK Fulbright Cyber Security Scholar Award.

REFERENCES

- [1] . 2009. Introduction to DNSCurve. <https://dnscurve.org/index.html> [Online] Accessed: 2019-05-29.
- [2] Fritz Alder, N Asokan, Arseny Kurnikov, Andrew Paverd, and Michael Steiner. 2019. S-FaaS: Trustworthy and Accountable Function-as-a-Service using Intel SGX. In *ACM Cloud Computing Security Workshop (CCSW '19)*.
- [3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. 2005. *DNS Security Introduction and Requirements*. Technical Report. <https://doi.org/10.17487/rfc4033>
- [4] ARM. 2009. ARM Security Technology - Building a Secure System using TrustZone Technology. <http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.prd29-genc-009492c/index.html> [Online] Accessed: 2019-05-29.
- [5] S. Bortzmeyer. 2016. *DNS Query Name Minimisation to Improve Privacy*. Technical Report. <https://doi.org/10.17487/RFC7816>
- [6] S Bortzmeyer. 2018. Encryption and authentication of the DNS resolver-to-authoritative communication. <https://tools.ietf.org/html/draft-bortzmeyer-dprive-resolver-to-auth-01>
- [7] Sergio Castillo-Perez and Joaquin Garcia-Alfaro. 2008. Anonymous Resolution of DNS Queries. Springer, Berlin, Heidelberg, 987–1000. https://doi.org/10.1007/978-3-540-88873-4_5
- [8] V.G. Cerf. 1991. *Guidelines for Internet Measurement Activities*. Technical Report. <https://doi.org/10.17487/rfc1262>
- [9] Cloudflare. -. DNS over TLS - Cloudflare Resolver. <https://1.1.1.1/dns/> [Online] Accessed: 2019-05-29.
- [10] Manuel Costa, Lawrence Esswood, Olga Ohrimenko, Felix Schuster, and Sameer Wagh. 2017. The pyramid scheme: Oblivious RAM for trusted processors. *arXiv preprint arXiv:1712.07882* (2017).
- [11] Victor Costan, Ilia Lebedev, and Srinivas Devadas. 2016. Sanctum: Minimal Hardware Extensions for Strong Software Isolation. , 857–874 pages. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/costan>
- [12] cs.nic. -. Knot Resolver. <https://www.knot-resolver.cz/> [Online] Accessed: 2019-05-29.
- [13] T. Dierks and E. Rescorla. 2008. *The Transport Layer Security (TLS) Protocol Version 1.2*. Technical Report. <https://doi.org/10.17487/rfc5246>
- [14] Huayi Duan, Cong Wang, Xingliang Yuan, Yajin Zhou, Qian Wang, and Kui Ren. 2017. LightBox: Full-stack Protected Stateful Middlebox at Lightning Speed. (Jun 2017). *arXiv:1706.06261* <http://arxiv.org/abs/1706.06261>
- [15] Annie Edmundson, Paul Schmitt, and Nick Feamster. 2018. ODNs: Oblivious DNS. <https://odns.cs.princeton.edu/> [Online] Accessed: 2019-05-29.
- [16] Hannes Federrath, Karl-Peter Fuchs, Dominik Herrmann, and Christopher Piosenky. 2011. Privacy-Preserving DNS: Analysis of Broadcast, Range Queries and Mix-Based Protection Methods. Springer, Berlin, Heidelberg, 665–683. https://doi.org/10.1007/978-3-642-23822-2_36
- [17] David Goltzsche, Signe Rusch, Manuel Nieke, Sebastien Vaucher, Nico Weichbrodt, Valerio Schiavoni, Pierre-Louis Aublin, Paolo Cosa, Christof Fetzer, Pascal Felber, Peter Pietzuch, and Rudiger Kapitza. 2018. EndBox: Scalable Middlebox Functions Using Client-Side Trusted Execution. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN '18)*. IEEE, 386–397. <https://doi.org/10.1109/DSN.2018.00048>
- [18] Google. 2018. DNS over TLS support in Android P Developer Preview. <https://security.googleblog.com/2018/04/dns-over-tls-support-in-android-p.html> [Online] Accessed: 2019-05-29.
- [19] Zi Hu, Liang Zhu, John Heidemann, Allison Mankin, Duane Wessels, and P Hoffman. 2016. Specification for DNS over Transport Layer Security (TLS). <https://doi.org/10.17487/RFC7858>
- [20] Thomas Knauth, Michael Steiner, Somnath Chakrabarti, Li Lei, Cedric Xing, and Mona Vij. 2018. Integrating Remote Attestation with Transport Layer Security. (Jan 2018). *arXiv:1801.05863* <http://arxiv.org/abs/1801.05863>
- [21] SPROUT Lab. 2019. PDOT Source Code. <https://github.com/sprout-uci/PDOT>
- [22] NLnet Labs. -. Stubby. <https://dnspriacy.org/wiki/display/DP/DNS+Privacy+Daemon+-+Stubby> [Online] Accessed: 2019-05-29.
- [23] NLnet Labs. -. Unbound. <https://nlnetlabs.nl/projects/unbound/about/> [Online] Accessed: 2019-05-29.
- [24] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B. Lee. 2015. Last-Level Cache Side-Channel Attacks are Practical. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 605–622. <https://doi.org/10.1109/SP.2015.43>
- [25] Y. Lu and G. Tsudik. 2010. Towards Plugging Privacy Leaks in the Domain Name System. In *2010 IEEE Tenth International Conference on Peer-to-Peer Computing (P2P)*. IEEE, 1–10. <https://doi.org/10.1109/P2P.2010.5569976>
- [26] Majestic. 2012. Majestic Million. <https://blog.majestic.com/development/majestic-million-csv-daily/>
- [27] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V. Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R. Savagaonkar. 2013. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy (HASP '13)*. ACM Press, New York, New York, USA, 1–1. <https://doi.org/10.1145/2487726.2488368>
- [28] Microsoft. 2017. Introducing Azure confidential computing. <https://azure.microsoft.com/en-us/blog/introducing-azure-confidential-computing/> [Online] Accessed: 2019-05-29.
- [29] P.V. Mockapetris. 1987. *Domain names - implementation and specification*. Technical Report. <https://doi.org/10.17487/rfc1035>
- [30] DNSCrypt Project. -. DNSCrypt. <https://dnscrypt.info/> [Online] Accessed: 2019-05-29.
- [31] Sajin Sasy, Sergey Gorbunov, and Christopher W Fletcher. 2017. ZeroTrace: Oblivious Memory Primitives from Intel SGX. *IACR Cryptology ePrint Archive* 2017 (2017), 549.
- [32] Ming-Wei Shih, Sangho Lee, Taesoo Kim, and Marcus Peinado. 2017. T-SGX: Eradicating Controlled-Channel Attacks Against Enclave Programs. In *NDSS Symposium*. <https://www.ndss-symposium.org/ndss2017/ndss-2017-programme/t-sgx-eradicating-controlled-channel-attacks-against-enclave-programs/>
- [33] Haya Shulman and Haya. 2014. Pretty Bad Privacy: Pitfalls of DNS Encryption. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society (WPES '14)*. ACM Press, New York, New York, USA, 191–200. <https://doi.org/10.1145/2665943.2665959>
- [34] Sandeep Tamrakar, Jian Liu, Andrew Paverd, Jan-Erik Ekberg, Benny Pinkas, and N. Asokan. 2017. The Circle Game: Scalable Private Membership Test Using Trusted Hardware. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security (ASIA CCS '17)*. <https://doi.org/10.1145/3052973.3053006>
- [35] Bohdan Trach, Alfred Krohmer, Franz Gregor, Sergei Arnautov, Pramod Bhatotia, and Christof Fetzer. 2018. ShieldBox: Secure Middleboxes using Shielded Execution. In *Proceedings of the Symposium on SDN Research (SOSR '18)*. ACM Press, New York, New York, USA, 1–14. <https://doi.org/10.1145/3185467.3185469>
- [36] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. 2015. Controlled-Channel Attacks: Deterministic Side Channels for Untrusted Operating Systems. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 640–656. <https://doi.org/10.1109/SP.2015.45>
- [37] Fangming Zhao, Yoshiaki Hori, and Kouichi Sakurai. 2007. Analysis of Privacy Disclosure in DNS Query. In *2007 International Conference on Multimedia and Ubiquitous Engineering (MUE '07)*. IEEE, 952–957. <https://doi.org/10.1109/MUE.2007.84>
- [38] Fangming Zhao, Yoshiaki Hori, and Kouichi Sakurai. 2007. Two-Servers PIR Based DNS Query Scheme with Privacy-Preserving. In *The 2007 International Conference on Intelligent Pervasive Computing (IPC '07)*. IEEE, 299–302. <https://doi.org/10.1109/IPC.2007.27>
- [39] Liang Zhu, Zi Hu, John Heidemann, Duane Wessels, Allison Mankin, and Nikita Somaiya. 2015. Connection-Oriented DNS to Improve Privacy and Security. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 171–186. <https://doi.org/10.1109/SP.2015.18>