

LightBox: Full-stack Protected Stateful Middlebox at Lightning Speed

Huayi Duan¹, Cong Wang¹, Xingliang Yuan², Yajin Zhou^{3,4}, Qian Wang⁵, and Kui Ren^{3,4}

¹City University of Hong Kong and City University of Hong Kong Shenzhen Research Institute; ²Monash University;

³College of Computer Science and Technology, School of Cyber Science and Technology, Zhejiang University;

⁴Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies;

⁵School of Cyber Science and Engineering, Wuhan University;

hy.duan@my.cityu.edu.hk, congwang@cityu.edu.hk, xingliang.yuan@monash.edu,

yajin_zhou@zju.edu.cn, qianwang@whu.edu.cn, kuiren@zju.edu.cn

ABSTRACT

Running off-site software middleboxes at third-party service providers has been a popular practice. However, routing large volumes of raw traffic, which may carry sensitive information, to a remote site for processing raises severe security concerns. Prior solutions often abstract away important factors pertinent to real-world deployment. In particular, they overlook the significance of metadata protection and stateful processing. Unprotected traffic metadata like low-level headers, size and count, can be exploited to learn supposedly encrypted application contents. Meanwhile, tracking the states of 100,000s of flows concurrently is often indispensable in production-level middleboxes deployed at real networks.

We present LightBox, the first system that can drive off-site middleboxes at near-native speed with stateful processing and the most comprehensive protection to date. Built upon commodity trusted hardware, Intel SGX, LightBox is the product of our systematic investigation of how to overcome the inherent limitations of secure enclaves using domain knowledge and customization. First, we introduce an elegant virtual network interface that allows convenient access to fully protected packets at line rate without leaving the enclave, as if from the trusted source network. Second, we provide complete flow state management for efficient stateful processing, by tailoring a set of data structures and algorithms optimized for the highly constrained enclave space. Extensive evaluations demonstrate that LightBox, with all security benefits, can achieve 10Gbps packet I/O, and that with case studies on three stateful middleboxes, it can operate at near-native speed.

CCS CONCEPTS

• **Networks** → **Middle boxes / network appliances; Network privacy and anonymity**; • **Security and privacy** → **Domain-specific security and privacy architectures**.

KEYWORDS

Intel SGX; stateful middleboxes; secure packet processing

ACM Reference Format:

Huayi Duan, Cong Wang, Xingliang Yuan, Yajin Zhou, Qian Wang, and Kui Ren. 2019. LightBox: Full-stack Protected Stateful Middlebox at Lightning Speed. In *2019 ACM SIGSAC Conference on Computer & Communications Security (CCS'19), November 11–15, 2019, London, UK*. ACM, NY, NY, USA, 17 pages. <https://doi.org/10.1145/3319535.3339814>

1 INTRODUCTION

Middleboxes underpin the infrastructure of modern networks, where they undertake critical network functions for performance, connectivity, and security [90]. Recently, a paradigm shift of migrating software middleboxes (aka virtual network functions) to professional service providers, e.g., public cloud, is taking place for the promising security, scalability and management benefits [8, 30, 76]. Its potential to enable a billion-dollar marketplace has already been widely recognized [21].

According to Zscaler [39], petabytes of traffic are now routed to its cloud-based security platform for middlebox processing every single day, and the number is still growing. Along with this seemingly unstoppable momentum comes an unprecedented security concern: how can end users be assured that their private information buried in the huge volumes of traffic, is not unauthorizedly leaked while being processed by the service provider (Fig. 1)? We are witnessing increasing and diversifying data breaches by service providers nowadays [53], yet embarrassingly we are also facing a daunting situation where full-scale traffic inspection seems mandatory to thwart stealthy threats [40]. In light of this, a reassuring solution should be capable of protecting sensitive traffic while retaining necessary middlebox functionality.

Over the past few years, a number of approaches have been proposed to address the problem above, and they can be categorized into two broadly defined classes: *software-centric* and *hardware-assisted*. The first line of solutions [3, 10, 26, 51, 77, 99] often rely on tailored cryptographic schemes. They are advantageous in providing provable security without hardware assumption, but often limited in functionality and sometimes in performance. The second line of solutions move middleboxes into a trusted execution environment, mostly Intel SGX enclave [56]. Hardware-assisted designs provide generally better functionality and performance than software-centric approaches. In this regard, efforts with particular focus on the modular design and programmability of secure

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '19, November 11–15, 2019, London, United Kingdom

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6747-9/19/11...\$15.00

<https://doi.org/10.1145/3319535.3339814>




Figure 1: Large volumes of traffic are redirected to service providers for middlebox processing, creating a unique vantage point for adversaries to learn sensitive information.

middleboxes [18, 35, 66, 86], deployment consideration [61], and code protection [66] have been actively made.

As with these designs, we prioritize the consideration of handling the intrinsic complexity and stringent performance requirements of middleboxes, and leverage the SGX enclave as a starting point to develop a secure middlebox system. We observe that while previous solutions have claimed the benefits and practicality of SGX-enabled design, they largely overlooked several key factors that are highly pertinent to the real-world deployment of off-site middleboxes. We have identified two important aspects and provided encouraging results, which we hope will lead the secure middlebox-as-a-service to the practical realm and stimulate its massive adoption.

1.1 Motivation

Necessity for metadata protection. Most existing designs only consider the protection of application payloads while redirecting traffic to remote middleboxes. Few of them protect traffic metadata, including *low-level packet headers, packet size, count and timestamps*. Such metadata is information-rich and highly exploitable. The importance of hiding communication metadata (e.g., IP addresses) has been reiterated in recent years [88, 89]. In the often cited example of Snowden Leaks, it is frustratedly put as “if you have enough metadata, you don’t really need content” [70].

By just exploiting the seemingly innocent *packet size, count and timing*, a variety of sophisticated traffic analysis attacks have been demonstrated: they can extract supposedly encrypted application contents such as website objects [91], VoIP conversations [96], streaming videos [68], instant messages [19], and user activities [16], by analyzing the distributions of metadata.

Such metadata may be obtained by an adversary who can sniff traffic anywhere on the transmission path. In fact, aggregating tremendous user traffic to the middlebox service provider creates a unique vantage point for traffic analysis, because of the much enlarged datasets for correlating information and statistical inference. We should thus protect not only application payloads but all aforementioned traffic metadata — what we dub *full-stack protection*.

Necessity for stateful middlebox. In contrast to L2 switches and L3 routers that process each packet independently, advanced middleboxes need to track various flow-level states to implement complex

Table 1: Functionality and security characterization of representative solutions for secure middleboxes.

	Function				Protection			
	Field	Op.	Stateful	Meta	HDR	P/L	Rule	State
<i>Software-centric</i>								
BlindBox [77]	P	PM	✗	✗	✗	✓	✓	N/A
YWLW16 [99]	P	PM	✗	✗	✗	✓	✓	N/A
SplitBox [3]	H+P	RM	✗	✗	✓	✓	✓	N/A
BlindIDS [10]	P	PM	✗	✗	✗	✓	✓	N/A
SPABox [26]	P	REX	✗	✗	✗	✓	✓	N/A
Embark [51]	H+P	RM	✓*	✗	✓	✓	✓	✓*
<i>Hardware-assisted</i>								
S-NFV [78]	N/A	N/A	N/A	✗	✗	✗	✗	✓
TrustedClick [18]	H+P	GN	✗	✗	✗	✓	✓	N/A
SGX-BOX [35]	P	GN	✓	✗	✗	✓	✓	✗
mbTLS [61]	P	GN	N/S	✗	✗	✓	✓	N/S
ShieldBox [86]	H+P	GN	✗	✗	✗	✓	✓	N/A
SafeBrick [66]	H+P	GN	✓*	✗	✓	✓	✓	✓*
LightBox	H+P	GN	✓	✓	✓	✓	✓	✓

Notations

- Field: which fields are processed, H (L2-L4 headers) and/or P (L4 payload).
- Op. (operation): PM (exact string pattern matching) \subset RM (range matching) \subset REX (regular expression matching) \subset GN (generic functions).
- Stateful: whether generic *flow-level* stateful processing is supported.
- Meta: packet size, count and timestamp.
- HDR: L2-L4 headers, e.g., ip address, port number, and TCP/IP flags.
- P/L: L4 payload, including all application content.
- Rule: middlebox processing rules, e.g., attacking signatures.
- State: *flow-level* states, e.g., connection status, statistics and stream buffers.
- N/A: the feature is not considered by design.
- N/S: the feature may be potentially supported, but not explicitly described.
- *: Embarks considers an ad hoc web proxy, but not generic stateful processing.
- *: SafeBrick considers a simple stateful firewall only.

functionality [45]. For example, intrusion detection systems typically keep per-flow stream buffers to eradicate cross-packet attack patterns [44, 81]; proxies and load balancers maintain front/backend connection states and packet pools to ensure end-to-end connectivity [65, 67]. Thus, supporting stateful processing is an essential functionality requirement in realistic middlebox products.

However, even with the power of trusted hardware, it is technically challenging to develop a secure yet efficient solution due to the unique features of stateful middleboxes. In particular, the per-flow states range from a few hundreds of bytes [47] to multiple kilobytes [45], and they need to stay tracked throughout the lifetime of flows or some expiration period. Moreover, it is not uncommon for production-level middleboxes to handle hundreds of thousands (or even more) of flows concurrently in real networks [20, 24, 44, 45]. The resulting gigabytes of runtime memory footprint are not easily manageable by any secure enclaves. Meanwhile, modern middleboxes feature packet processing delay within a few tens of microseconds [44, 54, 64] — a performance baseline that should be respected even if strong security guarantees are favored.

Characterization of prior arts. We summarize existing solutions and compile their functionality and security features in Table 1. We define “metadata” as the L2-L4 headers, packet size, count and timestamps. For better characterization, we separately label the headers as “HDR”, and the latter three as “Meta”. Some solutions [61, 77, 78] do not target the outsourcing scenario and have different security goals, but we include them anyway for completeness.

Regarding security, none of the existing systems considers full-stack protection. The closest to us are Embark [51], which applies deterministic encryption to each packet, and SafeBricks [66], which uses L2 secure tunneling to forward packets to the enclave. Both of them encrypt packets individually and thus may be vulnerable to traffic analysis attacks exploiting packet size and count.

Regarding stateful processing, SGX-BOX [35] allows inspection on traffic streams that are reassembled outside enclave, and SafeBricks implements a simple stateful firewall for testing. But none of them considers the high flow concurrency exhibited in real networks and the challenges therein. Embark supports an ad hoc web proxy that caches static HTTP contents, but it cannot support generic stateful functions with arbitrary operations over flow state. We defer more discussions on related work to Section 7.

Note that we do recognize the meaningful explorations and contributions made by these prior arts with different focuses. We believe that our efforts will make it more convincing and confident to deploy and operate secure middleboxes in practical settings.

1.2 Our Contribution

We design and build LightBox, the first SGX-enabled secure middlebox system that can drive off-site middleboxes at near-native speed with stateful processing and full-stack protection. By systematically tackling many well-known limitations of SGX, from the lack of system services including network I/O, trusted timing and synchronization, to the unacceptable overhead of secure memory oversubscription, we have provided affirmative and satisfactory answers to the following two major research questions.

1. How to securely forward raw packets to a remote enclave, without leaking their low-level metadata, while still making them conveniently accessible at line rate?

In pursuit of this goal and in reminiscent of the classic kernel-driven tun/tap tunnel devices, we have developed a virtual network interface (VIF) called `etap` (“enclave tap”). It allows access to packets without leaving the enclave, as if from the source network where the packets originate. For full packet protection, our design guarantees that the raw packets with L2-L4 headers are entirely delivered via a secure tunnel (e.g., TLS) terminated at a trusted enterprise gateway and the enclave. To frustrate traffic analysis, we pack the raw packets in a back-to-back manner and transmit them as continuous application payloads. As a result, the packet boundaries are obscured in the encrypted stream, so are the packet size and count. As a by-product of `etap`, we create a trusted clock for high-resolution and reliable timing inside the enclave.

We show how to progressively optimize the performance of `etap` with lock-free rings, cache line protection, and disciplined batching, so that it can catch up with the rate of physical network interfaces.

We also endeavor to improve the usability of `etap`, by further porting three networking frameworks on top of it: 1) an adaption layer of `libpcap` [33], 2) a lightweight TCP stream reassembly library [28], and 3) an advanced flow monitoring stack `mOS` [45]. These system efforts allow us to port or develop middleboxes that enjoy the security and performance benefits of `etap`, with little code modification. For example, a developer can write a middlebox in the `mOS` framework as usual, and then automatically turn it into a LightBox instance even without the knowledge of SGX.

2. How to enable the resource-demanding stateful middlebox processing within the highly constrained enclave space, without incurring unreasonably high overhead?

As mentioned before, stateful middleboxes in realistic settings have a large memory demand, which is at odds with the limited enclave memory supply. The naive paging approach for oversubscribing enclave memory incurs substantial overhead [2], which is intolerable to middleboxes as confirmed by our experiments. To reduce the enclave footprint, we propose to maintain only a small working set of states in the enclave, while keeping the vast remainder of them encrypted in untrusted memory, at the granularity of flow. While this general idea may seem natural at first sight, we show that the performance bottleneck can only be surmounted with carefully crafted data structures and algorithms.

Specifically, given the limited enclave resources, the data structures used to hold flow states must be very compact, and support efficient lookup, relocation, swapping and deletion of data items. To meet these requirements, we design a set of interlinked data structures. They separate the indexing and storage of flow states, enabling flexible lookup strategy. They also allow fast relocation of states inside enclave and swapping of states across enclave boundary, with very cheap pointer operations. Furthermore, after identifying the lookup procedure as the main roadblock on the critical path, we opt to the space-efficient cuckoo hashing for indexing flow states, and introduce a cache-friendly lookup algorithm to counteract the cache inefficiency of hashing-based scheme.

Efficiency aside, our design ensures the confidentiality, integrity, and freshness of the states throughout the management procedures.

Experiment. Extensive evaluations show that with our optimized designs, `etap` allows in-enclave packet I/O at 10Gbps rate with full-stack protection. We instantiate LightBox for three stateful middleboxes, which are arguably more complicated than any of those tested by prior arts, and evaluate each of them against the native version and a strawman variant solely relying on EPC paging. LightBox incurs negligible packet delay inflation to the native processing for the two middleboxes with 0.5KB and 5.5KB per-flow state, and moderate delay to the most complicated middlebox with 11.4KB per-flow state. It maintains constant performance and achieves multi-factor speedup over the strawman when tracking 100,000s of flows. The performance gap is widened as more flows are tracked and more severe paging penalty is imposed on the latter. Last but not the least, LightBox can achieve native speed on a real CAIDA trace for two of the three middleboxes under testing, and 2× speedup over the strawman for the other unoptimized one.

Lessons learned. From our experience in designing, building and evaluating LightBox, we have indeed challenged the common perception that one can run applications in secure enclaves, SGX in particular, at native speed almost without technical efforts but mundane code porting. Our results unveil that for realistic workloads that are security-critical and performance-sensitive, domain-specific design and optimization become a must to bypass the intricacies of secure enclaves. We hope to raise practitioners’ awareness that, as a particular lesson, the proper use of memory-efficient and cache-friendly data structures and algorithms will make a big difference to the performance of those enclave applications.

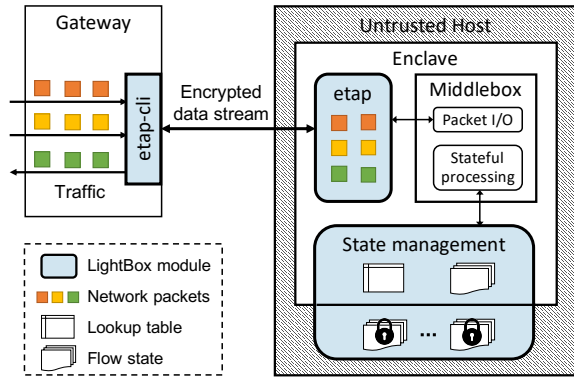


Figure 2: Overview of LightBox components.

2 OVERVIEW

2.1 Service Model

In a realistic service model, the enterprise redirects its traffic to the off-site middlebox hosted by the service provider for processing [41]. We assume that the middlebox code is not necessarily private and may be known to the service provider. This matches practical use cases where the source code is free to use, but only bespoke rule sets [82] are proprietary. We also consider a single middlebox. These simplifications allow us to concentrate on presenting the core designs of LightBox. Nonetheless, we stress that LightBox can be readily adapted to support service function chaining [83] and disjoint service providers [66], which mostly involves only changes to the service launching phase. We postpone the discussions of other service models to Appendix B.

Traffic forwarding. For ease of exposition, we consider the bounce model with one gateway [51, 76]: both inbound and outbound traffic is redirected from an enterprise gateway to the remote middlebox for processing and then bounced back. The other direct model, where traffic is routed from the source network to the remote middlebox and then directly to the next trusted hop, i.e., the gateway in the destination network [66], can be trivially supported by installing one `etap-cli` (see Section 2.3) on each gateway.

The communication endpoints themselves (e.g., a client in the enterprise network and an external server) may transmit data via a secure connection. To enable such already encrypted traffic to be processed by the middlebox, the gateway needs to intercept the secure connection and decrypt the traffic before redirection. We follow the common practice [35, 61, 66] to handle this issue. In particular, the gateway will receive the session keys from the endpoints to perform the interception, unbeknownst to the middlebox.

A dedicated high-speed connection will be typically established for traffic redirection [76]. Such services have been widely provisioned nowadays, for example AWS Direct Connect [5], Azure ExpressRoute [6], and Google Dedicated Interconnect [31]. The off-site middlebox, while being secured, should also be able to process packet at line rate to benefit from such dedicated links.

2.2 SGX Background

SGX introduces a trusted execution environment called enclave to shield code and data with on-chip security engines. It stands out

for the capability to run generic code at processor speed, with practically strong protection. Despite the benefits, it has several limitations. First, common system services cannot be directly used inside the enclave. Access to them requires expensive context switching to exit the enclave, typically via a secure API called OCALL. Second, memory access in the enclave incurs performance overhead. The protected memory region used by the enclave is called Enclave Page Cache (EPC). It has a conservative limit of 128MB in current product lines. Excessive memory usage in the enclave will trigger EPC paging, which can induce prohibitive performance penalties [2]. Besides, the cost of cache miss while accessing EPC is higher than normal, due to the cryptographic operations involved during data transferring between CPU cache and EPC. While such overhead may be negligible to certain applications, it becomes crucial to middleboxes with stringent performance requirements.

2.3 LightBox Overview

LightBox components. LightBox leverages an SGX enclave to shield the off-site middlebox. As shown in Fig. 2, a LightBox instance comprises two modules in addition to the middlebox itself: a virtual network interface `etap` and a state management module. The former is semantically equivalent to a physical network interface card (NIC), allowing packets I/O at line rate within the enclave. The latter provides automatic and efficient memory management of the large amount of flow states tracked by the middlebox.

The `etap` device is peered with one `etap-cli` program installed at the enterprise gateway. We establish a persistent secure channel between the two to tunnel the raw traffic, which is transparently encoded/decoded and encrypted/decrypted by `etap`. The middlebox and upper networking layers can directly access raw packets via `etap` without leaving the enclave.

The state management module maintains a small flow cache in the enclave, a large encrypted flow store in the untrusted memory, and an efficient lookup data structure in the enclave. The middlebox can lookup or remove state entries by providing flow identifiers. In case a state is not present in the cache but in the store, the module will automatically swap it with a cached entry.

Secure service launching. The enterprise needs to attest the integrity of the remotely deployed LightBox instance before launching the service. This is realized by the standard SGX attestation utility [1]. Specifically, the enterprise administrator can request a security measurement of the enclave signed by the CPU, and interact with Intel’s IAS API for verification. During attestation, a secure channel is established to pass configurations, e.g., middlebox processing rules, `etap` ring size and flow cache size, to the LightBox instance. Due to the space limit, we skip the verbose description here. We remark that for the considered service scenario where only two parties (the enterprise and the server provider) are involved, a basic attestation protocol between the two and Intel IAS is sufficient.

2.4 Adversary Model

In line with SGX’s security guarantee, we consider a powerful adversary. We assume that the adversary can gain full control over all user programs, OS and hypervisor, as well as all hardware components in the machine, with the exception of processor package and




Figure 3: Illustration of secure tunneling options.

memory bus. It can obtain a complete memory trace for any process, except those running in the enclave. It is also capable of observing network communications, modifying and dropping packets at will. In particular, the adversary can log all network traffic and conduct sophisticated inference to mine useful information. Our goal here is to thwart practical traffic analysis attacks targeting the *original packets* that are intended for processing at the off-site middleboxes.

Like many SGX applications [2, 37, 59, 63, 72], we consider side-channel attacks [32, 60, 79, 94, 98] out of scope. They can be orthogonally handled by corresponding countermeasures [13, 17, 34, 74, 79]. That said, we fully recognize the security benefits and limitations of SGX, and understand that they are still under rapid iteration [93]. It is yet by far, arguably, the most practical solution for a wide class of complicated tasks including secure middleboxes.

We do not deal with denial-of-service attacks. The middlebox code is assumed to be correct. Also, we assume that the enterprise gateway is always trusted and it does not have to be SGX-enabled.

3 THE ETAP DEVICE

The ultimate goal of *etap* is to enable in-enclave access to the packets intended for middlebox processing, *as if they were locally accessed from the trusted enterprise networks*. Towards this goal, we set forth the following design requirements.

- *Full-stack protection*: when the packets are transmitted in the untrusted networks, and when they traverse through the untrusted platform of the service provider, none of their metadata as defined in Section 1 is directly leaked.
- *Line-rate packet I/O*: *etaps* should deliver packets at a rate that can catch up with a physical network interface card (NIC), without capping the middlebox performance. A pragmatic performance target to shoot is 10Gbps.
- *High usability*: to use *etap*, we need to impose as few changes as possible to the secured middlebox. This implies that if certain network frameworks are used by the middlebox, they should be seamlessly usable inside the enclave too.

3.1 Overview

To achieve full-stack protection, an intuitive idea would be to securely *tunnel* the packets between the gateway and the enclave: the original packets are encapsulated and encrypted as the payloads of new packets, which contain non-sensitive header information (i.e., the IP addresses of the gateway and the middlebox server).

The naive way of encapsulating and encrypting packets individually, as used in L2 tunneling solution like IPSec, does not suffice for our purpose — it does not protect information pertaining to individual packets, including size, timestamp, and as a result, packet count. Padding each packet to the maximum size may hide exact




Figure 4: The architecture of *etap*.

packet size, but it incurs unnecessary bandwidth inflation, and still cannot hide the count and timestamps.

We thus consider encoding the packets as a continuous stream, which is treated as application payloads and transmitted via a secure channel (e.g., TLS). Such streaming design obfuscates packet boundaries, thus hiding the metadata we want to protect, as illustrated in Fig. 3. From a system perspective, the key to this approach is the VIF *tun/tap*¹ that can be used as an ordinary NIC to access the tunneled packets, as widely adopted by popular products like OpenVPN. While there are many userspace TLS suites and some of them even have handy SGX ports [4, 97, 100], the *tun/tap* device itself is canonically driven by the untrusted OS kernel. That is, even if we can terminate the secure channel inside the enclave, the packets are still exposed when accessed via the untrusted *tun/tap* interface.

This inspires us to develop *etap* (the “enclave tap”), which manages packets inside the enclave and enables direct access to them without exiting. The naming comes from the convention that *tap* is for L2 packets (more precisely, frames) while *tun* is for L3 packets, and we aim to protect the L2 header as well. From the middlebox’s point of view, accessing packets in the enclave via *etap* is identical to accessing them via a real NIC in the local enterprise networks.

3.2 Architecture

The major components of *etap* are depicted in Fig. 4. Every *etap* is peered with an *etap-cli* run by the gateway. The two share the same processing logic and since *etap-cli* operates as a normal program in the trusted gateway, we ignore a detailed description for it. A persistent connection will be established between the two for secure traffic tunneling. The *etap* peers will maintain a minimal traffic rate by injecting heartbeat packets to the tunnel.

At the heart of *etap* are two rings for queuing packet data: one for transmitting (TX) and the other for receiving (RX). A packet is described by a `pkt_info` structure, which stores in order the packet size, timestamp, and a buffer for raw packet data. Two additional data structures are used in preparing and parsing packets: a record buffer holds decrypted data and some auxiliary fields inside the enclave; a batch buffer stores multiple records outside the enclave.

The *etap* device is powered by two drivers. The core driver coordinates networking, encoding and cryptographic operations; it also maintains a trusted clock to overcome the lack of high-resolution timing inside the enclave. The poll driver is used by middleboxes to access packets. The two drivers source and sink the

¹<https://www.kernel.org/doc/Documentation/networking/tuntap.txt>

Algorithm 1: etap core driver's RX ring loop

```

1 ocall_fill_rx_bat_buf();
2 check_memory_safety(rx_bat_buf);
3 foreach enc_rec in rx_bat_buf do
4   rec_buf = decrypt(enc_rec);
5   finish_pending_partial_pkt(rec_buf);
6   while has_full_pkt(rec_buf) do
7     pkt_info = parse_next(rec_buf);
8     push_to_rx_etap_ring(pkt_info);
9   refresh_pending_partial_pkt(rec_buf)

```

two rings accordingly. We discuss the support of multiple RX rings for multi-threaded middleboxes in Appendix A.

Remark. The design of etap is agnostic to how the *real* networking outside the enclave is performed. For example, it can use standard kernel networking stack (this is currently used by us). For better efficiency, it can also use faster userspace networking frameworks based on DPDK [42] or netmap [69], as shown in Fig.6.

3.3 Drivers

Core driver. Upon initialization, the core driver takes care of necessary handshakes (via OCALL) for establishing the secure channel and stores the session keys inside the enclave. The packets intended for processing are pushed into the established secure connection in a back-to-back manner, forming a data stream at the application layer. At the transportation layer they are effectively organized into contiguous records (e.g., TLS records) of fixed size (e.g., 16KB for TLS), which then at the network layer are broken down into packets of maximum size. Each original packet is transmitted in the exact format of `pkt_info`, so the receiver will be able to recover from the continuous stream the original packet by first extracting its length, the timestamp, and then the raw packet data. The core driver is run by its own thread. The pseudo code in Alg. 1 outlines the the main RX loop. The TX side is similar and omitted here.

Trusted timing with etap clock. Middleboxes often demand reliable timing for packet timestamping, event scheduling, and performance monitoring. The timer should at least cope with the packet processing rate, i.e., at tens of microseconds. The SGX platform provides trusted relative time source [12], but its resolution is too low (at seconds) for our use case. Alternative approaches resort to system time provided by OS [66] and on-NIC PTP clock [86]. Yet, they both access time from untrusted sources, thus subject to adversarial manipulation. Another system [101] fetches time from a remote trusted website, and its resolution (at hundreds of milliseconds) is still unsatisfactory for middlebox systems.

We design a reliable clock by taking etap's architectural advantage. In particular, we treat `etap-cli` as a trusted time source to attach timestamps to the forwarded packets. The core driver can then maintain a clock (with proper offset) by updating it with the timestamp of each received packet. The resolution of the clock is determined by the packet rate, which in turn bounds the packet processing rate of the middlebox itself. Therefore, the clock should be sufficient for most timing tasks found in middleboxes. Furthermore, we collate the clock periodically with the round-trip delay




Figure 5: Performance of etap by applying three different synchronization mechanisms, without other optimizations suggested in Section 3.5.

estimated by the moderately low-frequency heartbeat packets sent from `etap-cli`, in a way similar to the NTP protocol [58]. Besides accuracy, such heartbeat packets additionally ensure that any adversarial delaying of packets, if it exceeds the collation period, will be detected when the packets are received by etap. We stress that it is still an open problem to provide trusted and high-resolution time for SGX applications [12, 13]. The proposed etap clock fits well for middlebox processing in our targeted high-speed networks.

Poll driver. The poll driver provides access to etap for upper layers. It supplies two basic operations, `read_pkt` to pop packets from RX ring, and `write_pkt` to push packets to TX ring. Unlike the core driver, the poll driver is run by the middlebox thread. It has two operation modes. In the default blocking mode, a packet is guaranteed to be read from or write to etap: in case the RX (resp. TX) ring is empty (resp. full), the poll driver will spin until the ring is ready. In the non-blocking mode, the driver returns immediately if the rings are not ready. This will allow the middlebox more CPU time for other tasks, e.g., processing cached events.

3.4 Security Analysis

The protection of application payloads in the traffic is obvious. We focus on the analysis of metadata. Also we discuss passive adversary only, because the active ones who attempt to modify any data will be detected by the employed authenticated encryption.

Metadata protection. Imagine an adversary located at the ingress point of the service provider's network, or one that has gained full privilege in the middlebox server. She can sniff the entire tunneling traffic trace between the etap peers. As illustrated in Fig. 3, however, the adversary is not able to infer the *packet boundaries* from the encrypted stream embodied as the encrypted payloads of observable packets, which have the maximum size most of the time. Therefore, she cannot learn the low-level headers, size and timestamps of the encapsulated individual packets in transmission. This also implies her inability to obtain the exact packet count, though this number is always bounded in a given period of time by the maximum and minimum possible packet size. Besides, the timestamp attached to the packets delivered by etap comes from the trusted clock, so it is invisible to the adversary.

As a result, a wide range of traffic analyses [16, 19, 68, 91, 96] that directly leverage the metadata will be thwarted, as no such information is available to the adversary.

Beyond metadata. Despite the aforementioned protection over metadata, we do not claim that our design defeats all possible inference attacks. For example, from the bursts in traffic the adversary may potentially learn the launching and termination of certain applications. It is also shown that some delicate analysis can glean meaningful information by looking at the total traffic volume and bursts [23]. A generic countermeasure would be to obfuscate exploitable traffic patterns by injecting calibrated noise [92]. We note that in-depth traffic analysis and mitigation is still a highly active research area [15, 38, 80]. In the context of middlebox outsourcing where an adversary may gain more visibility into the aggregated network traffic, our design can significantly raise the bar for realistic traffic analysis attacks.

3.5 Performance Boosting

While ensuring strong protection, `etap` is hardly useful if it cannot deliver packets at a practical rate. We therefore synergize several techniques to boost its performance.

Lock-free ring. The packet rings need to be synchronized between the two drivers of `etap`. We compare the performance of three approaches: a basic mutex (`sgx_thread_mutex_lock`), a spinlock without context switching (`sgx_thread_mutex_trylock`), and a classic single-producer-single-consumer lockless algorithm [50]. Our evaluation shows that the trusted synchronization primitives of SGX are too expensive for the use of `etap` (see Fig. 5), so we base further optimizations on the lock-free design.

Cache-friendly ring access. In the lock-free design, frequent updates on control variables will trigger a high cache miss rate, the penalty of which is amplified in the enclave. We adapt the cache-line protection technique [52] to our design to relieve this issue. It works by adding a set of new control variables local to the threads to reduce the contention on shared variables. Our evaluations show that this optimization results in a performance gain up to 31%.

Disciplined record batching. Recall that the core driver uses `bat_buf` to buffer the records. The buffer size has to be properly set for best performance. If too small, the overhead of `OCALL` cannot be well amortized. If too large, the core driver needs longer time to perform I/O: this would waste CPU time not only for the core driver that waits for I/O outside the enclave, but also for a fast poll driver that can easily drain or fill the ring. Through extensive experiments, we find a batch size around 10 to be a sweet spot that can deliver practically the best performance for different packet sizes in our settings (see Fig. 10).

3.6 Usability

A main thrust of `etap` is to provide convenient networking functions to in-enclave middleboxes, preferably without changing their legacy interfaces. On top of `etap`, we can port existing frameworks and build new ones. Here, we report our porting efforts of three of them, which greatly improve the usability of `etap`.

Compatibility with `libpcap`. Considering `libpcap` is widely used in networking frameworks and middleboxes for packet capturing [27, 49, 69, 81], we create an adaption layer that implements `libpcap` interfaces over `etap`, including the commonly used packet reading routines (e.g., `pcap_loop`, `pcap_next`), and filter routines




Figure 6: The networking stack enabled by `etap`.

(e.g., `pcap_compile`). This layer allows many legacy systems, including the ones discussed in Section 5, to transparently access protected raw packets inside the enclave.

TCP reassembly. This common function organizes the payloads of possibly out-of-order packets into streams for subsequent processing. To facilitate middleboxes demanding such functionality, we port a lightweight reassembly library `libntoh` [28] on top of `etap`. It exposes a set of APIs to create stream buffers for new flows, add new TCP segments, and flush the buffers with callback functions. It is used in one of our middlebox case-studies.

Advanced networking stack. We also manage to port an advanced networking stack called `mOS`, which allows for programming stateful flow monitoring middleboxes [45], into the enclave on top of `etap`. As a result, a middlebox built with `mOS` can automatically enjoy all security and performance benefits of `etap`, without the need for the middlebox developer to even have any knowledge of SGX. The porting is a non-trivial task as `mOS` has complicated TCP context and event handling, as well as more sophisticated payload reassembly logic than `libntoh`. Our current porting retains the core processing logic of `mOS` and only removes the threading features.

Note that the two stateful frameworks above track flow states themselves, so running them inside the enclave efficiently requires delicate state management, as described in the next section.

4 FLOW STATE MANAGEMENT

To get rid of the expensive application-agnostic EPC paging, a natural idea would be to carefully partition the working set of an SGX application into two parts, a small one that can fit in the enclave, and a large one that can securely reside in the untrusted main memory, while ensuring data swapping between the two in an on-demand manner. A related idea has been positively validated in a prior study [63]. But it still falls within the paradigm of paging, and reports a slowdown over native processing of several factors, even for moderate working sets of a few hundreds of MBs.

To effectively implement the idea above, we design a set of novel data structures specifically for managing flow states in stateful middleboxes. They are compact, such that collectively adding a few tens of MBs overhead to track one million flows concurrently. They are also interlinked, such that the data relocation and swapping involves only cheap pointer operations in addition to necessary data marshalling. To overcome the bottleneck of flow lookup, we further leverage the space-efficient cuckoo hashing to create a fast dual lookup algorithm. Altogether, our state management scheme

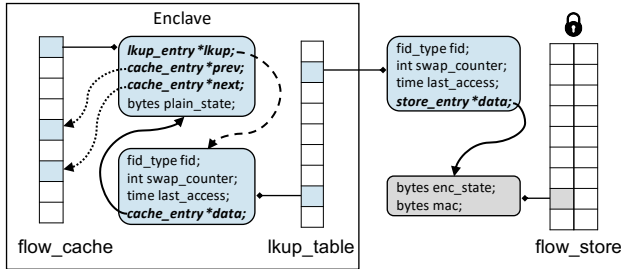


Figure 7: Data structures used in flow state management.

introduces small and nearly constant computation cost to stateful middlebox processing, even with 100,000s of concurrent flows.

Note that we focus on flow-level states, which are the major culprits that overwhelm memory. Other runtime states, such as global counters and pattern matching engines, do not grow with the number of flows, so we leave them in the enclave and handled by EPC paging whenever necessary. Our experiments confirm that the memory explosion caused by flow states is the main source of performance overhead.

4.1 Data Structures

The state management is centered around three abstract tables:

- `flow_cache`, which maintains the states of a fixed number of active flows *in the enclave*;
- `flow_store`, which keeps the encrypted states of inactive flows *in the untrusted memory*;
- `lkup_table`, which allows fast lookup of all flow states from *within the enclave*.

Among them, `flow_cache` has a fixed capacity, while `flow_store` and `lkup_table` can grow as more flows are tracked. Our design principle is to keep the data structures of `flow_cache` and `lkup_table` functional and minimal, so that they can scale to millions of concurrent flows. Figure 7 gives an illustration of them.

- The `cache_entry` holds raw state data. It keeps two pointers (dotted arrows) to implement the Least Recently Used (LRU) eviction policy, and links (dashed arrow) to a `lkup_entry`.
- The `store_entry` holds encrypted state data and authentication MAC. It is maintained in untrusted memory so does not consume enclave resources.
- The `lkup_entry` stores `fid`, a pointer (solid arrow) to either `cache_entry` or `store_entry`, and two small fields. The `fid` represents the conventional 5-tuple to identify flows. The `swap_count` serves as a monotonic counter to ensure the freshness of state; it is initialized to a random value and incremented by 1 on each encryption. The `last_access` assists flow expiration checking, it is updated with `etap` clock on each flow tracking. Note that the design of `lkup_entry` is independent of the underlying lookup structure, which for example can be plain arrays, search trees or hash tables.

The data structures above are succinct, making it efficient to handle high flow concurrency. Assume 8B (byte) pointer and 13B `fid`, then `cache_entry` uses 24B per *cached flow* and `lkup_entry` uses 33B per *tracked flow*. Assume 16K cache entries and full utilization of the underlying lookup structure, then tracking 1M flows requires only 33.8MB enclave memory besides the state data itself.

Algorithm 2: Fast flow tracking with dual lookup

Input: A `fid` extracted from input packet.

Output: The state of flow `fid`.

```

1 entry = flow_cache_cuckoo_lkup(fid);
2 if entry empty then // flow_cache miss
3   entry = flow_store_cuckoo_lkup(fid);
4   if entry empty then // flow_store miss
5     | entry = flow_store_alloc();
6     | check_memory_safety(entry);
7     | victim = drop_from_rear(flow_cache);
8     | victim = encrypt(victim);
9     | swap(entry.state, victim.state);
10  | entry = decrypt(entry);
11 Return entry.state;
```

4.2 Management Procedures

We refer to *flow tracking* as the process of finding the correct flow state on a given `fid`. It takes place in the early stage of the packet processing cycle. The identified state may be accessed anywhere and anytime afterwards [46, 47]. Thus, it should be pinned in the enclave immediately after flow tracking to avoid being accidentally paged out. The full flow tracking procedure is described in Alg. 2.

Initialization. For efficiency, we preallocate entries for all three components. During initialization, a random key is generated and stored inside the enclave for the required authenticated encryption.

Flow tracking. Given a `fid`, we first search through `lkup_table` to check if the flow has been tracked. If it is found in `flow_cache`, we will relocate it to the front of the cache by updating its logical position via the pointers, and return the raw state data. If it is found in `flow_store`, we will swap it with the LRU victim in `flow_cache`. In case of a new flow, an empty `store_entry` is created for the swapping. The entry swapping involves a series of strictly defined operations: 1) Checking memory safety of the candidate `store_entry`; 2) Encrypting the victim `cache_entry`; 3) Decrypting the `store_entry` to the just freed `flow_cache` cell; 4) Restoring the lookup consistency in the `lkup_entry`; 5) Moving the encrypted victim `cache_entry` to `store_entry`. At the end of flow tracking, the expected flow state will be cached in the enclave and returned to the middlebox.

Tracking termination. The tracking of a flow can be explicitly terminated (e.g., upon seeing FIN or RST flag). When this happens, the corresponding `lkup_entry` is removed and the `cache_entry` is nullified. This will not affect `flow_store`, as the flow has already been cached in the enclave.

Expiration checking. We periodically purge expired flow states to avoid performance degradation. The last access time field will be updated at the end of flow tracking for each packet using the `etap` clock. The checking routine will walk through the `lkup_table` and remove inactive entries in the tables.

4.3 Fast Flow Lookup

The fastest path in the flow tracking process above is indicated by `flow_cache` hit, where only a few pointers are updated to refresh LRU linkage. In case of `flow_cache` miss and `flow_store` hit, two memory copy (for swapping) and cryptographic operations are entailed. Due to the interlinked design, these operations have constant cost irrelevant to the number of tracked flows.

When encountering high flow concurrency, we found that the flow lookup sub-procedure becomes the main factor of performance slowdown, as confirmed by one of our tested middleboxes with an inefficient lookup design (PRADS, see Section 6.3). Given the constrained enclave resources, two requirements are therefore imposed on the underlying lookup structure: *search efficiency* and *space efficiency*.

Dual lookup design with cuckoo hashing. We recognize cuckoo hashing as the one to simultaneously achieve the two properties. It has guaranteed $O(1)$ lookup and superior space efficiency, e.g., 93% load factor with two hash functions and a bucket size of 4 [25]. One downside with hashing is their inherent cache-unfriendliness [36], which incurs a higher cache miss penalty in the enclave. Thus, while adopting cuckoo hashing, we still need a cache-aware design.

Our idea is to split `lkup_table` into a small table dedicated for `flow_cache`, and a much larger one for `flow_store`. The large one is searched only after a miss in the small one. The smaller table contains the same number of entries as `flow_cache` and has a fixed size that can well fit into a typical L3 cache (8MB). It is accessed on every packet and thus is likely to reside in L3 cache most of the time. Such a dual lookup design can perform especially well when the `flow_cache` miss rate is relatively low.

To validate the design, we evaluate the two lookup approaches with 1M flows, 512B states and `flow_cache` with 32K entries. As expected, Figure 8 shows that the lower the miss rate, the larger speedup the dual lookup achieves over the single lookup. Real-world traffic often exhibits temporal locality [11, 48]. We also estimate the miss rate of `flow_cache` over a real trace [9]. As shown in Fig. 9, the miss rate can be maintained well under 20% with 16K cache entries, confirming the temporal locality in the trace, hence the efficiency of the dual lookup design in practice.

4.4 Security of State Management

We show that the adversary can only gain little knowledge from the management procedures. It can neither manipulate the procedures to influence middlebox behavior. Therefore, the proposed management scheme retains the same security level as if it is not applied, i.e., when all states are handled by EPC paging.

We first analyze the adversary's view throughout the procedures. Among the three tables, `flow_cache` and `lkup_table` are always kept in the enclave, hence invisible to the adversary. Stored in untrusted memory, `flow_store` is fully disclosed. The adversary can obtain all `store_entry`'s, but never sees the state in clear text.

She will notice the creation of new flow state, but cannot link it to a previous one, even if the two have exactly the same content, because of the random initialization of the `swap_count`. Similarly, she is not able to track traffic patterns (e.g., packets coming in bursts) of a single flow, because the `swap_count` will increment




Figure 8: The speedup of dual lookup design over single lookup design.




Figure 9: The miss rate of `flow_cache` with varying size for a real network trace [9].

upon each swapping and produce different ciphertexts for the same flow state. In general, she cannot link any two `store_entry`'s.

The explicit termination of a flow is unknown to the adversary, as the procedure takes place entirely in the enclave. In contrast, she will notice state removal events during expiration checking. Yet, this information is useless as the entries are not linkable.

Now we consider an active adversary. Due to the authenticated encryption, any modification of `state_entry`'s is detectable. Malicious deletion of a `state_entry` will be also caught when it is supposed to be swapped into the enclave after a hit in `lkup_table`. She cannot inject a fake entry since `lkup_table` is inaccessible to her. Furthermore, the replay attack will be thwarted because `swap_count` keeps the freshness of the state.

5 INSTANTIATIONS OF LIGHTBOX

We have implemented a working prototype of LightBox and instantiated it for three case-study stateful middleboxes.²

5.1 Porting Middleboxes to SGX

A middlebox system should be first ported to the SGX enclave before it can enjoy the security and performance benefits of LightBox, as illustrated in Fig. 2. After that, the middlebox's original insecure I/O module will be seamlessly replaced with `etap` and the network frameworks stacked thereon; its flow state management procedures, including memory management, flow lookup and termination, will be changed to that of LightBox as well.

There are several ways to port a legacy middlebox. One is to build the middlebox with trusted LibOS [7, 87], which are pre-ported to SGX and support general system services within the enclave. Another more specialized approach is to identify only the necessary system services and customize a trusted shim layer for optimized performance and TCB size [55]. To prepare for our middlebox case-studies, we follow the second approach and implement a shim layer that supports the necessary system calls and `struct` definitions.

Some prior systems allow modular development of middleboxes that are automatically secured by SGX [35, 66, 86]. For middleboxes built this way, we can directly substitute their network I/O and flow state management with LightBox, augmenting them with full-stack protection and efficient stateful processing.

²<https://github.com/lightbox-impl/LightBox>




Figure 10: Performance of etap against varied batch size.




Figure 11: Performance of etap against varied ring size.




Figure 12: CPU usage of etap against middlebox throughput.




Figure 13: Performance of etap on real trace.

5.2 Middlebox Case Studies

We now introduce the three middleboxes we instantiated for LightBox. For discussions on the efforts of instantiating them with LightBox, we assume that they have already been ported to SGX. Both PRADS and lwIDS use `libpcr` for pattern matching, so we manually port it as a trusted library to be used within the enclave.

PRADS [27]. Capable of detecting network assets (e.g., OSes, devices) in packets against predefined fingerprints, PRADS has been widely used in academic research [29, 45, 47]. It uses `libpcap` for packet I/O, so its main packet loop can be directly replaced with the compatibility layer we built on etap (Section 3.6). We also adapt its own flow tracking logic to LightBox’s state management procedures without altering the original functionality. This affects about 200 lines of code (LoC) in the original PRADS project with 10K LoC.

lwIDS. Based on the tcp reassembly library `libntoh` [28], we built a `lightweight IDS` that can identify malicious patterns over reassembled data. Whenever the stream buffer is full or the flow is completed, the buffered content will be flushed and inspected against a set of patterns. Note that the packet I/O and main stream reassembly logic of lwIDS is handled by `libntoh` (3.8K LoC), which we have already ported on top of etap (Section 3.6). The effort of instantiating LightBox for lwIDS thus reduces to adjusting the state management module of `libntoh`, which amounts to a change of around 100 LoC.

mIDS. We design a more comprehensive middlebox, called mIDS, based on the mOS framework [45] and the pattern matching engine DFC [14]. Similar to lwIDS, mIDS will flush stream buffers for inspection upon overflow and flow completion; but to avoid consistent failure, it will also do the flushing and inspection when receiving out-of-order packets, as we found that the logic for handling such case is yet to be completed in current mOS code. Again, since we have ported mOS (26K LoC) with etap (Section 3.6), the remaining effort of instantiating LightBox for mIDS is to modify

the state management logic, resulting in 450 LoC change. Note that such effort is one-time: hereafter, we can instantiate any middlebox built with mOS without change.

6 EVALUATION

6.1 Methodology and Setup

Our evaluation comprises two main parts: in-enclave packet I/O, where we evaluate etap from various aspects and decide the practically optimal configurations (Section 6.2); middlebox performance, where we measure the efficiency of LightBox against a native and a strawman approach for the three case-study middleboxes (Section 6.3). We will also give discussions on experimental comparison between LightBox and previous systems (Section 6.4). We use a real SGX-enabled workstation with Intel E3-1505 v5 CPU and 16GB memory in the experiments. Equipped with 1Gbps NIC, the workstation is unfortunately incapable of reflecting etap’s real performance, so we prepare two experiment setups. In what follows, we will use K for thousand and M for million in the units.

Setup 1. The first setup is dedicated for evaluation on etap, where we run `etap-cli` and `etap` on the same standalone machine and let them communicate with the fast memory channel via kernel networking. Note that `etap-cli` needs no SGX support and runs as a normal user-land program. To reduce the side effect of running them on the same machine, we tame the kernel networking buffers such that they are kept small (500KB) but still performant. Our intent here is to demonstrate that etap can catch up with the rate of a real 10Gbps NICs in practical settings.

Setup 2. Deployed in a local 1Gbps LAN, the second setup is for evaluating middlebox performance. We use a separate machine as the *gateway* to run `etap-cli`, so it communicates with etap via the real link. The gateway machine also serves as the *server* to accept connections from clients (on other machines in the LAN). We then use `tcpkali` [71] to generate concurrent TCP connections transmitting random payloads from clients to the server; all ACK packets from the server to clients are filtered out. Our environment can afford up to 600K concurrent connections. We also obtain a real trace from CAIDA [9] for experiments; it is collected by monitors deployed at backbone networks. The trace is sanitized and contains only anonymized L3/L4 headers, so we pad them with random payloads to their original lengths specified in the header. We use the first 100M packets from the trace in our experiments.




Figure 14: Performance of PRADS under controlled settings.




Figure 16: PRADS on real trace.




Figure 15: Performance of lwIDS under controlled settings.




Figure 17: lwIDS on real trace.

6.2 In-enclave Packet I/O Performance

To evaluate `etap`, we create a bare middlebox which keeps reading packets from `etap` without further processing. It is referred to as `PktReader`. We keep a large memory pool (8GB) and feed packets to `etap-cli` directly from the pool.

Parameterized evaluation. We first investigate how batching size affects `etap` performance. The ring size is set as 1024. As shown in Fig. 10, the optimal size appears between 10 and 100 for all packet sizes. The throughput drops when the batching size becomes either too small or overly large, matching our expectation and analysis in Section 3.5. With a batching size of 10, `etap` can deliver small 64B (byte) packet at 7.4Gbps, and large 1024B packet at 12.4Gbps, which is comparable to advanced packet I/O framework on modern 10Gbps NIC [69]. We set 10 as the default batching size and use this configuration in all following experiments.

Shrinking `etap` ring is beneficial in that precious enclave resources can be saved for middlebox functions, and in the case of multi-threaded middleboxes, for efficiently supporting more RX rings. However, smaller ring size generally leads to lower I/O throughput. Figure 11 reports the results with varying ring sizes. As can be seen, the tipping point occurs around 256, where the throughput for all packet sizes begins to drop sharply as ring size decreases. Beyond that and up to 1024, the performance appears insensitive to ring size. We thus use 256 as the default ring size in all subsequent tests.

Resource consumption. The rings contribute to the major `etap` enclave memory consumption. One ring uses as small as 0.38MB as per the default configuration, and a working `etap` consumes merely 0.76MB. The core driver of `etap` is run by dedicated threads and we are also interested in its CPU consumption. The driver will spin in

the enclave if the rings are not available, since exiting enclave and sleeping outside is too costly. This implies that a slower middlebox thread will force the core driver to *waste* more CPU cycles in the enclave. To verify such effect, we tune `PkgReader` with different levels of complexity, and estimate the core driver’s CPU usage under varying middlebox speed. As expected, the results in Fig. 12 delineate a clear negative correlation between the CPU usage of `etap` and the performance of middlebox itself. With 70% utilization of a single core the core driver can handle packets at its full speed. Overall, we can see that an average commodity processor is more than enough for our target 10Gbps in-enclave packet I/O.

Performance on real trace. Figure 13 shows `etap`’s performance on the real CAIDA trace that has a mean packet size of 680B. We estimate the throughput for every 1M packets while replaying the trace to `etap-cli`. As shown, although there are small fluctuations overtime due to varying packet size, the throughput remains mostly within 11 – 12Gbps and 2 – 2.5Mpps. This further demonstrates `etap`’s practical I/O performance.

6.3 Middlebox Performance

We study the performance of the three middleboxes, each with three variants: the vanilla version (denoted as `Native`) running as a normal program; naive SGX port (denoted as `Strawman`) that uses `etap` and our ported `libntoh` and `mOS` for networking, but relies on EPC paging for however much enclave memory is needed; the `LightBox` instance as described in Section 5. It is worth noting that despite the name, the `Strawman` variants actually benefit a lot from `etap`’s efficiency. Our goal here is primarily to investigate the efficiency of our state management design.




Figure 18: Performance of mIDS under controlled settings.

We use the default configurations for all three middleboxes unless otherwise specified. For lwIDS we compile 10 pcre engines with random patterns for inspection; for mIDS we build the DFC engine with 3700 patterns extracted from Snort community ruleset. The flow state of PRADS, lwIDS, and mIDS has a size of 512B³, 5.5KB, and 11.4KB⁴, respectively; the latter two include stream reassembly buffer of size 4KB and 8KB. For LightBox variants, the number of entries of `flow_cache` is fixed to 32K, 8K and 4K for PRADS, lwIDS, and mIDS, respectively.

6.3.1 Controlled live traffic. To gain a better understanding of how stateful middleboxes behave in the highly constrained enclave space, we test them in controlled settings with varying number of concurrent TCP connections between clients and the server. We control the clients' traffic generation load such that the aggregated traffic rate at the server side remains roughly the same for different degrees of concurrency. By doing so the comparisons are made fair and meaningful. In addition, we start to collect data points only when all connections are established and stabilized. We measure the mean packet processing delay in microsecond (μ s) every 1M packets, and each reported data point is averaged over 100 runs.

PRADS. From Fig. 14, we can see that LightBox adds negligible overhead ($< 1\mu$ s) to native processing of PRADS regardless of the number of flows. In contrast, Strawman incurs significant and increasing overhead after 200K flows, due to the involvement of EPC paging. Interestingly, by comparing the subfigures it can also be seen that Strawman performs worse for smaller packets. This is because smaller packet leads to higher packet rate while saturating the link, which in turn implies higher page fault ratio. For 600K flows, LightBox attains $3.5\times - 30\times$ speedup over the Strawman.

lwIDS. Figure 15 presents similar results for lwIDS. Here, the performance of Strawman is further degraded, since lwIDS has larger flow state size than PRADS and its memory footprint exceeds 550MB even when tracking only 100K flows. For 64B packet, LightBox introduces $6 - 8\mu$ s packet delay ($4 - 5\times$ to native) because the state management dominates the whole processing; nonetheless, it still outperforms Strawman by $5 - 16\times$. For larger packets,

³PRADS has 124B flow state, which is too small under our current experiment settings. To better approximate realistic scenarios, we pad the flow state of PRADS to 512B with random bytes. No such padding is applied to lwIDS and mIDS.

⁴This size is resulted from the rearrangement of mOS's data structures pertaining to flow state. We merge all data structures into a single one to ease memory management.

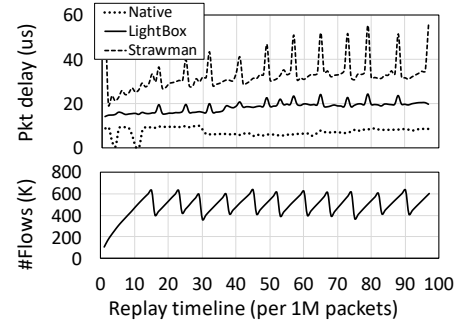


Figure 19: mIDS on real trace.

the network function itself becomes dominant and the overhead of LightBox over Native is reduced, as shown in Fig. 15 (b) and (c).

mIDS. Among the case-study middleboxes, mIDS is the most complicated one with the largest flow state. Here, our testbeds can scale to 300K concurrent connections. For each connection mIDS will track two flows, one for a direction, and allocate memory accordingly. But since we filter out the trivial ACK packets from the server to clients, we still count only one flow per connection. Figure 18 reveals that the performance of mIDS's three variants follows similar trends as in previous middleboxes: Native and LightBox are insensitive to the number of concurrent flows; conversely, the overhead of Strawman grows as more flows are tracked.

But in contrast to previous cases, now the overhead of LightBox over Native becomes notable. This is explained by mIDS's large flow state size, i.e., 11.4 KB, which leads to the substantial cost of encrypting/decrypting and copying states. Besides, we found that for each packet, in addition to its own flow, mIDS will also access the paired flow, *doubling the cost of our flow tracking design* (see Section 4.2). Nonetheless, we can see that the gap is closing towards larger packet size, as the network function processing itself weighs in. Later in this section we will discuss how to further improve our design to cope with large flow state and connection-based tracking.

6.3.2 Real trace. Now we investigate middlebox performance with respect to the real CAIDA trace. The trace is loaded by the gateway and replayed to the middlebox for processing. Again, we collect the data points every 1M packets. Packets of unsupported types are filtered out so only 97 data points are collected for each case. Since L2 headers are stripped in the CAIDA trace, we also adjust the packet parsing logic accordingly for the middleboxes. Yet another important factor for real trace is the flow timeout setting. We must carefully set the timeout so inactive flows are purged well in time, lest excessive flows overwhelm the testbeds. Here, we set the timeout for PRADS, lwIDS, and mIDS, to 60, 30, and 15 seconds, respectively. Table 2 reports the overall throughput of relaying the trace. Below we give more detailed analysis.

PRADS. Figure 16 shows that the packet delay of Strawman grows with the number of flows; it needs about 240μ s to process a packet when there are 1.6M flows. In comparison, LightBox maintains low and stable delay (around 6μ s) throughout the test. A bit surprisingly, it even edges over the native processing as more flows are tracked, attributed to an inefficient chained hashing design used in the native

Table 2: Throughput (Mbps) under CAIDA trace.

	Native	Strawman	LightBox
PRADS	429.24	67.399	928.06
lwIDS	689.11	182.57	685.36
mIDS	713.56	161.02	310.42

implementation. This highlights the importance of *efficient flow lookup in stateful middleboxes*.

lwIDS. Compared with PRADS, the number of concurrent flows tracked by lwIDS decreases, as shown in Fig. 17. This is due to the halved timeout and the more aggressive strategy we used for flow deletion: we remove a flow when a FIN or RST flag is received, and we do not handle TIME_WAIT event. It can be seen that with fewer flows, Strawman still incurs remarkable overhead, while the difference between LightBox and Native is indistinguishable.

mIDS. The case for mIDS is tricky. Its current implementation of flow timeout seems not to be fully working, so we replaced the related code with the logic of checking all flows for expiration every timeout interval. We also made some modifications to ensure that the packet formats and abnormal packets in the real trace can be properly processed. Figure 19 reports the test results. There is again a large gap between Strawman and Native. Yet, as in the controlled settings, there is some moderate gap between LightBox and Native, due to the large state and double flow tracking design.

6.3.3 Reflections on future improvement. Above results show that when the per-flow state size is not overly large, our current design, which treats the state as a whole chunk of raw data in an agnostic manner, suffices to achieve near-native performance for stateful processing. Otherwise, it may not be wise to manage the large state as a whole. We would expect more fine-grained partition and handling of the state to improve efficiency.

A promising direction is to separate the large stream buffer (e.g., 8KB in the case of mIDS) from the rest of the state (e.g., 3.4KB). The processing of each packet will only touch a small portion of the buffer, dispensing with the high overhead of encrypting, decrypting and copying the entire buffer. Only when it is time to flush the buffer for inspection do we need to load it into enclave in its entirety. This will significantly diminish the cost of the flow tracking routine. Note that here we should refrain from moving encrypted raw packets individually out from the enclave, as this will leak the packet size and count. A more secure way would be to divide the stream buffer into chunks of fixed length and handle packets in batches.

In a similar vein, for middleboxes that access the flow state of both connection directions on each packet, instead of treating both flows equally, we can manage only the necessary data fields of the paired flow, leaving the vast majority of its state untouched. This requires a slight redesign of our current data structures to support the effective linking of pairing flows.

A common theme in the fine-grained approaches suggested above is to reduce the amount of unnecessary data moved across the enclave boundary, and hence lessen the management overhead. We leave the detailed designs and exploration of potential trade-offs between security and efficiency as an interesting future work.

6.4 Comparison with Previous Systems

We now discuss previous secure middlebox systems built upon SGX regarding experiment settings and performance evaluation. They have been evaluated on various network functions, and they demonstrate that certain workloads can be run in the enclave with marginal performance overhead. This work differs from them in that it takes into consideration: 1) the complexity of stateful processing as seen in production-level middleboxes, and 2) the high flow concurrency encountered in deployed networks. These two features together pose unique challenges in operating stateful middleboxes in enclaves at a reasonable cost. To our best knowledge, such experiment settings have not been considered in prior works. Their experiments are mostly confined to small memory footprint within EPC limit (128MB), avoiding expensive EPC paging; but in our settings, the memory footprint of middleboxes can grow to multiple GBs. We discuss some representative systems below.

In ShieldBox [86], several stateless middleboxes, including some micro ones with simplistic functions, are evaluated. They process packets independently, without tracking any flow states. Since the middlebox memory footprint is always kept small, it is not very surprising to see that the performance of shielded middleboxes is close to native in virtually all test cases.

SGX-BOX [35] allows inspection over reassembled streams. However, it leaves flow state other than the stream buffer outside enclave unprotected, obviating the challenge of fully protecting stateful middlebox processing. Besides, it reports only preliminary evaluation that the overhead of inspecting a single stream is small.

Safebricks [66] presents a more diverse set of experiments. But the middleboxes under testing are still arguably too simple in reality. Unlike the stateful lwIDS and more advanced mIDS we have built, it uses a simple DPI application working on individual packets but not reassembled streams. The NAT and load balancer tested there are also stateless. It does include a stateful firewall, but the functionality seems basic, and more importantly, it was not evaluated against concurrent flows.

To sum up, due to the different focuses and methodology, it is not feasible to derive a direct and fair experimental comparison between LightBox and previous systems. We for the first time evaluate non-trivial stateful middleboxes under settings with high flow concurrency, and we hope that this work can invite more efforts on bridging the research-practice gap for secure middlebox systems. On the other hand, in addition to our intensive evaluations, it is also interesting and critical to experiment with other important aspects such as service function chaining, as done by previous works [66, 86]. Our current work has built solid a foundation for conducting further evaluations, and we leave them as future work.

7 RELATED WORK

Secure Middleboxes. BlindBox [77] is the first system that applies cryptographic protocols (i.e., searchable encryption and garbled circuit) to enable inspection on encrypted packet payloads. A list of follow-up designs are proposed. The work [99] emphasizes on the protection of middlebox rules and the support of more inspection rules. SPABox [26] and BlindIDS [10] put extra attention on reducing session setup cost. Besides payloads, privacy-preserving

packet header checking is also studied. The work [57] considers using heavy homomorphic encryption for generic network functions, which is mostly of theoretical interest only. SplitBox [3] employs a distributed model for a certain class of packet header processing with multi-party computation techniques. Embark [51] introduces a customized prefix-matching scheme; by integrating the technique from [77], it supports a wider class of middlebox functions. These software-centric solutions are often restricted in functionality (especially for stateful processing) and performance.

There are also several designs based on trusted hardware, i.e., Intel SGX. S-NFV [78] proposes to protect specifically the middlebox state, but not the entire middlebox processing over protected traffic. Trusted Click [18] and ShieldBox[86] port Click modular router [49] to the enclave. But they do not consider the protection of metadata, and lack support for stateful processing due to the limitation of Click that is inherently stateless. In a setting different from middlebox outsourcing, SGX-BOX [35] and mbTLS [62] enable middleboxes to intercept TLS connections and securely inspect traffic in the enclave, with primary focus on programmability and deployability, respectively. One note on SGX-BOX is that it employs the mOS framework for stateful processing outside enclave, so all states except the encrypted stream buffers are left unprotected and the correctness of stateful processing is not guaranteed. Targeting the same scenario as us, the latest work SafeBricks [66] pays extra attention to middlebox code protection, and applies IPSec for secure traffic tunneling. While protecting packet headers, it is still vulnerable to traffic analysis attacks leveraging packet size and count. None of these hardware-assisted solutions protects low-level traffic metadata as LightBox does, nor do they enable efficient stateful processing in real networks experiencing high flow concurrency.

Reducing overhead of SGX. One theme of our research is to minimize the performance overhead incurred by SGX while retaining security guarantees. A main strand of works approach this goal with switch-less ECALLs/OCALLs by avoiding expensive context switching [2, 95]. This approach is orthogonal to our designs. In fact, with the official release of such support [85], we can readily replace our normal use of OCALL with the switch-less version for even better efficiency. Eleos [63] introduces user-managed paging to alleviate the overhead of naive EPC paging. Its cache-store architecture is similar to ours, but for generality it entails coarse-grained data structure (e.g., page table) and complex procedures (e.g., address translation and page table walk), compared with our tailored and optimized designs. Therefore, it may not suit the performance-sensitive middlebox applications. To help grow EPC in the future, a recent study proposes to refine the underlying data structure for integrity checking [84]. Whether a large EPC can overcome SGX's current performance issue without enlarging the attack surface remains an open problem. Our state management design significantly increases the nominal secure memory that is usable by SGX applications. We hope that the proposed compact data structure and efficient lookup algorithm can provide helpful insights into future study towards this direction.

SGX-enabled systems. Many systems have been built with SGX, such as data analytics platforms [72, 75] and secure system services [2, 37]. Their application scenarios are different from ours,

and in particular, they do not face the challenge of operating stateful middlebox in real networks with strict performance requirements.

8 CONCLUSION

We present LightBox, an SGX-assisted secure middlebox system. While many researches have explored the possibilities of securing middleboxes with SGX and claimed its efficiency and practicality, this is the first work to ascertain the claims from a more pragmatic perspective. We bridge the research-practice gap by identifying two critical challenges, from both security and functionality aspect, and satisfactorily address them with domain knowledge and extensive customization. Among others, our first main technical contribution is an elegant in-enclave virtual network interface that is highly secure, efficient and usable; and our second main innovation is the flow state management scheme comprising data structures and algorithms optimized for the enclave space. They together build up a comprehensive solution for deploying off-site middleboxes with strong protection and stateful processing, at near-native speed. We hope that LightBox can push forward secure network function virtualization and middlebox-as-a-service to the practical realm, and that our work can invite more efforts on building hardware-assisted secure systems that are practical and usable.

ACKNOWLEDGMENTS

We thank our shepherd, Lorenzo De Carli, for helping us improve the paper, and the anonymous reviewers for their helpful comments. This work was supported in part by the Research Grants Council of Hong Kong under Grant CityU 11276816, Grant CityU 11212717, and Grant CityU C1008-16G, and by the National Natural Science Foundation of China (NSFC) under Grant 61572412 and Grant 61602396. Yajin Zhou was partially supported by the NSFC under Grant 61872438, and the Fundamental Research Funds for the Central Universities. Qian Wang was supported in part by the NSFC under Grant 61822207, and by the Equipment Pre-Research Joint Fund of the Ministry of Education of China (Youth Talent) under Grant 6141A02033327. Kui Ren was supported by the NSFC under Grant 61772236.

REFERENCES

- [1] Ittai Anati, Shay Gueron, Simon Johnson, and Vincent Scarlata. 2013. Innovative technology for CPU based attestation and sealing. In *Proc. of ACM HASP*.
- [2] Sergei Arnatov, Bohdan Trach, Franz Gregor, Thomas Knauth, Andre Martin, Christian Priebe, Joshua Lind, Divya Muthukumar, Dan O'Keeffe, Mark L. Stillwell, David Goltzsche, Dave Eyers, Rüdiger Kapitza, Peter Pietzuch, and Christof Fetzer. 2016. SCONE: Secure Linux Containers with Intel SGX. In *Proc. of USENIX OSDI*.
- [3] Hassan Jameel Asghar, Luca Melis, Cyril Soldani, Emiliano De Cristofaro, Mohamed Ali Kaafar, and Laurent Mathy. 2016. SplitBox: Toward Efficient Private Network Function Virtualization. In *Proc. of ACM HotMiddlebox*.
- [4] Pierre-Louis Aublin, Florian Kelbert, Dan O'Keeffe, Divya Muthukumar, Christian Priebe, Joshua Lind, Robert Krahn, Christof Fetzer, David Eyers, and Peter Pietzuch. 2017. TaLoS: Secure and transparent TLS termination inside SGX enclaves. *Imperial College London, Tech. Rep 5* (2017).
- [5] AWS. 2018. AWS Direct Connect. Online at: <https://aws.amazon.com/directconnect/>.
- [6] Azure. 2018. Azure ExpressRoute. Online at: <https://azure.microsoft.com/en-us/services/expressroute/>.
- [7] Andrew Baumann, Marcus Peinado, and Galen Hunt. 2014. Shielding applications from an untrusted cloud with Haven. In *Proc. of USENIX OSDI*.
- [8] Theophilus Benson, Aditya Akella, Anees Shaikh, and Sambit Sahu. 2011. Cloud-NaaS: a cloud networking platform for enterprise applications. In *Proc. of ACM SOCC*.

- [9] CAIDA. 2016. The CAIDA UCSD Anonymized Internet Traces 2016 - April 6th. Online at: http://www.caida.org/data/passive/passive_2016_dataset.xml.
- [10] Sébastien Canard, Aida Diop, Nizar Kheir, Marie Paindavoine, and Mohamed Sabt. 2017. BlindIDS: Market-Compliant and Privacy-Friendly Intrusion Detection System over Encrypted Traffic. In *Proc. of ACM AsiaCCS*.
- [11] Martin Casado, Teemu Koponen, Daekyeong Moon, and Scott Shenker. 2008. Rethinking Packet Forwarding Hardware. In *Proc. of ACM HotNets*.
- [12] Shanwei Cen and Bo Zhang. 2017. Trusted Time and Monotonic Counters with Intel Software Guard Extensions Platform Services. Online at: <https://software.intel.com/sites/default/files/managed/1b/a2/Intel-SGX-Platform-Services.pdf>.
- [13] Sanchuan Chen, Xiaokuan Zhang, Michael K Reiter, and Yingqian Zhang. 2017. Detecting privileged side-channel attacks in shielded execution with Déjà Vu. In *Proc. of ACM AsiaCCS*.
- [14] Byungkwon Choi, Jongwook Chae, Muhammad Jamshed, Kyoungsoo Park, and Dongsu Han. 2016. DFC: Accelerating string pattern matching for network applications. In *Proc. of USENIX NSDI*.
- [15] M. Conti, Q. Q. Li, A. Maragno, and R. Spolaor. 2018. The Dark Side(-Channel) of Mobile Devices: A Survey on Network Traffic Analysis. *IEEE Communications Surveys Tutorials* 20, 4 (2018), 2658–2713.
- [16] Mauro Conti, Luigi Vincenzo Mancini, Riccardo Spolaor, and Nino Vincenzo Verde. 2016. Analyzing android encrypted network traffic to identify user actions. *IEEE TIFS* 11, 1 (2016), 114–125.
- [17] Victor Costan, Ilya Lebedev, and Srinivas Devadas. 2016. Sanctum: Minimal hardware extensions for strong software isolation. In *Proc. of USENIX Security*.
- [18] Michael Coughlin, Eric Keller, and Eric Wustrow. 2017. Trusted Click: Overcoming Security Issues of NFV in the Cloud. In *Proc. of ACM SDN-NFV Security*.
- [19] Scott E. Coull and Kevin P. Dyer. 2014. Traffic Analysis of Encrypted Messaging Services: Apple iMessage and Beyond. *ACM SIGCOMM CCR* 44, 5 (2014).
- [20] Lorenzo De Carli, Robin Sommer, and Somesh Jha. 2014. Beyond pattern matching: A concurrency model for stateful deep packet inspection. In *Proc. of ACM CCS*.
- [21] Robert DeFrancesco. 2019. Securing The Cloud With Zscaler. Online at: <https://www.forbes.com/sites/robertdefrancesco/2019/03/06/securing-the-cloud-with-zscaler/#29281a785d52>.
- [22] Mihai Dobrescu, Norbert Egi, Katerina Argyraki, Byung-Gon Chun, Kevin Fall, Gianluca Iannaccone, Allan Knies, Maziar Manesh, and Sylvia Ratnasamy. 2009. RouteBricks: Exploiting Parallelism to Scale Software Routers. In *Proc. of ACM SOS*.
- [23] Kevin P Dyer, Scott E Coull, Thomas Ristenpart, and Thomas Shrimpton. 2012. Peek-a-boo, I still see you: Why efficient traffic analysis countermeasures fail. In *Proc. of IEEE S&P*.
- [24] Daniel E Eisenbud, Cheng Yi, Carlo Contavalli, Cody Smith, Roman Kononov, Eric Mann-Hielscher, Ardas Cilingiroglu, Bin Cheyney, Wentao Shang, and Jinnah Dylan Hosein. 2016. Maglev: A Fast and Reliable Software Network Load Balancer. In *Proc. of USENIX NSDI*.
- [25] Ulfar Erlingsson, Mark Manasse, and Frank McSherry. 2006. A cool and practical alternative to traditional hash tables. In *Proc. of 7th Workshop on Distributed Data and Structures (WDAS'06)*.
- [26] Jingyuan Fan, Chaowen Guan, Kui Ren, Yong Cui, and Chunming Qiao. 2017. SPABox: Safeguarding Privacy During Deep Packet Inspection at a MiddleBox. *IEEE/ACM Transactions on Networking* 25, 6 (2017), 3753–3766.
- [27] Edward Fjellskål. 2017. Passive Real-time Asset Detection System. Online at: <https://github.com/gamelinux/prads>.
- [28] Chema Garcia. 2018. libntoh. Online at: <https://github.com/sch3m4/libntoh>.
- [29] Aaron Gember-Jacobson, Raajay Viswanathan, Chaithan Prakash, Robert Grandl, Junaid Khalid, Sourav Das, and Aditya Akella. 2014. OpenNF: Enabling Innovation in Network Function Control. In *Proc. of ACM SIGCOMM*.
- [30] Glen Gibb, Hongyi Zeng, and Nick McKeown. 2012. Outsourcing network functionality. In *Proc. of ACM HotSDN*.
- [31] Google. 2018. Google Dedicated Interconnect. Online at: <https://cloud.google.com/interconnect/docs/concepts/dedicated-overview>.
- [32] Johannes Götzfried, Moritz Eckert, Sebastian Schinzel, and Tilo Müller. 2017. Cache attacks on Intel SGX. In *Proc. of ACM EuroSec*.
- [33] The Tcpdump Group. 2018. libpcap. Online at: <https://www.tcpdump.org>.
- [34] Daniel Gruss, Julian Lettner, Felix Schuster, Olya Ohrimenko, Istvan Haller, and Manuel Costa. 2017. Strong and Efficient Cache Side-Channel Protection using Hardware Transactional Memory. In *Proc. of USENIX Security*.
- [35] Juheng Han, Seongmin Kim, Jaehyeong Ha, and Dongsu Han. 2017. SGX-Box: Enabling Visibility on Encrypted Traffic Using a Secure Middlebox Module. In *Proc. of the First Asia-Pacific Workshop on Networking*.
- [36] Gregory L Heileman and Wenbin Luo. 2005. How Caching Affects Hashing. In *Proc. of ALENEX/ANALCO*.
- [37] Tyler Hunt, Zhiting Zhu, Yuanzhong Xu, Simon Peter, and Emmett Witchel. 2016. Ryoan: a distributed sandbox for untrusted computation on secret data. In *Proc. of USENIX OSDI*.
- [38] Mohsen Imani, Mohammad Saidur Rahman, and Matthew Wright. 2018. Adversarial Traces for Website Fingerprinting Defense. In *Proc. of ACM CCS*.
- [39] Zscaler Inc. 2019. Encrypted Traffic Dashboard. Online at: <https://www.zscaler.com/threatlab/encrypted-traffic-dashboard>.
- [40] Zscaler Inc. 2019. SSL Inspection. Online at: <https://www.zscaler.com/products/ssl-inspection>.
- [41] Zscaler Inc. 2019. Zscaler Architecture: Cloud from the beginning. Online at: <https://www.zscaler.com/products/cloud-architecture-security-as-a-service>.
- [42] Intel. 2017. Data plane development kit. Online at: <http://www.dpdk.org>.
- [43] Prerit Jain, Soham Desai, Seongmin Kim, Ming-Wei Shih, J Lee, Changho Choi, Youjung Shin, Taesoo Kim, Brent Byunghoon Kang, and Dongsu Han. 2016. Opensgx: An open platform for sgx research. In *Proc. of NDSS*.
- [44] Muhammad Asim Jamshed, Jihyung Lee, Sangwoo Moon, Insu Yun, Deokjin Kim, Sungryou Lee, Yung Yi, and Kyoungsoo Park. 2012. Kargus: a highly-scalable software-based intrusion detection system. In *Proc. of ACM CCS*.
- [45] Muhammad Asim Jamshed, YoungGyou Moon, Donghwi Kim, Dongsu Han, and Kyoungsoo Park. 2014. mOS: A Reusable Networking Stack for Flow Monitoring Middleboxes. In *Proc. of USENIX NSDI*.
- [46] Murad Kablan, Azzam Alsudais, Eric Keller, and Franck Le. 2017. Stateless Network Functions: Breaking the Tight Coupling of State and Processing. In *Proc. of USENIX NSDI*.
- [47] Junaid Khalid, Aaron Gember-Jacobson, Roney Michael, Anubhavnidhi Abhashkumar, and Aditya Akella. 2016. Paving the way for NFV: Simplifying middlebox modifications using StateAlyzr. In *Proc. of USENIX NSDI*.
- [48] Changhoon Kim, Matthew Caesar, Alexandre Gerber, and Jennifer Rexford. 2009. Revisiting route caching: The world should be flat. In *Proc. of International Conference on Passive and Active Network Measurement*.
- [49] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M Frans Kaashoek. 2000. The Click modular router. *ACM Transactions on Computer Systems (TOCS)* 18, 3 (2000), 263–297.
- [50] Leslie Lamport. 1977. Proving the correctness of multiprocess programs. *IEEE transactions on software engineering* 2 (1977), 125–143.
- [51] Chang Lan, Justine Sherry, Raluca Ada Popa, and Sylvia Ratnasamy. 2016. EM-BArk: Securely Outsourcing Middleboxes to the Cloud. In *Proc. of USENIX NSDI*.
- [52] Patrick PC Lee, Tian Bu, and Girish Chandranmenon. 2010. A lock-free, cache-efficient multi-core synchronization mechanism for line-rate network traffic monitoring. In *Proc. of IEEE Parallel & Distributed Processing (IPDPS)*.
- [53] Paige Leskin. 2018. The 21 scariest data breaches of 2018. Online at: <https://www.businessinsider.com/data-hacks-breaches-biggest-of-2018-2018-12>.
- [54] Hongda Li, Hongxin Hu, Guofei Gu, Gail-Joon Ahn, and Fuqiang Zhang. 2018. vNIDS: Towards Elastic Security with Safe and Efficient Virtualization of Network Intrusion Detection Systems. In *Proc. of ACM CCS*.
- [55] Joshua Lind, Christian Priebe, Divya Muthukumar, Dan O’Keeffe, Pierre-Louis Aublin, Florian Kelbert, Tobias Reiher, David Goltzsche, David Eysers, Rüdiger Kapitza, et al. 2017. Glamdring: automatic application partitioning for intel SGX. In *Proc. of USENIX ATC*.
- [56] 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 *Proc. of ACM HASP*.
- [57] Luca Melis, Hassan Jameel Asghar, Emiliano De Cristofaro, and Mohamed Ali Kaafar. 2016. Private Processing of Outsourced Network Functions: Feasibility and Constructions. In *Proc. of ACM SDN-NFV Security*.
- [58] David Mills. 1985. Network Time Protocol. Online at: <https://tools.ietf.org/html/rfc958>.
- [59] P. Mishra, R. Poddar, J. Chen, A. Chies, and R. A. Popa. 2018. Oblix: An Efficient Oblivious Search Index. In *Proc. of IEEE S&P*.
- [60] Ahmad Moghimi, Gorka Irazoqui, and Thomas Eisenbarth. 2017. Cachezoom: How SGX amplifies the power of cache attacks. In *Proc. of CHES*.
- [61] David Naylor, Richard Li, Christos Gkantsidis, Thomas Karagiannis, and Peter Steenkiste. 2017. And Then There Were More: Secure Communication for More Than Two Parties. In *Proc. of ACM CoNEXT*.
- [62] David Naylor, Richard Li, Christos Gkantsidis, Thomas Karagiannis, and Peter Steenkiste. 2017. And Then There Were More: Secure Communication for More Than Two Parties. In *Proc. of ACM CoNEXT*.
- [63] Meni Orenbach, Pavel Lifshits, Marina Minkin, and Mark Silberstein. 2017. Eleos: ExitLess OS Services for SGX Enclaves. In *Proc. of ACM EuroSys*.
- [64] Aurojit Panda, Sangjin Han, Keon Jang, Melvin Walls, Sylvia Ratnasamy, and Scott Shenker. 2016. NetBricks: Taking the V out of NFV. In *Proc. of USENIX OSDI*.
- [65] Parveen Patel, Deepak Bansal, Lihua Yuan, Ashwin Murthy, Albert Greenberg, David A. Maltz, Randy Kern, Hemant Kumar, Marios Zikos, Hongyu Wu, Changhoon Kim, and Naveen Karri. 2013. Ananta: Cloud Scale Load Balancing. In *Proc. of ACM SIGCOMM*.
- [66] Rishabh Poddar, Chang Lan, Raluca Ada Popa, and Sylvia Ratnasamy. 2018. SafeBricks: Shielding Network Functions in the Cloud. In *Proc. of USENIX NSDI*.
- [67] HAProxy Project. 2018. The reliable, high performance tcp/http load balancer. Online at: <http://www.haproxy.org/>.
- [68] Andrew Reed and Michael Kranch. 2017. Identifying HTTPS-Protected Netflix Videos in Real-Time. In *Proc. of ACM CODASPY*.

- [69] Luigi Rizzo. 2012. netmap: A Novel Framework for Fast Packet I/O. In *Proc. of USENIX ATC*.
- [70] Alan Rusbridger. 2013. The Snowden Leaks and the Public. *The New York Review of Books* (2013).
- [71] Satori. 2017. Fast multi-core TCP and WebSockets load generator. Online at: <https://github.com/machinezone/tcpkali>.
- [72] Felix Schuster, Manuel Costa, Cédric Fournet, Christos Gkantsidis, Marcus Peinado, Gloria Mainar-Ruiz, and Mark Russinovich. 2015. VC3: Trustworthy data analytics in the cloud using SGX. In *Proc. of IEEE S&P*.
- [73] Vyas Sekar, Norbert Egi, Sylvia Ratnasamy, Michael K Reiter, and Guangyu Shi. 2012. Design and implementation of a consolidated middlebox architecture. In *Proc. of USENIX NSDI*.
- [74] Jaebaek Seo, Byouyoung Lee, Seongmin Kim, Ming-Wei Shih, Insik Shin, Dongsu Han, and Taesoo Kim. 2017. SGX-Shield: Enabling address space layout randomization for SGX programs. In *Proc. of NDSS*.
- [75] Fahad Shaon, Murat Kantarcioglu, Zhiqiang Lin, and Latifur Khan. 2017. SGX-BigMatrix: A Practical Encrypted Data Analytic Framework With Trusted Processors. In *Proc. of ACM CCS*.
- [76] Justine Sherry, Shaddi Hasan, Colin Scott, Arvind Krishnamurthy, Sylvia Ratnasamy, and Vyas Sekar. 2012. Making middleboxes someone else's problem: network processing as a cloud service. In *Proc. of ACM SIGCOMM*.
- [77] Justine Sherry, Chang Lan, Raluca Ada Popa, and Sylvia Ratnasamy. 2015. Blind-Box: Deep Packet Inspection for Encrypted Traffic. In *Proc. of ACM SIGCOMM*.
- [78] Ming-Wei Shih, Mohan Kumar, Taesoo Kim, and Ada Gavrilovska. 2016. S-NFV: securing NFV states by using SGX. In *Proc. of ACM SDN-NFV Security*.
- [79] Shweta Shinde, Zheng Leong Chua, Viswesh Narayanan, and Prateek Saxena. 2016. Preventing page faults from telling your secrets. In *Proc. of AsiaCCS*. ACM.
- [80] Payap Sirinam, Mohsen Imani, Marc Juarez, and Matthew Wright. 2018. Deep Fingerprinting: Undermining Website Fingerprinting Defenses with Deep Learning. In *Proc. of ACM CCS*.
- [81] Snort. 2018. Network Intrusion Detection & Prevention System. Online at: <https://www.snort.org/>.
- [82] Snort. 2018. Rule Subscriptions. Online at: https://www.snort.org/products#rule_subscriptions.
- [83] Chen Sun, Jun Bi, Zhilong Zheng, Heng Yu, and Hongxin Hu. 2017. NFP: Enabling Network Function Parallelism in NFV. In *Proc. of ACM SIGCOMM*.
- [84] Meysam Taassori, Ali Shafiee, and Rajeev Balasubramonian. 2018. VAULT: Reducing Paging Overheads in SGX with Efficient Integrity Verification Structures. In *Proc. of ACM ASPLOS*.
- [85] Hongliang Tian, Qiong Zhang, Shoumeng Yan, Alex Rudnitsky, Liron Shacham, Ron Yariv, and Noam Milshten. 2018. Switchless Calls Made Practical in Intel SGX. In *Proc. of ACM SysTEX*.
- [86] Bohdan Trach, Alfred Krohmer, Franz Gregor, Sergei Arnaoutov, Pramod Bhatotia, and Christof Fetzer. 2018. ShieldBox: Secure Middleboxes using Shielded Execution. In *ACM Proc. of SOSR*.
- [87] Chia-che Tsai, Donald E Porter, and Mona Vij. 2017. Graphene-SGX: A Practical Library OS for Unmodified Applications on SGX. In *Proc. of USENIX ATC*.
- [88] Nirvan Tyagi, Yossi Gilad, Derek Leung, Matei Zaharia, and Nickolai Zeldovich. 2017. Stadium: A Distributed Metadata-Private Messaging System. In *Proc. of ACM SOSP*.
- [89] Jelle Van Den Hooff, David Lazar, Matei Zaharia, and Nickolai Zeldovich. 2015. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *Proc. of ACM SOSP*.
- [90] Michael Walfish, Jeremy Stribling, Maxwell Krohn, Hari Balakrishnan, Robert Morris, and Scott Shenker. 2004. Middleboxes No Longer Considered Harmful. In *Proc. of USENIX OSDI*.
- [91] Tao Wang, Xiang Cai, Rishab Nithyanand, Rob Johnson, and Ian Goldberg. 2014. Effective Attacks and Provable Defenses for Website Fingerprinting. In *Proc. of USENIX Security*.
- [92] Tao Wang and Ian Goldberg. 2017. Walkie-Talkie: An Efficient Defense Against Passive Website Fingerprinting Attacks. In *Proc. of USENIX Security*.
- [93] Wenhao Wang, Guoxing Chen, Xiaorui Pan, Yinqian Zhang, Xiaofeng Wang, Vincent Bindschaedler, Haixu Tang, and Carl A Gunter. 2017. Leaky cauldron on the dark land: Understanding memory side-channel hazards in sgx. In *Proc. of ACM CCS*.
- [94] Nico Weichbrodt, Anil Kurmus, Peter Pietzuch, and Rüdiger Kapitza. 2016. AsyncShock: Exploiting synchronisation bugs in Intel SGX enclaves. In *Proc. of ESORICS*.
- [95] Ofir Weisse, Valeria Bertacco, and Todd Austin. 2017. Regaining Lost Cycles with HotCalls: A Fast Interface for SGX Secure Enclaves. In *Proc. of ACM ISCA*.
- [96] Andrew M. White, Austin R. Matthews, Kevin Z. Snow, and Fabian Monrose. 2011. Phonotactic Reconstruction of Encrypted VoIP Conversations: Hookt on Fon-iks. In *Proc. of IEEE S&P*.
- [97] wolfSSL Inc. 2017. wolfSSL. Online at: <https://www.wolfssl.com/wolfssl-with-intel-sgx/>.
- [98] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. 2015. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *Proc. of*




Figure 20: Service function chain connected by etap's.

IEEE S&P.

- [99] Xingliang Yuan, Xinyu Wang, Jianxiong Lin, and Cong Wang. 2016. Privacy-preserving Deep Packet Inspection in Outsourced Middleboxes. In *Proc. of IEEE INFOCOM*.
- [100] Fan Zhang. 2017. mbedtls-SGX. Online at: <https://github.com/bl4ck5un/mbedtls-SGX>.
- [101] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. 2016. Towncrier: An authenticated data feed for smart contracts. In *Proc. of ACM CCS*.

A MULTI-THREADING SUPPORT

RSS Emulation for etap. Many middleboxes utilize multi-threading to achieve high throughput [22, 44, 45, 73]. The standard parallel architecture used by them relies on receiver-side scaling (RSS) or equivalent software approaches to distribute traffic into multiple queues by flows. Each flow is processed in its entirety by one single thread without affecting the others. We equip etap with an emulation of this NIC feature to cater for multi-threaded middleboxes.

With the emulation, multiple RX rings will be created by etap, and each middlebox thread is binded to one RX ring. The core driver will hash the 5-tuple to decide which ring to push a packet, and the poll driver will only read packets from the ring binded to the calling thread. As the number of rings increases, the size of each ring should be kept small to avoid excessive enclave memory consumption. Note that we have discussed practical ring size in Section 6.

Multi-threaded State Management. The RSS mechanism ensures that each flow is processed in isolation to others. For a multi-threaded middlebox, we assign each thread a separate set of `flow_cache`, `lkup_table`, and `flow_store`. There is no intersection between the sets, and thus all threads can perform flow tracking simultaneously without data racing. Note that compared to the single-threaded case, this partition scheme does not change memory usage in managing the same number of flows.

B EXTENSION OF SERVICE MODEL

To clearly lay out the core designs of LightBox, so far we have focused on a basic service model. That is, a single middlebox, and a single service provider hosting the middlebox service. Now we discuss how some other typical scenarios can be readily supported.

Service function chaining. Sometimes multiple logical middleboxes are chained together to process network traffic, which is commonly referred to as service function chaining [43, 83]. This service model is also considered in two recent systems for secure middlebox outsourcing. SafeBricks [66] chains the middleboxes within the same enclave, and isolates them by enforcing least privilege on each. In comparison, ShieldBox [86] chains the middleboxes with different enclaves run by different processes on the same physical machine. Both designs, however, run the chain on a single

physical machine that has a limited EPC size, and do not consider the resource-demanding stateful middlebox. Practical execution of a single stateful middlebox in the enclave is already a non-trivial task — what we strive to achieve in this paper — let alone running multiple enclaved stateful middleboxes on the same machine, where severe performance issue is almost inevitable.

To this end, we propose to drive each middlebox in the chain with a LightBox instance on a separate physical machine. Along the chain, one instance's `etap` will be simultaneously peered with previous and next instance's `etap` (or the `etap-cli` at the gateway). Now each `etap`'s core driver will effectively forward the encrypted traffic stream to the next `etap`. This way, each middlebox in the chain can access packet at line rate and run at its full speed. Note that the secure bootstrapping should be adjusted accordingly. In particular, the network administrator needs to attest each LightBox, and provision it with proper peer information.

Disjoint service providers. Middlebox outsourcing may span a disjoint set of service providers. A primary one may provide the networking and computing platform, yet others (e.g., professional cybersecurity companies) can provide bespoke middlebox functions and/or processing rules. Such service market segmentation calls for finer control over the composition of the security services.

The SGX attestation utility enables any participant of the joint service to attest enclaves on the primary service provider's platform. Therefore, they can securely provision their proprietary code/rule set to a trusted bootstrapping enclave. The code is then compiled in the bootstrapping enclave, and together with the rules, provisioned to LightBox enclave. Such on-the-fly compilation of private code in the enclave is first described in [72]. In [66], it is applied to bootstrap secure middleboxes with different network function vendors. We refer interested readers to [66] for more details.