

# Authenticated Key-Value Stores with Hardware Enclaves

Yuzhe (Richard) Tang  
Syracuse University  
Syracuse, NY USA  
ytang100@syr.edu

Ju Chen  
Syracuse University  
Syracuse, NY USA  
jchen133@syr.edu

Kai Li  
Syracuse University  
Syracuse, NY USA  
kli111@syr.edu

Jianliang Xu  
Hong Kong Baptist University  
Kowloon Tong, Hong Kong  
xujl@comp.hkbu.edu.hk

Qi Zhang  
IBM Research, TJ Watson  
Yorktown Heights, NY USA  
Q.Zhang@ibm.com

## ABSTRACT

Authenticated data storage on an untrusted platform is an important computing paradigm for cloud applications ranging from data outsourcing, to cryptocurrency and general transparency logs. These modern applications increasingly feature update-intensive workloads, whereas existing authenticated data structures (ADSs) designed with in-place updates are inefficient to handle such workloads. In this paper, we address this issue and propose a novel authenticated log-structured merge tree (eLSM) based key-value store by leveraging Intel SGX enclaves.

We present a system design that runs the code of eLSM store inside enclave. To circumvent the limited enclave memory (128 MB with the latest Intel CPUs), we propose to place the memory buffer of the eLSM store outside the enclave and protect the buffer using a new authenticated data structure by digesting individual LSM-tree levels. We design protocols to support query authentication in data integrity, completeness (under range queries), and freshness. The proof in our protocol is made small by including only the Merkle proofs at selective levels.

We implement eLSM on top of Google LevelDB and Facebook RocksDB with minimal code change and performance interference. We evaluate the performance of eLSM under the YCSB workload benchmark and show a performance advantage of up to 4.5X speedup.

## 1. INTRODUCTION

The practice of outsourcing data storage to the public cloud has been emerging and gaining popularity in security applications. For instance, using Amazon S3 to store the Bitcoin transactions (in cryptocurrency applications) or to host general transparency logs (e.g., in Google Certificate Transparency [1, 2]) can significantly bring down the operational cost of these security-centric schemes and are adopted in practice [3]. The modern security applications features user-generated content continuously generated in an intensive data-write stream (e.g., cryptocurrency transactions and user-requested certificates; see § 3.1 for details).

Data authenticity is the primary concern for storing and serving data on the untrusted cloud services which are constantly caught compromised in the real world. To guarantee data authenticity, a common approach is to employ the protocols of authenticated data

structures (ADSs) between an untrusted cloud server and trusted clients (i.e., data owner and query users). However, existing ADS schemes [4, 5, 6, 7, 8, 9, 10, 11, 12, 13] have several major limitations. First, they are designed based on update-in-place structures (i.e., requiring excessive communication and large proofs between the data owner and the cloud for updating the ADS), leading to known inefficiency problems in handling data updates [14, 15]. Second, existing schemes require the query users to verify the proof of results obtained from the cloud, which incurs high bandwidth and computation overheads at the query clients.

To address these limitations, in this paper, we leverage off-the-shelf hardware enclaves, in particular Intel Software Guard eXtension (SGX) [16], and propose a novel authenticated key-value store based on LSM trees. The motivation of our design is two-fold: 1. (Why LSM tree?) An LSM tree (log-structured merge tree) is a data structure that supports append-only writes and random-access data reads. Periodically, it conducts a batch operation, called *COMPACTION*, that reorganizes the data layout for better read performance in the future. By this design, an LSM tree has performance advantages in serving high-speed write streams and is widely adopted as the external-memory index structure in many modern storage systems including Google's BigTable [17]/LevelDB [18], Facebook RocksDB [19], Apache HBase [20], Apache Cassandra [21] (see § 2 for details). 2. (Why Intel SGX?) Without a trusted party, realizing an authenticated LSM tree on the cloud requires using costly cryptographic protocols, such as verifiable computations (VC [22, 23, 24, 25, 26]), to support verifiable *COMPACTION*.<sup>1</sup> With the advent of commercial trusted execution environments, notably Intel Software Guard eXtension or SGX [16], it is possible to build a trusted execution environment or enclave in proximity to the untrusted cloud platform. This makes bulk data transfer feasible and should be promising to support verifiable and efficient *COMPACTION* for authenticated LSM trees. In addition, the query users can be alleviated from the burden of result verification.

In our envisioned architecture, trusted cloud applications (e.g., a database server) run inside SGX enclaves and issue data read/write requests to our authenticated key-value store that is co-located in the cloud. In our system design, the code (namely the codebase of a vanilla LSM store<sup>2</sup>) is placed inside the enclave, which relies on the Intel SGX SDK [27] or an in-enclave Library OS (LibOS) [28, 29] to handle system calls. In terms of placing data (e.g., program states), a naive design is to store the data in the memory region inside the enclave. When handling data of Gigabytes, this design, however, imposes huge memory pressure inside enclave and would cause significant performance slowdown. To be more specific, the

<sup>1</sup>Another possible approach is to transfer the whole dataset back to trusted data owners over the Internet, which is also expensive.

<sup>2</sup>“LSM store” denotes the class of key-value stores designed based on LSM trees.

current family of Intel CPUs support 128 MB physical memory in the enclave, and when the enclave memory hosts more than 128 MB, it causes expensive enclave paging [30]. In our preliminary performance study, the overhead causes slowdown of more than two orders of magnitudes (See § 4.2 for details). More fundamentally, the slowdown is caused by the security needs to protect enclave memory. Even if Intel may remove the hard memory limits of 128 MB in future releases, putting the large data in a (larger) enclave still incurs the large slowdown by memory protection.

To circumvent the inefficiency, we propose to place the memory data outside the enclave. More precisely, among various memory data structures in an LSM store, we place the read buffer outside the enclave and leave other structures that often grow sublinearly with the data size inside the enclave, including index structures (e.g., a bloom filter), write buffer, etc. To ensure the integrity of the data placed outside the enclave, we propose an authenticated LSM tree, named by eLSM. eLSM builds a forest of Merkle trees, each digesting a “level” in an LSM tree (see the preliminary of an LSM tree in § 2). eLSM supports efficient reads and small-sized query proofs by presenting Merkle proofs at selective levels. We have proved the security of the query authentication schemes in eLSM (in § 5.3.1).

We also build the eLSM systems on Google LevelDB [18] and Facebook RocksDB [19]. In our systems, the eLSM Merkle proofs are embedded in individual data records in such a way that the proof of a query can be naturally constructed from the Merkle proofs embedded in the data records included in the query result. By this means, we minimize the code change needed in Google LevelDB, reducing performance interference at runtime. For RocksDB, the eLSM system is implemented as a middleware that does not require code change of the underlying RocksDB, but instead just relies on its callback interface [31]. More specifically, we implement authenticated COMPACT in some event handlers in the COMPACT path of RocksDB. With this, we believe the add-on design of eLSM is generally applicable to any LSM stores. By contrast, existing work in the field, notably Speicher [32], all requires significant code change of underlying LSM store (see § 7 for details).

We conduct a comprehensive performance study of eLSM under the YCSB workload benchmark [33]. The performance result shows that eLSM achieves lower operation latency than the baseline of update-in-place data structures by more than one order of magnitudes. Comparing with the eLSM design with memory buffers in the enclave, the memory placement outside the enclave achieves up to 4.5X speedup in most YCSB workloads.

The contributions made in this paper include the following:

1. This work addresses an emerging security workload, that is, supporting query authentication in the presence of frequent data updates. We propose a novel SGX-enabled authenticated key-value store.
2. We present the system designs of eLSM that are secure, efficient and generic. It places the memory data outside the enclave to circumvent the limited memory size in the enclave. It builds an authenticated LSM tree with small query proofs at selective tree levels. To the best of our knowledge, eLSM is the first data-authentication middleware on LSM stores, without any code change of the underlying store.
3. We implemented functional prototypes of eLSM on Google LevelDB and Facebook RocksDB. The code of our prototype is open-sourced [34]. We conducted a comprehensive performance study under the YCSB workload benchmark that shows up to 4.5X performance advantage.

## 2. PRELIMINARY: LSM TREE-BASED STORAGE SYSTEMS

This section presents the preliminaries of LSM trees and stores. Readers who are familiar with the details of LSM stores may skip this section. Additional background on Intel SGX is deferred to Appendix

§ A.

**Data structure:** A log-structured merge tree or LSM tree [35] is a data structure that organizes a dataset by so-called levels. A level is a collection of key-value records that are written in a time of proximity. Inside a level, it stores a sorted run of records, first ordered by data keys and then by time. Upon a write, an LSM tree stores it in the first level. The first level is thus a mutable data structure updated by individual data writes. Upon a read, it searches the tree by iterating through levels and finds the record that matches the queried data key. Periodically, normally at an offline time, the LSM tree runs a COMPACT operation that merges sorted runs in adjacent levels, to make room in lower levels<sup>3</sup> for upcoming writes and to facilitate reads in the future. An LSM tree consists of totally  $q$  levels, which are  $L_0, L_1, \dots$

**Storage systems:** LSM trees have been recently adopted in the design of many modern storage systems, including Google BigTable [17]/LevelDB [18], Apache HBase [20], Apache Cassandra [21], Facebook RocksDB [19], etc. In these systems, the LSM tree is adopted as an external-memory data structure that manages disk IOs. Specifically, the LSM tree buffers the first-level data in memory (stored in the MemTable structure) and, through the COMPACT operation, stores data at higher levels on disk. The data stored in MemTable is backed up by a log file on disk, named Write-Ahead Log or WAL. Data records are written to disk in a large data unit (e.g., several megabytes). Each data unit is persisted in a so-called SSTable file (Sorted Strings Table [17, 18, 19]). By this means, random-access writes are buffered in memory and the system causes only sequential writes to disk.

The LSM tree is also used as a primary index in the key-value store system. An SSTable is a file consisting of multiple data blocks. To support fast data reads in an SSTable, there is a compact B+ tree that indexes the different data blocks in the SSTable. In addition, a Bloom filter is built for each data block that indexes the records in the block. A Bloom filter can facilitate the case when the queried key is not found in the block.

In terms of performance, the LSM-tree based storage design represents a middle ground between the two classic designs, that is, the read-optimized update-in-place storage (e.g., B+ tree and many database indices [36, 37]) and the write-optimized log-structured storage (e.g., log-structured file systems [38]). On the one hand, an LSM tree (in an external memory model) serves data writes in an append-only fashion, in a way similar to log-structured file systems. On the other hand, it supports random-access reads without scanning the entire dataset, which is similar to update-in-place style B+ trees. An LSM tree reaps the benefits from both worlds, at the expense of assuming some offline hours to do the batched compaction operation.

## 3. RESEARCH FORMULATION

In this section, we present the research formulation by describing the motivating application and workloads, system model, security goals and design rationale of eLSM.

### 3.1 Motivating Scenarios

We present in depth a motivating scenario of our work, that is, general transparency. Other motivating applications, such as outsourced cloud databases, are explained in Appendix § B.

**Application scenario:** General transparency becomes a popular design paradigm for building trusted computing systems [39, 40], where internal events of a target system are recorded in and exposed to a publicly-auditable log. The transparency design is being adopted in many operational systems in the real world, including certificate transparency (or CT) [1, 41, 42], key transparency [2], and public Blockchains [43, 44]. In particular, CT exposes the certificates issued by known Certificate Authorities (CA) to the public

<sup>3</sup>Levels in an LSM tree are indexed and the levels with smaller indexes are lower levels.

in order for timely detection of mis-issued certificates [45, 46]. The Blockchain records buyer-seller transactions in a public ledger and allows anyone to audit the transaction history for assurance of no invalid transactions.

**The application workload** features 1) an intensive stream of small-data writes and 2) data reads that retrieve a selective part of the log for auditing and cause random-access to the storage medium. For instance, in Blockchain, transactions (i.e., small data writes) are continuously submitted by a large number of buyers/sellers and new Blockchain nodes who join the network as a simplified payment verification client (i.e., SPV client [47]) randomly accesses part of the transaction history (i.e., random-access reads). In the case of CT log, certificate registration requests (i.e., small data writes) are continuously submitted by a large number of domain owners (in the low-cost CA schemes such as Let’s Encrypt [48]). Also, the log is randomly accessed by web browsers running TLS handshakes which want to validate the certificates being used.

The write-intensive workload excludes conventional database storage engines that are designed based on update-in-place data structures (e.g., B+ trees). In update-in-place structures, a data update needs to overwrite the previous version of the record at the exact location where the record is stored. An update incurs lookups and random-accesses of the record’s previous location, leading to disk seeks and write amplification.

LSM stores lend themselves to serving the write-intensive workload, thanks to the design of append-only writes. Due to this reason, many LSM stores are adopted in practice. For cryptocurrency and Blockchain, Google’s LSM store LevelDB [18] is widely adopted in the software stack of many Blockchain clients including Bitcoin core [49], Ethereum [50], HyperLedger Fabric [51], multichain/stream [52], and other crypto-currencies. LevelDB [18] is used in Google CT systems [42] and RocksDB for the key transparency systems [53, 54, 2].

However, in these real-world uses of LSM stores, the handling of data integrity is unsatisfactory. On the one hand, data integrity is a critical security property in our target applications.<sup>4</sup> For instance, in CT, returning a revoked certificate (violating data freshness) may connect a user to an impersonator, leading to security breaches. On the other hand, most existing transparency systems protect data integrity by replicating data across many nodes. This design while decentralizes trusts incurs high overhead. Using SGX enclave is a much more lightweight approach. Thus, in this work, we are motivated to harden the data integrity in LSM stores by SGX enclaves.

### 3.2 System Model

In our system model, a trusted application runs inside the enclave on an untrusted cloud platform. The trusted application manages security-sensitive big-data and relies on a key-value store interface in enclave for data storage. The trusted application emits a read/write workload with characteristic described above (i.e., write-intensive workloads). The interface between the trusted application and data storage is a standard one in key-value store, described in Equation 1. Given a key-value record  $\langle k, v \rangle$ , a write request is  $PUT(k, v)$  and a read is  $v = GET(k)$ . Here, we assume the enclave runs a timestamp manager that assigns to a read/write operation a unique timestamp reflecting the real time.

$$\begin{aligned} ts &= PUT(k, v) \\ \langle k, v, ts \rangle &= GET(k, ts_q) \\ \{\langle k, v, ts \rangle\} &= SCAN(k_1, k_2, ts_q) \end{aligned} \quad (1)$$

The trusted application can run in multiple threads and issues the PUT/GET operations concurrently.

<sup>4</sup>By contrast, data confidentiality is less important in general transparency systems.

### 3.3 Security Goals

The security goal of this work is primarily on query authenticity, while we also address data confidentiality in § 5.6.2 and § 4.1. The query authenticity describes that given a dataset, whether the result of a read reflects the latest state of the dataset. To formally describe our security goal, we present the threats and security definition below.

**Threats:** The adversary in our model is the untrusted host outside the enclave. The adversarial host runs operating systems and the instance of LSM store. She can mount attacks to present forged query results to the enclave. Specifically, given a GET request, a malicious host can forge a fake result (breaking query-result integrity), or present a stale record (violating query freshness), or skip a legitimate record (violating query completeness). The definitions of these query-authenticity properties are described below.

**Security definitions:** Given a read  $\langle k, v, ts \rangle = GET(k, ts_q)$ , the query authenticity includes various correctness properties: 1) Query integrity is about whether the read-result  $\langle k, v, ts \rangle$  is a key-value record written by a legitimate write request before. If the read result is not written by the data owner, it violates the query integrity. 2) Query completeness is about whether a read result is complete. In case of point query (w.r.t. range query), the completeness is about the membership of a result and it prevents a legitimate record from being excluded in the result. For instance, if the store has a matching record to the read but it returns an empty result, the query completeness<sup>5</sup> is violated. 3) Query freshness states whether the result  $\langle k, v, ts \rangle$  has the largest timestamp (or is the latest) among all records of the queried key  $k$  and with a timestamp smaller/earlier than  $ts_q$ .

### 3.4 Design Motivation: Why LSM Tree based Digest Structure?

To support query authentication (with membership), the conventional approach is building a single Merkle tree over the entire dataset and updating the Merkle tree “in place” upon data updates. This update-in-place Merkle tree design is widely adopted in real-world systems such as digesting state in Ethereum [44].

To serve write-intensive workloads, however, the update-in-place digesting approach would incur high performance overhead. Briefly, with small data records (e.g., tens or hundreds of bytes as in a “tweets”), the hash digests which themselves are tens of bytes can grow comparable with the size of dataset itself. They need to be stored on a large storage medium such as disk. With digests stored on disk, the update-in-place digest structures cause random disk accesses and thus impose high overhead to the write path.

In this work, we present a new digest structure based on an LSM tree, called eLSM digests, where the digests are updated in an append-only fashion. The motivation of this design is based on the well-known fact [55, 56, 57, 58] that the append-only design of an LSM tree leads to performance advantages when serving write-intensive workloads. While existing LSM trees are applied to the data-access path in a storage system, our proposed eLSM presents a new paradigm that applies the LSM tree on the security path, particularly for efficient digest accesses.

## 4. eLSM-P1: A STRAWMAN DESIGN

In this section, we present a straightforward design of LSM store with SGX and name it by eLSM-P1. We then analyze the performance to motivate our primary design eLSM-P2 presented in the next section.

### 4.1 The eLSM-P1 Design

**Design space:** Recall that an LSM store is a user-space system that translates application-level data reads/writes to systems-level file reads/writes. To port an LSM store to SGX, in principle, one can place the code of the LSM store outside the enclave. That is, the LSM

<sup>5</sup>We may use the terms query authenticity and membership alternatively.

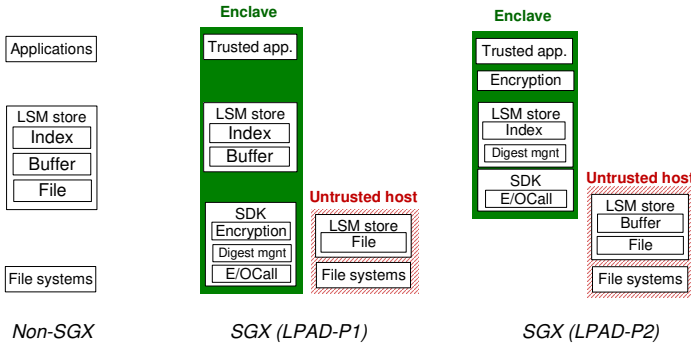


Figure 1: eLSM on SGX: System stack and designs

store runs in the untrusted world and the trusted application in enclave issues OCalls to send PUT/GET requests to the store. However, this code-outside-enclave design requires excluding “dynamically” updated memory data from enclave and incurs design complexity for data authentication. We dismiss this alternative design (more details in Appendix § D). In this work, we focus on building an LSM store with code inside an SGX enclave.

**A strawman:** We consider a strawman design, named eLSM-P1, that places the entire user-space codebase of an LSM store inside the enclave. The SSTable files are stored outside the enclave. The interaction between the enclave and the untrusted host occurs at the syscall levels, primarily for file management (e.g., `fwrite` and `fread`). Specifically, eLSM-P1 places outside the enclave the files at all LSM-tree levels, including the WAL file at level  $L_0$  and SSTables at levels  $L_{\geq 1}$ . Inside the enclave, it runs the codebase of an LSM store and stores the intermediate data including indices and data buffers. The system architecture of eLSM-P1 is depicted in Figure 1.

We implement a functional system prototype of eLSM-P1 on Google’s LevelDB [18]. To port Google LevelDB to SGX, we use Intel SGX SDK and modify the LevelDB codebase to call SDK’s syscalls.

**Security analysis:** eLSM-P1 provides security on data confidentiality and authenticity through SGX SDK’s data-protection mechanism. That is, SDK encrypts and digests the content of SSTable files stored outside the enclave, which ensures that an adversary host cannot extract information from the file content (in the ciphertext format) and forge a file block (whose integrity is verified by the hash proof of file block). Note that we do not address the confidentiality under side-channel attacks in SGX.

## 4.2 Performance Analysis of eLSM-P1

In this subsection, we first present a performance analysis of eLSM-P1 that identifies performance problem serving a large dataset. We then use this observation to motivate our next design, eLSM-P2.

**Observing eLSM-P1’s performance:** We first focus on analyzing the read path of eLSM-P1. Recall that on the read path of an LSM store, the CPU accesses disk-resident data by buffering it in memory. eLSM-P1 places the read buffers inside enclave, which may cause two performance problems: S1) The buffer in enclave incurs an extra data copy in the data-read path, that is, when the CPU accesses the data already buffered inside the untrusted memory, it creates a second data copy inside enclave. S2) When the in-enclave buffer grows large (e.g., beyond 128MB), it causes expensive enclave paging. For the common setting of disk-resident data, having a large read buffer is essential to the performance.

In order to quantify the performance slowdown caused by the in-enclave read buffer placement, we conduct a performance study based on our eLSM-P1 implementation. For comparison, we also implement the placement of read buffer outside enclave. This is done by allocating the read buffer of LSM store in the untrusted memory. In

the performance study, we store 5 GB dataset (larger than untrusted memory) and drive a read-only workload that consists of reads of randomly distributed data keys. We vary the read buffer size 4 MB to 2048 MB. The read latency is reported in Figure 2.

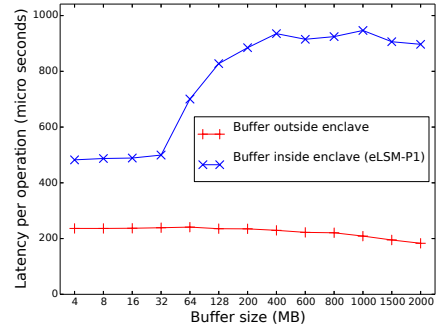


Figure 2: Placing read buffers inside/outside enclave and their performance difference

The results show that when the data buffer is small, the in-enclave buffer incurs 2X latency of the out-enclave buffer. This performance difference is caused by the extra data copy in enclave (S1). When the data buffer is larger beyond 64 MB, the latency of accessing in-enclave buffer grows significantly and it incurs 4.5X latency. The performance slowdown of the large in-enclave buffer is due to the expensive enclave paging. This performance characteristic prompts us to place the memory buffers of an LSM store outside the enclave.

**Motivating eLSM-P2 Design:** The above performance observation suggests a favorable design by placing the read buffer of an LSM store outside the enclave. We propose eLSM-P2 as a holistic system that materializes the design of placing read buffer outside the enclave. The eLSM-P2 system features two placement strategies, that is, placing read buffer outside enclave and placing code and other memory (meta)data inside enclave.

Concretely, in eLSM-P2, the code of an LSM store runs inside the enclave, and it switches the execution out of the enclave only when serving system calls (e.g., file reads/writes). The enclave code accesses the data on the read path outside enclave, which includes read buffers and disk files. Other memory data is placed inside enclave, including the write buffer (at level  $L_0$ ) and file indices (at levels  $L_{\geq 1}$ ). A common characteristic of data inside enclave is that they are meta-data in memory whose sizes are small enough (e.g., with file indices of a couple of megabytes, it can accommodate millions of files) and thus they can be safely placed in enclave without causing enclave paging. eLSM-P2 places only static memory data, that is, read buffers outside enclave, which significantly simplifies the complexity of a correct implementation. Table 1 summarizes design choices made in eLSM-P1 and eLSM-P2.

## 5. eLSM-P2 SYSTEM

Table 1: Design choices of eLSM-P1 and eLSM-P2

	Code placement	Data placement	Digest structure
eLSM-P1 (§ 4.1)	Inside enclave	Inside enclave	File granularity
eLSM-P2 (§ 5)	Inside enclave	Outside enclave	Record granularity

### 5.1 System Overview

Figure 3a depicts the system architecture of eLSM-P2 which runs the code for operations PUT,GET,COMPACTION inside enclave. The memory data including the buffer at Level  $L_0$  and file indices at Levels  $L_{\geq 1}$  are also placed inside the enclave. The read buffers and all SSTable files at Levels  $L_{\geq 1}$  are placed outside the enclave. The

WAL file is also stored outside enclave. The figure also illustrates the dataset in the LSM store, which we will use throughout this section to describe the details of eLSM-P2 system. In this example dataset, there is an LSM tree of three levels and six key-value records. Level  $L_1$  contains record  $\langle A, 9 \rangle$ , level  $L_2$  contains three records  $\langle T, 4 \rangle$ ,  $\langle Z, 7 \rangle$ ,  $\langle Z, 6 \rangle$  and level  $L_3$  contains four records  $\langle A, 2 \rangle$ ,  $\langle T, 0 \rangle$ ,  $\langle Y, 3 \rangle$ ,  $\langle Z, 1 \rangle$ . Here, we show the key-value record by its data key and timestamp. Record  $\langle T, 0 \rangle$  is of data key  $T$  and timestamp 0, which is the oldest record. For simplicity, the data value is omitted in this example.

**Protecting data outside enclave:** Because eLSM-P2 places outside enclave the data at non-zero levels, it entails data protection mechanisms. For data confidentiality, we require the data key in each record to be encrypted with deterministic encryption (DE), such that it can directly search the domain of ciphertext. We discuss the details of data confidentiality in § 5.6.2. For data authenticity, we build the eLSM digest structure to authenticate the data outside enclave, as described next.

## 5.2 Digest Structure

To digest an LSM tree, we propose a novel authenticated data structure, eLSM-P2 digests. There are two key designs: 1) eLSM-P2 builds a “forest” of Merkle trees, each digesting one LSM-tree level and each having its root stored in the enclave. 2) In a per-level Merkle tree, data records of the same key are digested in hash chains and records of different keys are digested in a Merkle tree. In particular, the hash chain is built in a temporal order where the chain header is the oldest record and the tail is the newest record. In Figure 3a, each of the three LSM tree levels is associated with a Merkle tree. Case 1): For a level of distinct data keys, such as level  $L_3$  in Figure 3a, it builds the leaf set of the Merkle tree directly on the data records. For instance,  $h_7 = H(\langle A, 2 \rangle)$  ( $H$  is a standard cryptographic hash algorithm with variable-length input) and  $h_6 = H(\langle T, 0 \rangle)$ . An intermediate node is the hash digest of the concatenation of the two children, for instance,  $h_8 = H(h_6 \| h_7)$ . Case 2) For a level that contains some records of the same keys, it constructs a hash chain over these records. For instance, level  $L_2$  contains two records of the same key,  $\langle Z, 7 \rangle$  and  $\langle Z, 6 \rangle$ . eLSM-P2 builds a *hash chain* on these two records, that is,  $h_4 = H(\langle Z, 7 \rangle \| H(\langle Z, 6 \rangle))$ . Then, it builds the Merkle tree over  $h_4$  (for records of key  $Z$ ) and  $h_2$  (for record  $\langle T, 4 \rangle$ ) for level  $L_2$ .

To materialize the eLSM-P2 digest structure, we present a simple storage design: Given a level  $L_i$  and its Merkle tree, each record at the level  $\langle k, v \rangle$  is augmented with its eLSM-P2 proof  $\pi$ , that is,  $\langle k, v \| \pi \rangle$ . Given a record, an eLSM-P2 proof is the set of Merkle tree nodes (or hashes) that surround the path from the leaf node of the record to the root node. For instance, in Figure 3a, the eLSM-P2 proof for record  $\langle A, 2 \rangle$  consists of hashes  $h_7$  and  $h_{11}$  (which are the siblings to nodes  $h_6$  and  $h_8$ ).

## 5.3 Read/Write Protocol

**Data read path** starts with the trusted application issuing a read operation,  $v = \text{GET}(k)$ . The enclave looks up its index to locate the target level and file, and it then notifies the eLSM-P2 store. **r1** The untrusted store serves the read operation on the target file and, in addition, runs algorithm  $\pi, v = \text{QUERY}_{\text{GET}}(m, k)$  to prepare a proof for authenticating the read result. The proof consists of Merkle authentication paths or Merkle proofs [59] at “relevant” LSM tree levels. Recall that a Merkle authentication path consists of the hashes surrounding the path from a leaf to the root in a Merkle tree and it can be used to verify the membership and non-membership of a record in a dataset. **r2** The eLSM-P2 store in the untrusted host then sends the result of  $\text{GET}(k)$  as well as the proofs ( $\pi$ ) to the enclave. The enclave verifies the authenticity of the read result, by running algorithm  $\text{Yes|No} = \text{VRFY}_{\text{GET}}(\pi, v)$ . The verification algorithm iterates through relevant levels, and, for each level, verifies the membership/non-membership of the queried data key ( $k$ ) using

the Merkle proof (in  $\pi$ ) and locally stored root hash.

A strawman of designing the eLSM-P2 proof is to scan all levels to prepare a proof (in algorithm  $\text{QUERY}$ ). We propose to reduce the proof size by including only Merkle proofs of the levels no higher than the level of the result record. This will allow algorithm  $\text{QUERY}$  to stop early when it reaches the first level, say  $L_i$ , that finds a matching record. The returned proof  $\pi = \pi_1, \dots, \pi_i$ , where  $\pi_1, \dots, \pi_{i-1}$  are the Merkle proofs for non-membership (there is not any matching record at levels  $L_1, L_2, \dots, L_{i-1}$ ).  $\pi_i$  is the Merkle proof for membership (there is a matching record at level  $L_i$ ). All Merkle proofs after level  $L_i$ , as will be seen, do not contain fresher records and are deliberately omitted. When there is no matching record,  $i = q$ .

*An example:* In Figure 3a, suppose the trusted application issues  $\text{GET}(Z)$  over dataset  $m$ . In step **r1**, the untrusted host serves  $\text{QUERY}_{\text{GET}}(m, Z)$  with authentic result  $\langle Z, 7 \rangle$  (the benign case).  $\langle Z, 7 \rangle$  is the newest record matching queried key  $Z$  and is located at level  $L_2$ . The proof is two Merkle authentication paths at levels  $L_1$  and  $L_2$ . Note that there is no need to include level  $L_3$  in the eLSM-P2 proof. Concretely, the proof at the first level is  $\langle A, 9 \rangle$  (denoted by  $\pi_1$ ). The proof at the second level is  $h_3, h_2$  (denoted by  $\pi_2$ ). Then in step **r2**, the enclave can verify the result authenticity in freshness and completeness based on the proof  $\pi = [\pi_1, \pi_2]$  (i.e., algorithm  $\text{VRFY}([\pi_1, \pi_2], \langle Z, 7 \rangle, [h_1, h_5])$ ).

Concretely, with the first-level proof  $\pi_1 = \langle A, 9 \rangle$ , the enclave verifies result authenticity by checking  $H(\pi_1) \stackrel{?}{=} h_5$ . If the  $\text{VRFY}$  algorithm runs through, it authenticates the fact that record  $\langle A, 9 \rangle$  is the only record at level  $L_1$ . From this, it can be derived that level  $L_1$  does not contain any record of key  $Z$  (i.e., the non-membership of a data key  $Z$  at level  $L_1$ ). With the second-level proof  $\pi_2 = h_3, h_2$ , the enclave verifies by checking  $H(h_2 \| H(\langle Z, 7 \rangle \| h_3)) \stackrel{?}{=} h_5$ . If successful, it authenticates the fact that a) record  $\langle Z, 7 \rangle$  is a valid record at level  $L_2$  (result integrity), b) record  $\langle Z, 7 \rangle$  is the newest record with key  $Z$  (result freshness). Fact b) is based on that there are no other records of key  $Z$  in the proof  $\pi_2$ . Based on these two proofs, one can establish that record  $\langle Z, 7 \rangle$  is the newest record in the dataset  $m$ .

Consider the malicious case when the untrusted host can return a stale record, say  $\langle Z, 6 \rangle$ , to the enclave. In this case, the malicious host can only present the following as a valid level- $L_2$  proof, that is,  $\pi'_2 = \langle Z, 7 \rangle, h_2$ . By this means, the enclave can verify the result integrity successfully by checking  $H(h_2 \| H(\langle Z, 7 \rangle \| H(\langle Z, 6 \rangle))) \stackrel{?}{=} h_5$ . However, as the newer result record  $\langle Z, 7 \rangle$  has to be included in the proof  $\pi_2$ , the enclave can detect that  $\langle Z, 6 \rangle$  is not the most fresh record (violating freshness).

**Data write path** starts with the trusted application issuing a write operation  $\text{PUT}(k, v)$ . To serve the write, the enclave maintains two in-enclave structures, a write buffer of level  $L_0$  and, for data recovery, a digest of the write-ahead log (WAL). Recall that a WAL stores recent data writes in temporal order and serves as the base to recover recent data in the case of fault. The storage of WAL is placed outside the enclave, while the enclave stores the hash digests of the WAL.

**w1** Serving the write  $\text{PUT}(k, v)$ , the enclave first assigns to the record to write the latest timestamp  $ts$ . It then writes to the memory buffer of level  $L_0$  inside enclave. Serving a timestamped write  $\text{PUT}(k, v, ts)$ , the enclave iteratively update its WAL digest by  $dig' = H(dig \| \langle k, v, ts \rangle)$ . **w2** When the write buffer at level  $L_0$  overflows, it is triggered to flush the content at Level  $L_0$  and to generate a file at Level  $L_1$ . In the system of an LSM store, the codebase for flush is shared with that for  $\text{COMPACTION}$ . **w3** The enclave switches out to append the write to the WAL in the untrusted domain. Enclave WAL can be extended to defend rollback attacks, which will be described in § 5.6.1.

*An example:* In Figure 3a, suppose the application calls  $\text{PUT}(Y)$ . The enclave assigns to the record the latest timestamp 10. It updates the WAL digest from  $dig$  to  $dig'$ , such that  $dig' = H(dig \| \langle Y, 10 \rangle)$

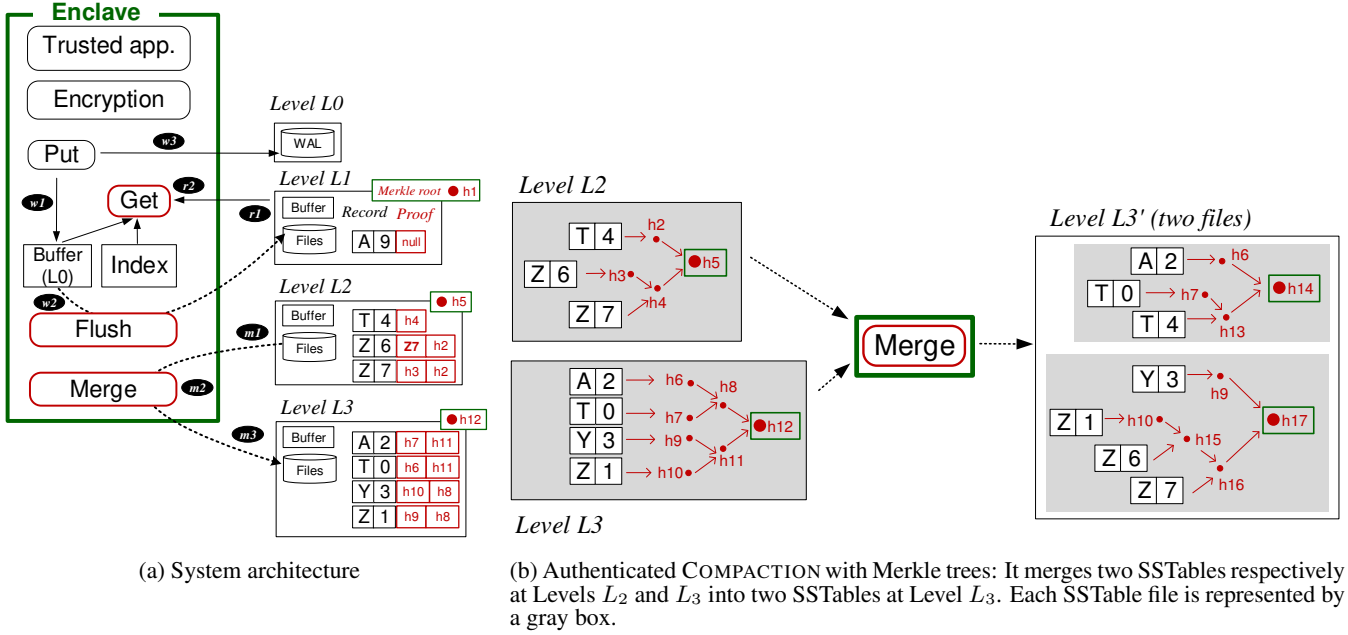


Figure 3: eLSM-P2 system with an LSM tree of three levels: A rectangle depicts data and a rounded rectangle depicts code. Shapes in red are where eLSM-P2 makes code change over the original LSM store. The root hashes (red dots in green boxes) are maintained with copies inside enclave.

(w1). The host appends the record to the WAL outside enclave (w3). If the buffer of Level  $L_0$  is overflowed by the new record, it will sort all records stored in  $L_0$ , and flush them to a new file at  $L_1$  (w2).

**COMPACTON path** starts with the trusted application in enclave issuing operation  $(L'_i, L'_{i+1}) = \text{COMPACTON}(L_i, L_{i+1})$ . For simplicity, we consider the most basic form of COMPACTON, namely, merging two adjacent levels. It is natural to extend it to more complicated cases such as merging more than two levels or merging subsets at the two levels. For the COMPACTON across two levels, eLSM-P2 carries out the computation inside enclave and only switches the execution outside enclave for file access. The process runs in the following steps: (m1) the enclave starts to issue OCalls to load all input files to untrusted memory so that the enclave can read the streams of data records (in their sorted order). (m2) The enclave then runs “authenticated COMPACTON” that merges input data at the two levels into one level. Internally, the enclave needs to verify the authenticity of input data, to conduct the actual computation for COMPACTON, to produce the digest of output data, and to generate the proofs embedded in the output data. We will describe in § 5.5.2 the detailed system design of the authenticated COMPACTON in enclave. (m3) The untrusted host makes effect of the COMPACTON by flushing merged data and proof to disk. The enclave updates the per-level digests by the newly produced ones.

*An example:* In Figure 3a, suppose the application calls  $\text{COMPACTON}(L_2, L_3)$ . In step (m1), the host loads the data at the two levels from disk to memory (in the untrusted world). In step (m2), the enclave verifies the data authenticity of input levels by reconstructing the Merkle tree at level  $L_2$  (and  $L_3$ ) and by checking if its root hash is equal with  $h_5$  (and  $h_{12}$ ). It will then merge the two levels’ data into one merged list, that is, from  $L_2 = [\langle T, 4 \rangle, \langle Z, 7 \rangle, \langle Z, 6 \rangle]$  and  $L_3 = [\langle A, 2 \rangle, \langle T, 0 \rangle, \langle Y, 3 \rangle, \langle Z, 1 \rangle]$  to output level  $L'_3 = [\langle A, 2 \rangle, \langle T, 4 \rangle, \langle T, 0 \rangle, \langle Y, 3 \rangle, \langle Z, 7 \rangle, \langle Z, 6 \rangle, \langle Z, 1 \rangle]$ . Meanwhile, it builds the Merkle tree over the output list, and based on it, generates the proofs embedded in data records. In step (m3), the digest of the new Merkle tree replaces that of level  $L_3$  (i.e.,  $h_{12}$ ).  $L_2$  becomes an empty list and its digest is updated as well.

### 5.3.1 Protocol Analysis

In this subsection we present the security analysis of the eLSM-P2 protocol. We first define the record freshness as below:

**DEFINITION 5.1. (Record freshness)** Given  $\langle k, v, ts \rangle = \text{GET}(k, m)$ , GET returns a fresh result if the record  $\langle k, v, ts \rangle$  is the newest one (i.e., with the largest timestamp) in the dataset  $m$  that matches data key  $k$ .

The protocol correctness is defined as below:

**DEFINITION 5.2. (Protocol correctness)** Given any state of dataset  $m$ , if a GET returns a correct and fresh record w.r.t. state  $m$ , the VRFY algorithm will complete successfully (i.e., returning Yes).

The protocol security is defined as below:

**THEOREM 5.3. (Protocol security)** Given a GET operation, if the enclave completes the VRFY algorithm successfully, the GET result is correct and fresh (as in Definition 5.1).

**Security analysis:** We prove the protocol security by contradiction. Assume there exists an adversary who can forge a proof on a stale (yet correct)<sup>6</sup> record on which the VRFY algorithm completes successfully. Logically, there are two cases for a stale record: 1) The record is stale at the level that contains it, 2) The record is fresh at the level that contains it but there are more fresh records at other levels. In the following, we prove the contradiction that the VRFY algorithm successfully completes on a stale record, respectively in the two cases.

For Case 1), suppose there is another record  $\langle k, v', ts' \rangle$  on the same level with result record  $\langle k, v, ts \rangle$  but fresher, that is,  $ts < ts'$ . The COMPACTON in enclave guarantees<sup>7</sup> that the two records are sorted by timestamps (note that they are of the same data key) and

<sup>6</sup>The record integrity or correctness can be easily authenticated by the cryptographic hashes.

<sup>7</sup>We assume program security in this work that the enclave program does not have exploitable bugs or run malware.

the eLSM-P2 builds a hash chain over them. More specifically,  $H(\langle k, v', ts' \rangle \| \dots \| H(\langle k, v, ts \rangle \| \dots))$ . In other words, the fresh record  $\langle k, v', ts' \rangle$  is among the neighbors of the path from the stale record  $\langle k, v, ts \rangle$  to the Merkle root node (think for an example,  $\langle Z, 7 \rangle$  is the neighbor of the path from  $\langle Z, 6 \rangle$  to root as in Figure 3b). In order for the adversary to forge a proof that can be verified successfully, she either breaks the security of cryptographic hashes (1a) or has to include *all* the neighbors of the path from itself to the root (1b). The former case (1a) is hard to any computationally bounded adversary (more formally, probabilistic polynomial time adversary). In the latter case (1b), the fresher record included in the neighbors is exposed to the enclave who will not pass the freshness check in the VRFY algorithm as it can simply detect that the result record is not fresh (e.g.,  $\langle Z, 7 \rangle$  has a larger timestamp and is fresher than  $\langle Z, 6 \rangle$ ). Therefore, the VRFY algorithm cannot complete successfully on a stale record. A contradiction is found in Case 1.

For Case 2), suppose the result record  $\langle k, v, ts \rangle$  resides at level  $L_i$  and there exists a fresher record  $\langle k, v', ts' \rangle$  stored at another level  $L_j$ . There are two cases: 2a)  $j < i$  and 2b)  $j > i$ . For Case 2a), successfully passing VRFY algorithm in eLSM-P2, it requires a non-membership proof of data key  $k$  at level  $L_j$ . This contradicts the assumption that  $L_j$  does contain a record of key  $k$ , namely  $\langle k, v', ts' \rangle$ . For Case 2b), eLSM-P2 does not require including any proof for the levels with index higher than  $i$  (recall that given a GET result residing at level  $L_i$ , an eLSM-P2 proof includes Merkle proofs for levels  $L_1, L_2, \dots, L_i$  but excludes the ones on any higher levels  $L_{i+1}, L_{i+2}, \dots$ ). The following lemma guarantees that those levels higher than  $L_i$  (i.e., with larger index values) cannot include any records fresher than  $\langle k, v, ts \rangle$ .

**LEMMA 5.4.** *In eLSM-P2 with the in-enclave COMPACTION, a record residing at a lower level must have a larger timestamp than any record of the same data key at a higher level. That is, given any two records of the same data key, say  $\langle k, v, ts \rangle$  residing at level  $L_i$  and  $\langle k, v', ts' \rangle$  at level  $L_{i'}$ , it holds that  $i < i'$  if and only if  $ts > ts'$ .*

For instance, in Figure 3a, at any level, say  $L_3$ , records are sorted, from key  $A$  to  $T$  to  $Z$ . Lemma 5.4 requires that in Figure 3a, an older record  $A$  with timestamp 2 is stored on a higher level  $L_3$  than the level a newer record  $\langle A, 9 \rangle$  is stored (which is level  $L_1$ ). To prove Lemma 5.4, the key intuition is that eLSM-P2 only allows moving key-value records from lower levels to higher levels, but not in the reversed order.

**Meta-data authenticity:** In our implementation, meta-data including Bloom filters, file indices, etc. are placed inside enclaves. The metadata authenticity is protected by the enclave.

## 5.4 Range Query and Deletes

**Processing Range Query:** eLSM-P2 supports range query processing with completeness. Given a queried key range, it iterates through *all* levels (unlike the case of exact-match query) and for each level, the untrusted world presents the range-query proof from the Merkle tree at the level.

Within one level, the range proof is constructed by treating the Merkle tree as a segment tree which is a classic data structure in computational geometry. A segment tree is essentially a full binary tree where each intermediate tree node represents a segment (or interval of data keys). Given a range query  $L$ , the segment tree can present  $2 \log L$  intermediate nodes, whose segments are union-ed to cover range  $L$  [60]. Based on this view, eLSM constructs a range proof by the sibling nodes of the segments covering the queried range.

For instance, suppose in Figure 3, a range query is  $\text{SCAN}([S, U])$  against Level  $L_3$ . The query proof will include 1) records which fall in or enclose the range, that is,  $\langle T, 0 \rangle, \langle Y, 3 \rangle, \langle Z, 1 \rangle$ , 2) the range proof that authenticates these records, that is,  $h_6$ . At level  $L_3$ , the segment union that covers the three records includes intermediate tree nodes  $h_{11}$  and  $h_7$ . Their siblings are  $h_6$  and  $h_8$ . While  $h_8$  can be reconstructed from  $\langle T, 0 \rangle$  and  $h_6$ , it can be omitted, leave the range

proof to be  $h_6$ . In this case, a query verifier can authenticate the query completeness, as records  $\langle T, 0 \rangle, \langle Y, 3 \rangle, \langle Z, 1 \rangle$  are authenticated and they consecutive in the leaf set of the Merkle tree.

**Security Analysis (Range Completeness):** The range-query proof authenticates query completeness. Informally, the security can be derived from the following facts: 1) The hash functions used in Merkle tree are collision resistant. 2) The data records in a level are sorted by data keys. 3) The records in the query proof form a range of data keys that cover the range in the query.

**Handling Deletes:** eLSM relies on the underlying LSM store to support delete operations. Briefly, a delete operation stores a so-called tombstone record in the LSM store, which triggers the next COMPACTION call to physically drop the records of the same key being merged. Before the COMPACTION, a GET operation may return the tombstone record which is interpreted by eLSM as “no record”.

## 5.5 System Implementation

We have implemented eLSM-P2 on Google LevelDB [18] and Facebook RocksDB [19]. For LevelDB, eLSM is implemented by directly changing the LevelDB codebase. This implementation approach has the advantage in performance. For RocksDB, eLSM is implemented as a RocksDB add-on, that is, without code changes to RocksDB. In this subsection, we present the LevelDB-based implementation in § 5.5.1, 5.5.2 and the RocksDB implementation in § 5.5.3.

In the protocol of eLSM-P2, a key-value record is stored with its proof, that is,  $\langle k, v \| \pi_i \rangle$ , where  $i$  is the index of the level where the record is currently located. To implement the embedded proof in an LSM store, it is required to add the code change in two paths, that is, a) the COMPACTION paths for updating records’ proof when they are merged to a different level (Note that the proof is sensitive to which level the record is located), and b) the GET path where the proof is used to authenticate the record membership/non-membership in a level. Next, we present the code change on the GET path in § 5.5.1, and then present that on the COMPACTION path in § 5.5.2.

### 5.5.1 Query Verification in Read Path

The code change on the GET path serves two purposes: First, for the data stored in the untrusted world, eLSM-P2 needs to present non-membership proofs for the levels that do not have a matching record. Instead of returning null for non-membership levels, eLSM-P2 returns the two neighboring records whose keys are smaller and larger than the queried key. For instance, in Figure 3a, when it queries  $\text{GET}(k = B)$  on level  $L_3$ , it returns records  $\langle A, 2 \rangle$  and  $\langle T, 0 \rangle$  (with their eLSM-P2 proofs).

Second, inside the enclave, eLSM-P2 requires implementing the VRFY algorithm to verify the query authenticity given a proof. Note that the proof can be directly extracted from the result record returned from the untrusted part of eLSM-P2.

**Support mmap reads:** LevelDB supports the data reads in two ways: Read through a user-space read buffer and read through mmap’ed files. eLSM implementation over LevelDB supports both read paths in LevelDB. On the buffer-based read path, eLSM-P2 allows the enclave code to access the buffer outside the enclave and let the untrusted code manage the buffer for eviction. On the mmap-based read path, it switches out of the enclave, upon opening a file, to mmap the file to the untrusted memory. Then the enclave code directly accesses the mmap’ed file in the untrusted memory.

### 5.5.2 Authenticated COMPACTION

Recall that eLSM-P2 runs the COMPACTION computation inside the enclave by accessing data from the untrusted world. For data authenticity, the in-enclave COMPACTION needs to be augmented with three extra steps: a) Before the COMPACTION, the enclave authenticates the input data read from the untrusted world. b) After the COMPACTION, the enclave needs to digest the output data and store the digest in the enclave. c) In the end, it generates from the

output Merkle tree the proofs embedded in individual data records.

**Integrating the COMPACTION authentication** with the underlying LSM store (i.e., LevelDB and RocksDB) is realized by event handling. Concretely, the compaction process in an LSM store triggers various internal events. Two events are of particular interest to our implementation, that is, `Filter()` and `OnTableFileCreated()`. Event `Filter()` occurs whenever the compaction produces a data record. Event `OnTableFileCreated()` occurs whenever the compaction (or other procedure like flush) produces a new file on disk. In RocksDB, these two events are exposed in its callback functions (as will be described in details). In LevelDB, we have to modify the LevelDB codebase to expose the events (i.e., adding hooks) to applications. The pseudo-code that implements authenticated COMPACTION by the two callbacks is depicted in Figure 4. eLSM-P2 runs the codebase of LSM store and the implemented callback functions in enclave. Steps a) and b): In `Filter()`, it constructs two types of Merkle trees, respectively for the input files and the output file. Given a key-value record, it parses the file name and level information from the value field and then updates the input Merkle tree at the corresponding level (Line 20 in Figure 4). It also builds an output Merkle tree for the level the compaction output is located (Line 21). When the compaction finishes, it checks the equality of the Merkle root hash stored in enclave and the root hashes of the input Merkle trees reconstructed. If the equality check passes, the Merkle root hash for the output file takes effect (Line 36-39). Step c): In `OnTableFileCreated()`, it embeds the Merkle proof in the individual data records in the output file.

**Merkle tree construction:** Constructing a Merkle tree is a primitive operation in eLSM-P2, specifically for authenticating the input and output of a COMPACTION. The Merkle tree in our scheme is constructed incrementally with the arrival of records in the input/output data stream. That is, when a data record arrives, the enclave checks if the record shares the data key with the previous record. If this is the case, the enclave builds a hash chain on the previous record. Otherwise, it will incrementally build a Merkle tree. For instance, in Figure 3b, upon reading the output record  $\langle T, 4 \rangle$ , the enclave checks if the previous record (i.e.,  $\langle T, 0 \rangle$ ) has the same data key. Since the two records are of the same key, it then builds a hash chain over them. In the next iteration when the output record  $\langle Y, 3 \rangle$  arrives, the previous record (i.e.,  $\langle T, 4 \rangle$ ) has a different data key, it then treats as a leaf of the Merkle tree the hash of previous record, namely  $h_3$ .

Our implementation retains the COMPACTION features in the vanilla LSM store, including versioning policies, tombstone delete, etc.

**Multi-threading:** eLSM-P2 supports concurrent operations in a multi-threaded enclave. For concurrent reads/writes in MemTable at Level  $L_0$ , eLSM-P2 implementation relies on the synchronization support (i.e., mutex and condition variables) in enclave provided by Intel SGX SDK [61]. Concurrent reads on upper levels  $L_{\geq 1}$  can be processed in parallel without synchronization.

Concurrent COMPACTION with reads/writes needs to be synchronized, as is in LevelDB. For instance, when an SSTable produced by a COMPACTION replaces an old SSTable at the same level, the replacement operation needs to be synchronized with any pending read on the old file. This type of synchronization in eLSM-P2 is realized by checks during query verification VRFY, such that a read before/after the file replacement can only be verified successfully with the Merkle hash of the old/new SSTable file. To implement this, the file replacement outside enclave synchronously calls into the enclave (through a blocking ECall) in order to update the file hash maintained inside the enclave. When a GET operation returns (which is potentially executed in parallel with a COMPACTION), the enclave verifies its result against the latest file hash (updated by the COMPACTION) in the enclave. We use an in-enclave mutex to guard the file hash under concurrent hash

```

1 Record prev_r; Hash prev_h;
2 MHT[] input_mhts; MHT output_mht;
3
4 void MHT_add(Record r, MHT mht) {
5     if (prev_r == NULL) h=H(r);
6     else if (r.datakey != prev_r.datakey) {
7         h=H(r);
8         mht.add_leaf(prev_h);
9     } else h=H(prev_h||r);
10    prev_r = r; prev_h = h;}
11
12 void vanilla_compaction(SSTable La, SSTable Lb,
13     int filter(Record a)
14     int onTableFileCreated(SSTable f));
15
16 void auth_filter(Record r) {
17     MHT_add(r, input_mhts[parseFile(r)]);
18     MHT_add(r, output_mht);}
19
20 int auth_onTableFileCreated(SSTable f){
21     SSTable f_new;
22     for(record r : f){
23         r.addEmbeddedProof(output_mht);
24         f_new.write(r);}
25     f = f_new;}
26
27 void auth_compaction(SSTable La, SSTable Lb){
28     vanilla_compaction(La, Lb,
29         auth_filter, auth_onTableFileCreated);
30     //check digest equality on input and then sign
31     if(input_mhts[La].roothash == La.roothash &&
32        input_mhts[Lb].roothash == Lb.roothash)
33        Lb.roothash = output_mht.roothash;}

```

Figure 4: Authenticated COMPACTION in eLSM

updates (from file replacement in a COMPACTION) and hash reads (for verifying a GET operation).

### 5.5.3 Implementation as a RocksDB Add-on

The eLSM-P2 implementation on LevelDB requires the change of LevelDB’s codebase. A more modular approach (with benefits in easy maintenance and deployment, etc.) is to implement eLSM-P2 without the code change of underlying LSM stores. For this reason, we present the second implementation on RocksDB.

Comparing with LevelDB, the system of RocksDB exposes callbacks through which application programs can listen to and handle RocksDB’s internal events. Different RocksDB events exposed occur at the granularity of level, file, record, etc. The callback API is similar to stored procedures (supported in commercial database systems) and is widely available in other LSM stores (e.g., HBase Coprocessor [62]).

The RocksDB-based eLSM is implemented as a series of event handlers. 1) As previously described, the authenticated COMPACTION is implemented as handlers for events `Filter` and `OnTableFileCreated`. The former event is triggered every time the underlying RocksDB encounters a key-value record during a COMPACTION, and the latter is when a new file is written to the disk. 2) The embedded proof on RocksDB is extended to cover not only the Merkle proof at the current file, but also all Merkle proofs at the current level and previous non-hit levels. The Merkle proofs at non-hit levels are for the non-membership. 3) To implement the authenticated flush, it wraps the code for digesting a MemTable in the iterator (i.e., `next()`) exposed by a pluggable MemTable [63].

## 5.6 Discussion

### 5.6.1 Freshness and Rollback Attacks

In a rollback attack, the untrusted host can replace the authenticated data storage with an older but also authenticated version. To detect and defend the rollback attack, we can harden the security of eLSM systems (both eLSM-P1 and eLSM-P2) with a trusted monotonic counter. A trusted monotonic counter provides the state freshness across power cycles. In reality, one can build

a trusted monotonic counter on TPM chips (e.g., the SDK `sgx_create_monotonic_counter` service [64] or others [65, 66] based on Intel Management Engine (ME) [67]) or multiple remote enclaves (e.g., as in ROTE [66]). Given a trusted monotonic counter, the eLSM would periodically write the hash of current dataset across all levels to the counter. This hash can be implemented by the static hash of all non-zero levels (updated only when it merges) and the hash of WAL file. By this means, one can guarantee the freshness of entire dataset across power cycles against rollback attacks. The dataset-wide freshness can derive the query freshness.

To reduce the performance impact, one can allocate a write buffer for logging the monotonic counter. The size of the write buffer is tunable by system administrator.

### 5.6.2 Data Confidentiality

eLSM can be extended to support data confidentiality for applications of the need (e.g., outsourcing sensitive data to the cloud). This is handled by building an encryption layer on eLSM. While the values can be encrypted using the standard semantically-secure schemes (e.g., AES), the data keys need to be encrypted in a searchable fashion. We follow the common approach of using searchable encryption schemes [68], which is adopted in CryptDB and other searchable databases [69, 70]. In our design, the key of a data record is deterministic-encrypted (DE) [71, 72]. The use of DE allows for querying the ciphertext in the untrusted world. In our system implementation, we used the Intel SDK function `SGX_RIJNDAEL128GCM_ENCRYPT` [27]. For range query, one can use Order-Preserving Encryption (OPE) [73, 74] to encrypt the data keys. In existing work, data confidentiality in encrypted key-value stores is treated in an ad-hoc fashion (e.g., Speicher [32] is an encrypted LSM store where the correlation of block index and query is disclosed) or by more formal security protocols, e.g., Oblivious RAM, at a higher cost [75].

## 5.7 Case Studies

Our eLSM can be integrated with existing systems for general transparency and cryptocurrency. Here we focus on describing the case of integrating eLSM with Google’s Certificate Transparency (CT) system [42].

A vanilla CT protocol involves with three actors, log servers, log auditors and log monitors. Log servers collect newly issued certificates (from the certificates authorities or CAs). A log auditor running along with a web browser needs to validate the certificate being used by the browser. Given a certificate, the log auditor queries the log server for a proof of inclusion (or membership) of the certificate in the CT log. A log monitor run by a domain owner needs to monitor all certificates and to detect any certificates misissued under her domain name. A log monitor continuously send queries to the log server and downloads all certificates.

With eLSM, one can build a trustworthy and efficient CT log server, which provides query authenticity and fork prevention, yet without the overhead incurred in the vanilla CT log server. More specifically, it eliminates the cost of running multiple log servers and does not need to gossip. The eLSM scheme can enable lightweight log monitors who only download the certificates of their own domain names, resulting low and sublinear bandwidth. In addition, with eLSM and a secure enclave, one can safely delegate the log-monitoring logic to the enclave, which relieves low-power clients (like web browser running on mobile phones) from the cumbersome auditing work.

We have built a prototype of eLSM-based log server. In our implementation, we download certificates from the Google CT log server<sup>8</sup> and store them in eLSM by key-value records: The hostname of a certificate is used as the data key and the certificate itself (more specifically, the hash of the certificate) is the data value. eLSM receives GET queries from a log auditor and continuously monitors

<sup>8</sup><https://ct.googleapis.com/pilot>

the certificate of interest (e.g., of a particular hostname). eLSM provides query authenticity and freshness which is essential for trusted log auditing.

## 6. PERFORMANCE EVALUATION

In this section, we present the performance evaluation of eLSM-P2 and eLSM-P1 using YCSB benchmarking tools [33].

### 6.1 Experiment Design and Setup

**Setup:** In our experiments, we use a laptop equipped with an SGX CPU. Specifically, the hardware specs include an Intel 8-core i7-6820HK CPU of 2.70 GHz with 8 MB cache, a 16 GB RAM and 1 TB disk.

We run our experiments in the YCSB framework [33]. We use YCSB to both generate the workload and to execute the experiments. YCSB framework works in two phases: the load phase when it initializes the system by populating the dataset, and the evaluation phase when it drives the target workload to the system and measures the performance. When initializing each experiment, we typically scan the loaded dataset so that it is loaded in the untrusted memory. By this means, we mainly consider the setting of memory-resident data with size ranging from several hundreds of megabytes to four gigabytes (that is, millions of records, each with a 16-byte key and 100-byte value by default). With such small record size, four gigabytes is the maximal data size that we can tolerate in experiment time.

We port the open-source LevelDB-YCSB adapter [76] to the SGX architecture. This is done by running the YCSB platform in the untrusted world and wrap each PUT/GET request as an ECall (as in SGX SDK) into the enclave. Inside the enclave, we run the YCSB measurement code that measures various performance metrics.

**Baseline: Eleos:** We implement a baseline of an in-memory data store. In this in-memory store, the entire dataset is stored in enclave as a sorted array. To make data update efficient, we leave 30% of the array space empty to accommodate data insertions without moving existing data. For implementation, we use Eleos [30], a state-of-the-art virtual memory management engine in enclave without calling expensive enclave paging. Their approach, briefly, is to monitor all memory references and to relocate data between enclave and untrusted memory. We implement a sorted array in enclave linked with Eleos. The array serves data reads with binary search and is updated “in place”. For fair comparison, the data in Eleos is persisted to disk periodically. This is done by maintaining a write buffer storing recent data updates and switching out enclave (through an OCall) for data persistence on disk.

### 6.2 Overall Performance under YCSB (Macro-benchmark)

We present the performance result of eLSM under YCSB. In this set of experiments, we vary the workloads in terms read-write ratio, key distribution, etc. and evaluate eLSM performance. The purpose is to present a holistic view regarding the performance of eLSM-P1 and eLSM-P2.

To conduct the experiments, we fix the initial dataset at 3 GB. In the evaluation phase, we drive millions of operations to the key-value store for performance measurement. We use the uniform distribution to generate the dataset and queries. We turn on the default compaction strategy. By this means, we conduct a series of experiments with varying the read-write ratio of the workload. Each experiment is run for three times and the average performance metric and standard deviation are reported. The operation latency of eLSM-P1 and eLSM-P2 under varying read-write ratios are shown in Figure 5a. The result shows that eLSM-P2 outperforms eLSM-P1 in most workloads except for a small set of write-only workloads. Specifically, as the workload becomes more read intensive, eLSM-P2 (with the `mmap` configuration as described in § 5.5.1) has its operation latency decreased. This performance characteristic is due to that

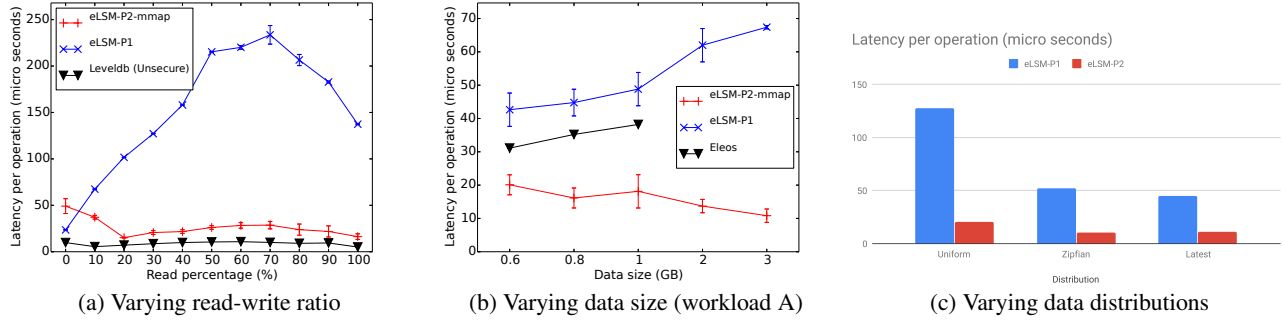


Figure 5: Performance of eLSM-P2 and eLSM-P1 under YCSB workloads

eLSM-P2 has to cause disk IO for data persistence on the write path while on the read path it can read the memory (through the `mmap` files). As the workload transitions from writes to reads, eLSM-P1’s latency first increases and then decreases near the end. The increase of latency is caused by overflowing the enclave memory (of 128 MB) and enclave paging, as will be validated in other experiments (e.g., in Figure 6a). In addition, compared with the ideal approach (running an unsecured LevelDB), the slowdown caused by eLSM-P2 ranges between  $1.5\times$  and  $4X$ .

Comparing eLSM-P2 and eLSM-P1, when the workload is write-only, eLSM-P1 is faster. For most workloads, eLSM-P2 has a smaller operation latency than eLSM-P1, and the performance discrepancy reaches the highest when the workload consists of 70% reads (Note the uniform key distribution in this workload). In this setting, eLSM-P2 achieves  $4.5X$  performance speedup comparing eLSM-P1. This performance result clearly supports the design tradeoff made in eLSM-P1 and eLSM-P2, where eLSM-P2 optimizes the read path by placing the read buffer outside enclave and avoiding enclave paging, which inevitably causes the write overhead, including authenticating `COMPACTION` and embedding eLSM-P2 proofs in the software layer. eLSM-P1 does not have such write overhead (data security is provided by the hardware-level memory protection in `SGX`). From the performance result, it can be seen that the eLSM-P2’s design to trade off write performance for read is worthwhile, as the majority of workloads favors eLSM-P2.

The second experiment is to report the operation latency under varying data sizes. We initialize the system with data of varying sizes from 0.6 GB to 3 GB. In the evaluation phase, we drive into the system YCSB workload A which consists of 50% reads and 50% writes with data keys generated following a Zipfian distribution. We measure the operation latency for eLSM-P2 (in `mmap` configuration), eLSM-P1 and the baseline of Eleos. The result is shown in Figure 5b. With the increasing data sizes, Eleos can scale only to 1 GB data which is limited by their open-source project [77, 30]. The discrepancy between the latency of eLSM-P2 and eLSM-P1 increases, which reaches a maximal of  $7X$  difference when the data size is 3 GB.

Figure 5c presents the operation latency when the workload is generated with different key distributions. YCSB provides three common distribution for generating data keys, that is, Uniform, Zipfian and Latest. Among the three, Latest is the key distribution that has the best temporal locality (as it tends to read the latest inserted records), leading to smaller working sets. In this experiment, we use an initial dataset of 3 GB. In general, eLSM-P2 is less sensitive to key distribution than eLSM-P1. Under the uniform distribution, eLSM-P1 causes the highest operation latency. Because the working set size is the largest when data keys are generated uniformly, and it causes the highest memory pressure in enclave in eLSM-P1’s design.

### 6.3 Read Performance

This set of experiments evaluate the read performance of eLSM-P2 and eLSM-P1. We initialize the storage system with datasets of varying sizes from 8 MB to 3 GB. We then drive a read-only workload of one million GET requests and measure the latency of these requests. In addition to Eleos, we consider another baseline that was mentioned in our initial performance study (in Figure 2). The baseline places the read buffers outside the enclave (similar to eLSM-P2) but does not have the data-authentication measures including the eLSM Merkle trees and proofs (unlike eLSM-P2). This baseline is thus unsecured but it serves to show the ideal performance. We use the `mmap` configuration in eLSM-P2 which allows for accessing files pinned in memory.

The read latency is presented in Figure 6a. When the data size is smaller than the enclave memory size (i.e., smaller than 128 MB), eLSM-P1 and the baseline of Eleos perform better than eLSM-P2 because eLSM-P2 incurs proof and verification in the software layer. When the data size grows beyond the enclave memory (i.e., larger than 128 MB), eLSM-P2 outperforms eLSM-P1 and Eleos. Again, Eleos limits the data scalability to 1 GB data. The longer latency in Eleos may be caused by the overhead for runtime monitoring and extra data copy in enclave memory. eLSM-P2 generally keeps the read latency constant with the increasing data size.

On the read path, eLSM-P2 can support both `mmap` files and a user-space buffer. Note that eLSM-P1 cannot support `mmap` files as such files must reside in the kernel-space memory outside the enclave. We compare the performance of `mmap` read and buffer-based read in eLSM-P2. The result is in Figure 6b. As data grows, the performance advantages of `mmap` configuration in eLSM-P2 become clearer. At the largest data scale tested, eLSM-P2 with `mmap` achieves  $5X$  speedup of read latency comparing with the read buffer configuration.

We also compare eLSM-P2 (the buffer configuration) with eLSM-P1 under varying buffer sizes. Both eLSM-P2 (buffer) and eLSM-P1 support user-space memory buffers and this experiment is intended for a fairer comparison. In the experiment, we fix the data size to be 2 GB and vary the buffer size from 32 MB to 2 GB. The result is in Figure 6c. It can be seen that eLSM-P2 stays constant with increasing buffer size (Note that the data size is fixed and this is a different setting from the previous experiment in Figure 6b) and eLSM-P1’s read latency increases sharply around 128 MB data size. In general, eLSM-P2 (buffer) achieves  $1.6X \sim 2.3X$  speedups against eLSM-P1.

### 6.4 Write Performance

This set of experiments measure the write performance of eLSM-P2 and eLSM-P1. In the experiments, we initialize the system with dataset of varying sizes from 0.2 GB to 4 GB. In the evaluation phase, we measure the write latency. For eLSM-P2 and eLSM-P1, we consider the configurations with and without `COMPACTION`. We also consider the baseline of Eleos.

The write performance with `COMPACTION` is reported in Figure 7a.

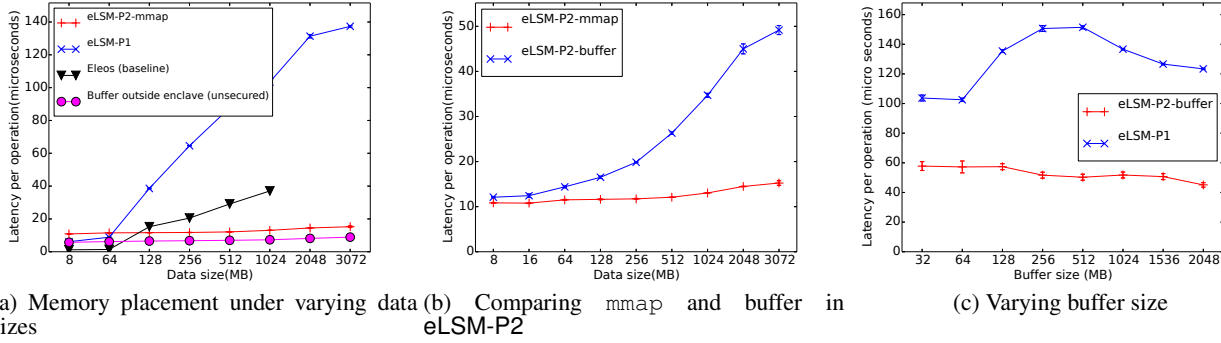


Figure 6: Read performance of eLSM-P2, eLSM-P1 and the Eleos baseline

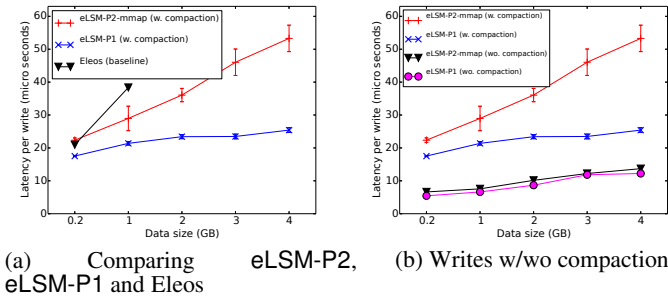


Figure 7: Write performance with/without COMPACTION in eLSM-P2 and eLSM-P1

We report as the write latency the average execution time of running an PUT request plus the time for COMPACTON amortized to the individual PUT. Among the three approaches, eLSM-P1 is the fastest on the write path. With COMPACTON turned on, eLSM-P2’s write latency is about  $1.3X \sim 2.3X$  of that in eLSM-P1. Eleos is the slowest and can only scale to the data size of 1 GB. Because in both eLSM-P2 and eLSM-P1, the memory footprint in the write path is limited to a small write buffer (of 4 MB), whereas the Eleos is an update-in-place structure with the working set size equal to the data size. Even without hardware-level enclave paging (as is optimized out by Eleos), the update-in-place structure still incurs a large amount of memory copies across the enclave boundary, leading to a higher write latency than eLSM-P2 and eLSM-P1.

We also compare the write performance with and without COMPACTON. The result is presented in Figure 7b. It can be seen that turning on COMPACTON, it causes  $2X \sim 4X$  slowdown on the write path. In both cases, namely with and without COMPACTON, eLSM-P2 has a higher write latency than eLSM-P1. This is due to the overhead caused by building the embedded proof in eLSM-P2.

## 7. RELATED WORK

**Software Systems on SGX:** Since the release of Intel SGX, there has been a large body of research on building software systems on SGX. The existing works have addressed the in-enclave support for OS (e.g., Haven [78], Graphene-SGX [29], SCONE [79], and Panoply [80]), security applications [81], side-channel attacks [82, 83, 84, 85, 86] and defenses [87, 88, 89], databases [90, 91, 92, 93, 94], etc.

In particular, there is a line of research on building key-value stores in enclave. Pesos [95] supports the secure hosting of a key-value store where the enclave enforces fine-grained access-control policies.

SecureKeeper [96] secures ZooKeeper style coordination services hosted in the public cloud. Their approach is to confine the computation of user-provided data in enclave and to encrypt data in transit. Eleos [30] supports in-memory key-value stores and particularly addresses the problem of fitting in enclave data larger than 128 MB. They provide a virtual-memory abstraction in enclave and optimize out the expensive enclave paging by monitoring user-space memory accesses and by relocating data dynamically between the enclave memory and untrusted memory. HardIDX [97] is a secure index in enclave that seals external data using authenticated encryption [72]. Concerto [98] supports concurrent key-value stores with consistency guarantees. It is based on a novel design to check strong consistency by leveraging homomorphic secure hash with enclave. Shieldstore [99] supports in-memory key-value stores by placing data outside enclave and by running inside enclave an engine for record-grained encryption and integrity-checking. Shieldstore does not support data persistence or LSM tree.

**LSM Storage Systems:** bLSM [55] optimizes the LSM tree performance by row-based data storage and fine-grained compaction. Prior work [56] minimizes the write amplification under the skewed key access pattern. Pebble [57] reduces the write amplification by organizing storage layout in skip lists and avoiding data rewriting in the same level. Accordion [100] applies the principle of LSM tree to the memory management in order to solve the write amplification caused by frequent compaction. Compaction is critical to the performance of LSM tree storage. Existing work studies distributed compaction management in a cluster setting [58]. Beyond disk storage, the LSM tree has been applied for main-memory databases with high compression rate [101], on non-volatile memory [102], and for spatial databases in the AsterixDB project [103]. Concurrency of the LSM tree is studied in cLSM [104] that supports snapshot scan, conditional update, and concurrent COMPACTON.

**Comparison with Speicher:** Speicher [32], published recently, is a secure system of LSM store with Intel SGX. It places inside the enclave the data keys in MemTable and meta data in other levels to enable search. Data values are stored outside the enclave. Speicher presents a series of performance optimization for efficient IO in SGX.

Our eLSM is developed independently with Speicher. Comparing with it, eLSM has technical distinctions listed as following: 1) eLSM supports efficient data reads (GET) with early stops. As mentioned in § 5.3, eLSM stops processing a GET as early as the first hit is found, that is, without scanning all the levels. This leads to significant performance improvement especially for workloads exhibiting temporal locality (recall Figure 5c). Also, the capability is much needed in the context of incremental log monitoring where only the recent records are of interest (see the case study in § 5.7). In addition, we present non-trivial security analysis of this capability (i.e., Theorem 5.3 and Lemma 5.4). By contrast, existing work including Speicher requires each GET to iterate through all levels. 2)

We present a modular implementation of eLSM without changing the code of RocksDB. This is realized by the eLSM design in embedding proofs in data records and leveraging the callback hooks in RocksDB (§ 5.5.3). By contrast, existing work including Speicher requires code modification of the underlying LSM stores [32]. 3) We have integrated eLSM in Google’s certificate transparency for a real-world case study (§ 5.7). We also open-source the code of eLSM [34].

At last, Speicher supports encrypted data keys in the untrusted world. To search the encrypted LSM store, it needs to “correlate the block” index of an SSTable with the queried key. As the block index being accessed is leaked to the untrusted world, its correlation to data key may lead to information leakage and break the query confidentiality. This security implication of this leakage however is not well documented in the paper.

## 8. CONCLUSION

This paper presents a novel SGX-based LSM key-value store to address emerging application needs of authenticating data under frequent updates. The proposed eLSM-P2 system places memory buffers outside enclave for efficiency. It authenticates the data with small proofs at selective levels of an LSM tree. Implementation on LevelDB and RocksDB is presented. The performance studies with YCSB workloads show a 4.5X speedup of eLSM-P2 over the baseline of eLSM-P1.

## 9. REFERENCES

- [1] “Certificate transparency, the internet standards.” [Online]. Available: <https://tools.ietf.org/html/rfc6962>
- [2] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman, “CONIKS: bringing key transparency to end users,” in *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12-14, 2015.*, J. Jung and T. Holz, Eds. USENIX Association, 2015, pp. 383–398. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/melara>
- [3] “Introducing oak, a free and open certificate transparency log (let’s encrypt).” [Online]. Available: <https://bit.ly/2HtgDAO>
- [4] R. Tamassia, “Authenticated data structures,” in *Algorithms - ESA 2003, 11th Annual European Symposium, Budapest, Hungary, September 16-19, 2003, Proceedings, 2003*, pp. 2–5. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-39658-1\\_2](http://dx.doi.org/10.1007/978-3-540-39658-1_2)
- [5] C. U. Martel, G. Nuckolls, P. T. Devanbu, M. Gertz, A. Kwong, and S. G. Stubblebine, “A general model for authenticated data structures,” *Algorithmica*, vol. 39, no. 1, pp. 21–41, 2004.
- [6] C. Papamanthou, R. Tamassia, and N. Triandopoulos, “Authenticated hash tables,” in *Proceedings of the 2008 ACM Conference on Computer and Communications Security, CCS 2008, Alexandria, Virginia, USA, October 27-31, 2008*, P. Ning, P. F. Syverson, and S. Jha, Eds. ACM, 2008, pp. 437–448. [Online]. Available: <http://doi.acm.org/10.1145/1455770.1455826>
- [7] Y. Zhang, J. Katz, and C. Papamanthou, “Integridb: Verifiable SQL for outsourced databases,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, I. Ray, N. Li, and C. Kruegel, Eds. ACM, 2015, pp. 1480–1491. [Online]. Available: <http://doi.acm.org/10.1145/2810103.2813711>
- [8] P. Devanbu, M. Gertz, C. Martel, and S. G. Stubblebine, “Authentic data publication over the internet,” *Journal of Computer Security*, vol. 11, p. 2003, 2003.
- [9] H. Pang and K.-L. Tan, “Authenticating query results in edge computing,” in *Proceedings of the 20th International Conference on Data Engineering*, ser. ICDE ’04. Washington, DC, USA: IEEE Computer Society, 2004, pp. 560–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=977401.978163>
- [10] F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin, “Dynamic authenticated index structures for outsourced databases,” in *SIGMOD Conference, 2006*, pp. 121–132.
- [11] S. Papadopoulos, Y. Yang, and D. Papadias, “Cads: Continuous authentication on data streams,” in *VLDB, 2007*, pp. 135–146.
- [12] Y. Yang, D. Papadias, S. Papadopoulos, and P. Kalnis, “Authenticated join processing in outsourced databases,” in *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2009, Providence, Rhode Island, USA, June 29 - July 2, 2009*, U. Çetintemel, S. B. Zdonik, D. Kossmann, and N. Tatbul, Eds. ACM, 2009, pp. 5–18. [Online]. Available: <http://doi.acm.org/10.1145/1559845.1559849>
- [13] Y. Yang, S. Papadopoulos, D. Papadias, and G. Kollios, “Authenticated indexing for outsourced spatial databases,” *VLDB J.*, vol. 18, no. 3, pp. 631–648, 2009. [Online]. Available: <http://dx.doi.org/10.1007/s00778-008-0113-2>
- [14] C. Papamanthou, E. Shi, R. Tamassia, and K. Yi, “Streaming authenticated data structures,” in *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques*,

- Athens, Greece, May 26-30, 2013. *Proceedings*, ser. Lecture Notes in Computer Science, T. Johansson and P. Q. Nguyen, Eds., vol. 7881. Springer, 2013, pp. 353–370. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-38348-9\\_22](http://dx.doi.org/10.1007/978-3-642-38348-9_22)
- [15] Y. Qian, Y. Zhang, X. Chen, and C. Papamanthou, “Streaming authenticated data structures: Abstraction and implementation,” in *Proceedings of the 6th edition of the ACM Workshop on Cloud Computing Security, CCSW '14, Scottsdale, Arizona, USA, November 7, 2014*, G. Ahn, A. Oprea, and R. Safavi-Naini, Eds. ACM, 2014, pp. 129–139. [Online]. Available: <http://doi.acm.org/10.1145/2664168.2664177>
- [16] “Intel Corp. Software Guard Extensions programming reference.” [Online]. Available: <https://goo.gl/Ka3pnU>
- [17] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. Gruber, “Bigtable: A distributed storage system for structured data (awarded best paper!),” in *OSDI, 2006*, pp. 205–218.
- [18] “Google LevelDB.” [Online]. Available: <http://code.google.com/p/leveldb/>
- [19] “Facebook RocksDB.” [Online]. Available: <http://rocksdb.org/>
- [20] “Apache HBase.” [Online]. Available: <http://hbase.apache.org/>
- [21] “Apache Cassandra.” [Online]. Available: <http://cassandra.apache.org/>
- [22] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, “Proof verification and the hardness of approximation problems,” *J. ACM*, vol. 45, no. 3, pp. 501–555, 1998. [Online]. Available: <http://doi.acm.org/10.1145/278298.278306>
- [23] B. Parno, J. Howell, C. Gentry, and M. Raykova, “Pinocchio: Nearly practical verifiable computation,” in *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*. IEEE Computer Society, 2013, pp. 238–252. [Online]. Available: <http://dx.doi.org/10.1109/SP.2013.47>
- [24] S. T. V. Setty, B. Braun, V. Vu, A. J. Blumberg, B. Parno, and M. Walfish, “Resolving the conflict between generality and plausibility in verified computation,” in *Eighth Eurosys Conference 2013, EuroSys '13, Prague, Czech Republic, April 14-17, 2013*, 2013, pp. 71–84. [Online]. Available: <http://doi.acm.org/10.1145/2465351.2465359>
- [25] B. Braun, A. J. Feldman, Z. Ren, S. T. V. Setty, A. J. Blumberg, and M. Walfish, “Verifying computations with state,” in *ACM SIGOPS 24th Symposium on Operating Systems Principles, SOSP '13, Farmington, PA, USA, November 3-6, 2013*, 2013, pp. 341–357. [Online]. Available: <http://doi.acm.org/10.1145/2517349.2522733>
- [26] R. S. Wahby, S. T. V. Setty, Z. Ren, A. J. Blumberg, and M. Walfish, “Efficient RAM and control flow in verifiable outsourced computation,” in *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2014*, 2015. [Online]. Available: <http://www.internet-society.org/doc/efficient-ram-and-control-flow-verifiable-outsourced-computation>
- [27] “Intel Software Guard Extensions (Intel SGX) SDK.” [Online]. Available: <https://software.intel.com/en-us/sgx-sdk/download>
- [28] A. Baumann, M. Peinado, and G. C. Hunt, “Shielding applications from an untrusted cloud with haven,” *ACM Trans. Comput. Syst.*, vol. 33, no. 3, pp. 8:1–8:26, 2015. [Online]. Available: <http://doi.acm.org/10.1145/2799647>
- [29] C. Tsai, D. E. Porter, and M. Vij, “Graphene-sgx: A practical library OS for unmodified applications on SGX,” in *2017 USENIX Annual Technical Conference, USENIX ATC 2017, Santa Clara, CA, USA, July 12-14, 2017*. USENIX Association, 2017, pp. 645–658. [Online]. Available: <https://www.usenix.org/conference/atc17/technical-sessions/presentation/tsai>
- [30] M. Orenbach, P. Lifshits, M. Minkin, and M. Silberstein, “Eleos: Exitless OS services for SGX enclaves,” in *Proceedings of the Twelfth European Conference on Computer Systems, EuroSys 2017, Belgrade, Serbia, April 23-26, 2017*, G. Alonso, R. Bianchini, and M. Vukolic, Eds. ACM, 2017, pp. 238–253. [Online]. Available: <http://doi.acm.org/10.1145/3064176.3064219>
- [31] “Compaction filter.” [Online]. Available: <https://github.com/facebook/rocksdb/wiki/Compaction-Filter>
- [32] M. Bailieu, J. Thalheim, P. Bhatotia, C. Fetzer, M. Honda, and K. Vaswani, “SPEICHER: securing lsm-based key-value stores using shielded execution,” in *17th USENIX Conference on File and Storage Technologies, FAST 2019, Boston, MA, February 25-28, 2019.*, 2019, pp. 173–190. [Online]. Available: <https://www.usenix.org/conference/fast19/presentation/bailieu>
- [33] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, “Benchmarking cloud serving systems with ycsb,” in *SoCC, 2010*, pp. 143–154.
- [34] “elms source code.” [Online]. Available: <https://github.com/anonymous-submission2/eLSM-P2>
- [35] P. E. O’Neil, E. Cheng, D. Gawlick, and E. J. O’Neil, “The log-structured merge-tree (lsm-tree),” *Acta Inf.*, vol. 33, no. 4, pp. 351–385, 1996.
- [36] R. Elmasri and S. B. Navathe, *Fundamentals of Database Systems, 2nd Edition*. Benjamin/Cummings, 1994.
- [37] A. Silberschatz, H. F. Korth, and S. Sudarshan, *Database System Concepts, 5th Edition*. McGraw-Hill Book Company, 2005.
- [38] M. Rosenblum, *The Design and Implementation of a Log-Structured File-System*. Kluwer, 1995.
- [39] “General transparency and the trillion project.” [Online]. Available: <https://www.certificate-transparency.org/general-transparency>
- [40] “Security through transparency.” [Online]. Available: <https://security.googleblog.com/2017/01/security-through-transparency.html>
- [41] “Certificate transparency.” [Online]. Available: <https://www.certificate-transparency.org/>
- [42] “Auditing for TLS certificates.” [Online]. Available: <https://github.com/google/certificate-transparency>
- [43] “Bitcoin.” [Online]. Available: <https://bitcoin.org/en/>
- [44] “Ethereum project.” [Online]. Available: <https://www.ethereum.org/>
- [45] “Comodo: Comodo fraud incident 2011-03-23.” [Online]. Available: [https://www.comodo.com/ComodoFraud-Incident-2011-03-23.html\(March2011\)](https://www.comodo.com/ComodoFraud-Incident-2011-03-23.html(March2011))
- [46] “Microsoft: Erroneous verisign-issued digital certificates pose spoofing hazard.” [Online]. Available: [https://technet.microsoft.com/library/security/ms01-017\(Mar2001\)](https://technet.microsoft.com/library/security/ms01-017(Mar2001))
- [47] “Simplified payment verification.” [Online]. Available: <http://docs.electrum.org/en/latest/spv.html>
- [48] “Let’s Encrypt - Free SSL/TLS Certificates.” [Online]. Available: <https://letsencrypt.org>
- [49] “LevelDB optimization for Bitcoin.” [Online]. Available: <https://github.com/bitcoin/bitcoin/blob/v0.8.0/doc/release-notes.txt>
- [50] “Diving into ethereums world state.” [Online]. Available: <http://bit.ly/2m4ZS70>
- [51] “LevelDB for state database in HyperLedger.” [Online]. Available: <https://hyperledger-fabric.readthedocs.io/en/>

- release-1.2/ledger.html#state-database
- [52] “Multichain data streams.” [Online]. Available: <https://www.multichain.com/developers/data-streams/>
- [53] “Google key-transparency project: A transparent and secure way to look up public keys.” [Online]. Available: <https://github.com/google/keytransparency>
- [54] “Google trillion project.” [Online]. Available: <https://github.com/google/trillian>
- [55] R. Sears and R. Ramakrishnan, “blsm: a general purpose log structured merge tree,” in *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2012, Scottsdale, AZ, USA, May 20-24, 2012*, K. S. Candan, Y. Chen, R. T. Snodgrass, L. Gravano, and A. Fuxman, Eds. ACM, 2012, pp. 217–228. [Online]. Available: <http://doi.acm.org/10.1145/2213836.2213862>
- [56] C. M. Jermaine, E. Omiecinski, and W. G. Yee, “The partitioned exponential file for database storage management,” *VLDB J.*, vol. 16, no. 4, pp. 417–437, 2007. [Online]. Available: <http://dx.doi.org/10.1007/s00778-005-0171-7>
- [57] P. Raju, R. Kadekodi, V. Chidambaram, and I. Abraham, “Pebblesdb: Building key-value stores using fragmented log-structured merge trees,” in *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*. ACM, 2017, pp. 497–514. [Online]. Available: <http://doi.acm.org/10.1145/3132747.3132765>
- [58] M. Y. Ahmad and B. Kemme, “Compaction management in distributed key-value datastores,” *PVLDB*, vol. 8, no. 8, pp. 850–861, 2015. [Online]. Available: <http://www.vldb.org/pvldb/vol8/p850-ahmad.pdf>
- [59] R. C. Merkle, “Protocols for public key cryptosystems,” in *IEEE Symposium on Security and Privacy*, 1980, pp. 122–134.
- [60] C. Zheng, G. Shen, S. Li, and S. Shenker, “Distributed segment tree: Support of range query and cover query over DHT,” in *5th International workshop on Peer-To-Peer Systems, IPTPS 2006, Santa Barbara, CA, USA, February 27-28, 2006*, 2006. [Online]. Available: <http://www.iptps.org/papers-2006/Zheng-DST06.pdf>
- [61] “Intel SGX synchronization primitives.” [Online]. Available: <https://tools.ietf.org/html/rfc6962#section-7.2>
- [62] “HBase Coprocessor.” [Online]. Available: [https://blogs.apache.org/hbase/entry/coprocessor\\_introduction](https://blogs.apache.org/hbase/entry/coprocessor_introduction)
- [63] “RocksDB Hooks.” [Online]. Available: <https://github.com/facebook/rocksdb/wiki/rocksdb-basics>
- [64] “Intel SGX SDK’s sgx\_create\_monotonic\_counter.” [Online]. Available: <https://software.intel.com/en-us/node/709160>
- [65] R. Strackx and F. Piessens, “Ariadne: A minimal approach to state continuity,” in *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, T. Holz and S. Savage, Eds. USENIX Association, 2016, pp. 875–892. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/strackx>
- [66] S. Matetic, M. Ahmed, K. Kostianinen, A. Dhar, D. Sommer, A. Gervais, A. Juels, and S. Capkun, “ROTE: rollback protection for trusted execution,” in *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017.*, E. Kirda and T. Ristenpart, Eds. USENIX Association, 2017, pp. 1289–1306. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/matetic>
- [67] “Intel Management Engine ME.” [Online]. Available: <https://intel.ly/2k8bUMi>
- [68] R. Curtmola, J. A. Garay, S. Kamara, and R. Ostrovsky, “Searchable symmetric encryption: improved definitions and efficient constructions,” in *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, October 30 - November 3, 2006*, A. Juels, R. N. Wright, and S. D. C. di Vimercati, Eds. ACM, 2006, pp. 79–88. [Online]. Available: <http://doi.acm.org/10.1145/1180405.1180417>
- [69] R. A. Popa, C. M. S. Redfield, N. Zeldovich, and H. Balakrishnan, “Cryptdb: protecting confidentiality with encrypted query processing,” in *Proceedings of the 23rd ACM Symposium on Operating Systems Principles 2011, SOSOP 2011, Cascais, Portugal, October 23-26, 2011*, T. Wobber and P. Druschel, Eds. ACM, 2011, pp. 85–100. [Online]. Available: <http://doi.acm.org/10.1145/2043556.2043566>
- [70] S. Tu, M. F. Kaashoek, S. Madden, and N. Zeldovich, “Processing analytical queries over encrypted data,” *PVLDB*, vol. 6, no. 5, pp. 289–300, 2013. [Online]. Available: <http://www.vldb.org/pvldb/vol6/p289-tu.pdf>
- [71] M. Bellare, A. Boldyreva, and A. O’Neill, “Deterministic and efficiently searchable encryption,” in *CRYPTO*, ser. Lecture Notes in Computer Science, vol. 4622. Springer, 2007, pp. 535–552.
- [72] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*. Chapman and Hall/CRC Press, 2007.
- [73] A. Boldyreva, N. Chenette, Y. Lee, and A. O’Neill, “Order-preserving symmetric encryption,” in *Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings*, 2009, pp. 224–241. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-01001-9\\_13](http://dx.doi.org/10.1007/978-3-642-01001-9_13)
- [74] R. A. Popa, F. H. Li, and N. Zeldovich, “An ideal-security protocol for order-preserving encoding,” in *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*. IEEE Computer Society, 2013, pp. 463–477. [Online]. Available: <http://dx.doi.org/10.1109/SP.2013.38>
- [75] A. Ahmad, K. Kim, M. I. Sarfaraz, and B. Lee, “OBLIVATE: A data oblivious filesystem for intel SGX,” in *NDSS*. The Internet Society, 2018.
- [76] “YCSB LevelDB adaptor.” [Online]. Available: <https://github.com/jtsui/ycsb-leveldb>
- [77] “Exitless services for SGX enclaves.” [Online]. Available: <https://github.com/acsl-technion/eleos>
- [78] A. Baumann, M. Peinado, and G. C. Hunt, “Shielding applications from an untrusted cloud with haven,” in *11th USENIX Symposium on Operating Systems Design and Implementation, OSDI ’14, Broomfield, CO, USA, October 6-8, 2014.*, 2014, pp. 267–283. [Online]. Available: <https://www.usenix.org/conference/osdi14/technical-sessions/presentation/baumann>
- [79] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O’Keeffe, M. Stillwell, D. Goltzsche, D. M. Eyers, R. Kapitza, P. R. Pietzuch, and C. Fetzer, “SCONE: secure linux containers with intel SGX,” in *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016.*, K. Keeton and T. Roscoe, Eds. USENIX Association, 2016, pp. 689–703. [Online]. Available: <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/arnautov>
- [80] S. Shinde, D. Le Tien, S. Tople, and P. Saxena, “Panoply: Low-tcb linux applications withSGXenclaves,” 2017.
- [81] J. Lind, C. Priebe, D. Muthukumaran, D. O’Keeffe, P. Aublin, F. Kelbert, T. Reiher, D. Goltzsche, D. M. Eyers, R. Kapitza, C. Fetzer, and P. R. Pietzuch, “Glamdring: Automatic

- application partitioning for intel SGX,” in *2017 USENIX Annual Technical Conference, USENIX ATC 2017, Santa Clara, CA, USA, July 12-14, 2017*. USENIX Association, 2017, pp. 285–298. [Online]. Available: <https://www.usenix.org/conference/atc17/technical-sessions/presentation/lin>
- [82] Y. Xu, W. Cui, and M. Peinado, “Controlled-channel attacks: Deterministic side channels for untrusted operating systems,” in *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. IEEE Computer Society, 2015, pp. 640–656. [Online]. Available: <http://dx.doi.org/10.1109/SP.2015.45>
- [83] J. V. Bulck, N. Weichbrodt, R. Kapitza, F. Piessens, and R. Strackx, “Telling your secrets without page faults: Stealthy page table-based attacks on enclaved execution,” in *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, E. Kirda and T. Ristenpart, Eds. USENIX Association, 2017, pp. 1041–1056. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/van-bulck>
- [84] M. Hähnel, W. Cui, and M. Peinado, “High-resolution side channels for untrusted operating systems,” in *2017 USENIX Annual Technical Conference, USENIX ATC 2017, Santa Clara, CA, USA, July 12-14, 2017*. USENIX Association, 2017, pp. 299–312. [Online]. Available: <https://www.usenix.org/conference/atc17/technical-sessions/presentation/hahnel>
- [85] F. Brasser, U. Müller, A. Dmitrienko, K. Kostiaainen, S. Capkun, and A. Sadeghi, “Software grand exposure: SGX cache attacks are practical,” *CoRR*, vol. abs/1702.07521, 2017. [Online]. Available: <http://arxiv.org/abs/1702.07521>
- [86] S. Lee, M. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado, “Inferring fine-grained control flow inside SGX enclaves with branch shadowing,” in *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, E. Kirda and T. Ristenpart, Eds. USENIX Association, 2017, pp. 557–574. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/lee-sangho>
- [87] M.-W. Shih, S. Lee, T. Kim, and M. Peinado, “T-sgx: Eradicating controlled-channel attacks against enclave programs,” in *NDSS Symposium 2017 in San Diego, California*.
- [88] S. Shinde, Z. L. Chua, V. Narayanan, and P. Saxena, “Preventing page faults from telling your secrets,” in *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2016, Xi’an, China, May 30 - June 3, 2016*, X. Chen, X. Wang, and X. Huang, Eds. ACM, 2016, pp. 317–328. [Online]. Available: <http://doi.acm.org/10.1145/2897845.2897885>
- [89] S. Chen, X. Zhang, M. K. Reiter, and Y. Zhang, “Detecting privileged side-channel attacks in shielded execution with déjà vu,” in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2017, Abu Dhabi, United Arab Emirates, April 2-6, 2017*, R. Karri, O. Sinanoglu, A. Sadeghi, and X. Yi, Eds. ACM, 2017, pp. 7–18. [Online]. Available: <http://doi.acm.org/10.1145/3052973.3053007>
- [90] F. Schuster, M. Costa, C. Fournet, C. Gkantsidis, M. Peinado, G. Mainar-Ruiz, and M. Russinovich, “VC3: trustworthy data analytics in the cloud using SGX,” in *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. IEEE Computer Society, 2015, pp. 38–54. [Online]. Available: <https://doi.org/10.1109/SP.2015.10>
- [91] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa, “Oblivious multi-party machine learning on trusted processors,” in *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016*, T. Holz and S. Savage, Eds. USENIX Association, 2016, pp. 619–636. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/ohrimenko>
- [92] C. Priebe, K. Vaswani, and M. Costa, “Enclavedb - a secure database using sgx.” IEEE, May 2018. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/enclavedb-a-secure-database-using-sgx/>
- [93] W. Zheng, A. Dave, J. G. Beekman, R. A. Popa, J. E. Gonzalez, and I. Stoica, “Opaque: An oblivious and encrypted distributed analytics platform,” in *14th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2017, Boston, MA, USA, March 27-29, 2017*, 2017, pp. 283–298. [Online]. Available: <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/zheng>
- [94] S. Bajaj and R. Sion, “Correctdb: SQL engine with practical query authentication,” *PVLDB*, vol. 6, no. 7, pp. 529–540, 2013. [Online]. Available: <http://www.vldb.org/pvldb/vol6/p529-bajaj.pdf>
- [95] R. Krahn, B. Trach, A. Vahldiek-Oberwagner, T. Knauth, P. Bhatotia, and C. Fetzer, “Pesos: policy enhanced secure object store,” in *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, R. Oliveira, P. Felber, and Y. C. Hu, Eds. ACM, 2018, pp. 25:1–25:17. [Online]. Available: <http://doi.acm.org/10.1145/3190508.3190518>
- [96] S. Brenner, C. Wulf, D. Goltzsche, N. Weichbrodt, M. Lorenz, C. Fetzer, P. R. Pietzuch, and R. Kapitza, “SecureKeeper: Confidential ZooKeeper using Intel SGX,” in *Proceedings of the 17th International Middleware Conference, Trento, Italy, December 12 - 16, 2016*. ACM, 2016, p. 14. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2988350>
- [97] B. Fuhry, R. Bahmani, F. Brasser, F. Hahn, F. Kerschbaum, and A. Sadeghi, “Hardidx: Practical and secure index with SGX,” in *Data and Applications Security and Privacy XXXI - 31st Annual IFIP WG 11.3 Conference, DBSec 2017, Philadelphia, PA, USA, July 19-21, 2017, Proceedings*, ser. Lecture Notes in Computer Science, G. Livraga and S. Zhu, Eds., vol. 10359. Springer, 2017, pp. 386–408. [Online]. Available: [https://doi.org/10.1007/978-3-319-61176-1\\_22](https://doi.org/10.1007/978-3-319-61176-1_22)
- [98] A. Arasu, K. Eguro, R. Kaushik, D. Kossmann, P. Meng, V. Pandey, and R. Ramamurthy, “Concerto: A high concurrency key-value store with integrity,” in *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May 14-19, 2017*, S. Salihoglu, W. Zhou, R. Chirkova, J. Yang, and D. Suciu, Eds. ACM, 2017, pp. 251–266. [Online]. Available: <http://doi.acm.org/10.1145/3035918.3064030>
- [99] T. Kim, J. Park, J. Woo, S. Jeon, and J. Huh, “Shieldstore: Shielded in-memory key-value storage with SGX,” in *Proceedings of the Fourteenth EuroSys Conference 2019, Dresden, Germany, March 25-28, 2019*, 2019, pp. 14:1–14:15. [Online]. Available: <https://doi.org/10.1145/3302424.3303951>
- [100] E. Bortnikov, A. Braginsky, E. Hillel, I. Keidar, and G. Sheffi, “Accordion: Better memory organization for LSM key-value stores,” *PVLDB*, vol. 11, no. 12, pp. 1863–1875, 2018. [Online]. Available: <http://www.vldb.org/pvldb/vol11/p1863-bortnikov.pdf>
- [101] H. Zhang, D. G. Andersen, A. Pavlo, M. Kaminsky, L. Ma, and R. Shen, “Reducing the storage overhead of main-memory OLTP databases with hybrid indexes,” in *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, 2016, pp.

1567–1581. [Online]. Available: <http://doi.acm.org/10.1145/2882903.2915222>

- [102] L. Mármol, S. Sundaraman, N. Talagala, and R. Rangaswami, “NVMKV: A scalable, lightweight, ftl-aware key-value store,” in *2015 USENIX Annual Technical Conference, USENIX ATC '15, July 8-10, Santa Clara, CA, USA*, S. Lu and E. Riedel, Eds. USENIX Association, 2015, pp. 207–219. [Online]. Available: <https://www.usenix.org/conference/atc15/technical-session/presentation/marmol>
- [103] S. Alsubaiee, A. Behm, V. R. Borkar, Z. Heilbron, Y. Kim, M. J. Carey, M. Dreseler, and C. Li, “Storage management in asterixdb,” *PVLDB*, vol. 7, no. 10, pp. 841–852, 2014. [Online]. Available: <http://www.vldb.org/pvldb/vol7/p841-alsubaiee.pdf>
- [104] G. Golan-Gueta, E. Bortnikov, E. Hillel, and I. Keidar, “Scaling concurrent log-structured data stores,” in *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, 2015, pp. 32:1–32:14. [Online]. Available: <http://doi.acm.org/10.1145/2741948.2741973>
- [105] I. Anati, S. Gueron, S. P. Johnson, and V. R. Scarlata, “Innovative technology for cpu based attestation and sealing.”
- [106] “Intel TXT.” [Online]. Available: [http://www.intel.com/technology/security/downloads/TrustedExec\\_Overview.pdf](http://www.intel.com/technology/security/downloads/TrustedExec_Overview.pdf)
- [107] “TPM.” [Online]. Available: <http://www.trustedcomputinggroup.org/tpm-main-specification/>
- [108] “ARM TrustZone.” [Online]. Available: <https://www.arm.com/products/security-on-arm/trustzone>
- [109] “IBM SCPU.” [Online]. Available: <http://www-03.ibm.com/security/cryptocards/>
- [110] “Litecoin.” [Online]. Available: <https://litecoin.org/>
- [111] “Google Chrome.” [Online]. Available: <https://www.google.com/chrome/>
- [112] “LevelDB in CT log server.” [Online]. Available: <https://github.com/google/certificate-transparency/blob/5fce65cb60cfe7808afc98de23c7dd5ddbfa1509/docs/DesignDoc.md>
- [113] “Working with IndexedDB.” [Online]. Available: <https://developers.google.com/web/ilt/pwa/working-with-indexeddb>
- [114] “Introducing the browser-based blockchain.” [Online]. Available: <https://medium.com/nimiq-network/introducing-the-browser-based-blockchain-63d408add368>
- [115] B. Schneier, *Data and Goliath: The Hidden Battles to Capture Your Data and Control Your World*, 1st ed. W. W. Norton & Company, 2015.

## APPENDIX

### A. ADDITIONAL PRELIMINARIES

#### A.1 Software Guard eXtension (SGX)

Intel SGX is a security-oriented x86-64 ISA extension on the Intel Skylake CPU, released in 2015. SGX provides a “security-isolated world” for trustworthy program execution on an otherwise untrusted hardware platform. At the hardware level, the SGX’s trusted world or enclave includes a tamper-proof CPU which automatically encrypts memory accesses upon cache write-backs. Programs executed outside the enclave trying to access enclave memory only get to see the ciphertext and cannot succeed. At the software level, the SGX enclave includes only some unprivileged program and excludes any OS kernel code, by explicitly prohibiting system services (e.g., system calls) inside the enclave.

To use the technology, a client initializes an enclave by uploading her program to the server host and uses SGX’s seal and attestation mechanism [105] to verify the correct setup of the enclave environment (i.e., the binding between the client’s program and a

genuine CPU supporting SGX). During the program execution, the enclave can be entered and exited proactively (by SGX instructions, e.g., `EENTER` and `EEXIT`) or passively (by interrupts or traps). These world-switch events trigger the context saving/reloading in both hardware and software levels. Comparing prior TEE solutions [106, 107, 108, 109], SGX uniquely supports multi-core concurrent execution, dynamic paging, and interrupted execution.

The software built on SGX relies on an underlying “RPC” mechanism to switch the execution between the enclave and untrusted world. In SGX SDK, such RPC mechanism is supported by `ECall/OCall` where in the case of an `ECall` (`OCall`), the untrusted host (enclave) switches to enclave (the untrusted host) in order to call a function there. Alternatively, one can use the LibOS supports in enclave to load and run unmodified application software in enclave. There are existing research prototypes for enclave LibOS, such as Haven [28], Graphene-SGX [29], and SCONE [79].

**SGX storage and costs:** In SGX, there is a reserved region of main memory where the CPU stores data encrypted and authenticated. This CPU-protected region of memory, called PRM, is small and in current Intel CPU, the size is limited to 128 MB. The SGX architecture supports a virtual memory in enclave of arbitrary size. When the data stored in enclave virtual memory is larger than 128 MB, it causes the event of “enclave paging”, that is, the hardware will transfer the data between untrusted main memory and the protected memory region (PRM). The enclave paging is expensive as the process involves asynchronous enclave exit (i.e., `AEX`), OS-managed page table, and SGX instructions (i.e., `EWB`) to evict a victim page to remap the requested page. This enclave paging event can easily become the system bottleneck when the application has a working set larger than 128 MB.

#### A.2 Merkle Hash Trees

A Merkle Hash Tree (or MHT) [59] is a method of digesting and authenticating a dataset. Specifically, given an array of data records, a Merkle tree’s leaf set consists of the hashes of all data records. The non-leaf node in a Merkle tree is the hash of the concatenation of the children of this node. The root hash of the Merkle tree digests the entire data array with records in the fixed order. Merkle trees are often used in constructing a proof system between a verifier and a prover. Given a query over a dataset, a Merkle proof consists of all hashes of the tree nodes surrounding the path from the leaf (that matches the query) to the root. Hence, it is also called Merkle authentication path. When the Merkle tree is built over a sorted dataset, the Merkle proof can be used to authenticate the membership and non-membership of a query result in the dataset. In particular, when there is no matching record to the query, the Merkle proof for non-membership consists of the two authentication paths respectively for the two records immediately larger and smaller than the queried data.

### B. DETAILED TARGET SCENARIOS AND WORKLOADS

In practice, the LSM based key-value stores are being used increasingly in security sensitive applications. As an example, LevelDB serves as the data storage solution in security-critical scenarios including certificate transparency log [1, 2], Blockchain/cryptocurrency [43, 44, 110], web browser storage [111], etc. In the following, we give a more detailed description of these scenarios before summarizing the common traits in these target applications.

**In cryptocurrency and Blockchain applications,** LSM stores serve as the ledger storage. For instance, LevelDB is adopted in Bitcoin [49], Ethereum [50], HyperLedger [51], multichain/stream [52], etc. These are user-space storage systems that translate application-level reads/writes (on individual data records) to lower-level file reads/writes. With an LSM tree, they do so

by translating random writes to sequential file writes, while still supporting efficient processing of random-access reads. Specifically, an LSM tree represents a dataset  $m$  by a series of levels  $L_0, L_1, \dots$  where each level consists of key-value records sorted by data keys. Upon a write operation, the LSM tree only updates the first level  $L_0$  in memory. When the memory buffer at  $L_0$  is full, it flushes all buffered recent writes (sorted in order) to the disk, creating a new file in  $L_1$ . Periodically, the LSM tree will merge multiple levels into one, by accessing all files/records of involved levels in batch. By this means, on the write path, it only cause sequential file access, which reduces disk seeks.

**In certificate-transparency log (CT log [1]),** LSM stores serve as a publicly verifiable storage of certificates [112]. A certificate is stored in the CT log with its identity as the key and the certificate as the value. It serves the workloads of data updates by inserting or revoking certificates, and the read queries to retrieve certificates by identity. The data integrity of a certificate is critical to the security in CT, as using a revoked certificate may lead to the misuse of a stolen key and facilitating impersonation attacks.

**Web browser storage** is one of the canonical applications of LSM stores. LevelDB is primarily used as the storage backend in Google Chrome, responsible for storing a large amount of structured data [113]. The data stored in LevelDB in Chrome can be personal files and information, used in Browser extensions. For instance, a Chrome extension for supporting Blockchain (i.e., Nimiq [114]), LevelDB stores the Blockchain data through the web interface. In these applications, serving the fresh data with integrity is security critical (as mentioned above).

**User-generated content:** For an example, suppose a Twitter-like startup company wants to outsource the storage of tweets to a third-party cloud (e.g., for economic reasons). The application-level web server accepting user requests will generate two types of queries for the outsourced storage underneath: 1) writes of new tweets as social users post them, 2) reads of recent tweets to analyze trendy stories and to present them in the “Moments” page (on a Twitter-like website). This persistence workload features a stream of small writes that arrive at a high rate and reads that need to be served in a real-time fashion. An LSM tree is suited to serve this workload – with the first level ( $L_0$ ) of constant size and a constant number of levels, the LSM tree can serve intensive writes with small memory and bound the latency of a read in quasi-linear time. The data authenticity in this Twitter application means that social users will neither be fooled by a fake post nor miss their friends’ newest update.

**Streaming data analytics:** For another example, many big-data applications, such as real-time monitoring and analytics (e.g., Spark), feature an intensive volume of data reads and writes as a typical “big-data” attitude is to collect everything first and to process it later [115]. They require to minimize the time between data collection and processing for real-time analytics. An LSM tree is well suited to serve big-data storage as its log-structured design allows to “collect everything” at high write throughput and its merge design allows to “process” the data as soon as possible.

We summarize the common traits in the above application scenarios and formulate our work in terms of target workloads (§ 3.1) and security requirements (§ 5.3.1).

### B.0.1 Target Workloads

In our target applications, the workload features intensive write streams and random-access reads. For instance, Blockchain applications feature an intensive stream of incoming cryptocurrency transactions and random-access queries to selectively download relevant blocks/transactions when new Blockchain nodes join the network (as in lightweight SPV nodes [47]). CT log’s read/write workloads features an intensive stream of updates, especially in a large-scale setting such as key transparency scheme (CONIKS [2]). CT log serves random-access reads to selectively retrieve the certificate of interest.

In general, our target workload is characterized as below: Our data model is a key-value store, where each record contains a data key through which the record is accessed. On the read path, the workload supports read queries where each read request specifies the data key and retrieves the associated value. The read queries cause random access as the data key can be arbitrarily distributed (C1).

On the write path, there is an *intensive* stream of data writes that need to be persisted in real time. The stream consists of small random-access writes where keys are arbitrarily distributed (C1). To be specific, a data write manifests in an insertion of a key-value record, where keys can be arbitrarily distributed. The target workload features an intensive write stream, requiring the underlying storage system to provide high throughput and low latency.

This target workload excludes many database storage engines that are designed based on update-in-place data structures (e.g., B+ trees). In update-in-place structures, a data update needs to overwrite the previous version of the record at the exact location where the record is stored. An update incurs lookups and random-accesses of the record’s previous location, leading to disk seeks and write amplification. The LSM stores are suited for the target workload of intensive small writes (C2), due to its append-only design that persists the writes in their arrival order.

## C. WRITE BUFFER PLACEMENT: MEASUREMENT RESULTS

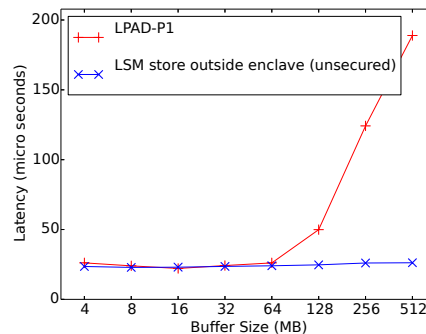


Figure 8: Disk writes: Placing data outside

We conduct a preliminary performance study of placing the write buffer inside enclave (i.e., eLSM-P1) in comparison with that outside enclave. The latter is implemented by the enclave issuing data writes to the untrusted memory (even without security protection) and flushing the write buffer with a world switch. In the performance result, it can be seen that 1) for an unsecured LSM store, a small write buffer results in a similar performance with a large write buffer. 2) With a small buffer, placing the buffer inside enclave and outside enclave result in the similar performance. In other words, placing the write buffer outside enclave does not improve the performance, yet incurs much higher implementation complexity.

## D. ALTERNATIVE DESIGN: TCB-MINIMAL CODE PLACEMENT

To decide which part of the codebase can be placed outside enclave, we take the following view: The codebase of a storage system consists of “computation”-oriented code and data-access oriented code. In an LSM store, there are two major computations: the merge operation and the versioning (the latter is about selecting and comparing timestamps among different records of the same key). When placing the computation-oriented code outside enclave, it entails very expensive cryptographic protocols/schemes to ensure the computation security. For instance, to ensure the integrity of computation result, it requires running verifiable computation

protocols [22] in the untrusted world, whereas the state of the art systems for verifiable computation [23, 24, 25, 26] cause performance slowdown (comparing unsecured computation) by multiple orders of magnitudes. Because of this, we have to place the computation-oriented code inside the enclave.

For the data-access code, it can be placed outside enclave, while imposing affordable security overhead. Concretely, one can build an authenticated data structure (ADS) on the untrusted data and attach a  $O(\log N)$  proof to individual read/write operations for their authenticity.

Running the code outside enclave leads to more OCalls (one Put/Get, at least one OCall) than that of running the code inside enclave. For the former, one Put/Get operation causes at least one OCall, while for the latter, it causes a OCall only when it flushes or misses a read buffer (which can be amortized to multiple Put/Get operations). Due to this reason, we think the alternative design by placing code outside enclave is also unfavorable from the performance perspective.<sup>9</sup>

---

<sup>9</sup>Building an LSM tree outside enclave also require non-trivial engineering efforts implementing the digest structure.