

Confidential Consortium Framework: Secure Multiparty Applications with Confidentiality, Integrity, and High Availability

Heidi Howard*
Azure Research, Microsoft

Fritz Alder†
imec-DistriNet, KU Leuven
Belgium

Edward Ashton
Azure Research, Microsoft

Amaury Chamayou
Azure Research, Microsoft

Sylvan Clebsch
Azure Research, Microsoft

Manuel Costa
Azure Research, Microsoft

Antoine Delignat-Lavaud
Azure Research, Microsoft

Cédric Fournet
Azure Research, Microsoft

Andrew Jeffery†
University of Cambridge
UK

Matthew Kerner
Microsoft

Fotios Kounelis†
Imperial College London
UK

Markus A. Kuppe
Microsoft Research

Julien Maffre
Azure Research, Microsoft

Mark Russinovich
Microsoft

Christoph M. Wintersteiger
Azure Research, Microsoft

ABSTRACT

Confidentiality, integrity protection, and high availability, abbreviated to CIA, are essential properties for trustworthy data systems. The rise of cloud computing and the growing demand for multiparty applications however means that building modern CIA systems is more challenging than ever. In response, we present the Confidential Consortium Framework (CCF), a general-purpose foundation for developing secure stateful CIA applications. CCF combines centralized compute with decentralized trust, supporting deployment on untrusted cloud infrastructure and transparent governance by mutually untrusted parties.

CCF leverages hardware-based trusted execution environments for remotely verifiable confidentiality and code integrity. This is coupled with state machine replication backed by an auditable immutable ledger for data integrity and high availability. CCF enables each service to bring its own application logic, custom multiparty governance model, and deployment scenario, decoupling the operators of nodes from the consortium that governs them. CCF is open-source and available now at <https://github.com/microsoft/CCF>.

1 INTRODUCTION & MOTIVATION

Together data confidentiality, integrity protection, and high availability form the CIA triad, a classic formulation in information security [129]. We begin by examining each property, with a particular focus on running trustworthy multiparty applications on untrusted infrastructure.

Data confidentiality. Organizations are responsible for ensuring the privacy of personal data. Increasingly this responsibility is codified in law, and the penalties for failing to fulfill this obligation can be substantial, for instance up to 4% of turnover in the

case of GDPR [97]. Even where data is not personal, companies may wish to keep data confidential to protect intellectual property, for competitive advantage, or to maintain systems security, for instance when storing secrets. Encryption at rest and in-flight are well-established approaches to achieving confidentiality, but confidentiality during execution is more challenging. Moreover, encryption alone does not fully solve the problem of confidentiality. Instead, it reduces the problem of protecting arbitrary data into the problem of protecting keys, which in turn need to be managed, stored, and released according to some well-defined policy.

Integrity protection. Not only are organizations responsible for maintaining data confidentiality, they are also obligated to protect data in their care from unauthorized or unintended modification. In fact, integrity protection of the code that accesses data is often a prerequisite of data confidentiality. Code integrity together with transparency [34] allows parties who share data to agree on how data can be used, for instance, a bank can comply with anti-money laundering legislation by executing queries on behalf of a government without disclosing all customer information.

Cloud computing enables efficient scaling of applications with cost proportional to use and a low barrier to entry [9], but the mass migration to the cloud means that the trusted computing base (TCB) of our systems is growing over time [103]. Data integrity and confidentiality are vital when leveraging untrusted cloud infrastructure [112] but even more difficult to ensure remotely. As a result, some highly sensitive applications, such as medical, financial, or governmental services, are unable to migrate to the public cloud. This challenging situation leaves us with the following research question: *Is it possible to remove cloud providers from the TCB of multiparty applications and still enable developers to leverage the cloud's compute and storage capabilities?* This requirement is particularly important with the increasing demand for multiparty scenarios: the

*Corresponding author's email: heidi.howard@microsoft.com

†Work done while at Microsoft

integration of data systems between parties who wish to compute over shared data but do not fully trust one another. Collaborating and combining data from multiple sources increases its value and enables novel use cases [103]. However, this makes confidentiality and integrity more difficult, as we must reason about the access rights and requirements of multiple distinct participants.

High availability. As modern society increasingly depends on digital infrastructure, it is vital that applications are dependable and highly available. Even if the necessary consistency and cost trade-offs are made, it is not possible to guarantee 100% availability of digital infrastructure, and so applications should be resilient to failures expected during normal operation.

CIA triad. We need a principled approach to developing CIA applications that is also highly pragmatic, supporting a wide range of stateful applications and modern deployment scenarios, in particular, delegation to untrusted cloud infrastructure and governance by multiple untrusted parties.

In this paper, we present the Confidential Consortium Framework, also known as CCF, which combines centralized cloud compute with decentralized trust (section 2). CCF leverages cloud-based trusted execution environments for remotely attestable confidentiality and integrity (section 3). This is coupled with state machine replication and a transactional key-value store, backed by an immutable ledger for high availability (section 4) and auditing. CCF decouples the operators of CCF nodes and the controllers of the system (the consortium members) (section 5). CCF is highly flexible, allowing developers to bring their own application logic, and their own multiparty governance model for highly tunable oversight.

As we examine in section 8, we are not the first to consider data confidentiality, integrity protection, or high availability in the context of cloud computing or multiparty collaboration. Whereas most prior systems offer either a data storage solution (in the form of a ledger, database, or key-value store) or a secure execution solution (relying instead on a secondary storage system) in isolation, CCF offers an end-to-end solution, supporting both execution and storage. CCF has a unique auditable governance model, designed to operate between untrusted environments and parties, codified in a programmable contract. Moreover, CCF occupies a security versus usability sweet spot, offering a relatively small trusted computing base and a simple yet flexible programming model. Last but not least, CCF is trusted in production with services such as Azure Confidential Ledger [82] and Azure Managed CCF [114] relying on features such as reconfiguration, disaster recovery, snapshotting, indexing, and live code updates, evidencing the importance of a general-purpose and self-contained design.

2 CCF OVERVIEW

A CCF service, as shown in Figure 1, consists of a collection of CCF nodes. These nodes are deployed by one or more operators, such as cloud providers, and managed by the members of a consortium. Users invoke endpoints provided by the application logic on the CCF nodes. The application logic executes the operation corresponding to the endpoint, updating the state of the service, and responds to the user.

Endpoints vary depending on the application, for instance, a distributed logging application might simply provide `write_message`




Figure 1: Overview of a CCF service. Users, operators & consortium members interact with the three CCF nodes.




Figure 2: Key components of each CCF node.

and `read_message` endpoints, which take a message ID and take/return a corresponding message.

On the other hand, consider a banking application, managed by a consortium of financial institutions. Endpoints such as `credit`, `debit`, and `transfer`, could take an account ID (or IDs) and an amount in USD and return with success (or an error such as insufficient funds), and the new account balance. Further endpoints might include `apply_interest`, which updates all account balances from a given bank accordingly, or `audit`, which is available only to a financial regulator, and returns the names of account holders whose total funds exceed some threshold.

The application-specific logic is responsible only for implementing the application itself, naming the endpoints which can be invoked and how they transform a user request to interact with the service state. CCF implements all else necessary to deploy a confidential, integrity-protected, and highly available stateful service. The programming model of CCF aims to be as easy to use as today’s functions-as-a-service platforms whilst providing stronger guarantees.

CCF utilizes hardware-protected trusted execution environments (TEEs), also known as enclaves. Whilst offerings vary, in general, TEEs provide a secure and isolated execution space that guarantees that code and data are protected with respect to confidentiality and integrity. In this paper, the TEE we will be focusing on is Intel SGX (refer to section 8 for a discussion of other TEE platforms). Each CCF node has its own TEE. The framework code running on CCF nodes is divided between the TEE and the untrusted host. The application logic is always executed inside the TEE. This forms the basis of CCF’s confidentiality and integrity guarantees.

Table 1: Sample of the cryptographic keys utilized by nodes in CCF.

Key (Type)	Purpose	Lifetime & Storage
Service Identity (Asymmetric)	Public-key certificate (X.509) used for TLS server authentication (subsection 3.1) & receipt verification (subsection 3.5).	Generated when a CCF service is started, either for the first time or following disaster recovery (subsection 5.2), the associated private key is shared with all trusted CCF nodes and kept only in enclave memory.
Node Identities (Asymmetric)	Public-key certificates (X.509) used for node authentication.	Generated separately by each CCF node when it joins the CCF service, the associated private key is stored in enclave memory and never shared.
Ledger Secret (Symmetric)	Used to encrypt updates to private maps in the ledger.	Shared between all trusted CCF nodes, the encrypted ledger secret is recorded in the underlying key-value store (subsection 5.2).

The application logic accesses data via a key-value store interface provided by CCF. Each endpoint invocation executes transactionally over the key-value store, potentially reading and writing from multiple keys. CCF records the result of these transactions in an append-only ledger, replicated to the other CCF nodes and to persistent storage, forming the basis of CCF’s high availability guarantees.

Threat Model. Regarding safety properties such as data confidentiality, integrity, and correctness, we assume that the hosts running the CCF nodes, and by extension their operators, cannot be trusted. We also do not trust users. Attestation refers to the process by which a host can produce a verifiable proof that it has a TEE and of what code is running inside the TEE. We assume that once their attestation has been verified, the TEEs on CCF nodes can be trusted. Verification of an attestation requires trusting the hardware and the hardware manufacturer. We assume that the code running inside the TEEs, CCF itself and the application logic, has been implemented correctly. We do not trust other applications running on the same host, the operating systems of either the guest or host, the hypervisor, or BIOS. The persistent storage is outside the trust boundary and thus could be modified or rolled back by a malicious host.

Provided that the consortium consists of multiple independent parties, together with a suitable constitution (defined in subsection 5.1), then we do not trust individual consortium members. Whilst we assume that CCF nodes have TEEs, we do not assume that consortium members or users have their own TEEs.

We make the standard cryptographic assumptions such as the existence of collision-resistant hash functions. Table 1 gives an overview of CCF’s main cryptographic keys. CCF also provides rekeying and certificate expiry/renewal, though the implementation details [85] are beyond the scope of this paper. In the case of catastrophic failure, CCF’s disaster recovery procedure (described in subsection 5.2) can perform a best-effort recovery from a single copy of the ledger in persistent storage. This recovered service presents a changed service identity, so the recovery is detectable by users (Table 1).

Regarding liveness, we assume that the operators are not deliberately attacking the availability (or performance) of CCF. However, the CCF nodes are subject to the usual failures such as crash faults and network delays. We do not describe any specific measures to prevent denial-of-service attacks, but standard techniques can be applied.

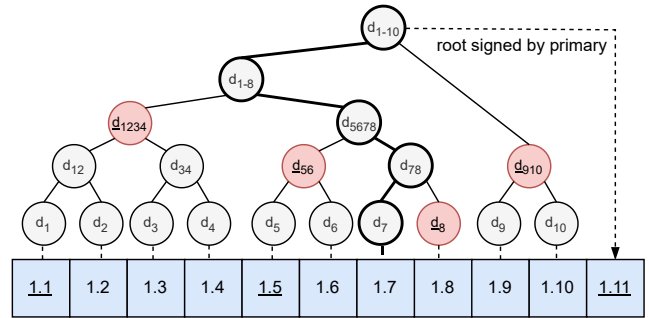


Figure 3: Example of a CCF ledger. Each box represents a transaction and the transaction ID (view & sequence number). The underlined transaction IDs are signature transactions, which each contain a Merkle root over the ledger prefix.

3 INTRA-NODE ARCHITECTURE

In this section we examine the architecture of a static single-node CCF network; we leave the multi-node architecture (section 4) and governance (section 5) to future sections. The key components of a CCF node are shown in Figure 2.

3.1 Application logic

Users connect to the CCF node over TLS, with the connection terminating inside the TEE and the CCF *service certificate* as a root of trust. Users may then invoke endpoints, provided by the application logic, using the HTTP REST API. Each endpoint invocation is executed against the latest version of the key-value store. Once completed, a transaction comprised of the request’s write-set is appended to the end of the ledger and the user receives a response to their request. This response includes the ID assigned to the transaction. This *transaction ID* is an ordered pair of two natural numbers, the view and the sequence number (the index of the transaction in the ledger). We will discuss views later (section 4) but for now we simply assume that transaction IDs are unique to each transaction.

Each CCF endpoint declares how callers should be authenticated. Each invocation is first checked by CCF against these declared policies and the application logic is only called if the caller passes the checks. The application logic can then implement its own authorization logic based on these authenticated claims.



Figure 4: Transaction statuses and transitions. Committed and Invalid states are final, once observed the alternative final status will never be observed (except after disaster recovery).

3.2 Integrity protected ledger

As illustrated in Figure 3, CCF maintains a *Merkle tree* [78], over the ledger. Each leaf vertex of the Merkle tree corresponds to a transaction in the ledger and each interior vertex is the digest of its two children. The Merkle tree is extended to the right as each new transaction is appended to the ledger. The root of the tree, known as the *Merkle root*, is a cryptographic commitment to the whole ledger content. The CCF node periodically signs the Merkle root, and appends this to the ledger in the form of a *signature transaction*. The signature transactions provide integrity protection and thus a transaction can only be considered committed once it has been followed by a signature transaction. The logical ledger is then divided into multiple physical ledger files, each terminating with a signature transaction, as it is written to persistent storage by the host.

In addition to the endpoints exposed by the application logic, CCF also provides its own built-in endpoints common to every service. One such endpoint, `tx`, returns the transaction status (see Figure 4) of a transaction given its transaction ID. Recall that CCF replies to the user as soon as its request has finished executing and includes the associated transaction ID in the response. This `tx` endpoint thus provides one mechanism for the user to check whether the resulting transaction has been committed to the ledger.

3.3 Transactional key-value store

The underlying key-value store in CCF consists of a set of maps, each of which is a named collection of key-value pairs of a given type. Each map may either be private, meaning its updates are encrypted before leaving the TEE and being appended to the ledger, or public, in which case their updates are written to the ledger in plain text. Each transaction in the ledger includes a set of updates, each either a write-to or a removal-of a single key, to be applied atomically to the maps. These updates are subdivided into updates to public maps (unencrypted) and updates to private maps (encrypted).

The current state of the maps is kept in memory on the CCF node and the key-value store is utilized both by the application logic and by CCF itself. For instance, the map `ccf.internal.signatures` stores the signed Merkle root from the signature transaction. CCF’s internal and governance maps are public, enabling auditing without ledger decryption.

3.4 Read-only & historical transactions

Not all transactions update the key-value store, and as such CCF provides a fast path for read-only transactions. Read-only transactions are executed against the latest version of the key-value store

but are not then recorded on the ledger. The user receives a response from CCF as per usual, except the transaction ID included is the ID of the last transaction applied to the key-value store instead of a new one.

The application logic can also choose to query the historical state of the ledger over a range of sequence numbers. Since historical queries may require fetching and processing many transactions from persistent storage, they can take some time. To accelerate historical range queries, CCF enables applications to define their own indexing strategy within the application logic. The indexer on the CCF node pre-processes in-order each transaction in the ledger as it is committed and stores the results for future use. Alternatively, this can also be done lazily when a historical query is received. The indexer’s state is kept in-memory for quick access but can be offloaded to persistent storage if needed.

As an example, an application’s indexing strategy could record for each key every transaction ID that writes to the key. Such an index could be utilized by our banking example from earlier to implement a `get_statement` endpoint that returns all recent credits/debits given an account ID.

3.5 Verifiable receipts

Receipts are signed statements, associated with a given transaction, that can be verified and audited offline. Like the key-value store, receipts are utilized both internally by CCF, to ensure the validity of snapshots (subsection 4.4), and externally by users. Receipts can be leveraged either to audit CCF or to prove to a third party that a transaction was committed in a particular context and at a specified position in the ledger. One way users can generate a receipt is by invoking the built-in `receipt` endpoint with the transaction ID.

At its core, the receipt includes a signature over the Merkle root (from a subsequent signature transaction) and the sequence of vertices on the path from the transaction leaf to the root, known as a *Merkle proof*. In Figure 3, the Merkle proof for transaction 1.7 is $[(right, d_8), (left, d_{56}), (left, d_{1234}), (right, d_{910})]$. The Merkle proof can be used to verify that the transaction leaf, which includes the transaction ID and a digest of the write-set, was in the ledger at the time of the signature. The application logic may also choose to attach arbitrary claims to a transaction and thus its receipt to make them verifiable offline.

4 INTER-NODE ARCHITECTURE

Last section we focused on a single-node CCF network. This section describes how we extend that to a multi-node network, providing high availability in the presence of crash faults. Moreover, multi-node CCF networks also increase the throughput of read transactions, which can be handled by any node. CCF includes its own consensus layer for replicating transactions in the ledger, inspired by Raft [94] and adapted for a fault model more suited to trusted execution. This protocol requires n CCF nodes to tolerate $f = \lfloor \frac{n-1}{2} \rfloor$ crash faults. To build confidence in our approach, we used TLA+ [68] to formally specify and model check CCF’s consensus layer including reconfiguration [88].




Figure 5: Examples of five CCF ledgers. Underlined transactions IDs are signature transactions and the thick boxes highlight the committed transactions.

Table 2: Possible votes and primaries during an election, based on the ledgers shown in Figure 5 (left). For each possible candidate (row), the table shows which nodes might vote for it and thus whether it could win the election.

candidate	n_0	n_1	n_2	n_3	n_4	could win?
n_0	✓	✗	✗	✗	✗	✗
n_1	✓	✓	✗	✗	✗	✗
n_2	✓	✓	✓	✓	✓	✓
n_3	✓	✓	✗	✓	✓	✓
n_4	✓	✓	✗	✓	✓	✓

4.1 Transaction replication

Time is divided into a series of *views*. During most views, one of the nodes is the *primary* and the remaining nodes are the *backups*. The primary executes the user requests, appends the resulting transaction to its ledger, and replies to the users. The primary regularly replicates its ledger onto the other nodes using the integrity-protected `append_entries` remote procedure call (RPC).

As before, the primary regularly appends signature transactions. The primary tracks the matched sequence number on each node and once a signature transaction has been copied onto $\lceil \frac{n-1}{2} \rceil$ other nodes then it considers the signature transaction and all previous transactions as committed and updates its *commit sequence number* accordingly. This updated commit sequence number is included in future `append_entries` so that the backups can update their local commit sequence number.

Recall that each transaction ID uniquely identifies a transaction. Each `append_entries` includes a set of transactions and the previous transaction ID. The backup checks that the previous transaction ID matches its ledger before appending the new transaction. This check ensures that if any two ledgers contain a transaction with the same ID then the ledgers up to and including that transaction are identical. Figure 5 shows an example of the ledgers from five nodes, where node n_2 is the primary in view 3. Note that the last committed transaction is always a signature transaction and that all committed transactions are present in at least a majority of ledgers.

4.2 Primary elections

Eventually, one or more nodes will fail. If the failed node (or nodes) are backups, and $\lceil \frac{n-1}{2} \rceil$ backups remain then the primary will




Figure 6: Node states & transitions: PENDING, TRUSTED including PRIMARY, CANDIDATE, & BACKUP, RETIRING, and RETIRED.

continue as before. If the primary fails then a new primary will be elected, as illustrated by Figure 6.

Note that each node keeps track of its perspective on the current view, increasing it accordingly as messages are exchanged. Each RPC includes the sender's view, and the receiver will update its view if the sender's view is greater before processing the message. If the sender's view is less than the receiver's view then the receiver will not process the message. Instead, it will reply negatively to the message including its view so that the sender can update its own view. When a node increases its view, it becomes a backup if it is not already one. The primary sends regular heartbeats, in the form of empty `append_entries` RPCs to its backups. If a backup does not hear from a primary within its election timeout, it becomes a *candidate*, increments its view, and dispatches a `request_vote` RPC to all nodes including the view and sequence number of its last signature transaction. After checking the candidate's view, the recipient of the `request_vote` RPC checks if the candidate's ledger is at least as up-to-date as its own, and if so, provided it has not voted for another candidate this view, it will vote for the candidate. In doing so, the backup also resets its own election timeout.

A candidate's ledger is at least as up-to-date as the voter's ledger if (i) the view of the candidate's last signature transaction is greater than the view of the voter's last signature transaction, or (ii) the views are equal and the sequence number of the candidate's last signature transaction is greater or equal to the sequence number of the voter's last signature transaction.

An example of this voting criteria is shown in Table 2. n_0 cannot become primary as the other nodes would not vote for it and n_1 cannot become primary as only one other node (n_0) would vote for it and n_2 could become primary as all nodes would vote for it. n_3 and n_4 could become primary as n_0 and n_1 could vote for them, and they could vote for each other.

If a candidate receives votes from $\lceil \frac{n-1}{2} \rceil$ other nodes, then it becomes a primary. It discards any transactions after the last signature transaction and begins sending `append_entries` to its backups. The new view will begin with a signature transaction and the primary can update its commit sequence number to this transaction as soon as it has been replicated to the majority of ledgers. The new primary tracks the sequence number of the next transaction to send to each backup, starting with the first signature transaction of the next view. The ledgers of the backups might have fallen behind or even diverged from that of the new primary. Recall that when a backup receives an `append_entries` it checks that it has

the previous transaction ID. If not, it responds negatively to the primary including the sequence number of what it believes to be the latest common point. The primary decrements its sequence number for that primary, utilizing the information provided by the backup, and tries again until the `append_entries` is successful. If a backup has conflicting transactions then they are deleted before the new transactions are appended. In other words, the primary’s ledger is considered the ground truth when there are conflicts.

Consider the example on the right of Figure 5. Following the situation on the left of Figure 5, n_4 became the primary in view four. After rolling back transaction 3.5, n_4 added signature transaction 4.5 and appended it to n_0 , n_1 & n_2 . n_4 also appended transactions 4.6, 4.7, and signature transaction 4.8 to its ledger but has so far only replicated these to n_1 .

If a candidate’s election timer expires without it either stepping up to a primary or down to a backup then it restarts the election by incrementing its view once more and sending new `request_vote` RPCs. It might be the case that a new primary is elected whilst the previous primary is still alive. In this scenario, the new primary will eventually learn of the new view and will step down to a backup in the process. Multiple backups can become candidates at a similar time, splitting the votes. To reduce the likelihood of this, the election timer is randomly chosen from a configurable range. The primary also keeps track of the last time it received an `append_entries` response from each backup, and it steps down if it does not hear from at least $\lceil \frac{n-1}{2} \rceil$ backups within a specified time window. This ensures that a primary that cannot make progress, for instance due to a partial network partition, steps down cleanly, allowing another node to become primary, and does not produce an arbitrarily long ledger suffix which will never be committed.

4.3 User interaction

Users can connect to any CCF node, and backup nodes can forward user requests to the primary. When a node fails, the users connected to it will retry with other nodes. Read-only requests, including receipt requests, can be served by any CCF node and do not need to be forwarded to the primary. However, to ensure that CCF provides session consistency to its users, once a backup has forwarded a request to a primary it will forward all subsequent requests on the same TLS session, including reads, to the same primary, and terminate the session if this forwarding is not possible due to a primary change.

Note that each node maintains a view history, containing the start index for each view, according to its ledger. This history can help a node determine transaction status (Figure 4), for instance, a transaction has status `Invalid` if a greater view started at a smaller sequence number.

4.4 Atomic reconfiguration

When a node fails, the CCF network can tolerate one fewer fault than it could before. To restore the fault tolerance of the CCF network, the failed node should be replaced with a new node. This process, known as *reconfiguration*, allows an arbitrary transition from any node configuration to any other node configuration. For

Table 3: Sample of CCF’s built-in maps. Some names have been edited for brevity and all are prefixed with `public:ccf.gov` or `public:ccf.internal`.

Name	Purpose
<code>users.certs</code>	User certificates
<code>members.certs</code>	Consortium member certificates
<code>nodes.info</code>	Node identity certificates & properties
<code>nodes.code_ids</code>	Code versions that are allowed to join
<code>service.info</code>	Service identity certificate & service status
<code>signatures</code>	Merkle roots and signatures
<code>tree</code>	Merkle tree for historical receipts
<code>history</code>	Governance operations signed by members
<code>constitution</code>	CCF service constitution
<code>modules</code>	JavaScript application logic
<code>endpoints</code>	JavaScript endpoints
<code>proposals</code>	Governance proposals
<code>proposals_info</code>	Status of governance proposals & ballots
<code>ledger_secret</code>	Encrypted ledger secret
<code>recovery_shares</code>	Encrypted shares to recover ledger secret
<code>members_keys</code>	Members’ public encryption keys

instance, any number of nodes can be added, removed, or a combination of both, atomically changing the total number of nodes and thus changing the fault tolerance accordingly.

Reconfiguration is achieved by a reconfiguration transaction. Reconfiguration transactions are similar to signature transactions, in that they contain updates to built-in maps and affect the operation of a CCF node directly, distinguishing them from user-defined updates to application tables. For reference, Table 3 provides a sample of some built-in maps used by CCF.

A reconfiguration transaction includes writes to `nodes.info` which (among other things) maps each node’s ID to a node status (Figure 6). The current configuration is the set of nodes whose status is `TRUSTED`. Reconfiguration transactions can add a node by changing its status from `PENDING` to `TRUSTED` and remove a node by changing its status to `RETIRED`. The question of how nodes are first added to the map as `PENDING` and removed from the map is addressed in section 5.

Each node maintains a sorted list of *active configurations*, where each configuration contains the transaction ID of the reconfiguration transaction and the associated set of nodes. At the head of the list is the current configuration, followed by any pending configurations. Winning an election or committing a transaction requires a majority quorum in each active configuration, therefore `request_vote` and `append_entries` are sent by the candidate or primary to all nodes across the active configurations. As soon as a node appends a reconfiguration transaction to its ledger, regardless of whether the node is a primary or backup, that configuration is added to the active configurations. The reconfiguration transaction, like user transactions, is considered committed by the primary only when a subsequent signature transaction has been committed on all active configurations. When a node learns that a reconfiguration transaction has been committed, all earlier configurations are removed from the active configurations. A newly added node

begins participating in consensus, including potentially starting an election, as soon as it appends to its ledger the first signature transaction following the reconfiguration transaction that added it.

The process described thus far requires new nodes to replay all transactions in the history of the CCF service in order to become up to date. In practice, nodes can begin from a snapshot and use the consensus layer to simply learn the transactions since the last snapshot.

Recall that nodes add new configurations to the set of active configurations as soon as the associated reconfiguration transaction is appended to their ledger, they do not wait for the reconfiguration transaction to be committed or even followed by a signature transaction. Uncommitted transactions however can be rolled back following a view change. When this occurs, the rolled-back reconfigurations are removed from the set of active configurations. Note that the set of active configurations will always contain at least one configuration, the current configuration, which is either the initial configuration or the last committed reconfiguration, which cannot be rolled back.

4.5 Node retirement

The process of retiring nodes includes an additional step. The reconfiguration transaction first updates the node's status to `RETIRING`. As soon as the reconfiguration transaction is committed, the previous configurations are removed as before and the retired node is no longer part of the configuration. Before we shut down the retiring node, the primary must ensure that all future primaries also know that the node has been successfully retired. Therefore, once this reconfiguration transaction has been committed, the primary adds a second reconfiguration transaction to update the node's state to `RETIRED`. Once this second transaction is committed, then the retired node can be shut down. If there is a view change during reconfiguration, the new primary must commit the first reconfiguration transaction before appending the second reconfiguration transaction.

A primary may commit a reconfiguration transaction that retires itself. In this case, it should stop adding new transactions after the signature transaction which commits its retirement. Once the first reconfiguration transaction is committed, it should then stop sending heartbeats but remain online, replicating its ledger and voting for new primaries, whilst a new primary establishes itself. As before, the retiring node can be shut down once the new primary commits the reconfiguration transaction which sets its state to `RETIRED`.

5 MULTIPARTY GOVERNANCE

CCF is designed to support long-running applications and thus must remain available in the face of typical management operations such as node replacement, live code updates, changes to the set of users, and more. Such actions must be done with great care to preserve CCF's strong confidentiality and integrity properties. Moreover, since CCF services can be managed by multiple parties, these parties should have an opportunity to sign off on management operations before they are executed.

Table 4: Sample of the governance actions defined in CCF's default constitution [87].

Governance action	Purpose
<code>set_user</code>	Add user
<code>set_member</code>	Add consortium member
<code>set_js_app</code>	Update JS application logic
<code>add_node_code</code>	Allow a new code ID
<code>transition_node_to_trusted</code>	Allow node to join
<code>set_constitution</code>	Update constitution
<code>transition_service_to_open</code>	Allow user requests
<code>set_recovery_threshold</code>	Update recovery threshold

5.1 Governance proposals & ballots

At a high level, CCF's approach to multiparty governance is to require consortium members to oversee governance operations using governance proposals and ballots, which are processed by the programmable constitution.

The *constitution* is a contract between the consortium members describing all the available governance actions and the associated voting criteria. The constitution, encoded in JavaScript, is provided to a CCF service at start-up. The constitution defines a `resolve` function, which takes a governance proposal (defined below) and votes by consortium members, and determines if the proposal has been accepted. The constitution also defines `apply`, which takes an accepted proposal and executes the governance actions within it to modify the key-value store. The default constitution [87] accepts a governance proposal once it has received votes from a strict majority of consortium members. Alternative constitutions could give consortium members different voting power, for instance giving a few consortium members, who are the main parties running the application, veto power. Alternatively, a constitution could vary the number of votes needed depending on the action, for instance giving one consortium member, who is also an operator, the unilateral power to add or remove nodes. The constitution can also be updated using a governance proposal (if permitted by the current constitution).

A *governance proposal* consists of a set of actions, defined in the constitution, and a *ballot* is a vote for or against a particular proposal, conditional on the proposal itself and the current state of the key-value store. Similar to user transactions, governance proposals and ballots are recorded on the underlying CCF ledger. However, unlike user transactions, they always originate from a request signed by a consortium member. This signature is also stored on the ledger.

Table 4 shows a sample of the governance actions defined in the `apply` function of the default constitution and Listing 1 shows the implementation of one such action. This action, `add_node_code`, can be used when updating the code inside the TEE (CCF itself or its C++ application logic), and adds the given code ID to the set of trusted code IDs. This is done by writing the code ID to the nodes `code_ids` map (listed in Table 3) in the key-value store. That map is then read from when a new node attempts to join the service, to establish whether they are running a trusted version of the code.

```
1 ["add_node_code", new Action(  
2   function (args) {  
3     checkType(args.code_id, "string", "code_id") },  
4   function (args, proposalId) {  
5     ccf.kv["public:ccf.gov.nodes.code_ids"].set(args.code_id, "AllowedToJoin");  
6     invalidateOtherOpenProposals(proposalId)}}]
```

Listing 1: An example action from the apply function in CCF’s default constitution. Adapted for brevity.

Governance proposals have the power to change many fundamental properties of the service, not least modifying the constitution itself. Any errors, deliberate or otherwise, are therefore potentially damaging. Consortium members should examine proposals carefully before submitting their ballot to ensure they do not leak confidential state. Proposals are encoded as succinct JSON documents so that they are easy to inspect offline. Note that the default constitution [87] provided by CCF should be suitable for most deployments.

5.2 Disaster recovery protocol

Governance and reconfiguration require consensus and thus majority agreement. If more than a majority of nodes fail simultaneously then CCF cannot make progress. In this case, the disaster recovery protocol can enable the service to restart. Note that disaster recovery restores the CCF service to a safe state, even from just a single copy of the ledger files, but it is not guaranteed to preserve all committed transactions. This best-effort approach is necessary when any consensus-based service wishes to recover from the persistent storage of just a minority of nodes. The newly recovered service will have a new service identity (defined in Table 1), making it clear to users that a disaster recovery has occurred, and thus that the service may have been rolled back to a previous state.

During disaster recovery, a node or set of nodes is started in recovery mode with a set of ledger files from the previous CCF service (and optionally, snapshot files to speed things up). The public parts of transactions are restored, and further nodes may join the network. A consortium member will submit a proposal to open the new service, specifying the old and new service identities to ensure this proposal is bound to a precise recovery, and once the proposal receives sufficient approving ballots, the service is opened.

In order to decrypt the updates to private maps, consortium members must submit their recovery shares to the CCF service. Recall that updates to private maps are encrypted using the ledger secret, as defined in Table 1. The ledger secret is recorded in the ledger, encrypted using the *ledger secret wrapping key*. The ledger secret wrapping key is itself stored in the ledger after being split into n shares, known as *recovery shares*, where n can be up to the number of members in the consortium. Each recovery share is first encrypted using one of the consortium member’s public encryption keys. The ledger secret wrapping key is divided into k -of- n shares [108], such that n recovery shares are created, less than k shares leak nothing about the ledger secret, and any k shares suffice to recover the ledger secret. Once these shares are submitted by the consortium members, the ledger secret can be reconstructed in TEE, and the previous ledger’s private state decrypted. Once

this is complete, the CCF service is fully recovered. The parameter k , known as the recovery threshold, can be configured through governance. Table 3 includes a summary of the built-in maps utilized by CCF’s disaster recovery procedure.

6 DISCUSSION

We set out to investigate whether it is possible to develop a pragmatic solution for deploying confidential, integrity protected, and highly available multiparty applications on untrusted cloud infrastructure. We now evaluate CCF along these dimensions.

6.1 Confidentiality

Keeping data confidential requires protection during every stage of a CCF service. Firstly, data is encrypted in transit between users and CCF nodes with a TLS connection, using the service identity certificate (defined in Table 1), and terminating inside the TEE. Data is encrypted during execution and in memory by the TEEs. Data is encrypted at rest (on disk) and in transit between CCF nodes as the transactions in the ledger are encrypted with the ledger secret (Table 1). The service certificate private key and ledger secret are not shared with new CCF nodes until their attestation has been verified. Whilst the application logic has read-write access to the key-value store (except CCF’s read-only internal/governance maps), the application logic can limit which users can invoke specific endpoints using user authentication policies. The application logic and set of authorized users can only be modified via a governance proposal. There is a fundamental tradeoff between confidentiality and auditability; some transactions, including all governance operations, are written unencrypted to permit offline audit, while application data which does not require such audit can be encrypted by the ledger secret.

6.2 Integrity

Whilst the ledger secret is the primary mechanism for providing confidentiality, it is the signature transactions (subsection 3.2) that provide integrity protection to the ledger whilst outside the TEEs. The Merkle tree constructed for the signature transactions is again utilized to provide Merkle proofs for receipts (subsection 3.5). Integrity protection with signature transactions ensures that a malicious party cannot modify the ledger undetected whilst it is in persistent storage, however, the ledger could be rolled back to a previously valid prefix. Notably, CCF’s design avoids the need to depend upon a dedicated rollback protection mechanism. Existing rollback protection approaches include SGX trusted monotonic counters [24] (which have since been deprecated [53]), using TPMs [53, 75], or a service based on replication across TEEs such as ROTE [77], LCM [21], Tiks in Engraft [126], NARRATOR [90] or Nimble [4]. Rollback protection is also closely related to forking protection mechanisms and likewise, CCF avoids the need for them by design. The architecture of CCF treats nodes as ephemeral. When a node crashes, it cannot simply resume from persistent storage as we cannot be sure that its state is up-to-date. Instead, the CCF node rejoins the service via reconfiguration (subsection 4.4) and establishes a fresh node identity. This stricter fault model as well as the need for integrity protection are the motivation behind CCF’s

choice of a custom consensus protocol over a classic approach such as Multi-Paxos [67] or Raft [94].

CCF adopts a transparent approach to management where governance proposals and ballots are recorded in public maps and signed by consortium members to enable auditability and disaster recovery (subsection 5.1).

6.3 High availability

CCF achieves high availability through replication across CCF nodes (section 4). The consensus layer in CCF uses strict majority quorums thus it can continue to operate even when up to a minority of nodes fail simultaneously. Users can simply connect to any CCF node, requests are handled locally if they are read-only, including historical transactions and receipt generation, and forwarded to the primary otherwise (subsection 4.3). If the CCF node that a user is connected to fails, then the user can retry with another node.

Failed nodes can be replaced by reconfiguration to restore the fault tolerance of the service (subsection 4.4). For instance, a CCF service with five nodes can tolerate two faults, if one node fails then the service can only tolerate one further fault. Reconfiguring to remove the failed node and add a new healthy node restores the fault tolerance of the service back to two faults. In addition to recovering from failures, reconfiguration can also be used proactively to migrate the service, to change the replication factor, or for code updates, either to CCF itself or to the application logic. CCF's consensus layer supports arbitrary reconfiguration, moving from any set of nodes to any other set of nodes, without stopping the service. CCF also achieves reconfiguration in just a single transaction, with a second transaction only needed to confirm node removal. The classic approaches to reconfiguration in Raft either requires nodes to be added or removed one at a time [38, 93], or always requires two transactions [94]. Our design is particularly useful as it enables CCF services to minimize the window of vulnerability until the usual degree of resilience is restored. If more than a majority of nodes are lost, CCF can recover from persistent storage through its disaster recovery procedure (subsection 5.2). During disaster recovery, the service may be reverted to a prior state (a suffix of the ledger may be lost).

The constitution (subsection 5.1) defines what groups of consortium members must vote in favor of a proposal for it to be accepted, for instance, a majority quorum. Votes are recorded and calculated using the ledger and thus consortium members do not need to be online simultaneously to propose and accept a proposal. Proposing and voting for a proposal can be done on the fly and does not require stopping the service. If a sufficient number of consortium members lose their keys, then governance operations and disaster recovery will no longer be possible.

6.4 Pragmatism

CCF is designed to make the deployment of trustworthy multiparty applications on untrusted infrastructure as simple as possible. CCF supports custom application logic that can be written in either JavaScript/TypeScript or C++, providing greater flexibility compared to permissioned blockchains which often support DSLs such as Solidity, Move [17], or Daml. The programming model is simple, application logic exposes endpoints that are invoked by users,

and state changes are recorded by a transaction over the built-in data store. Transactions are applied to the key-value store atomically and in isolation thus the application will never observe partial transaction execution. Application logic is not required to be deterministic, it may be executed multiple times but its associated transaction is applied exactly once. The constitution is also fully programmable, allowing the consortium members granular control over governance operations. The constitution, written in JavaScript, is provided at start-up and can be updated by a governance proposal (if supported by the current constitution).

Many aspects of CIA support in CCF are configurable, allowing CCF to adapt to different deployment scenarios. Maps in the key-value store can be either private (encrypted on the ledger) or public (plain-text). Public maps are not confidential but can be audited without decrypting the ledger, and thus can be audited offline. Application logic maps are private by default, and public maps are also supported. CCF's built-in maps are public for transparency. CCF can be deployed on hosts without TEEs using *virtual* mode. This is useful for development and, in practice, for replication if confidentiality and integrity are not needed. CCF in virtual mode still provides highly available state machine replication with advanced features such as atomic reconfiguration, indexing, snapshots, historical transactions, disaster recovery, receipts, and multiparty governance. Likewise, CCF can be deployed on a single node if high availability is not needed.

Users of CCF can determine on a per-request basis the consistency/durability guarantee required. CCF provides two levels of guarantee: local execution & global commit. Users receive a reply with a transaction ID as soon as it has been executed locally. The user can determine after having observed the response if it wishes to proceed without checking for commit. Likewise, receipts can be expensive to retrieve and thus are issued on demand, at the user's request, again using the transaction ID. There is optional support for user request signing, via the same mechanism that consortium members sign governance operations. CCF also allows applications to define their own indexing strategy.

CCF currently supports Intel SGX, and it is designed with a view to supporting more TEE platforms, for instance, by minimizing dependencies on features specific to SGX (or specific versions of SGX) e.g. sealing [48], trusted time and monotonic counters [24], or ECALL/OCALL interface provided by SGX.

7 IMPLEMENTATION & EVALUATION

CCF is written in C++20 and built on the OpenEnclave SDK [37]. The host and the TEE communicate via a pair of lock-free multi-producer single-consumer ringbuffers to minimize the expensive transitions to/from the TEE [10, 49, 96, 128]. CCF uses a host-side thread, a TEE-side thread, and optionally, a TEE-side thread pool with a configurable number of worker threads. CCF's REST API supports HTTP/1.1 and HTTP/2 with a custom HTTP header in responses for the transaction ID. Built-in endpoints use JSON and COSE payloads. CCF supports TLS 1.2 & 1.3 using OpenSSL 1.1.1 [95]. Request signing by consortium members, and optionally by users, supports HTTP signatures [23] and COSE Sign1 [105].




Figure 7: Impact of the number of CCF nodes on writes (left) and reads (center) and the impact of read/write request ratio on single node throughput (right).

Table 5: Throughput (tx/s) for writes/reads on a five node service. Virtual mode refers to CCF running without SGX.

	SGX	Virtual
C++	64.8 K / 881 K	118 K / 1.24 M
JS	15.7 K / 90.7 K	33.7 K / 219 K

User authentication is implemented using JWT (JSON Web Tokens) [57] and by X.509 certificates [18]. Our Merkle tree implementation [81] uses SHA-256 [91]. The ledger secret uses AES256-GCM [104] to encrypt the transactions on the ledger. AES256-GCM is also used to encrypt any persistent storage utilized by the indexer. Diffie-Hellman key exchange [35] is used for node-to-node message headers and message forwarding. The encryption of ledger secrets uses RSA OAEP [16]. The map implementation in CCF is based on the Compressed Hash-Array Mapped Prefix-tree (CHAMP) [116]. CCF’s JavaScript engine is QuickJS [15].

Experiment Setup. Unless specified otherwise, our setup was as follows. We used DC16s_v3 (16 cores, 128 GB memory) VMs [86] in the UK South region of Azure, running Ubuntu 20.04. Our Ice Lake CPUs support Intel SGX2 [55]. Our C++ application logic implements a simple logging application, where messages with corresponding identifiers are posted, and later retrieved with read-only transactions. Messages are private and 20 characters each. We use four VMs: three CCF nodes, configured with 10 worker threads, with one user VM, simulating sufficient users (closed loop with up to 1k concurrent requests) to stress test CCF. To measure the performance of CCF itself, instead of the optional node-to-node forwarding logic, the user directly writes to the primary.

Figure 7 shows that CCF consistently achieves a throughput of at least 65 thousand requests per second. Increasing the ratio of reads to writes slightly increases performance. This effect is most notable for the 100% read workload, where all user requests are handled directly by the CCF node that receives them since only writes must be handled by the primary. We also observed similar performance using public maps instead of private ones.

Table 5 illustrates the significant performance difference between the C++ implementation of the logging application and the JavaScript implementation. Table 5 also demonstrates the significant overhead of Intel SGX. We hope that future TEEs will substantially reduce this overhead such that performance is comparable to that of




Figure 8: Impact of the signatures on response time (left & center) and write throughput (right).

non-confidential offerings. This is supported by early benchmarks of AMD SEV-SNP which report a performance overhead of just 2–8% [3].

Figure 8 investigates the performance impact of generating signature transactions. To measure the overhead of signatures, the signature transaction frequency has been set to every 100 transactions and most other sources of latency variance removed, for instance, we use one node and one user. We then measure the *response time*, the time for the user to receive a positive response from the CCF node, including a transaction ID that can then be used to check for commit. We observe that the response time of write requests is 1.2–1.3 ms, with the exception that approximately every 100 requests, a signature is triggered, and the response time increases to 2.3 ms. The graph on the right shows how the throughput of writes is impacted by the frequency of signature transactions, demonstrating the tradeoff between time to commit (which requires a signature), and overall throughput.

Figure 9 shows the impact of a primary failure and subsequent re-configuration on a CCF service with three initial nodes $\{n_0, n_1, n_2\}$, and three consortium members $\{m_0, m_1, m_2\}$ using the default constitution. Listing 2 also shows five governance transactions from the ledger in Figure 9. The graph shows two users, one sending writes to n_0 , the initial primary, and the other sending reads to n_1 , one of the two backups. At A, we kill n_0 , writes stop immediately as the service is without a primary and the reads continue (in fact their throughput increases as the backup is no longer handling messages from the primary). n_2 is then elected primary and writes resume. During this time our test infrastructure, simulating an operator, detects the failure and begins preparing a new node, n_3 , by copying the snapshots onto the new host, n_3 tries to send a join message to the primary and succeeds at B. At C, our test infrastructure starts a governance proposal (p_3 in Listing 2) as consortium member m_0 to add n_3 to the network and remove n_0 . Consortium members m_0 and m_1 then submit ballots to support the proposal p_3 . By D the governance proposal has been accepted and n_3 will now be added to the current configuration at the consensus layer. At E, the reconfiguration has been completed, the fault tolerance of the service has been restored, and n_0 can be retired by the operator.

8 RELATED WORK

We divide the state-of-the-art into techniques that do not leverage trusted execution (subsection 8.1) and those that do (subsection 8.2).




Figure 9: Impact of primary failure (A) and subsequent node replacement (B-E) on the availability of reads and writes.

```
Bal = "export function vote (proposal, proposer_id) {return true}"
txid 3.198408: // Point B on Figure 9
  map public:ccf.gov.nodes.info:
    n3: {status: Pending}
txid 3.207341: // Point C on Figure 9
  map public:ccf.gov.proposals:
    p3: {actions: [{name: transition_node_to_trusted, args: {
      node_id: n3}}, {name: remove_node, args: {node_id: n0}}]}
  map public:ccf.gov.proposals.info:
    p3: {ballots: {}, proposer_id: m0, state: Open}
txid 3.208098:
  map public:ccf.gov.proposals.info:
    p3: {ballots: {m0: Bal}, state: Open}
txid 3.209096: // Point D on Figure 9
  map public:ccf.gov.proposals.info:
    p3: {ballots: {m1: Bal, m0: Bal}, final_votes: {m1: true, m0:
      true}, state: Accepted}
  map public:ccf.gov.nodes.info:
    n0: {status: Retiring}
    n3: {status: Trusted}
txid 3.210844: // Point E on Figure 9
  map public:ccf.gov.nodes.info:
    n0: {status: Retired}
```

Listing 2: Sample of the key updates from the ledger in Figure 9, where n_0 is replaced with n_3 . Adapted for brevity.

8.1 Approaches without TEEs

Raft [94] does not provide confidentiality or integrity on untrusted infrastructure as the protocol only targets high availability. The system model of Raft assumes that nodes can only fail by crashing, and thus does not handle other faults, for instance, nodes subverting the consensus protocol to reorder transactions or fork the ledger. The same is true of other crash fault-tolerant (CFT) consensus algorithms [67] and atomic broadcast protocols [58].

TEEs can help systems to achieve integrity and confidentiality, however, simply executing a CFT protocol within TEEs does not necessarily provide the desired CIA properties. For instance, Engraft [126] uses model checking to prove that Raft within TEEs does not provide integrity. A key concern is that the persistent storage is outside the trust boundary and thus under the control of the untrusted host, which can truncate the ledger to violate Raft’s linearizability guarantees. Moreover, CCF also provides features, such as receipts and multiparty governance, which are beyond the scope of Raft.

Unlike their CFT counterparts, Byzantine fault-tolerant (BFT) protocols, can tolerate arbitrary faults including malicious nodes. However, safety is only guaranteed by BFT protocols provided that a threshold of participants, often two-thirds, are non-faulty. This limitation is fundamental to BFT protocols, yet it is unrealistic for cloud environments where the parties involved are not

all running their own nodes. BFT protocols are also typically less available than CFT protocols, as they require at least two-thirds of nodes to be available to make progress and some BFT protocols also do not support reconfiguration, which is key for restoring fault tolerance after a node failure. BFT protocols are computationally expensive relative to CFT protocols and BFT protocols also require the re-execution of user requests across nodes, so the application logic must be deterministic, making it difficult to support general-purpose programming languages. Distributed ledger technologies and permissioned blockchains [27, 30, 46] often utilize BFT state machine replication to provide distributed multiparty applications with integrity provided that a threshold of participants are non-faulty. Many blockchains struggle with poor performance [41], high energy footprints [98], and typically they do not provide confidentiality, although zero knowledge proofs [40] may be used to avoid recording private data on blockchains.

Fully homomorphic encryption, arbitrary computation over encrypted data, is an active research area, but it is not yet a practical option [76] for most applications. Secure Multiparty Computation provides both confidentiality and integrity but can be expensive and is highly application-specific [127].

8.2 TEE-based approaches

At the time of writing, there are various TEE platforms, at differing stages of development. These include ARM CCA [8] with TrustZone [7] and Realms [71], Intel SGX [31] and TDX [51], RISC-V Keystone [69], and AMD SEV [61], SEV-ES [60], and SEV-SNP [50]. Like most systems discussed in this section, CCF opted for Intel SGX due to its availability at the time of development, SDK support [37, 54, 119], and its small TCB relative to VM-based approaches. Early Intel SGX offerings had highly limited enclave memory (128 MB) but this is no longer the case [36], for instance, the 3rd Gen Intel Xeon Scalable Processors support up to 512 GB [52]. In January 2022, Intel reaffirmed its commitment to SGX on its server-grade processors [101]. There are two general approaches to building trustworthy data systems leveraging TEEs: lift-and-shift and enclave-aware.

Lift-and-shift approaches focus on porting existing applications to a TEE with minimal modifications. Examples include Haven [13], Graphene [120, 121], SCONE [10], Occlum [110], Ryoan [44], Rattel [32], S-FaaS [1], and Parma [56]. Cloud providers are beginning to offer lift-&-shift as a service, in the form of Confidential VMs or Confidential containers [99]. Kata containers [29] has experimental support for AMD SEV-SNP [28] and lift-and-shift is already being utilized, for instance, by Constellation [118]. Lift-and-shift is easy to adopt for stateless applications but has a large trusted computing base (TCB), violating the principle of least privilege, and is more challenging for applications requiring durability.

Enclave-aware approaches are new system designs with a focus on dividing code between the trusted enclave and the untrusted host. This approach has a higher barrier to entry compared to lift-&-shift solutions but also a smaller TCB. Some systems in this space are highly application specific [22, 70, 106, 111]. Others are more application agnostic, for instance, recent efforts to put privacy into blockchains with off-chain processing on TEEs [2, 26, 47, 66, 131], a TEE-based file system [65], or Glamdring [73] which allows

Table 6: Comparison of data systems utilizing TEEs, ordered by publication year. (C) Confidentiality, (E) Execution Integrity, (D) Data Integrity, (A) High Availability, (R) Reconfiguration, (U) Untrusted infrastructure, (OS) Open Source.

	C	E	D	A	R	U	OS
Hybster [14]		✓	✓	✓		✓	
Fabric PC [20]	✓	✓	✓	✓	✓		✓
Pesos [64]	✓		✓	✓		✓	
EnclaveDB [100]	✓	✓	✓			✓	
VeritasDB [113]			✓			✓	
Ekiden [26]	✓	✓	✓	✓	✓		
ShieldStore [63]	✓		✓			✓	✓
Speicher [12]	✓		✓			✓	
Azure SQL AE [5]	✓			✓	✓	✓	
Avocado [11]	✓		✓	✓	✓	✓	✓
Enclave [117]	✓					✓	
FastVer [6]			✓			✓	
HotStuff-M/VABA-M [132]			✓	✓		✓	
Precursor [79]	✓		✓			✓	
(Chained) Damysus [33]			✓	✓		✓	✓
Engraff [126]	✓	✓	✓	✓		✓	✓
FlexiTrust [43]			✓	✓		✓	
Treaty [39]	✓		✓			✓	✓
Data Station [130]	✓	✓	✓			✓	
CCF	✓	✓	✓	✓	✓	✓	✓

developers to annotate sensitive data and automatically partitions the application between the enclave and host.

Table 6 compares CCF to similar TEE-based approaches, along the dimensions of data confidentiality, integrity protection (divided into execution and data), high availability, support for reconfiguration, and support for untrusted infrastructure. All approaches listed are based on Intel SGX, except for Data station [130] which opted for AMD SEV.

Avocado [11] is a confidential, integrity-protected key-value store, built for Intel SGX using SCONE [10]. Unlike CCF, Avocado is purely in-memory with a node replacement protocol based on Hermes [62] and does not support disaster recovery. Avocado uses multi-writer ABD [74] for replication instead of a distributed consensus algorithm. Avocado supports sharding on individual operations since Avocado does not support transactions so no cross-shard transaction support is needed. Like many systems utilizing TEEs, its architecture is heavily influenced by the limited enclave memory in early SGX offerings and the high cost of enclave page cache paging. Similarly, Treaty [39] is a confidential, distributed key-value store built on SGX. The distribution layer is built on top of Speicher [12] and like Avocado, its networking stack is based on eRPC [59]. Treaty uses two-phase commit for replication instead of a consensus algorithm such as Raft. Treaty supports sharding but does not support server-side application logic, so users submit interactive transactions directly. Protection against rollback attacks is achieved using a trusted counter. FastVer [6] adds integrity protection to the Faster [25] key-value store, using a combination of trusted execution, Merkle trees, and deferred memory verification.

Whilst FastVer does not have many of the features of CCF such as confidentiality or custom application logic, it does achieve outstanding performance. FlexiTrust [43] is motivated by the claim that current TEE protocols have three issues compared to classic BFT protocols: liveness attacks, rollback attacks, and sequential consensus. FlexiTrust tries to address these. It is a BFT protocol for TEEs that uses $3f + 1$ nodes but only one phase (usually CFT protocols require one phase but BFT protocols require two) and a TEE is only used by the primary. FlexiTrust provides no confidentiality guarantees but does perform well (up to 200K txns/sec). Hyperledger Fabric [46] is a permissioned blockchain, whose nodes are divided into peers, which execute the chaincode (smart contract/app logic), and the orderers, which use Raft (implementation from etcd) to order transactions on the ledger. Hyperledger Fabric Private Chaincodes (PC) [20, 45] is a proposed extension to Fabric that uses Intel SGX for private execution by modifying the peers to execute chaincode inside an enclave. With Fabric PC, like Ekiden [26], the consensus layer remains outside the TEEs.

We have released a technical report [102] describing an earlier iteration of CCF and a subsequent paper on a technique for providing accountability following TEE compromise [109].

9 CONCLUSION & FUTURE WORK

CCF is a general-purpose platform for constructing confidential, integrity protected & highly available services, with support for delegation to untrusted cloud infrastructure and governance by mutually untrusted parties. CCF is a production quality, open source project that we believe has substantial potential as a foundation for new systems research into trustworthy computing, and it is already being built upon [66, 80, 83, 84]. CCF offers an all-in-one platform for deploying trustworthy multiparty applications on untrusted infrastructure. Moreover, aspects of CCF’s design, such as its disaster recovery procedure, atomic reconfiguration, and governance model, could also be useful techniques in their own right.

CCF is by no means the final word on confidential computing. Our trusted computing base still includes the hardware manufacturer and research is ongoing into verifying the correctness of the code and protecting against side-channel attacks [42, 92]. Intel SGX has also been the subject of various exploits [19, 89, 107, 115, 122–125], capable of compromising the confidentiality and/or integrity of execution. We have applied the available mitigations, however, further exploits may be uncovered. SGX also has a significant performance overhead, as shown in Table 5, and is not universally available. In response, we have recently added experimental support for another trusted execution platform, AMD SEV-SNP, and we aim to support future offerings such as Intel TDX. CCF’s threat model is a significant improvement over the current state of the art and a substantial step in the right direction. The removal of cloud providers from the trusted computing base opens up a myriad of new possibilities for reimagining multiparty distributed systems that leverage the cloud’s compute capability whilst minimizing trust in cloud providers.

ACKNOWLEDGMENTS

We are sincerely grateful to the following people for their contributions to CCF at various stages of the project: Christine Avanesians,

Dominic Ayre, Jacqueline Aziz, Theodore Butler, Miguel Castro, Mahati Chamarthy, Joonwon Choi, Renato Golin, Istvan Haller, Syed Hamza, Sid Krishna, Sangeeth Kumar, Paul Liétar, Thomas Moscibroda, Kartik Nayak, Olga Ohrimenko, Joe Powell, Jonathan Protzenko, Maik Riechert, Takuro Sato, Felix Schuster, Roy Schuster, Alex Shamis, Lyndon Shi, Ross P. Smith, Cale Teeter, Shoko Ueda, and Olga Vrousitou. We would also like to thank the contributors to Open Enclave [37], Confidential Containers [56], and smalloc [72]. Finally, we would also like to thank the anonymous VDLB reviewers for their insightful feedback.

REFERENCES

- [1] Fritz Alder, N. Asokan, Arseny Kurnikov, Andrew Paverd, and Michael Steiner. 2019. S-FaaS: Trustworthy and Accountable Function-as-a-Service Using Intel SGX. In *Proceedings of the 2019 ACM SIGSAC Conference on Cloud Computing Security Workshop* (London, United Kingdom) (CCSW'19). Association for Computing Machinery, New York, NY, USA, 185–199. <https://doi.org/10.1145/3338466.3358916>
- [2] Enterprise Ethereum Alliance. 2021. Enterprise Ethereum Alliance Off-Chain Trusted Compute Specification v1.1. <https://entethalliance.github.io/trusted-computing/spec.html> [Last accessed: 2023-Oct-06].
- [3] AMD. 2021. Microsoft Azure Confidential Computing Powered by 3rd Gen EPYC CPUs. <https://community.amd.com/t5/business/microsoft-azure-confidential-computing-powered-by-3rd-gen-epyc/ba-p/497796> [Last accessed: 2023-Oct-06].
- [4] Sebastian Angel, Aditya Basu, Weidong Cui, Trent Jaeger, Stella Lau, Srinath Setty, and Sudheesh Singanamalla. 2023. Nimble: Rollback Protection for Confidential Cloud Services. In *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI '23)*. USENIX Association, Boston, MA, 193–208. <https://www.usenix.org/system/files/osdi23-angel.pdf>
- [5] Panagiotis Antonopoulos, Arvind Arasu, Kunal D. Singh, Ken Eguro, Nitish Gupta, Rajat Jain, Raghav Kaushik, Hanuma Kodavalla, Donald Kossmann, Nikolas Ogg, Ravi Ramamurthy, Jakub Szymaszek, Jeffrey Trimmer, Kapil Vaswani, Ramarathnam Venkatesan, and Mike Zwilling. 2020. Azure SQL Database Always Encrypted. In *Proceedings of the 2020 International Conference on Management of Data* (Portland, OR, USA) (SIGMOD '20). Association for Computing Machinery, New York, NY, USA, 1511–1525. <https://doi.org/10.1145/3318464.3386141>
- [6] Arvind Arasu, Badrish Chandramouli, Johannes Gehrke, Esha Ghosh, Donald Kossmann, Jonathan Protzenko, Ravi Ramamurthy, Tahina Ramananand, Aseem Rastogi, Srinath Setty, Nikhil Swamy, Alexander van Renen, and Min Xu. 2021. FastVer: Making Data Integrity a Commodity. In *Proceedings of the 2021 International Conference on Management of Data* (Virtual Event, China) (SIGMOD '21). Association for Computing Machinery, New York, NY, USA, 89–101. <https://doi.org/10.1145/3448016.3457312>
- [7] ARM. 2013. GlobalPlatform based Trusted Execution Environment and TrustZone Ready - The foundations for trusted services. https://community.arm.com/cfs-file/_key/telligent-evolution-components-attachments/01-2142-00-00-00-00-51-36/GlobalPlatform-based-Trusted-Execution-Environment-and-TrustZone-R.pdf [Last accessed: 2023-Oct-06].
- [8] ARM. 2023. Learn the architecture - Introducing Arm Confidential Compute Architecture - Issue 2.0. <https://developer.arm.com/documentation/den0125/latest> [Last accessed: 2023-Oct-06].
- [9] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. 2010. A View of Cloud Computing. *Commun. ACM* 53, 4 (apr 2010), 50–58. <https://doi.org/10.1145/1721654.1721672>
- [10] Sergei Arnaudov, Bohdan Trach, Franz Gregor, Thomas Knauth, Andre Martin, Christian Priebe, Joshua Lind, Divya Muthukumaran, Dan O’Keeffe, Mark L. Stillwell, David Goltzsche, David Eyers, Rüdiger Kapitza, Peter Pietzuch, and Christof Fetzer. 2016. SCONe: Secure Linux Containers with Intel SGX. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation* (Savannah, GA, USA) (OSDI'16). USENIX Association, USA, 689–703. <https://www.usenix.org/system/files/conference/osdi16/osdi16-arnaudov.pdf>
- [11] Maurice Bailleu, Dimitra Giantsidi, Vasilis Gavrielatos, Do Le Quoc, Vijay Nagarajan, and Pramod Bhatotia. 2021. Avocado: A Secure In-Memory Distributed Storage System. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, USA, 65–79. <https://www.usenix.org/system/files/atc21-bailleu.pdf>
- [12] Maurice Bailleu, Jörg Thalheim, Pramod Bhatotia, Christof Fetzer, Michio Honda, and Kapil Vaswani. 2019. Speicher: Securing LSM-Based Key-Value Stores Using Shielded Execution. In *Proceedings of the 17th USENIX Conference on File and Storage Technologies* (Boston, MA, USA) (FAST'19). USENIX Association, USA, 173–190. <https://www.usenix.org/system/files/fast19-bailleu.pdf>
- [13] Andrew Baumann, Marcus Peinado, and Galen Hunt. 2015. Shielding Applications from an Untrusted Cloud with Haven. *ACM Trans. Comput. Syst.* 33, 3, Article 8 (aug 2015), 26 pages. <https://doi.org/10.1145/2799647>
- [14] Johannes Behl, Tobias Distler, and Rüdiger Kapitza. 2017. Hybrids on Steroids: SGX-Based High Performance BFT. In *Proceedings of the Twelfth European Conference on Computer Systems* (Belgrade, Serbia) (EuroSys '17). Association for Computing Machinery, New York, NY, USA, 222–237. <https://doi.org/10.1145/3064176.3064213>
- [15] bellard. 2021. QuickJS. <https://bellard.org/quickjs/> [Last accessed: 2023-Oct-06].
- [16] Mihir Bellare and Phillip Rogaway. 1995. Optimal asymmetric encryption - How to Encrypt with RSA. In *Advances in Cryptology - EUROCRYPT '94*, Alfredo De Santis (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 92–111. <https://doi.org/10.1007/BFb0053428>
- [17] Sam Blackshear, Evan Cheng, David L. Dill, Victor Gao, Ben Maurer, Todd Nowacki, Alistair Pott, Shaz Qadeer, Rain, Dario Russi, Stephane Sezer, Tim Zekian, and Runtian Zhou. 2019. Move: A Language With Programmable Resources. Technical Report. Calibra. <https://developers.diem.com/papers/diem-move-a-language-with-programmable-resources/2019-06-18.pdf> [Last accessed: 2023-Oct-06].
- [18] Sharon Boeyen, Stefan Santesson, Tim Polk, Russ Housley, Stephen Farrell, and David Cooper. 2008. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280. <https://doi.org/10.17487/RFC5280>
- [19] Pietro Borrello, Andreas Kogler, Martin Schwarzl, Moritz Lipp, Daniel Gruss, and Michael Schwarz. 2022. \mathcal{A} EPIC Leak: Architecturally Leaking Uninitialized Data from the Microarchitecture. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 3917–3934. <https://www.usenix.org/system/files/sec22-borrello.pdf>
- [20] Marcus Brandenburger, Christian Cachin, Rüdiger Kapitza, and Alessandro Sorniotti. 2018. Blockchain and Trusted Computing: Problems, Pitfalls, and a Solution for Hyperledger Fabric. arXiv:1805.08541 [cs.DC] <https://arxiv.org/abs/1805.08541>
- [21] Marcus Brandenburger, Christian Cachin, Matthias Lorenz, and Rüdiger Kapitza. 2017. Rollback and Forking Detection for Trusted Execution Environments Using Lightweight Collective Memory. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE Computer Society, USA, 157–168. <https://doi.org/10.1109/DSN.2017.45>
- [22] Stefan Brenner, Colin Wulf, David Goltzsche, Nico Weichbrodt, Matthias Lorenz, Christof Fetzer, Peter Pietzuch, and Rüdiger Kapitza. 2016. SecureKeeper: Confidential ZooKeeper Using Intel SGX. In *Proceedings of the 17th International Middleware Conference* (Trento, Italy) (Middleware '16). Association for Computing Machinery, New York, NY, USA, Article 14, 13 pages. <https://doi.org/10.1145/2988336.2988350>
- [23] Mark Cavage and Manu Sporny. 2019. *Signing HTTP Messages*. Internet-Draft draft-cavage-http-signatures-12. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-cavage-http-signatures/12/> Work in Progress.
- [24] Shanwei Cen and Bo Zhang. 2017. *Trusted Time and Monotonic Counters with Intel Software Guard Extensions Platform Services*. Technical Report. Intel. https://community.intel.com/legacyfs/online/drupal_files/managed/1b/a2/Intel-SGX-Platform-Services.pdf [Last accessed: 2023-Oct-06].
- [25] Badrish Chandramouli, Guna Prasaad, Donald Kossmann, Justin Levandoski, James Hunter, and Mike Barnett. 2018. FASTER: A Concurrent Key-Value Store with In-Place Updates. In *Proceedings of the 2018 International Conference on Management of Data* (Houston, TX, USA) (SIGMOD '18). Association for Computing Machinery, New York, NY, USA, 275–290. <https://doi.org/10.1145/3183713.3196898>
- [26] Raymond Cheng, Fan Zhang, Jernej Kos, Warren He, Nicholas Hynes, Noah Johnson, Ari Juels, Andrew Miller, and Dawn Song. 2019. Ekiden: A Platform for Confidentiality-Preserving, Trustworthy, and Performant Smart Contracts. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE Computer Society, USA, 185–200. <https://doi.org/10.1109/EuroSP.2019.00023>
- [27] ConsenSys. 2022. Quorum. <https://github.com/ConsenSys/quorum> [Last accessed: 2023-Oct-06].
- [28] Kata Containers. 2021. [RFC] [WIP] Confidential Computing Enablement. <https://github.com/kata-containers/kata-containers/issues/1332> [Last accessed: 2023-Oct-06].
- [29] Kata Containers. 2022. Kata Containers. <https://github.com/kata-containers> [Last accessed: 2023-Oct-06].
- [30] Corda. 2022. Corda. <https://github.com/corda/corda> [Last accessed: 2023-Oct-06].
- [31] Victor Costan and Srinivas Devadas. 2016. Intel SGX Explained. Cryptology ePrint Archive, Paper 2016/086. <https://eprint.iacr.org/2016/086>
- [32] Jinhua Cui, Shweta Shinde, Satyaki Sen, Prateek Saxena, and Pinghai Yuan. 2022. Dynamic Binary Translation for SGX Enclaves. *ACM Trans. Priv. Secur.* 25, 4, Article 32 (jul 2022), 40 pages. <https://doi.org/10.1145/3532862>
- [33] Jérémie Decouchant, David Kozhaya, Vincent Rahli, and Jiangshan Yu. 2022. DAMYSUS: Streamlined BFT Consensus Leveraging Trusted Components.

- In *Proceedings of the Seventeenth European Conference on Computer Systems* (Rennes, France) (*EuroSys '22*). Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3492321.3519568>
- [34] Antoine Delignat-Lavaud, Cédric Fournet, Kapil Vaswani, Sylvan Clebsch, Maik Riechert, Manuel Costa, and Mark Russinovich. 2023. Why Should I Trust Your Code? Confidential Computing Enables Users to Authenticate Code Running in TEEs, but Users Also Need Evidence This Code is Trustworthy. *Queue* 21, 4 (sep 2023), 94–122. <https://doi.org/10.1145/3623460>
- [35] W. Diffie and M. Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory* 22, 6 (1976), 644–654. <https://doi.org/10.1109/TIT.1976.1055638>
- [36] Muhammad El-Hindi, Tobias Ziegler, Matthias Heinrich, Adrian Lutsch, Zheguang Zhao, and Carsten Binnig. 2022. Benchmarking the Second Generation of Intel SGX Hardware. In *Data Management on New Hardware* (Philadelphia, PA, USA) (*DaMoN'22*). Association for Computing Machinery, New York, NY, USA, Article 5, 8 pages. <https://doi.org/10.1145/3533737.3535098>
- [37] Open Enclave. 2022. Open Enclave SDK. <https://github.com/openenclave/openenclave> [Last accessed: 2023-Oct-06].
- [38] etcd v3.5. 2021. Operations Guide: Runtime reconfiguration - Cluster reconfiguration operations. <https://etcd.io/docs/v3.1/op-guide/runtime-configuration/#cluster-reconfiguration-operations> [Last accessed: 2023-Oct-06].
- [39] Dimitra Giantsidi, Maurice Bailleu, Natacha Crooks, and Pramod Bhatotia. 2022. Treaty: Secure Distributed Transactions. In *2022 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE Computer Society, USA, 14–27. <https://doi.org/10.1109/DSN53405.2022.00015>
- [40] S. Goldwasser, S. Micali, and C. Rackoff. 1989. The Knowledge Complexity of Interactive Proof Systems. *SIAM J. Comput.* 18, 1 (feb 1989), 186–208. <https://doi.org/10.1137/0218012>
- [41] Vincent Gramoli, Rachid Guerraoui, Andrei Lebedev, Chris Natoli, and Gauthier Voron. 2023. Diablo: A Benchmark Suite for Blockchains. In *Proceedings of the Eighteenth European Conference on Computer Systems* (Rome, Italy) (*EuroSys '23*). Association for Computing Machinery, New York, NY, USA, 540–556. <https://doi.org/10.1145/3552326.3567482>
- [42] Daniel Gruss, Julian Lettner, Felix Schuster, Olga Ohrimenko, Istvan Haller, and Manuel Costa. 2017. Strong and Efficient Cache Side-Channel Protection Using Hardware Transactional Memory. In *Proceedings of the 26th USENIX Conference on Security Symposium* (Vancouver, BC, Canada) (*SEC'17*). USENIX Association, USA, 217–233. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-gruss.pdf>
- [43] Suyash Gupta, Sajjad Rahnama, Shubham Pandey, Natacha Crooks, and Mohammad Sadoghi. 2023. Dissecting BFT Consensus: In Trusted Components We Trust!. In *Proceedings of the Eighteenth European Conference on Computer Systems* (Rome, Italy) (*EuroSys '23*). Association for Computing Machinery, New York, NY, USA, 521–539. <https://doi.org/10.1145/3552326.3587455>
- [44] Tyler Hunt, Zhiting Zhu, Yuanzhong Xu, Simon Peter, and Emmett Witchel. 2016. Ryoan: A Distributed Sandbox for Untrusted Computation on Secret Data. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation* (Savannah, GA, USA) (*OSDI'16*). USENIX Association, USA, 533–549. <https://www.usenix.org/system/files/conference/osdi16/osdi16-hunt.pdf>
- [45] Hyperledger. 2021. Fabric Private Chaincode RPC 1.0. <https://github.com/hyperledger/fabric-rfcs/blob/main/text/0000-fabric-private-chaincode-1.0.md> [Last accessed: 2023-Oct-06].
- [46] Hyperledger. 2022. Hyperledger Fabric. <https://github.com/hyperledger/fabric> [Last accessed: 2023-Oct-06].
- [47] Hyperledger. 2022. Hyperledger Fabric Private Chaincode. <https://github.com/hyperledger/fabric-private-chaincode> [Last accessed: 2023-Oct-06].
- [48] Intel. 2016. Introduction to Intel SGX Sealing. <https://www.intel.com/content/www/us/en/developer/articles/technical/introduction-to-intel-sgx-sealing.html> [Last accessed: 2023-Oct-06].
- [49] Intel. 2018. Performance Considerations for Intel Software Guard Extensions (Intel SGX) Applications. https://community.intel.com/legacyfs/online/drupal_files/managed/09/37/Intel-SGX-Performance-Considerations.pdf [Last accessed: 2023-Oct-06].
- [50] Intel. 2020. AMD SEV-SNP: Strengthening VM isolation with integrity protection and more. <https://www.amd.com/system/files/TechDocs/SEV-SNP-strengthening-vm-isolation-with-integrity-protection-and-more.pdf> [Last accessed: 2023-Oct-06].
- [51] Intel. 2021. Intel Trust Domain Extensions - White Paper. <https://cdrdrv2.intel.com/v1/dl/getContent/690419> [Last accessed: 2023-Oct-06].
- [52] Intel. 2021. Product Brief: 3rd Gen Intel Xeon Scalable Processors. <https://www.intel.com/content/www/us/en/products/docs/processors/xeon/3rd-gen-xeon-scalable-processors-brief.html> [Last accessed: 2023-Oct-06].
- [53] Intel. 2021. Unable to find Alternatives to Monotonic Counter Application Programming Interfaces (APIs) in Intel Software Guard Extensions (Intel SGX) for Linux* to Prevent Sealing Rollback Attacks. <https://www.intel.com/content/www/us/en/support/articles/000057968/software/intel-security-products.html> Article ID: 000057968 [Last accessed: 2023-Oct-06].
- [54] Intel. 2022. Intel Software Guard Extensions (Intel SGX) SDK for Linux OS - Developer Reference. https://download.01.org/intel-sgx/latest/linux-latest/docs/Intel_SGX_Developer_Reference_Linux_2.17_Open_Source.pdf Revision 2.17 [Last accessed: 2023-Oct-06].
- [55] Intel. 2022. Which Platforms Support Intel Software Guard Extensions (Intel SGX) SGX2? <https://www.intel.com/content/www/us/en/support/articles/000058764/software/intel-security-products.html> Article ID: 000058764 [Last accessed: 2023-Oct-06].
- [56] Matthew A. Johnson, Stavros Volos, Ken Gordon, Sean T. Allen, Christoph M. Wintersteiger, Sylvan Clebsch, John Starks, and Manuel Costa. 2023. Parma: Confidential Containers via Attested Execution Policies. [arXiv:2302.03976 \[cs.CR\]](https://arxiv.org/abs/2302.03976) <https://arxiv.org/abs/2302.03976>
- [57] Michael Jones, John Bradley, and Nat Sakimura. 2015. JSON Web Token (JWT). RFC 7519. <https://doi.org/10.17487/RFC7519>
- [58] Flavio P. Junqueira, Benjamin C. Reed, and Marco Serafini. 2011. Zab: High-Performance Broadcast for Primary-Backup Systems. In *Proceedings of the 2011 IEEE/IFIP 41st International Conference on Dependable Systems & Networks (DSN '11)*. IEEE Computer Society, USA, 245 – 256. <https://doi.org/10.1109/DSN.2011.5958223>
- [59] Anuj Kalia, Michael Kaminsky, and David G. Andersen. 2019. Datacenter RPCs Can Be General and Fast. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) (*NSDI'19*). USENIX Association, USA, 1–16. <https://www.usenix.org/system/files/nsdi19-kalia.pdf>
- [60] David Kaplan. 2017. Protecting VM Register State with SEV-ES. <https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/white-papers/Protecting-VM-Register-State-with-SEV-ES.pdf> [Last accessed: 2023-Oct-06].
- [61] David Kaplan, Jeremy Powell, and Tom Woller. 2021. AMD Memory Encryption. <https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/white-papers/memory-encryption-white-paper.pdf> [Last accessed: 2023-Oct-06].
- [62] Antonios Katsarakis, Vasilis Gavrielatos, M.R. Siavash Katebzadeh, Arpit Joshi, Aleksandar Dragojevic, Boris Grot, and Vijay Nagarajan. 2020. Hermes: A Fast, Fault-Tolerant and Linearizable Replication Protocol. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (*ASPLOS '20*). Association for Computing Machinery, New York, NY, USA, 201–217. <https://doi.org/10.1145/3373376.3378496>
- [63] Taehoon Kim, Joongun Park, Jaewook Woo, Seungheun Jeon, and Jaehyuk Huh. 2019. ShieldStore: Shielded In-Memory Key-Value Storage with SGX. In *Proceedings of the Fourteenth EuroSys Conference 2019* (Dresden, Germany) (*EuroSys '19*). Association for Computing Machinery, New York, NY, USA, Article 14, 15 pages. <https://doi.org/10.1145/3302424.3303951>
- [64] Robert Krahn, Bohdan Trach, Anjo Vahldiek-Oberwagner, Thomas Knauth, Pramod Bhatotia, and Christof Fetzer. 2018. Pesos: Policy Enhanced Secure Object Store. In *Proceedings of the Thirteenth EuroSys Conference* (Porto, Portugal) (*EuroSys '18*). Association for Computing Machinery, New York, NY, USA, Article 25, 17 pages. <https://doi.org/10.1145/3190508.3190518>
- [65] Sandeep Kumar and Smruti R. Sarangi. 2021. SecureFS: A Secure File System for Intel SGX. In *Proceedings of the 24th International Symposium on Research in Attacks, Intrusions and Defenses* (San Sebastian, Spain) (*RAID '21*). Association for Computing Machinery, New York, NY, USA, 91–102. <https://doi.org/10.1145/3471621.3471840>
- [66] Hyperledger Labs. 2022. Hyperledger Private Data Objects. <https://github.com/hyperledger-labs/private-data-objects>
- [67] Leslie Lamport. 1998. The Part-Time Parliament. *ACM Trans. Comput. Syst.* 16, 2 (may 1998), 133 – 169. <https://doi.org/10.1145/279227.279229>
- [68] Leslie Lamport. 2002. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Professional, Pearson Education, USA. <https://lamport.azurewebsites.net/tla/book.html> [Last accessed: 2023-Oct-06].
- [69] Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. 2020. Keystone: An Open Framework for Architecting Trusted Execution Environments. In *Proceedings of the Fifteenth European Conference on Computer Systems* (Heraklion, Greece) (*EuroSys '20*). Association for Computing Machinery, New York, NY, USA, Article 38, 16 pages. <https://doi.org/10.1145/3342195.3387532>
- [70] Mingyu Li, Jinhao Zhu, Tianxu Zhang, Cheng Tan, Yubin Xia, Sebastian Angel, and Haibo Chen. 2021. Bringing Decentralized Search to Decentralized Services. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI'21)*. USENIX Association, USA, 331–347. <https://www.usenix.org/system/files/osdi21-li.pdf>
- [71] Xupeng Li, Xuheng Li, Christoffer Dall, Ronghui Gu, Jason Nieh, Yousuf Sait, and Gareth Stockwell. 2022. Design and Verification of the Arm Confidential Compute Architecture. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI'22)*. USENIX Association, Carlsbad, CA, 465–484.

- <https://www.usenix.org/system/files/osdi22-li.pdf>
- [72] Paul Liétar, Theodore Butler, Sylvan Clebsch, Sophia Drossopoulou, Juliana Franco, Matthew J. Parkinson, Alex Shamis, Christoph M. Wintersteiger, and David Chisnall. 2019. Smmalloc: A Message Passing Allocator. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on Memory Management* (Phoenix, AZ, USA) (*ISMM 2019*). Association for Computing Machinery, New York, NY, USA, 122–135. <https://doi.org/10.1145/3315573.3329980>
- [73] Joshua Lind, Christian Priebe, Divya Muthukumar, Dan O’Keeffe, Pierre-Louis Aublin, Florian Kelbert, Tobias Reiher, David Goltzsche, David Eyers, Rüdiger Kapitza, Christof Fetzter, and Peter Pietzuch. 2017. Glamdring: Automatic Application Partitioning for Intel SGX. In *Proceedings of the 2017 USENIX Conference on Usenix Annual Technical Conference* (Santa Clara, CA, USA) (*USENIX ATC ’17*). USENIX Association, USA, 285–298. <https://www.usenix.org/system/files/conference/atc17/atc17-lind.pdf>
- [74] N. A. Lynch and A. A. Shvartsman. 1997. Robust Emulation of Shared Memory Using Dynamic Quorum-Acknowledged Broadcasts. In *Proceedings of the 27th International Symposium on Fault-Tolerant Computing (FTCS ’97)* (FTCS ’97). IEEE Computer Society, USA, 272. <https://doi.org/10.1109/FTCS.1997.614100>
- [75] André Martin, Cong Lian, Franz Gregor, Robert Krahn, Valerio Schiavoni, Pascal Felber, and Christof Fetzter. 2021. ADAM-CS: Advanced Asynchronous Monotonic Counter Service. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE Computer Society, USA, 426–437. <https://doi.org/10.1109/DSN48987.2021.00053>
- [76] Paulo Martins, Leonel Sousa, and Artur Mariano. 2017. A Survey on Fully Homomorphic Encryption: An Engineering Perspective. *ACM Comput. Surv.* 50, 6, Article 83 (dec 2017), 33 pages. <https://doi.org/10.1145/3124441>
- [77] Sinisa Matetic, Mansoor Ahmed, Kari Kostianen, Aritra Dhar, David Sommer, Arthur Gervais, Ari Juels, and Srđjan Capkun. 2017. ROTE: Rollback Protection for Trusted Execution. In *Proceedings of the 26th USENIX Conference on Security Symposium* (Vancouver, BC, Canada) (*SEC’17*). USENIX Association, USA, 1289–1306. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-matetic.pdf>
- [78] Ralph C. Merkle. 1987. A Digital Signature Based on a Conventional Encryption Function. In *A Conference on the Theory and Applications of Cryptographic Techniques on Advances in Cryptology (CRYPTO ’87)*. Springer-Verlag, Berlin, Heidelberg, 369–378. https://doi.org/10.1007/3-540-48184-2_32
- [79] Ines Messadi, Shivananda Neumann, Nico Weichbrodt, Lennart Almstedt, Mohammad Mahhouk, and Rüdiger Kapitza. 2021. Precursor: A Fast, Client-Centric and Trusted Key-Value Store Using RDMA and Intel SGX. In *Proceedings of the 22nd International Middleware Conference* (Québec city, Canada) (*Middleware ’21*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3464298.3476129>
- [80] Microsoft. 2022. LSKV. <https://github.com/microsoft/LSKV> [Last accessed: 2023-Oct-06].
- [81] Microsoft. 2022. Merklecpp. <https://github.com/microsoft/merklecpp> [Last accessed: 2023-Oct-06].
- [82] Microsoft. 2022. Microsoft Azure confidential ledger. <https://learn.microsoft.com/en-us/azure/confidential-ledger/overview> [Last accessed: 2023-Oct-06].
- [83] Microsoft. 2022. scitt-ccf-ledger. <https://github.com/microsoft/scitt-ccf-ledger> [Last accessed: 2023-Oct-06].
- [84] Microsoft. 2022. W3C DID for Confidential Consortium Framework. <https://github.com/microsoft/did-ccf> [Last accessed: 2023-Oct-06].
- [85] Microsoft. 2023. Confidential Consortium Framework Documentation. <https://microsoft.github.io/CCF/main/> [Last accessed: 2023-Oct-06].
- [86] Microsoft. 2023. DCsv3 and DCdsv3-series. <https://learn.microsoft.com/en-us/azure/virtual-machines/dcv3-series> [Last accessed: 2023-Oct-06].
- [87] Microsoft. 2023. Default constitution for CCF. <https://github.com/microsoft/CCF/tree/main/samples/constitutions/default> [Last accessed: 2023-Oct-06].
- [88] Microsoft. 2023. TLA+ Specifications for the Confidential Consortium Framework. <https://github.com/microsoft/CCF/tree/main/ta> [Last accessed: 2023-Oct-06].
- [89] Kit Murdock, David Oswald, Flavio D. Garcia, Jo Van Bulck, Daniel Gruss, and Frank Piessens. 2020. Stundervolt: Software-based Fault Injection Attacks against Intel SGX. In *Proceedings of the 41st IEEE Symposium on Security and Privacy (S&P’20)*. IEEE Computer Society, USA, 1466–1482. <https://doi.org/10.1109/SP40000.2020.00057>
- [90] Jianyu Niu, Wei Peng, Xiaokuan Zhang, and Yinqian Zhang. 2022. NARRATOR: Secure and Practical State Continuity for Trusted Execution in the Cloud. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) (*CCS ’22*). Association for Computing Machinery, New York, NY, USA, 2385–2399. <https://doi.org/10.1145/3548606.3560620>
- [91] National Institute of Standards and Technology (NIST). 2015. FIPS 180-4: Secure Hash Standard. <http://dx.doi.org/10.6028/NIST.FIPS.180-4>
- [92] Oleksii Oleksenko, Bohdan Trach, Robert Krahn, Andre Martin, Christof Fetzter, and Mark Silberstein. 2018. Varys: Protecting SGX Enclaves from Practical Side-Channel Attacks. In *Proceedings of the 2018 USENIX Conference on Usenix Annual Technical Conference* (Boston, MA, USA) (*USENIX ATC ’18*). USENIX Association, USA, 227–239. <https://www.usenix.org/system/files/conference/atc18/atc18-oleksenko.pdf>
- [93] Diego Ongaro. 2014. *Consensus: Bridging Theory and Practice*. Ph.D. Dissertation. Stanford. <https://web.stanford.edu/~ouster/cgi-bin/papers/OngaroPhD.pdf> [Last accessed: 2023-Oct-06].
- [94] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference* (Philadelphia, PA) (*USENIX ATC ’14*). USENIX Association, USA, 305–320. <https://www.usenix.org/system/files/conference/atc14/atc14-paper-ongaro.pdf>
- [95] OpenSSL. 2023. OpenSSL homepage. <https://www.openssl.org/> [Last accessed: 2023-Oct-06].
- [96] Meni Orenbach, Pavel Lifshits, Marina Minkin, and Mark Silberstein. 2017. Eleos: ExitLess OS Services for SGX Enclaves. In *Proceedings of the Twelfth European Conference on Computer Systems* (Belgrade, Serbia) (*EuroSys ’17*). Association for Computing Machinery, New York, NY, USA, 238–253. <https://doi.org/10.1145/3064176.3064219>
- [97] European Parliament. 2016. General Data Protection Regulation (GDPR) 2016/679. <https://eur-lex.europa.eu/eli/reg/2016/679/oj> [Last accessed: 2023-Oct-06].
- [98] Moritz Platt, Johannes Sedlmeir, Daniel Platt, Jiahua Xu, Paolo Tasca, Nikhil Vadgama, and Juan Ignacio Ibañez. 2021. The Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work. In *2021 IEEE 21st International Conference on Software Quality, Reliability and Security Companion (QRS-C)*. IEEE Computer Society, USA, 1135–1144. <https://doi.org/10.1109/QRS-C55045.2021.00168>
- [99] Peter Pogorski. 2023. Announcing public preview of confidential containers on Azure Container Instances. <https://techcommunity.microsoft.com/t5/apps-on-azure-blog/announcing-public-preview-of-confidential-containers-on-azure/ba-p/3755623> [Last accessed: 2023-Oct-06].
- [100] Christian Priebe, Kapil Vaswani, and Manuel Costa. 2018. EnclaveDB: A Secure Database Using SGX. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, USA, 264–278. <https://doi.org/10.1109/SP.2018.00025>
- [101] Anil Rao. 2022. Rising to the Challenge - Data Security with Intel Confidential Computing. <https://community.intel.com/t5/Blogs/Products-and-Solutions/Security/Rising-to-the-Challenge-Data-Security-with-Intel-Confidential/post/1353141> [Last accessed: 2023-Oct-06].
- [102] Mark Russinovich, Edward Ashton, Christine Avanesians, Miguel Castro, Amaury Chamayou, Sylvan Clebsch, Manuel Costa, Cédric Fournet, Matthew Kerner, Sid Krishna, Julien Maffre, Thomas Moscibroda, Kartik Nayak, Olya Ohrimenko, Felix Schuster, Roy Schwartz, Alex Shamis, Olga Vrousgou, and Christoph M. Wintersteiger. 2019. *CCF: A Framework for Building Confidential Verifiable Replicated Services*. Technical Report MSR-TR-2019-16. Microsoft. <https://www.microsoft.com/en-us/research/publication/ccf-a-framework-for-building-confidential-verifiable-replicated-services/> [Last accessed: 2023-Oct-06].
- [103] Mark Russinovich, Manuel Costa, Cédric Fournet, David Chisnall, Antoine Delignat-Lavaud, Sylvan Clebsch, Kapil Vaswani, and Vikas Bhatia. 2021. Toward Confidential Cloud Computing. *Commun. ACM* 64, 6 (May 2021), 54–61. <https://doi.org/10.1145/3453930>
- [104] Joseph A. Salowey, David McGrew, and Abhijit Choudhury. 2008. AES Galois Counter Mode (GCM) Cipher Suites for TLS. RFC 5288. <https://doi.org/10.17487/RFC5288>
- [105] J. Schaad. 2017. *CBOR Object Signing and Encryption (COSE)*. RFC 8852. RFC Editor. <https://doi.org/10.17487/RFC8152>
- [106] Felix Schuster, Manuel Costa, Cédric Fournet, Christos Gkantsidis, Marcus Peinado, Gloria Mainar-Ruiz, and Mark Russinovich. 2015. VC3: Trustworthy Data Analytics in the Cloud Using SGX. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy (SP ’15)*. IEEE Computer Society, USA, 38–54. <https://doi.org/10.1109/SP.2015.10>
- [107] Michael Schwarz, Samuel Weiser, Daniel Gruss, Clémentine Maurice, and Stefan Mangard. 2017. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment, Michalis Polychronakis and Michael Meier (Eds.)*. Springer International Publishing, Cham, 3–24. https://doi.org/10.1007/978-3-319-60876-1_1
- [108] Adi Shamir. 1979. How to Share a Secret. *Commun. ACM* 22, 11 (nov 1979), 612–613. <https://doi.org/10.1145/359168.359176>
- [109] Alex Shamis, Peter Pietzuch, Burcu Canakci, Miguel Castro, Cedric Fournet, Edward Ashton, Amaury Chamayou, Sylvan Clebsch, Antoine Delignat-Lavaud, Matthew Kerner, Julien Maffre, Olga Vrousgou, Christoph M. Wintersteiger, Manuel Costa, and Mark Russinovich. 2022. IA-CCF: Individual Accountability for Permissioned Ledgers. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 467–491. <https://www.usenix.org/system/files/nsdi22-paper-shamis.pdf>
- [110] Youren Shen, Hongliang Tian, Yu Chen, Kang Chen, Runji Wang, Yi Xu, Yubin Xia, and Shoumeng Yan. 2020. Occlum: Secure and Efficient Multitasking Inside a Single Enclave of Intel SGX. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (*ASPLoS ’20*). Association for Computing

- Machinery, New York, NY, USA, 955–970. <https://doi.org/10.1145/3373376.3378469>
- [111] Signal. 2022. Private Contact Discovery Service (Public Archive). <https://github.com/signalapp/ContactDiscoveryService/> [Last accessed: 2023-Oct-06].
- [112] Jatinder Singh, Jennifer Cobbe, Do Le Quoc, and Zahra Tarkhani. 2021. Enclaves in the Clouds: Legal Considerations and Broader Implications. *Commun. ACM* 64, 5 (apr 2021), 42–51. <https://doi.org/10.1145/3447543>
- [113] Rohit Sinha and Mihai Christodorescu. 2018. VeritasDB: High Throughput Key-Value Store with Integrity. *Cryptology ePrint Archive, Paper 2018/251*. <https://eprint.iacr.org/2018/251>
- [114] Shubhra Sinha. 2022. Microsoft introduces preview of Azure Managed Confidential Consortium Framework. <https://techcommunity.microsoft.com/t5/azure-confidential-computing/microsoft-introduces-preview-of-azure-managed-confidential/ba-p/3648986> [Last accessed: 2023-Oct-06].
- [115] Dimitrios Skarlatos, Mengjia Yan, Bhargava Gopireddy, Read Sprabery, Josep Torrellas, and Christopher W. Fletcher. 2019. MicroScope: Enabling Microarchitectural Replay Attacks. In *Proceedings of the 46th International Symposium on Computer Architecture (Phoenix, Arizona) (ISCA '19)*. Association for Computing Machinery, New York, NY, USA, 318–331. <https://doi.org/10.1145/3307650.3322228>
- [116] Michael J. Steindorfer and Jurgen J. Vinju. 2016. Fast and Lean Immutable Multi-Maps on the JVM based on Heterogeneous Hash-Array Mapped Tries. *CoRR abs/1608.01036* (2016). [arXiv:1608.01036](http://arxiv.org/abs/1608.01036) <http://arxiv.org/abs/1608.01036>
- [117] Yuanyuan Sun, Sheng Wang, Huorong Li, and Feifei Li. 2021. Building Enclave-Native Storage Engines for Practical Encrypted Databases. *Proc. VLDB Endow.* 14, 6 (apr 2021), 1019–1032. <https://doi.org/10.14778/3447689.3447705>
- [118] Edgeless Systems. 2022. Constellation: Always Encrypted Kubernetes. <https://github.com/edgeless/constellation> [Last accessed: 2023-Oct-06].
- [119] Edgeless Systems. 2022. EGo. <https://github.com/edgeless/ego> [Last accessed: 2023-Oct-06].
- [120] Chia-Che Tsai, Kumar Saurabh Arora, Nehal Bandi, Bhushan Jain, William Jannen, Jitin John, Harry A. Kalodner, Vrushali Kulkarni, Daniela Oliveira, and Donald E. Porter. 2014. Cooperation and Security Isolation of Library OSes for Multi-Process Applications. In *Proceedings of the Ninth European Conference on Computer Systems (Amsterdam, The Netherlands) (EuroSys '14)*. Association for Computing Machinery, New York, NY, USA, Article 9, 14 pages. <https://doi.org/10.1145/2592798.2592812>
- [121] Chia-Che Tsai, Donald E. Porter, and Mona Vij. 2017. Graphene-SGX: A Practical Library OS for Unmodified Applications on SGX. In *2017 USENIX Annual Technical Conference (Santa Clara, CA, USA) (USENIX ATC '17)*. USENIX Association, USA, 645–658. <https://www.usenix.org/system/files/conference/atc17/atc17-tsai.pdf>
- [122] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. 2018. Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient out-of-Order Execution. In *Proceedings of the 27th USENIX Conference on Security Symposium* (Baltimore, MD, USA) (SEC'18). USENIX Association, USA, 991–1008. https://www.usenix.org/system/files/conference/usenixsecurity18/sec18-van_bulck.pdf
- [123] Jo Van Bulck, Daniel Moghimi, Michael Schwarz, Moritz Lippi, Marina Minkin, Daniel Genkin, Yuval Yarom, Berk Sunar, Daniel Gruss, and Frank Piessens. 2020. LVI: Hijacking Transient Execution through Microarchitectural Load Value Injection. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, USA, 54–72. <https://doi.org/10.1109/SP40000.2020.00089>
- [124] Stephan van Schaik, Andrew Kwong, Daniel Genkin, and Yuval Yarom. 2020. SGAXe: How SGX Fails in Practice. <https://sgaxe.com/files/SGAXe.pdf> [Last accessed: 2023-Oct-06].
- [125] Stephan van Schaik, Alex Seto, Thomas Yurek, Adam Batori, Bader Albassam, Christina Garman, Daniel Genkin, Andrew Miller, Eyal Ronen, and Yuval Yarom. 2022. SoK: SGX.Fail: How Stuff Get eXposed. <https://sgx.fail>.
- [126] Weili Wang, Sen Deng, Jianyu Niu, Michael K. Reiter, and Yinqian Zhang. 2022. ENGRAFT: Enclave-Guarded Raft on Byzantine Faulty Nodes. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (Los Angeles, CA, USA) (CCS '22)*. Association for Computing Machinery, New York, NY, USA, 2841–2855. <https://doi.org/10.1145/3548606.3560639>
- [127] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. 2017. Global-Scale Secure Multiparty Computation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (Dallas, Texas, USA) (CCS '17)*. Association for Computing Machinery, New York, NY, USA, 39–56. <https://doi.org/10.1145/3133956.3133979>
- [128] Ofir Weisse, Valeria Bertacco, and Todd Austin. 2017. Regaining Lost Cycles with HotCalls: A Fast Interface for SGX Secure Enclaves. *SIGARCH Comput. Archit. News* 45, 2 (jun 2017), 81–93. <https://doi.org/10.1145/3140659.3080208>
- [129] Michael E Whitman and Herbert J Mattord. 2011. *Principles of information security, 4th edition*. Course Technology Inc, USA.
- [130] Siyuan Xia, Zhiru Zhu, Chris Zhu, Jinjin Zhao, Kyle Chard, Aaron J. Elmore, Ian Foster, Michael Franklin, Sanjay Krishnan, and Raul Castro Fernandez. 2022. Data Station: Delegated, Trustworthy, and Auditable Computation to Enable Data-Sharing Consortia with a Data Escrow. *Proc. VLDB Endow.* 15, 11 (sep 2022), 3172–3185. <https://doi.org/10.14778/3551793.3551861>
- [131] Yang Xiao, Ning Zhang, Jin Li, Wenjing Lou, and Y. Thomas Hou. 2020. PrivacyGuard: Enforcing Private Data Usage Control with Blockchain and Attested Off-Chain Contract Execution. In *Computer Security - ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part II* (Guildford, United Kingdom). Springer-Verlag, Berlin, Heidelberg, 610–629. https://doi.org/10.1007/978-3-030-59013-0_30
- [132] Sravya Yandamuri, Ittai Abraham, Kartik Nayak, and Michael Reiter. 2021. Brief Announcement: Communication-Efficient BFT Using Small Trusted Hardware to Tolerate Minority Corruption. In *35th International Symposium on Distributed Computing (DISC 2021) (Leibniz International Proceedings in Informatics (LIPIcs))*, Seth Gilbert (Ed.), Vol. 209. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 62:1–62:4. <https://doi.org/10.4230/LIPIcs.DISC.2021.62>