

# Autonomous Membership Service for Enclave Applications

Hung Dang, Ee-Chien Chang  
School of Computing, National University of Singapore  
{hungdang,change}@comp.nus.edu.sg

## Abstract

Trusted Execution Environment, or enclave, promises to protect data confidentiality and execution integrity of an outsourced computation on an untrusted host. Extending the protection to distributed applications that run on physically separated hosts, however, remains non-trivial. For instance, the current enclave provisioning model hinders elasticity of cloud applications. Furthermore, it remains unclear how an enclave process could verify if there exists another concurrently running enclave process instantiated using the same codebase, or count a number of such processes. In this paper, we seek an autonomous membership service for enclave applications. The application owner only needs to partake in instantiating the very first process of the application, whereas all subsequent process commission and decommission will be administered by existing and active processes of that very application. To achieve both safety and liveness, our protocol design admits unjust excommunication of a non-faulty process from the membership group. We implement the proposed membership service in a system called AMES. Our experimental study shows that AMES incurs an overhead of 5% – 16% compared to vanilla enclave execution.

## 1. INTRODUCTION

Intel Software Guard Extensions (SGX) [21] protects data confidentiality and execution integrity of outsourced computations by offering a hardware-protected trusted execution environment (TEE), or *enclave*, that guards the user’s security-critical application code and its data against untrusted system software including the operating system (OS). The CPU prevents any non-enclave code from reading or modifying enclave memory at runtime. SGX architecture relies on the OS to instantiate enclave. A user can verify if a specific enclave is correctly instantiated and running at a remote host via a *remote attestation* protocol [9], which also establishes a secure channel via which the user can provision application secrets, if any, to the enclave.

While SGX allows users to harden security of software applications on an untrusted host (or server) [10, 14], extending its protection to distributed applications running on a number of hosts is non-trivial. More specifically, if the application involves a secret that must not be revealed to the untrusted host, the secret provisioning requires remote attestation by the application owner. In the context of cloud service, this hinders the elasticity of cloud applications, for the cloud service providers cannot commission additional enclave processes dynamically at its sole discretion. In another example, it is unclear how an enclave process could verify if there exists another concurrently running enclave process instantiated using the same codebase [1], or count the number

of such processes, in an autonomic manner. This capability is desired in implementing singleton applications (i.e., only a single instance of the application can be created and run) and decentralized floating licensing (i.e., floating licensing without the central license server).

Let us use an example of a *forking attack* [12] on application that maintains persistent state to illustrate the challenges of extending TEE protections to a distributed environment. At a high level, the adversary (e.g., cloud service provider) instantiates multiple enclave processes of the same application and feeds them with different, potentially conflicting, inputs so as to violate the expected behaviours of the victim application. More specifically, consider an application **Prog** that serves read and write requests of a stateless client. The adversary runs two independent enclave processes (or processes for short) of **Prog**, denoted by  $p_1$  and  $p_2$ , while keeping them oblivious to each other’s existence. Since the adversary handles all I/O operations to the processes, the client is only presented with a simple API for read and write requests, and is unaware of the underlying provisioning of the processes. At the beginning, both processes assume an initial state  $s_0$ . The client issues two sequential requests, the first writes data **dat** to **Prog**’s memory, changing its state to  $s_1$ , and the second reads from **Prog** its latest written data. From the application logic, the client should receive **dat** as a response to the second request. Nevertheless, if the adversary routes the write request to  $p_1$  and read request to  $p_2$ , the client will receive an obsolete (but authenticated) response drawn from the stale state  $s_0$ . This attack may wreak severe havoc in financial applications wherein an account balance should reflect at all time a correct transaction history.

The example described earlier motivates the need of a *membership service* for enclave applications in a distributed environment. The membership service of an application **Prog** not only facilitates seamless commissioning of **Prog**’s enclave processes, but also monitors and keeps track of the status of all **Prog**’s active processes, ensuring that their collective existence adheres to the Service Level Agreement (SLA). In a straight forward approach, one can implement the membership service by maintaining a special *directory server* that keeps track of all active processes, and excommunicates a process that it suspects of being faulty from the membership group. The commissioning of a new process is augmented such that the newly instantiated process remains *nonoperational* until its request to join the existing membership group is approved by the directory server. Nonetheless, this approach is not desirable, for the directory server becomes an obvious single-point-of-failure. A logical attempt to avoid the single-point-of-failure would be to replicate the directory server into a distributed fault-tolerant directory service comprising of a cluster of replicas

coordinating their operations via a consensus protocol such as Paxos [19]. However, it is challenging to protect these replicas themselves against forking attack [12], wherein the adversary manage to splits the directory service replicas into two or more “cliques”, each of which is oblivious to the others’ existence. Such adversary can then selectively route communications of the existing processes, or join request of the newly instantiated processes to different clique of directory service replicas, thereby evading the control mechanism that enforces the SLA.

In this paper, we seek a membership service that is autonomous. In particular, once an application `Prog` has been successfully bootstrapped, all subsequent process commission and decommission of `Prog` are administered by its own active processes. This self-administering is attained by designating one of `Prog`’s active processes at a *leader*, while the other assume the role of follower. The leader is tasked to monitor the status of all follower processes by exchanging heartbeat messages. On the one hand, the process commission requires acknowledgement from the leader and a quorum of existing followers before the new process becomes fully functional. On the other hand, any process that is suspected of having halted is excommunicated from the membership group. Our protocol design necessarily admits unjust suspicion and excommunication of a non-faulty process, in order to sidestep the difficulty of solving the membership problem in an asynchronous environment when faults may be present [8]. To prevent the non-faulty yet excommunicated process from disrupting the membership service, its future communication to other processes of `Prog` is discarded, rendering it isolated. The isolated process eventually halts by committing suicide, fulfilling the excommunication.

We consider a powerful adversary that controls the OS (or a hypervisor) of the hosts on which the enclave processes are instantiated. The adversary can schedule or corrupt all non-enclave processes, and restart the processes at will. Further, it can control all network communications between the processes regardless of whether they are instantiated on the same or across different hosts. The adversary can modify, reorder or delay the network messages arbitrarily. Nonetheless, we assume that the adversary cannot break enclave protections, and is limited by the constraints of the cryptographic methods employed.

We implement the proposed solution that enables autonomous membership service for SGX enclaves in a system called AMES. The application owner only needs to partake in the commissioning of the first enclave process (i.e., conducting remote attestation and provision application secret to the enclave). All subsequent process commissioning (or dismantling) during the lifetime of the application will be handled automatically and securely. To facilitate the incorporation of AMES logic to an enclave application `Prog`, we provide application developers with a AMES library that envelops the membership service’s operations in a control thread, separating them from the original operational logic of `Prog` that is to be implemented in a worker thread (or threads).

In summary, we make the following contributions in this paper.

- We propose a new approach to enable autonomous membership service for enclave-based applications, which allows seamless commissioning and dismantling of processes over the lifetime of the application.

- We provide a AMES library that encapsulates the membership service’s operations, which can be easily incorporated to any enclave-based application.
- We conduct empirical evaluation of the overhead incurred by AMES on both our local cluster with 35 servers and the Google Cloud Platform (GCP) nodes spanning across 4 regions. The experimental results show that AMES incurs an overhead of 5% – 16% over vanilla enclave execution.

The rest of this paper is organised as follows. Section 2 provides backgrounds on key properties of Intel SGX and distributed consensus protocols. Section 3 presents our problem statement, covering system and adversarial models, design goals, and challenges. We propose our autonomous membership service and its rationale in Section 4, before discussing the implementation details in Section 5. Section 6 reports our empirical evaluation, while Section 7 reviews the related work. Finally, Section 8 concludes our work.

## 2. PRELIMINARIES

In this section, we first describe key characteristics of Intel SGX [21]. We then provide a brief overview of consensus protocols.

### 2.1 Intel SGX

**Enclave Execution.** Intel SGX [21] is a set of CPU extensions that are available on Intel CPUs starting from the Skylake microarchitecture [3], capable of providing hardware-protected TEE (or *enclave*) for generic computations. It enables a host to instantiate one or multiple enclaves simultaneously. An enclave is associated with a CPU-guarded address space which is accessible only by the enclave code; the CPU blocks any non-enclave code’s attempt to access the enclave memory. This effectively isolates the enclave from other enclaves concurrently running on the same host, from the OS, and from other user processes. Memory pages can be swapped out of the enclave memory, but they are encrypted using the processor’s key prior to leaving the enclave.

Enclaves cannot directly execute OS-provided services such as I/O. In order to access those services, enclaves have to employ OCalls (calls executed by the enclave code to transfer the control to non-enclave code) and ECalls (API for untrusted applications to transfer control back to the enclave). These ECalls and OCalls constitute the enclave boundary interface, enabling a communication between the enclave code and the untrusted application to service OS-provided functions. Clearly, enclave developer should design such interface with caution, for ECalls exposed to the untrusted application may open up an attack surface to the protected execution environment.

**Attestation.** Intel SGX allows a validator to verify if a specific enclave has been properly instantiated with the correct code [9] via its attestation mechanisms. These mechanisms also enable the validator and the attesting enclave to establish a secure, authenticated channel via which sensitive data can be communicated.

If the attestation is carried out between two enclaves instantiated on the same platform (or host), the mechanism in use is referred to as *local attestation*. Once the attesting enclave has been initiated, the CPU computes its *measurement*

(i.e., the hash of its initial state). Next, the CPU produces a message authentication code (MAC) of such measurement using a key that is accessible only by the validating enclave. The measurement and the MAC are then sent to the validating enclave for verification. Alternatively, if the validator is a remote party, the CPU issues a *remote attestation* by signing the measurement with its private key under the Enhance Privacy ID (EPID) scheme [18] [9]. The remote party obtaining the attestation then relies on the Intel’s Attestation Service (IAS) to verify the signature contained in the attestation [4], and then check the measurement value against a known value.

**Data sealing.** Enclaves can persist their private state to non-volatile memory via the data sealing mechanism. The enclave first requests the CPU for an enclave-specific key bound to its measurement. It then encrypts its private state using the requested key before storing it on persistent storage. The mechanism ensures that the sealed data can only be accessed by the enclave that sealed it. Nevertheless, enclave recovery using sealed data is susceptible to rollback attacks wherein an attacker (e.g., the malicious OS) provides the enclave with properly sealed but stale data [20].

## 2.2 Consensus Protocols

A rich literature of *consensus protocols* has been devoted to address a variety of fault tolerance problems in distributed systems [23, 13]. Consensus algorithms can be broadly categorized by the failure model that they assume. On the one hand, the crash failure model assumes that a faulty process may stop processing any message, and it does not resume [23]. On the other hand, the Byzantine failure model may allow for arbitrary faults. A process experiencing Byzantine failure may deviate from its expected behaviors in any manner, it may equivocate (i.e., sending contradictory messages to other processes), or it may intentionally delay its activity for any period of time [13].

Consensus algorithms aim to achieve *safety* and *liveness* in the presence of failures. Safety requires that non-faulty processes reach an agreement and never return conflicting results for the same request, whereas liveness means that these processes eventually agree on a value. In the following, we give an overview of Raft [23], a consensus protocol that assumes crash failure model and ensures safety regardless of timing. Nonetheless, its liveness depends on timing (e.g., the communication channels between processes are partially synchronous, wherein messages are delivered within an unknown but finite bound).

**Raft Consensus Protocol.** Raft assumes a system of  $n$  deterministic processes (or servers), among which at most  $f = \frac{n-1}{2}$  could be faulty. A faulty process fails by crashing. Each process stores a log that contains a series of commands, or events of interest. Raft ensures that logs of non-faulty processes contain the same sequence of commands.

A process can be in one of the three roles, namely *leader*, *follower* and *candidate*. Raft divides time into *terms* numbered with consecutive integers. In each term, one process is elected as the leader, and all other processes are followers. The leader exchanges periodic heartbeats with all followers to maintain its authority. If a follower does not receive any heartbeat after an *election timeout* period, it assumes that the leader has crashed. The follower then increases its term, changes its role to candidate and begins an election.

The candidate becomes the leader if it receives votes from a majority of processes. We refer readers to [23] for further details on the leader election and its election criteria.

During normal operation, the followers are passive, only responding to requests from the leader and candidate. The leader receives all the commands (e.g., requests from the clients), and replicates the commands on the followers. The leader first appends a command as a new entry which is uniquely identified by the leader’s *term* and an index to its log. Next, it broadcasts the entry to all of the followers. The followers append the received entry to their logs, and acknowledge the receipt to the leader. Once the leader receives the acknowledgement from a majority of the processes (i.e.,  $f + 1$  or more processes), it *commits* the entry (i.e., execute the command contained in the entry), and all preceding entries in its log. The leader records the highest index it knows to be committed, and includes this index in subsequent messages to the followers, thereby informing them on the committed entries.

Although Raft is designed for crash failure model, there exists attempt that deploys Raft in a Byzantine setting [7]. This is achieved through the use of TEEs. More specifically, by running the consensus protocol inside an enclave that offers attested execution, one can restrict adversarial behaviours of the faulty processes, thereby reducing the threat model from Byzantine fault tolerance to crash fault tolerance, to which Raft applies.

## 3. THE PROBLEM

In this section, we first define our system model in Section 3.1. Next, we present the system goals that we seek to achieve in Section 3.2. We then describe in Section 3.3 the adversary model against which our proposed autonomous membership service is designed. Finally, we detail the challenges of enabling autonomous membership service for enclave processes in Section 3.4.

### 3.1 System Model

We consider a typical outsourced computation model consisting of a cloud service provider (CSP) and its users. The CSP operates a set of machines (or hosts) that are equipped with Intel SGX-enabled processors [21]. A user (i.e., application owner) uploads her code to the cloud, and executes her application on the cloud’s infrastructure. The application involves secrets that the user wishes to keep private from the CSP. This can be achieved by running the application code (or a critical portion of its code) inside an SGX enclave. In particular, the code uploaded to the cloud does not contain any secret. The SGX architecture then relies on the CSP’s hypervisor or OS to instantiate an enclave process using the application code. Subsequently, the enclave process attests itself to the application owner so as to convince the latter that it is properly instantiated. Upon successful attestation, the secret materials are securely provisioned to the enclave.

The application code can be split into security-critical and non-critical components. The critical component is to be run in an enclave (or enclaves). The CSP may spawn multiple instances of the same application. Each application instance comprises both the critical and non-critical components, and they are bound together. If the critical component is divided into a set of enclaves, we assume that they are also uniquely bound together. In another word, the service provider can-

not mix-and-match a component of one instance with that of another instance without breaking the function of the application. Consequently, it suffices to identify an application instance by its enclave (or one of its enclaves). Without loss of generality, we assume throughout the rest of the paper that the application runs entirely in a single enclave, and abuse the language to refer to each instance as a *process*.

**Functional vs. Nonfunctional process.** We assume that the application involves secret materials that are concealed from the CSP or any untrustworthy party, and that such secret is necessary for the application’s operation (e.g., encryption/decryption keys). While the CSP has a sole discretion in instantiating a new process using the application code, such a process is initially *nonoperational*. Only until it is provisioned with the aforementioned application secret does it become *operational*. Put differently, although the CSP facilitates the commissioning of a new process, it cannot alone complete such a commission.

**Communication.** The processes communicate only by passing messages along a fixed set of channels facilitated by CSP. Each pair of processes is connected via a reliable, authenticated point-to-point communication channel that does not drop any message. These channels, however, are subject to a delivery scheduler controlled by the CSP. This allows the CSP to impose arbitrary delay on any message at any channel of its choice. When a channel is not intervened by a malicious delivery scheduler, it is *synchronous*, in a sense that all messages sent via that channel are delivered within a finite delay  $\Delta$  known to all processes. The malicious scheduler can postpone the delivery of any message for an arbitrary duration, causing a *communication fault*, wherein messages communicated via that channel are not delivered within the delay  $\Delta$ . Finally, we do not assume any global clock, or bounds on relative local clock speed and execution speed of each process.

**Faults.** We call a process *faulty* if it has crashed or halted. A process is *correct*, otherwise. Under the communication model described earlier, correct processes may be *isolated* from other correct processes. Let us consider an application **Prog**. We denote by  $n_s(\text{Prog})$  the number of **Prog**’s existing processes at a given moment  $s$ . We deem a process  $p$  isolated if  $p$  cannot communicate with a quorum of  $\lfloor \frac{n_s(\text{Prog})}{2} \rfloor$  processes in a synchronous manner. We call a process *active* if it is correct and not isolated.

Let  $f_s(\text{Prog})$  and  $i_s(\text{Prog})$  be the number of faulty and isolated processes among the  $n_s(\text{Prog})$  processes, respectively. We deem the application **Prog** to be in *anarchy* at the given moment  $s$  if:

$$f_s(\text{Prog}) + i_s(\text{Prog}) > \lfloor \frac{n_s(\text{Prog}) - 1}{2} \rfloor \quad (1)$$

Hereafter, when it is clear from the context, we omit  $s$  and **Prog** from the notations. Figure 1 depicts different faulty scenarios **Prog**’s processes may experience, and state where **Prog** is in anarchy.

### 3.2 System Goals

To allow for elasticity, the CSP needs to commission or decommission processes dynamically overtime. Nonetheless, this may expose the system to forking attacks, as described earlier in Section 1. To address this problem, we seek an

autonomous membership service that protects the application against such forking attacks. The membership service polices the commissioning of new processes, and excommunication of faulty or isolated processes from the system, subject to the SLA concerning the upkeep of the application. More specifically, while the CSP can instantiate a new process at will, such a process remains nonoperational until its request to join the existing group of processes is approved by the membership service. In addition, the membership service actively monitors the status of each member process, and excommunicates a process that it suspects of being faulty or isolated from the group.

To render the membership service autonomous, our approach is to augment each process of an application **Prog** with a membership list that keeps track of **Prog**’s active processes, and ensure that the membership lists of all active processes converge without relying on manual intervention of the application owner, or a trusted third party:

- If a process becomes faulty, it is required that the faulty process is eventually suspected of having halted, and excommunicated from the membership list of all active processes. Communication from the excommunicated process are discarded.
- A newly instantiated process must join the system via the membership service before it is rendered functional. The membership service ensures that this new process is added to the membership list of all existing active processes, and the former is brought up-to-date with the current system membership list.
- Except for the instantiation of the first process, the application owner does not involve in operation of the membership service. All subsequent process attestation and secret provisioning are handled by active processes in the system.

### 3.3 Adversary Model

We consider a powerful adversary that possesses all capabilities of the CSP. In practice, such adversary can be an errant insider having complete access to the cloud infrastructure, or an attacker exploiting vulnerabilities in the cloud’s software stack. The adversary is interested in conducting a forking attack on the victim application, attempting to launch an arbitrarily large number of enclave instances in violation of the SLA, or to isolate the concurrently running instances from each other. Nonetheless, we assume that the adversary does not seek to undermine the application’s availability (e.g., by decommissioning all of its instances).

The adversary has complete control over the cloud’s OS and/or hypervisor. It can schedule or corrupt all non-enclave processes, and restart enclave instances at will. Moreover, it can control all network communications between the enclave instances regardless of whether they are instantiated on the same or across different hosts. More specifically, the adversary can read, modify, reorder, delay or even drop the messages sent by the enclaves arbitrarily.

Nonetheless, it cannot compromise protection mechanisms of the SGX processors. That is, it can neither access the enclave runtime memory nor learn the application secrets that are protected by the enclave execution, and does not have any access to the processor’s private keys that are used for attestation and data sealing functions. Finally, we assume

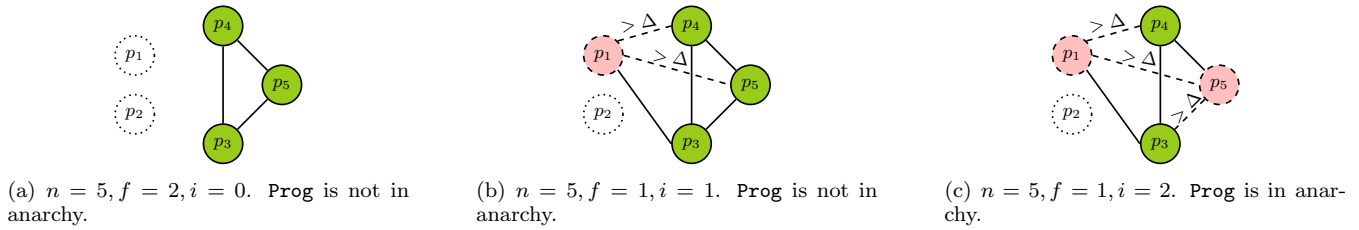


Figure 1: Illustrations of faulty, isolated and active processes. They are depicted in dotted, dashed and solid circles, respectively.

that the adversary is computationally bounded, and cannot break standard cryptographic assumptions.

We assume the application code running inside the enclave is secure. We do not consider side-channel attacks against the enclave execution [26] and DoS attacks against the system.

### 3.4 Challenges

The first challenge in enabling autonomous membership service stems from the need of injecting application secret to the newly instantiated process. Since the untrusted OS is in charge of enclave instantiation, any secret material necessitated for the operations of the enclave process cannot be contained in the enclave code, but should rather be provisioned to the enclave only after it has been properly instantiated. In a typical pipeline of enclave instantiation, this secret provisioning is realised via the remote attestation procedure carried out by the application owner. Nonetheless, it is unreasonable to assume the involvement of the application owner in the instantiating of every enclave process, especially when processes may be commissioned dynamically to meet the elasticity requirement of the application. Ideally, the application owner only needs to involve in the instantiating of the first enclave instance during the application provisioning process, and does not have to take part in the upkeep of her application.

The second challenge lies in ensuring the security of the membership service itself, especially against forking attack [12] or logical partition. To further elaborate on this challenge, we describe in the following two straw-man solutions and point out their weaknesses. The first solution assumes a trusted *directory server* keeping track of all running processes of the application, thereby offering the membership service. The second solution attempts to implement the *directory server* as a replicated state machine so as to avoid single-point-of-failure.

**Directory server.** A trusted directory server is in charge of managing and attesting the correct instantiating of application processes, and keeping track of all such running processes. More specifically, the user uploads the expected measurement of her application enclave process (i.e., the hash of its initial state) along with its secret to the directory server via a secure channel. The directory server then performs a remote attestation with the newly instantiated enclave process on behalf of the user, and provides it with the application secret upon successful attestation. It maintains a list of concurrently running processes, and refuses the instantiating of the new process if it violates the SLA.

The directory server determines the aliveness of a process in its process list by listening to a **heartbeat** message sent by

that process, to which it responds with an acknowledgement **ack**. If it does not receive the **heartbeat** from the process after a certain time-out period, it would assume that such process has crashed, removing such process from its process list. It is important to note that the excommunication of a process may be due to communication fault, rather than the said process having halted (i.e., unjust excommunication). In order to keep the logic of the directory server simple, we can enforce the unjust excommunication by forcing the excommunicated process to halt. This can be achieved by programming a process to *commit suicide* if it does not receive an **ack** from the directory server after a configurable **time-to-live** period, and to restart its **time-to-live** timer every time it receives an **ack**. Put differently, the **ack** from the directory server acts as a *lease* that allows the receiving process to continue operating for a fixed period of time until it expires [16].

This approach, nevertheless, is not desired, for the directory server becomes a single-point-of-failure, or target for attacks. For example, if the directory server is implemented using TEE, it is subject to rollback attack [20] in which the adversary may deceive it to resume operations using stale process list, violating the integrity and consistency of the directory server’s operations.

**Replicated directory service.** One approach to sidestep the aforementioned single-point-of-failure issue is by replicating the directory server. In particular, the directory service now comprises of a set of replicas that leverage fault-tolerant protocol to implement a replicated state machine, ensuring consistency of the service’s operation despite individual replica failures [25]. Nonetheless, it is challenging to protect these replicas against forking attack [12], wherein the adversary manage to splits the directory service replicas into two “cliques”, each of which is unaware of the other. Such attack allows the adversary to run twice as many number of processes as what would otherwise be allowed by the SLA. Addressing this issue, then, requires a membership service offered to the directory service replicas themselves.

## 4. OUR DESIGN

In this section, we present our system AMES (Autonomous Membership Service). We first give an overview of AMES’s design. We later discuss in details the main protocols that comprise the autonomous membership service. We remark that although our discussion assumes processes are provisioned as SGX enclaves, the technique underlying AMES is applicable to other Trusted Execution Environment platforms that support similar features to SGX enclaves.

## 4.1 Overview

To enable the autonomous membership service for an application `Prog`, AMES augments `Prog`'s processes to equip them with the following parameters:

- `peerList`: a list that records all active processes of `Prog`.
- term number: a value that represents the current term (which we shall define in the following). The term number increases monotonically.
- replicated log: a log that records messages the process exchange during the operation of the membership service. A log entry is indexed using the term number of the term at which it is added, together with sequence number that increases monotonically over time across terms.
- `commitIndex`: the index of the latest entry in the replicated log that is known to be committed. This value increases monotonically over time.
- heartbeat timeout `T`: a randomized timeout used in AMES's `heartbeat` mechanism, unique for each process.
- candidature timeout: the timeout used during the leader election. This timeout is the same for all processes.

While a process can always derive `peerList` from its replicated log, and one may argue that they can be merged into a single data structure, we find that treating them as two separate entities keep our exposition clearer.

AMES requires the involvement of the application owner during the application bootstrap so as to attest the correct instantiating of the first process, and to provide it with the application secret (Section 4.2). Once the first instance is fully operational, the application owner may go offline until the time when she wishes to shutdown the application. All subsequent process instantiating (or dismantling) over the lifetime of the application will be policed by AMES. To this end, `Prog` needs to incorporate in its implementation AMES's protocols for managing the processes' `peerLists` and ensuring that they converge. We describe these protocols in the remaining of this section, deferring discussion on the implementation details to Section 5.

**Workflow.** Let  $n$  denote the number of processes of `Prog` at a given time. AMES first elects one process to be the *leader*, while the other  $[n - 1]$  processes are *followers*. A leader is associated with a *term*, which is a time period of arbitrary length in which the leader retains its authority. The leader relies on a `heartbeat` mechanism to monitor the status of the processes in the system, and to maintain its authority over them (Section 4.3). In normal operation (i.e., the leader is active), the followers need not exchange `heartbeat` messages among themselves, but simply respond to messages from the leader.

The commission of a new process requires an "approval" from the membership service via a `NEWPROCESS` protocol orchestrated by the leader. More specifically, the leader attests the new process, and checks against the SLA if its commission is eligible (e.g., if the quota of concurrently running processes has not been consumed). The leader then

provisions the application secret to the new process, rendering it functional, and informs all followers to update their `peerLists` accordingly (Section 4.4). If the leader suspects a follower process  $p_h$  of having halted (e.g., it did not receive a response from  $p_h$  for the `heartbeat` message it has sent to  $p_h$  after a configurable heartbeat timeout  $T_h$  which is unique to  $p_h$ ), it drives the other followers to expel  $p_h$  from their `peerList` via the `EXPELPROCESS` protocol (Section 4.5). Finally, when the user decides to shutdown the application, she sends the shutdown request to the leader, who then coordinates the termination of all active processes via a `SHUTDOWN` protocol, at the end of which they all cease to operate (Section 4.6).

It is important to note that AMES admits unjust excommunication of a active process whose connection with the leader is severed. To prevent the excommunicated process from disrupting the membership service, its future communication (i.e., after the point of excommunication) will be discarded, and the process will eventually halt.

The workflow described above has not accounted for leader failure. If the leader halts or becomes isolated, followers will eventually detect such failure, and engage in a leader election protocol to elect a new leader. In particular, if a follower does not receive a `heartbeat` message from the leader after its heartbeat timeout, it will assume that the current leader is no longer viable, and attempt to claim the leadership for itself (Section 4.3). When a new leader is elected, the new term starts. During the period where a process attempting to claim the leadership, it is referred to as a *candidate*. More than one candidates may compete for the leadership at the same time. The leader election guarantees that only one candidate succeeds, and all active processes in the system recognize the new leader as legitimate.

## 4.2 Application Bootstrap

In AMES, application bootstrap refers to the instantiating of the very first process of the application. This procedure is initiated by the application owner, who instructs the cloud service provider to instantiate an enclave process using the application code. Since the cloud's OS/hypervisor is not trusted, the instantiating has to be attested by the application owner, and application secret is only provisioned to the enclave after a successful attestation.

AMES benefits directly from Intel SGX's remote attestation protocol [9] in implementing the above mentioned attestation and secret provisioning. We highlight key characteristics of the protocol, and refer the readers to [9] for further details. The SGX architecture features a special enclave, called *Quoting Enclave*, to facilitate the remote attestation. After being instantiated, the attesting enclave invokes the SGX's `EREPOR`T instruction and specifies the Quoting Enclave as its target. The `EREPOR`T creates a structure, called `report`, that contains all necessary information to verify the correctness of the attesting enclave's instantiating, and produces a message authentication code (MAC) tag for the report. The MAC is computed using a "report key" that is associated only with the attesting enclave and its intended target (i.e., Quoting Enclave in this case). The Quoting Enclave verifies the `report`, and subsequently replaces its MAC with a signature signed using a CPU's private key under the EPID group signature scheme [18], producing a remote attestation `quote`<sup>1</sup>. As Intel does not publish the group public

<sup>1</sup>EPID allows a signer to sign objects with his private key

key associated with the CPUs’ private key, the verifier has to rely on the Intel Attestation Service to validate the remote attestation **quote**.

It is worth noting that SGX remote attestation mechanism allows the verifier to include a challenge in her attestation requests, and the attesting enclave to include a manifest that contains response to the verifier’s challenge and an ephemerally generated public key to the signed structure **report**, which is then carried over to the **quote** [9]. The challenge and response ensures freshness of the **quote**, avoiding replay attack by the malicious OS. The ephemeral public key allows the verifier to securely provision application secrets to the enclave. Consequently, upon successful application bootstrap, the first application process becomes fully functional, containing the secrets necessitated for the application’s operation.

### 4.3 HEARTBEAT EXCHANGE and Leader Election

During normal operation, the leader uses **heartbeat** mechanism to maintain its leadership and to monitor status of follower processes. A **heartbeat** is essentially a message sent by the leader, and acknowledged with a corresponding **ack** message by a follower, via an authenticated channel. A **heartbeat** always contains the term number and **commitIndex** of the leader (the latter is the index of the latest entry in the leader’s replicated log that has been committed). The follower relies on the leader’s **commitIndex** to detect if its replicated log and state are up-to-date with those of the leader. If a follower finds that its log misses some entries, or its state is out of date, it fetches the missing log entries and commit them, bringing its state up-to-date. The **heartbeat** can also contain information regarding the commission (ex-communication) of a new (existing) follower process, which we shall elaborate in the following subsections.

When a new process becomes functional, it initially assumes the role of a follower. The follower process  $p_i$  and the leader process  $p_L$  establish a randomized timeout  $T_i$ , which is unique to each follower process, chosen from a fixed interval. Via the **heartbeat** mechanism,  $p_L$  can identify if it has become isolated, and halts if it is so. On the other hand, if  $p_L$  is active but cannot receive an **ack** message from  $p_i$  over a period of  $T_i$ , it suspects the latter of having halted, and triggers the **EXPELPROCESS** protocol to inform all other followers of  $p_i$  failure. Alternatively, if the follower  $p_i$  is active but cannot receive a **heartbeat** from the leader over a period of  $T_i$ , it changes its role to *candidate*, increases its term number and broadcasts a **RequestVote** message to all existing processes in its **peerList** to claim the leadership.

The **RequestVote** message from  $p_i$  takes a form of  $(p_i, \text{newTerm}, \text{lastLogIndex})$ , where **newTerm** is its current term number (which has just been increased), and **lastLogIndex** is the index of the last entry in its replicated log. Upon receiving **RequestVote** from  $p_i$ , a receiver process  $p_r$  acts as follows:

- If  $p_i$  is not in  $p_r$ ’s **peerList**, it ignores  $p_i$ ’s **RequestVote**.
- If  $p_i$  appears in its **peerList**, but  $p_r$  has recently received a message from the leader  $p_L$  (i.e., its timeout  $T_r$

without uniquely identifying the signer or linking his different signatures. This is achieved by assigning each signer to a “group” such that the group’s public key can be used to verify all signatures produced by the group member.

has not expired),  $p_r$  queues  $p_i$ ’s **RequestVote** if it has not queued any other **RequestVote**. If  $p_r$  has queued a **RequestVote** from another candidate  $p_c$ ,  $p_r$  determines the dominating candidate among  $p_i$  and  $p_c$ <sup>2</sup>, keeping only the **RequestVote** of the dominating candidate, and informing the dominated candidate on the existence of the dominating one. Once its **heartbeat** timeout  $T_r$  expires,  $p_r$  check if it is dominated by the candidate whose **RequestVote** is currently kept in its queue. If so, it grants its vote to the candidate; otherwise it discards such a **RequestVote** and broadcasts its own. On the other hand, if  $p_r$  receives the next **heartbeat** from  $p_L$  before  $T_r$  expires, it resets  $T_r$  and discards any **RequestVote** in its queue.

- If  $p_i$  is in its **peerList**, but it has voted for another candidate (potentially itself) that is not dominated by  $p_i$ ,  $p_r$  informs  $p_i$  on such candidate.

If  $p_i$  collects a quorum of votes from a *majority* of the processes in its **peerList** for its **RequestVote** before its **heartbeat** timeout expires, it assumes the role of the leader. It then drives the processes to expel  $p_L$  from their **peerList**. While waiting for the votes, if  $p_i$  is informed of another candidate that outweighs itself,  $p_i$  resets its **heartbeat** timeout and waits for the other candidate to pronounce its leadership. When  $p_i$  is informed of a new leader whose term is as large as its current term via a **heartbeat** from the latter, it acknowledges the leader’s **heartbeat**, and returns its role to follower. If  $p_i$  fails to collect a quorum of responses from a majority of the processes, or a **heartbeat** from the new leader within its **heartbeat** timeout, it increases its term and triggers a new election with a new **RequestVote**. Finally, if  $p_i$  is still assuming the candidate role (i.e., it can neither claim the leadership nor return to the follower role) when its candidature timeout expires, it commits suicide.

*Security Arguments:* The leader election in AMES follows Raft’s leader election [23] to certain extent. The key difference here is that a candidate  $p_i$  shall halt (by committing suicide) if it cannot collect a quorum of responses or a **heartbeat** from the new leader within the candidature timeout. Here, we reason about the correctness of this design when the system is not in anarchy.

The fact that  $p_i$  issuing its **RequestVote** implies either the leader  $p_L$  has become faulty or isolated, or the connection between  $p_i$  and the active  $p_L$  has been severed.

If  $p_L$  is isolated, it will detect its isolation by itself and halts, thereby becoming faulty. When  $p_L$  is indeed faulty<sup>3</sup>, other processes will either grant  $p_i$  their vote or inform it on the dominating candidate. If  $p_i$  is active (i.e., it is not isolated), it will have received enough responses to either claim the leadership, and thereby orchestrating the removal of  $p_L$ , or to wait to be contacted by the new leader. On the other hand, if  $p_L$  is active, it would have simultaneously driven other processes to expel  $p_i$  from their **peerList**, preventing them from responding to  $p_i$ ’s **RequestVote**. This, in turn, prevents  $p_i$  from being able to collect a quorum of

<sup>2</sup>We say a candidate  $p_c$  dominates (or outweighs) another candidate  $p_e$ , if it has a more up-to-date log, which is determined by the index of last entries in the logs [23].

<sup>3</sup>In this case, no new process can be commissioned, and no process failures can be detected during the period when the system is in the leader election stage.

responses for its `RequestVote`, which subsequently leads to  $p_i$ 's committing suicide.

*Remark.* While our design admits a potential unjust excommunication of an active process whose connection to the active leader is severed, this is necessary to sidestep the difficulty of solving membership problem in an asynchronous environment when faults may be present [15, 24]. By requiring the active yet excommunicated process to commit suicide, AMES preserves the consistency of the leader's expelling decision and safety of the membership service.

#### 4.4 NEWPROCESS Protocol

After the application has been bootstrapped and as long as it is not in anarchy, AMES enables the application to commission new processes in an autonomous manner, without relying on manual intervention of the application owner, or a trusted third party. This section describes how AMES commissions a new process  $p_j$  and ensures that the `peerLists` of all active processes are updated accordingly and that they converge. Let  $n$  be the number of existing processes at the beginning of  $p_j$ 's commission, and assume that the SLA allows  $n > 2$ .

**Case  $n = 1$ .** The only running and functional process at the beginning of  $p_j$ 's commission is the leader  $p_L$ . The interaction between the two processes happen in two phases. In the first phase,  $p_j$  contacts  $p_L$  for attestation. The results of this attestation are two-fold. First,  $p_L$  is convinced that  $p_j$  is correctly instantiated. Second, the two processes establish a secure and authenticated communication channel. In the second phase,  $p_L$  updates its `peerList` to include  $p_j$ , provisions the application secret as well as its current `peerList` to  $p_j$  via the secure channel and establishes a heartbeat timeout  $T_j$ . Likewise,  $p_j$ 's `peerList` now contains itself and  $p_L$ . At this point,  $p_j$  becomes fully functional, and assumes a role of a follower.

*Security Arguments:* The interaction between  $p_L$  and  $p_j$  described earlier follows the workflow of a 2-phase-commit (2PC) protocol [11], with  $p_L$  assuming the role of the coordinator. The attestation in the first phase acts as a prepare phase in the 2PC protocol, whereas the provisioning of the application secret to  $p_j$  and updating of `peerLists` mirror the commit phase of the 2PC protocol. Provided that the application is not in anarchy (i.e.,  $p_L$  remains active), our NEWPROCESS protocol does not suffer from blocking, and the `peerLists` of  $p_L$  and  $p_j$  converge.

**Case  $n = 2$ .** Let us call the two existing processes  $p_L$  and  $p_2$ , with  $p_L$  being the leader. The commissioning of  $p_j$  happens in two phase. In the first phase,  $p_L$  engages in a remote attestation with  $p_j$ , and at the same time queries  $p_2$  (this query is piggybacked on a `heartbeat` sent to  $p_2$ ) if it is ready to update its `peerList` to contain  $p_j$ . In the second phase, if the attestation checks out, and  $p_2$  responds with a "yes" vote (the vote is piggybacked on the `ack` message corresponding to the `heartbeat` carrying the query),  $p_L$  updates its `peerList` to include  $p_j$ , transfers the application secret as well as the current `peerList` to  $p_j$ , and establishes a heartbeat timeout  $T_j$  with  $p_j$ .  $p_L$  also informs  $p_2$  to update its `peerList` accordingly.  $p_j$  now becomes fully functional, and assume a role of a follower.

*Security Arguments:* Similar to the case where  $n = 1$ , the

interaction between  $p_L$ ,  $p_2$ , and  $p_j$  described above follows the workflow of a 2PC protocol [11], with  $p_L$  assuming the role of the coordinator. Provided that the application is not in anarchy (i.e., both existing processes are active), our NEWPROCESS protocol does not suffer from blocking, and the `peerLists` of all processes converge at the end of  $p_j$ 's commission.

**Case  $n > 2$ .** The leader  $p_L$  first checks if the commission of  $p_j$  is eligible according to the SLA. If so, it engages in a remote attestation with  $p_j$ . Next,  $p_L$  follows Raft [23] to informing all other followers about the commission of  $p_j$ :

1.  $p_L$  puts an entry  $\langle \text{add}, p_j \rangle$  to its log.
2.  $p_L$  broadcasts  $\langle \text{add}, p_j \rangle$  to all the followers.
3. Upon receiving  $\langle \text{add}, p_j \rangle$ , a follower  $p_i$  acknowledges the receipt by sending  $\langle \text{ack}_i, \text{add}, p_j \rangle$  to  $p_L$ , and puts  $\langle \text{add}, p_j \rangle$  to its log.
4. Once  $p_L$  confirms that  $\langle \text{add}, p_j \rangle$  has been replicated on a majority of the processes, its adds  $p_j$  to its `peerList`, establishes a heartbeat timeout  $T_j$  with  $p_j$ , and transfers to  $p_j$  the application secret and its `peerList`. In another word,  $p_L$  "commits" the entry  $\langle \text{add}, p_j \rangle$ .
5.  $p_L$  announces the commit of  $\langle \text{add}, p_j \rangle$  in the next message it exchanges with the followers. Once the followers see the commit, they add  $p_j$  to their `peerList`.

*Security Arguments:* Provided that the application is not in anarchy, the `peerLists` of all active processes converge. This follows directly from the security guarantees of Raft [23], which ensures protocol safety as long as any majority of the processes are active (i.e., they are operational and can communicate with each other without communication fault).

#### 4.5 EXPELPROCESS Protocol

When the leader  $p_L$  suspects that a process  $p_h$  has become faulty (e.g., it does not receive a response for the `heartbeat` sent to  $p_h$  after the timeout  $T_h$  bound to  $p_h$ ), it drives all other followers to expel  $p_h$  from their `peerList`.

1.  $p_L$  puts an entry  $\langle \text{Expel}, p_h \rangle$  to its log.
2.  $p_L$  broadcasts  $\langle \text{Expel}, p_h \rangle$  to all the followers.
3. A follower  $p_i$  acknowledges the receipt by sending  $\langle \text{ack}_i, \text{Expel}, p_h \rangle$  to  $p_L$ , and puts  $\langle \text{Expel}, p_h \rangle$  to its log.
4. Once  $p_L$  confirms that  $\langle \text{Expel}, p_h \rangle$  has been replicated on a majority of the processes (*excluding*  $p_h$ ), its expels  $p_h$  from its `peerList`. In another word,  $p_L$  commits the entry  $\langle \text{Expel}, p_h \rangle$ .
5.  $p_L$  announces the commit of  $\langle \text{Expel}, p_h \rangle$  in the next message it exchanges with the followers. Once the followers see the commit, they expel  $p_h$  from their `peerList`.

*Security Arguments:* Since the EXPELPROCESS protocol is essentially an instance of Raft [23], it follows from the security guarantees of Raft that the `peerLists` of all processes at the end of the EXPELPROCESS protocol converge, so long



as the application is not in anarchy. It is worth noting that AMES admits unjust excommunication of a correct but potentially isolated process or an active process whose connection to the leader  $p_L$  is severed. To prevent the excommunicated process from disrupting the membership service, all its future communication will be discarded. This is enforced by having other processes not responding to message from a process that is not in their `peerList`. This would also lead the unjustly excommunicated process to halting. In particular, the unjustly excommunicated process will trigger a leader election, for it has not received any `heartbeat` from the leader  $p_L$ . Nonetheless, it cannot obtain the required quorum of responses from other processes for its `RequestVote`, thereby committing suicide as mentioned in Section 4.3.

## 4.6 SHUTDOWN Protocol

When the application owner wishes to shutdown all processes of the application, she sends the shutdown request to the leader  $p_L$ , who then coordinates the termination of all active processes as follows:

1.  $p_L$  puts an entry `<preShutdown>` to its log. It seals the application data, if any, to persistent storage using SGX's data sealing mechanism.
2.  $p_L$  broadcasts `<preShutdown>` to all the followers.
3. A follower  $p_i$  puts the received `<preShutdown>` to its log. Next, it seals the application data, if any, to persistent storage, and responds to the leader by sending `<acki, preShutdown>`.
4. Once  $p_L$  confirms that `<preShutdown>` has been replicated on a majority of the processes, it broadcasts `<commitShutdown>` to all the followers, thereby announcing the committing of the shutdown request. Finally, it informs the application owner that shutdown request has been served, and halts.
5. Once the followers see the `<commitShutdown>`, they halt.

*Security Arguments:* Similar to the EXPELPROCESS protocol, the SHUTDOWN protocol is an instance of Raft [23]. It follows from the security guarantees of Raft that so long as the application is not in anarchy, all active processes shall respond to the shutdown request, and halt. Even if the application is in anarchy, as long as the leader commits suicide, all other processes will eventually halt (as we discuss in the next subsection), thereby also bringing the shutdown request into effect.

## 4.7 Anarchy

The application is in anarchy if the total number of faulty and isolated processes (denoted by  $f$  and  $i$  respectively) exceeds  $\lfloor \frac{n-1}{2} \rfloor$ . When the application falls into anarchy, the status of the current leader  $p_L$  becomes critical. If  $p_L$  remains active, the application may still be operational. On the other hand, if it is either faulty or isolated, all existing processes will eventually halt, thereby shutdown the entire application. We remark that this only affects availability of the application, not the security of the membership service. More specifically, even if the application is in anarchy,

the `peerList` of active processes do not diverge. In the following, we examine the situations where the application is in anarchy and how the status of the leader  $p_L$  affects the operation of the application.

**Case ( $f > \lfloor \frac{n-1}{2} \rfloor$ ):** If the leader  $p_L$  is among the  $f$  faulty processes, the remaining correct processes will trigger leader election. Nonetheless, since  $f > \lfloor \frac{n-1}{2} \rfloor$ , none of them will be able to obtain a sufficient quorum of responses for their `RequestVote`. Consequently, all correct process shall eventually halt by committing suicide, which effectively shuts the application down in its entirety.

If  $p_L$  is active, it will attempt to trigger the EXPELPROCESS to excommunicate faulty processes. Nonetheless, it cannot collect sufficient acknowledgement after step (3) of the EXPELPROCESS protocol, and hence shall identify itself as being isolated, thereby halting by committing suicide. This further increases  $f$ , and  $p_L$  now has become faulty. Subsequently, the remaining correct processes will trigger the leader election, and eventually halt for not being able to collect a sufficient quorum of responses to their `RequestVote`, shutting down the application.

**Case ( $i > \lfloor \frac{n-1}{2} \rfloor$ ):** If  $p_L$  identifies itself as being isolated (i.e.,  $p_L$  is among the  $i$  isolated processes), it will halt and thereby becomes faulty. Once  $p_L$  is faulty, the remaining isolated processes will trigger the leader election, and will eventually halt for not being able to collect a sufficient quorum of responses to their `RequestVote`. In another words, the  $i$  isolated processes will eventually become faulty, rendering  $f > \lfloor \frac{n-1}{2} \rfloor$  and causing the application to shut down, as discussed earlier.

If  $p_L$  is active, and can successfully drives the followers to complete all NEWPROCESS and EXPELPROCESS protocols that it triggers, the application remains operational. We note that whenever there is a new process added to the system membership list, or an existing process is excluded, the value of  $n$  changes, which may moves the application out of or into anarchy.

**Case ( $f < \lfloor \frac{n-1}{2} \rfloor \wedge i < \lfloor \frac{n-1}{2} \rfloor \wedge (f + i) > \lfloor \frac{n-1}{2} \rfloor$ ):** If  $p_L$  is isolated, it will halt and thereby become faulty. Once  $p_L$  is faulty, the remaining isolated processes will trigger the leader election. Nonetheless, being isolated, they cannot collect a sufficient quorum of responses to their `RequestVote`. Consequently, they will halt, and thus become faulty. At this point, the total number of faulty processes exceeds  $\lfloor \frac{n-1}{2} \rfloor$ , causing the application to shut down.

If  $p_L$  is active, and can successfully drives the correct followers to exclude the faulty processes from their `peerLists` via the EXPELPROCESS protocol, the application continues to be operational. Similar to the previous case where  $i > \lfloor \frac{n-1}{2} \rfloor$ , each process commission or excommunication changes the value of  $n$ , which may move the application out of or into anarchy.

## 4.8 Timeout Configurations

It should be clear from our security arguments thus far that safety of AMES does not depend on timing. In another words, `peerLists` of all active process converge even if they do not share any global clock, and their execution speed may differ. Nonetheless, its liveness (i.e., the service being able to make progress) is inevitably dependent upon timing, especially HEARTBEAT\_EXCHANGE and Leader Elec-

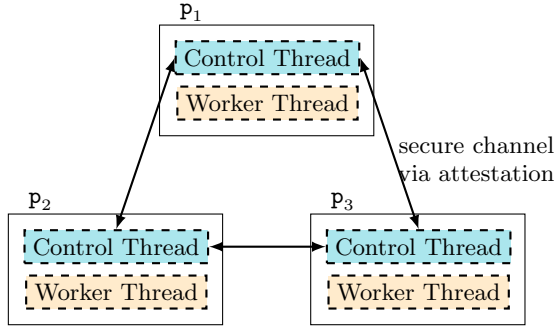


Figure 2: A process in AMES comprises a control thread that is responsible for partaking in the autonomous membership service, and a worker thread that is responsible for handling the application’s operational logic.

tion protocols. A necessary condition for liveness in AMES is:

$$\text{broadcastTime} \ll \text{heartbeat timeout} \leq \text{candidature timeout} \ll \text{EIATime}$$

where `broadcastTime` is the the average time required for a process to send messages (in parallel) to other processes and receive their responses; `heartbeat timeout` and `candidature timeout` are the timeouts described in Section 4.1; and `EIATime` are events’ inter-arrival time, with events being a new process instantiating, or failure/dismantling of existing processes in the membership cluster. Ideally, the timeouts should be an order of magnitude larger than the `broadcastTime`, and `EIATime` should be a few orders of magnitude larger than the timeouts. It is also reasonable to configure the `heartbeat timeout` to be less than or equal to the `candidature timeout`. This setting allows the candidate to retry a leader election with a new `RequestVote` at a higher term, in case there is a split vote and the current term ends with no leader being elected.

While `broadcastTime` and `EIATime` depends on the underlying communication system and dynamic workloads of the application, respectively, the timeouts can be tuned at our discretion. As we show in our empirical evaluation, `broadcastTime` ranges from `10ms` to `60ms` for processes located across different geographical regions. We expect typical `EIATime` in typical applications to be in the neighborhood of tens of seconds. Consequently, the timeouts could be set in the range of `250 – 500ms`. Section 6 elaborates on the effect of timeout configuration on the performance of AMES.

## 5. IMPLEMENTATION

To facilitate the incorporation of AMES’s logic in the implementation of an enclave application `Prog`, we provide application developers with a AMES library. The library envelops the membership service’s operations in a *control thread*, separating them from the original operational logic of `Prog` that is to be implemented in a *worker thread* (or threads). As a result, each `Prog`’s enclave process comprises one control thread, and one or more worker threads (as depicted in Figure 2). Our AMES library implementation contains approximately 3,000 lines of C++ code, and builds on the Raft implementation by RethinkDB [5].

**Control thread vs worker thread.** The control thread

is responsible for exchanging `heartbeat` messages, conducting remote attestation and secure channel establishment in the `NEWPROCESS` protocol, and handling messages that are exchanged in other protocols of AMES. It is worth noting that the message exchange of the control thread entails context switching of the enclave. To reduce context switching overhead, our implementation piggybacks control thread’s messages on the worker thread’s interface with the untrusted application (e.g., its `ECall` and `OCall`) whenever necessary.

Recall that both the control thread and the worker thread belong to the same enclave, the control thread can seal the application data handled by the worker thread to persistent storage, if need be, in the `SHUTDOWN` protocol. In addition, the control thread handles the process’s suicide by destroying its own enclave.

**Reducing IAS access overhead.** In the `NEWPROCESS` protocol, the newly instantiated process has to attest itself to the leader, so as to convince the latter of its correct instantiation. If the new process is not located in the same physical host with the leader, they have to engage in a remote attestation procedure. In such a situation, the leader (i.e., remote verifier) cannot verify the attestation of an attesting process locally, but has to rely on the Intel’s Attestation Service (IAS) to verify the signature contained in the attestation [4]. Contacting the IAS in every run of `NEWPROCESS` protocol would incur a significant overhead that is unfavorable for the application performance.

We make an observation that the multitenancy nature of cloud computing allows for a remedy of the aforementioned IAS access overhead. In particular, it is often the case that the CSP provisions many processes (either of the same or different applications) on the same physical host. Therefore, we can reduce the IAS access overhead by introducing a `HostAttestationDelegate` (HAD) enclave at each physical host, and routing the remote attestation between two processes instantiated on two separate physical hosts via their prospective HAD enclaves.

Let us consider two physical hosts  $h_1$  and  $h_2$ , and denote their HAD enclaves by `HAD_E1` and `HAD_E2`, respectively, as depicted in Figure 3. These two HAD enclaves engage in a conventional remote attestation, and rely on IAS to verify the attestation. Once `HAD_E1` and `HAD_E2` have completed the attestation and established a secure channel, any pair of processes  $p_1^A$  and  $p_2^A$  of an application `A` provisioned on  $h_1$  and  $h_2$ , respectively, can remotely attest one another without having to invoke the IAS. In particular, to attest its correct instantiating to  $p_2^A$ ,  $p_1^A$  first leverages the local attestation procedure to prove its correctness to `HAD_E1` which then conveys this attestation result to `HAD_E2` via the secure channel. Finally, `HAD_E2` locally attests to  $p_2^A$  that it is correctly instantiated, gaining the latter’s trust, and then conveys  $p_1^A$ ’s attestation result to  $p_2^A$ . This approach essentially amortizes the IAS access overhead incurred by the attestation of HAD enclaves, and allows two processes  $p_1^A$  and  $p_2^A$  of an application `A` provisioned on two separate hosts  $h_1$  and  $h_2$  to conduct remote attestation without contacting the IAS.

## 6. EVALUATION

In this section, we report our experimental study of AMES. Our experiments are conducted in two different settings. The first setting is based on an in-house (local) cluster consisting of 35 servers, each equipped with Intel Xeon E3-

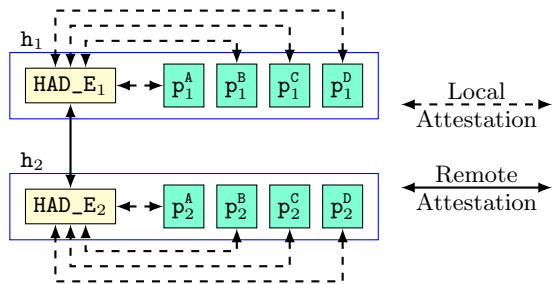


Figure 3: AMES introduces `HostAttestationDelegate` enclave at each host, and leverages them to reduce IAS access overhead incurred in the `NEWPROCESS` protocol.

Table 1: Latency (ms) between different regions on GCP.

Zone	us-west1	us-west2	us-east1	us-east4
us-west1	0.0	24.7	66.7	59.0
us-west2	24.7	0.0	62.9	60.5
us-east1	66.7	62.9	0.0	12.7
us-east4	59.1	60.4	12.7	0.0

1240 2.1GHz CPUs, 32GB RAM and 2TB hard drive. The latency between any two servers are roughly  $0.13ms$ . In this setting, we run each process on a separate server. The second setting is based on Google Cloud Platform (GCP), in which we use a separate instance provisioned with 2 vCPUs and 8GB RAM for each process. The instances on GCP span across 4 regions, namely Oregon (us-west1), Los Angeles (us-west2), South Carolina (us-east1) and North Virginia (us-east4). The average latency between these regions observed in our experiment is detailed in Table 1.

The trusted code base in our experiment is implemented using Intel SGX SDK [2]. For the experiments that we run on GCP, we configured the SDK to run in simulation mode, as SGX is not available on GCP instances. We measured the latency of each SGX operation on our local cluster’s CPU with SGX Enabled BIOS support, and injected it to the simulation. We observe that public key operations are expensive: signing and signature verification take roughly  $450\mu s$  and  $844\mu s$ , respectively. Context switching and symmetric key operations take less than  $5\mu s$ . Remote attestation, which involves access to the IAS, takes about  $250ms$  on average. Unless otherwise stated, the results reported in the following are averaged over 20 independent runs.

## 6.1 AMES Overhead

We evaluate overhead incurred by AMES over vanilla SGX execution using five benchmarks (i.e., *mcf*, *deepsjeng*, *leela*, *exchange2*, and *xz*) selected from SPEC CPU2017 [6] on our local cluster. We leverage Intel SGX SDK [2] to port the benchmarks into enclave execution, enveloping their operational logic in a worker thread. We then inject a control thread that is responsible to implement the membership service to the enclave using our AMES library.

Figure 4 reports the running time of each AMES-enabled benchmarks under different range of heartbeat timeouts normalized against their own vanilla SGX execution’s running time. For a given range  $[a, b]$ , a heartbeat timeout is chosen uniformly at random between  $a$  and  $b$ , inclusively. The overhead ranges from 5% to 16%, with smaller timeout range leads to higher overhead. This is so because smaller

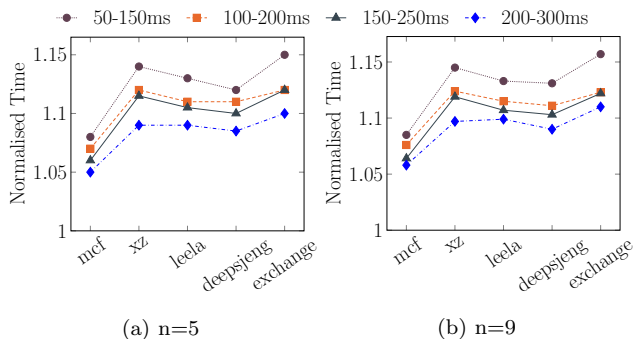


Figure 4: AMES’s overhead compared to vanilla SGX execution. The running time of each benchmark is normalized against its own vanilla SGX execution’s.

heartbeat timeouts result in higher frequency of the leader contacting the follower processes to maintain its authority, which in turns leads to higher communication overhead and control switching. We also observe that the overhead is more evident when benchmark workload is computation intensive and requires few control switching for I/O (e.g., for the heartbeat timeout range of  $[50 - 100](ms)$ , *mcf* witnesses an overhead of 8%, whereas *exchange* sports a 15% overhead. This makes sense because our implementation piggybacks AMES control thread’s messages on the worker thread’s interface with the untrusted application (e.g., its `ECall` and `OCall`) whenever necessary. Thus, the more interactions the worker thread has with the untrusted application over the course of the execution, the fewer control switching the process has to trigger exclusively for exchanging the control thread’s messages, and hence lower the overhead. Comparing across figure 4a and figure 4b, we can see that while the overhead increases when there are more processes in the system, the difference is insignificant. In particular, when  $n$  increases from 5 to 9, the difference in the normalized running time is less than 1% for all benchmarks in our experiment.

## 6.2 Leader Election

In the second set of experiments, we study the effect of timeout configurations on AMES leader election. We measure the leader election latency as a period from when we crash the leader of a cluster of  $n$  servers (with  $n$  ranges from 5 to 33) to the moment when the new leader announces its authority. We experiment with different range of heartbeat timeouts and candidature timeout. For a given range  $[a, b]$ , a heartbeat timeout is chosen uniformly at random between  $a$  and  $b$ , inclusively, while the candidature timeout is set to  $5 \times b$ .

Figure 5 reports the latency under difference timeout settings on our local cluster and GCP. On our local cluster, the smaller the timeout range, the lower the leader election latency is. However, this phenomenon is not observed in our experiments on the GCP. More specifically, with the timeout range of  $[50 - 100](ms)$ , the leader election latency on GCP ranges from  $405 - 552ms$ , while the timeout range of  $[200 - 300](ms)$  yields the lower leader election latency ranging from  $325 - 458ms$ .

A careful analysis of the difference between our local cluster and GCP confirms the necessity of proper timeout configurations (see Section 4.8) with respect to communication

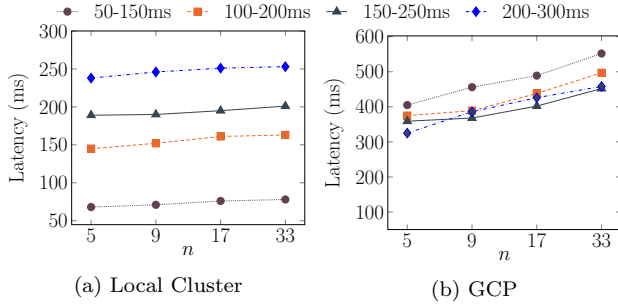


Figure 5: Leader election latency.

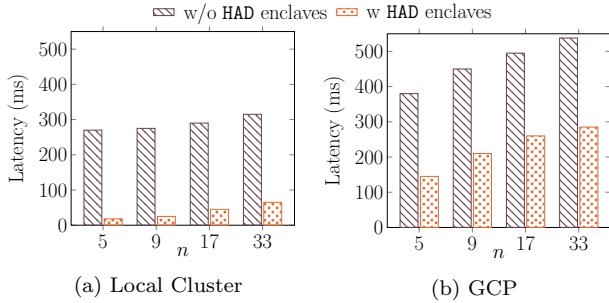


Figure 6: NEWPROCESS's latency.

latency (or `broadcastTime`). Recall that on our local cluster, the latency between any two servers is roughly  $0.13ms$ , whereas that of GCP ranges from  $24ms$  to  $66ms$ , depends on the geolocation of the servers. When the heartbeat and candidature timeout is too small compared to the communication latency, the candidates face difficulty in collecting sufficient responses to its `RequestVote` in time, entering the new terms and triggering new leader election prematurely. This leads to unnecessary leader elections, and exacerbates the system availability.

### 6.3 NEWPROCESS Protocol

We now examine the cost of the NEWPROCESS protocol with and without `HostAttestationDelegate` (HAD) enclave. Recall that by incorporating HAD enclave, the leader and the new process can engage in a remote attestation without contacting the IAS.

Figure 6 depicts the latency of NEWPROCESS protocol with respect to different number of existing processes in the system at the beginning of the new process's commission. The latency of NEWPROCESS protocol is measured as the duration between the new process  $p_j$  signaling its join request till  $p_j$  receiving the application secret and becoming fully functional.

On our local cluster (Figure 6a), NEWPROCESS without HAD enclaves takes from  $270 - 325ms$ , whereas incorporating HAD enclaves reduces such a cost to  $18 - 65ms$ . This difference is primarily due to the cost of IAS access, which is approximately  $250ms$  in our experiments. A similar result is observed on GCP experiments (Figure 6b), wherein incorporating HAD enclaves significantly reduces the cost of the NEWPROCESS protocol. Compared to the experiments on our local cluster, the NEWPROCESS latency observed on GCP is higher (e.g.,  $145 - 285ms$  compared to  $18 - 65ms$  with HAD enclaves). This is so because the cost of the underlying consensus protocol itself, especially the broadcast time during the consensus run, differ significantly between




Figure 7: EXPELPROCESS's latency.

the two settings.

### 6.4 EXPELPROCESS Protocol

Finally, we examine the cost of the EXPELPROCESS protocol by crashing a process  $p_h$ , and measure the latency from such the crash till the time when  $p_h$  is expelled from the leader's `peerList`. On our local cluster, the heartbeat timeout is chosen randomly in the range of  $[25 - 50]ms$ , whereas the corresponding range on GCP is  $[200 - 300]ms$ . Figure 7 plots the latency of the EXPELPROCESS protocol observed on our local cluster and GCP, with respect to different  $n$ .

Similar to NEWPROCESS protocol, the latency observed on GCP is significantly higher than that observed in our local cluster. We attribute this to the difference in the timeout configurations, which in turn is influenced by the broadcast time of each setting.

## 7. RELATED WORK

**Membership service.** Birman et al. [24] formulated the problem of group membership in asynchronous environments. The authors acknowledge that the system's liveness entails the possibility of erroneous failure detection, and the need of suppressing communication from a process that is unjustly excommunicated. Our autonomous membership service follows this observation, fulfilling the excommunication by having an unjustly excommunicated process committing suicide. Amir et al. [8] proposed a membership protocol that relies on atomic broadcast, whereas Mishra et al. [22] leveraged a multicast facility that preserves the partial order of messages exchanged among the communicating processes to construct a membership protocol. Our approach, on the other hand, does not make any assumption on the underlying broadcast/multicast primitives. Instead, it builds on a well-known consensus protocol, namely Raft [23], to implement the membership service.

**Live migration and replications of SGX enclaves.** Gu et al. [17] present technique for live enclave migration on the cloud. In contrary to our membership service, their system model allows only one instance of the application to be functional at any time. ReplicaTEE [25] attempts to enable seamless replication of SGX enclaves by relying on a distributed directory service layer that handles provisioning of enclave processes, and keeps track of the number of existing processes. Nevertheless, it is unclear how the directory service's replicas are protected against forking attack, wherein the adversary splits the directory service replicas into two cliques, each of which is fully operational but unaware of the other. This attack effectively allows the adversary to replicate more enclave processes than what would otherwise be allowed by the SLA.

## 8. CONCLUSION

We have described an autonomous membership service for enclave-based applications. While the application bootstrap necessitates the involvement of the application owner, all subsequent process commission and decommission are administered the existing and active processes of the application. The proposed membership service necessarily admits unjust excommunication of a non-faulty process from the membership group, which is then fulfilled by the unjustly excommunicated process committing suicide. The experimental study shows that AMES incurs 5% – 16% overhead compared to vanilla enclave execution.

## 9. REFERENCES

- [1] Ensuring only a single instance of enclave. <https://software.intel.com/en-us/forums/intel-software-guard-extensions-intel-sgx/topic/709552>.
- [2] Intel SGX SDK for Linux. <https://github.com/01org/linux-sgx>.
- [3] Intel skylake processors. <https://ark.intel.com/products/codename/37572/Skylake>.
- [4] Intel software guard extensions: Intel attestation service api. <https://software.intel.com/sites/default/files/managed/7e/3b/ias-api-spec.pdf>.
- [5] Rethinkdb. <https://github.com/rethinkdb/rethinkdb>.
- [6] SPEC CPU2017 Benchmarks. <https://www.spec.org/cpu2017/Docs/overview.html>.
- [7] The Coco Framework. <http://aka.ms/cocopaper>.
- [8] Yair Amir, Danny Dolev, Shlomo Kramer, and Dalia Malki. Membership algorithms for multicast communication groups. In *International Workshop on Distributed Algorithms*, pages 292–312. Springer, 1992.
- [9] Ittai Anati, Shay Gueron, Simon Johnson, and Vincent Scarlata. Innovative technology for cpu based attestation and sealing. In *Proceedings of the 2nd international workshop on hardware and architectural support for security and privacy*, volume 13. ACM New York, NY, USA, 2013.
- [10] Andrew Baumann, Marcus Peinado, and Galen Hunt. Shielding applications from an untrusted cloud with haven. In *OSDI*, 2014.
- [11] Philip A Bernstein, Vassos Hadzilacos, and Nathan Goodman. Concurrency control and recovery in database systems. 1987.
- [12] Marcus Brandenburger, Christian Cachin, Matthias Lorenz, and Rüdiger Kapitza. Rollback and forking detection for trusted execution environments using lightweight collective memory. In *Dependable Systems and Networks (DSN), 2017 47th Annual IEEE/IFIP International Conference on*, pages 157–168. IEEE, 2017.
- [13] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
- [14] Hung Dang and Ee-Chien Chang. Privacy-preserving data deduplication on trusted processors. In *Cloud Computing (CLOUD), 2017 IEEE 10th International Conference on*, pages 66–73. IEEE, 2017.
- [15] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. Technical report, MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR COMPUTER SCIENCE, 1982.
- [16] Cary Gray and David Cheriton. *Leases: An efficient fault-tolerant mechanism for distributed file cache consistency*, volume 23. ACM, 1989.
- [17] Jinyu Gu, Zhichao Hua, Yubin Xia, Haibo Chen, Binyu Zang, Haibing Guan, and Jinming Li. Secure live migration of sgx enclaves on untrusted cloud. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 225–236. IEEE, 2017.
- [18] Simon Johnson, Vinnie Scarlata, Carlos Rozas, Ernie Brickell, and Frank Mckeen. Intel software guard extensions: Epid provisioning and attestation services. *White Paper*, 1:1–10, 2016.
- [19] Leslie Lamport et al. Paxos made simple. *ACM Sigact News*, 32(4):18–25, 2001.
- [20] Sinisa Matetic, Mansoor Ahmed, Kari Kostiaainen, Aritra Dhar, David Sommer, Arthur Gervais, Ari Juels, and Srdjan Capkun. Rote: Rollback protection for trusted execution. *IACR Cryptology ePrint Archive*, 2017:48, 2017.
- [21] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R Savagaonkar. Innovative instructions and software model for isolated execution. *HASP@ ISCA*, 10, 2013.
- [22] Shivakant Mishra, Larry L Peterson, and Richard D Schlichting. A membership protocol based on partial order. In *Dependable Computing for Critical Applications 2*, pages 309–331. Springer, 1992.
- [23] Diego Ongaro and John K Ousterhout. In search of an understandable consensus algorithm. In *USENIX Annual Technical Conference*, pages 305–319, 2014.
- [24] Aleta M Ricciardi and Kenneth P Birman. Process membership in asynchronous environments. Technical report, Cornell University, 1993.
- [25] Claudio Soriente, Ghassan Karame, Wenting Li, and Sergey Fedorov. Replicatee: Enabling seamless replication of sgx enclaves in the cloud. *arXiv preprint arXiv:1809.05027*, 2018.
- [26] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 640–656. IEEE, 2015.