Trustworthy confidential virtual machines for the masses

Anna Galanou*
TU Dresden
Germany
anna.galanou@tu-dresden.de

Khushboo Bindlish
DFINITY Foundation
Switzerland
khushboo.bindlish@dfinity.org

Luca Preibsch
Friedrich-Alexander-Universität
Erlangen-Nürnberg
Germany
luca.preibsch@fau.de

Yvonne-Anne Pignolet
DFINITY Foundation
Switzerland
yvonneanne@dfinity.org

Christof Fetzer
TU Dresden
Germany
christof.fetzer@tu-dresden.de

Rüdiger Kapitza
DFINITY Foundation
Switzerland
Friedrich-Alexander-Universität
Erlangen-Nürnberg
Germany
ruediger.kapitza@fau.de

ABSTRACT

Confidential computing alleviates the concerns of distrustful customers by removing the cloud provider from their trusted computing base and resolves their disincentive to migrate their workloads to the cloud. This is facilitated by new hardware extensions, like AMD's SEV Secure Nested Paging (SEV-SNP), which can run a whole virtual machine with confidentiality and integrity protection against a potentially malicious hypervisor owned by an untrusted cloud provider. However, the assurance of such protection to either the service providers deploying sensitive workloads or the end-users passing sensitive data to services requires sending proof to the interested parties. Service providers can retrieve such proof by performing remote attestation while end-users have typically no means to acquire this proof or validate its correctness and therefore have to rely on the trustworthiness of the service providers.

In this paper, we present *Revelio*, an approach that features two main contributions: i) it allows confidential virtual machine (VM)-based workloads to be designed and deployed in a way that disallows any tampering even by the service providers and ii) it empowers users to easily validate their integrity. In particular, we focus on web-facing workloads, protect them leveraging SEV-SNP, and enable end-users to remotely attest them seamlessly each time a new web session is established. To highlight the benefits of *Revelio*, we discuss how a standalone stateful VM that hosts an open-source collaboration office suite can be secured and present a replicated protocol proxy that enables commodity users to securely access the Internet Computer, a decentralized blockchain infrastructure.

CCS CONCEPTS

- **Security and privacy** → *Virtualization and security*; **Distributed systems security**; **Web application security**.

---

*Research was conducted while working at DFINITY Foundation

KEYWORDS

Confidential Computing, TEEs, AMD SEV-SNP, Attestation, TLS

1 INTRODUCTION

Responding to cloud customers' concerns over the security of their workloads, cloud providers are currently rolling out offerings with novel security extensions enabled, facilitating the use of hardware-isolated compartments, called *trusted execution environments (TEEs)*. By migrating their security-sensitive workloads inside TEEs, customers are ensured that their data remain shielded from any unauthorised access of privileged system software layers, such as the hypervisor, and the cloud-managing personnel. In recent years, VM-based TEEs, where entire VMs can be isolated, have gained a lot of traction from the industry because of their seamless deployability, leveraging technologies like AMD's SEV-SNP [2], Intel's Trust Domain Extensions (TDX) [23] and ARM's Confidential Compute Architecture (CCA) [5].

Although confidential computing is considered a big step forward in terms of cloud security, its benefits cannot be reaped to a maximum degree without additional effort. In particular, end-users of VM-based TEEs are unaware if a service they access is secured via trusted execution or served by a commodity environment. They could be assured of these guarantees via remote attestation, which essentially provides proof for the authenticity of the underlying hardware as well as for the integrity of the software loaded inside the TEE. However, to our knowledge, there is currently no standardised, widely established approach to provide remote attestation as a service to end-users besides the option to expose directly the hardware manufacturer-provided APIs that are typically accessible only to the VM owner/service provider.

When it comes to verifying the integrity of the TEE state, the end-users need to be aware of what constitutes an acceptable state rather than just depending on a Certificate Authority (CA) to validate the TEE authenticity. The TEE state is reflected on the attestation report, which includes a cryptographic hash over the VM's initial memory context

right before its launch. This hash is usually compared against a pre-computed expected value and upon a successful match, the validating party can be reassured that the VM has been initialised in a known good state. This is sufficient for service owners who intend to verify that a VM runs within a TEE as expected and they have full control and knowledge over the provided VM's image.

However, for end-users the situation is different and they are largely in the dark regarding the implementation and configuration of the services they access. Besides that, even if they had access to the VM's source code and therefore be in the position to validate the hash of the VM's initial state against an expected value, they would inherently have to trust the service provider not to tamper with the VM after it has been measured, via management APIs (e.g., ssh) or directly boot it with a malicious kernel for instance. The latter scenario is a real threat since the VM's initial state typically only includes the virtual firmware used to boot the machine and nothing more. Therefore, the measurement included in the attestation report is not indicative of either the operating system or the root filesystem's state that the VM booted with. Consequently, by solely having access to a remote attestation report, which contains an initial VM context, end-users cannot really verify if their data is securely handled.

The implicit trust in the service provider cannot be avoided even if end-to-end encryption is in place, as is in the case of email services (e.g. ProtonMail), collaboration suites (e.g. CryptPad), cloud storage services (e.g. Tresorit), search engines (e.g. DuckDuckGo) etc. Even though this technique provides a high level of protection for user data and their privacy, such guarantees are ensured as long as the endpoints are not compromised. If the server-side software, for instance, has vulnerabilities, an attacker can manage to execute code remotely and retrieve login credentials, or any other sensitive information. Therefore, verifying the state of such applications and enlightening the end-users about them is of significant value.

In this paper, we present *Revelio*; an approach that enables end-users to easily, systematically, and deterministically validate a VM and its hosted services, so that they no longer need to entrust service providers with the integrity and confidentiality of their data. Service providers manage the VM and its services, but are averted from either tampering with an attested VM or manipulating it during booting. In particular, we describe how to reduce their privileges to a level, where they can only perform denial of service to the VM but cannot access users' data or tamper with it. The end-users on the other hand are enabled to retrieve a complete fingerprint of the VM's initial state and attest it. Since only a small fraction of the users will have the technical skills or the time to determine if a VM and its offered services are properly implemented and securely configured, we discuss how the assessment can be delegated to a third party. This allows less technically savvy users to benefit from knowledgeable users or institutions that can perform the validation of the system on their behalf. We detail the steps on how this can be achieved,

especially for web-based services, and show how commodity users who are entirely unaware of trusted execution implicitly profit from our system.

Revelio's contributions can be summarised as follows:

- We present an approach that enables seamless and automated remote attestation when accessing a *Revelio VM*-hosted web-facing service via a web extension. Depending on the user and the service provider, the validation can be performed by a third party, such as an auditing company, a community using a blockchain infrastructure, or the end-users themselves.
- We introduce and implement a systematic way to disallow service providers from accessing or even altering users' data. To achieve that, we securely configure the VM network-wise, render its root filesystem, containing the services of interest, immutable and extend its initial cryptographic fingerprint to cover the necessary state so that end-users receive a representative report of the attested VM.
- We highlight how Revelio can address a wide variety of use cases, as well as elaborate on how it can protect a stateful open-source collaboration office suite and a protocol translation proxy that enables access to the Internet Computer, a decentralised blockchain infrastructure.

We implemented Revelio leveraging AMD's SEV-SNP [2], however, upcoming VM-based TEEs, such as TDX [23] and ARM's CCA [5] can also be alternatives for our approach. Therefore, Revelio can be deployed in a hardware-agnostic fashion, as long the TEE follows the VM model.

2 BACKGROUND

This section provides background information on the AMD SEV-SNP technology, a boot method of the confidential VMs, as well as the basic workflow of Secure Socket Layer (SSL) certificate generation for web servers. While the first two build the essential basis for Revelio, the latter part is needed to be adapted so that the control of service providers over the Revelio VMs can be reduced.

2.1 AMD Secure Encrypted Virtualization (SEV) Technology

Revelio is relying on AMD SEV-SNP technology to guarantee the confidentiality and integrity of VM's memory and hence of the services running on it. SEV uses a unique memory encryption key for each VM and tags it with an address space identifier, preventing cross-TEE attacks and unauthorised usage inside the processor. The encryption keys are generated by a firmware running on a dedicated processor called, AMD Secure Processor (AMD-SP). The encryption is transparent to the hypervisor, performed inside dedicated hardware in the on-die memory controller that encrypts data when it is written to DRAM and decrypts it when read.

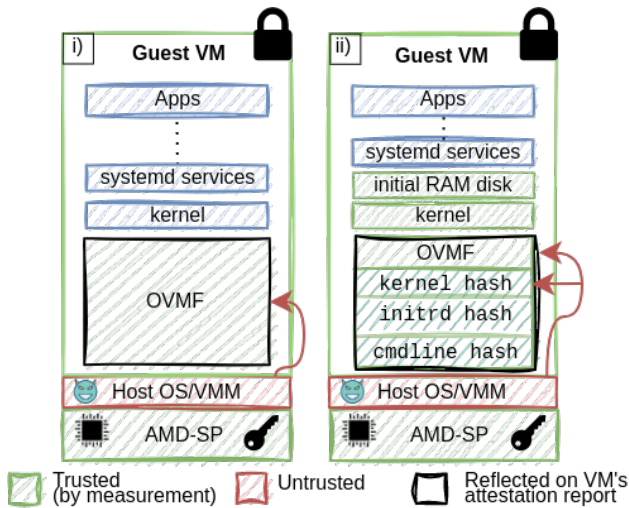*2.1.1 Remote Attestation on SEV-SNP.* Besides the VM protection, SEV-SNP provides proof to the user about the correct

Figure 1: Direct boot (i) vs measured direct boot (ii).

VM deployment and the authenticity of the SEV hardware by performing remote attestation. The provided proof is called an attestation report and holds several pieces of information, among which is the VM's measurement; a cryptographic hash of the components that form the initial VM's state and is being taken by the trusted firmware running on AMD-SP. It also contains the unique identifier of the processor, called Chip ID, as well as the TCB version, which refers to the version number of the SNP firmware. To allow the authentication of an SEV platform and enable the remote attestation protocol, SEV signs the attestation report with the Versioned Chip Endorsement Key (VCEK), which is unique to the local platform's processor and takes into consideration the TCB version. While the data contained in the attestation report are fixed during VM's lifetime, there is a protected path between the AMD-SP and the VM so that arbitrary data can be added to the AMD-SP-signed attestation report via the field `REPORT_DATA` and be cryptographically linked to it.

2.1.2 Measured direct boot. After the VM has finished loading, AMD-SP finalises its hash measurement, which typically corresponds only to the first virtual firmware volume and nothing else. To implement Revelio, an approach is required that imprints the exact initial state including any persistent state, used to bring up the VM, in the attestation report. This is achieved by an approach called *measured direct boot* [30], implemented as a series of patches to the hypervisor, QEMU, and to the virtual firmware, OVMF. The modified OVMF package creates space in the firmware binary (Fig. 1) so that it can store a table with the hashes of the kernel, initial RAM disk (initrd) and the kernel command line. When QEMU boots a VM, it hashes each of these blobs and injects the hashes into the special table. These components are passed by the hypervisor as arguments via the `fw_config` interface. OVMF measures each component upon reception and compares the measurement to the ones placed into the designated

table. In case any of the measurements do not match, the boot fails. The hypervisor is untrusted, but the firmware is measured and reflected on the attestation report, hence the hashes injected by the hypervisor can be verified by anyone performing remote attestation and having access to all sources of the VM.

2.1.3 Securely saving persistent data. Besides remote attestation, SEV-SNP provides the possibility to generate additional key material for data sealing that can be derived from different sources depending on their use case. In the most basic case, a VM can request a sealing key derived from VM's measurement, which essentially makes it accessible only by a VM with an identical cryptographic fingerprint. The sealing key is exchanged via a trusted path between the AMD-SP and the VM and stays protected from any component outside of the TCB.

2.2 Certificate generation and Certificate Authorities

One of Revelio's key contributions is establishing trust to web-facing services with which end-users interact. End-users typically have no evidence to assess the security guarantees services promise to offer. Nowadays, the vast majority of these services use the hypertext transfer protocol secure (HTTPS); a secure version of HTTP, the primary protocol used to send data between a web browser and a website. HTTPS uses an encryption protocol to encrypt communications, called Transport Layer Security (TLS), the successor protocol of Secure Socket Layer (SSL). Coupled with a Public Key Infrastructure (PKI), SSL/TLS provides verifiable identities via certificate chains and private communication. Revelio leverages this to provide a binding between the web interface of the service and its cryptographic fingerprint imprinted on the attestation report.

When a browser connects to a SSL-secured site it will first retrieve its SSL certificate and check that it has not expired, it has been issued by a trusted CA and that it is being used by the intended website. An SSL certificate is essentially a signed attestation binding a subject to a public key and can be issued by CAs, who in turn have their own certificates terminating at a small set of self-signed root certificates. In order to obtain and install a certificate, a server operator is used to manually prove control of the domain name and complete a payment transaction to a CA. Nowadays, CAs like Let's Encrypt [1], offer domain-validated certificates through a standard protocol at no cost to server operators.

Let's Encrypt is the first browser-trusted CA that established a standard protocol, called Automatic Certificate Management Environment (ACME), to automate identity validation and certificate issuance via clients, like certbot [39], without intervention from web server operators or CA staff. To obtain a certificate for the domain, the user can pass to the client a Certificate Signing Request (CSR) [21] and ask Let's Encrypt CA to issue a certificate with a specified public key. The CSR includes a signature by the private key corresponding to the public key in the CSR, among other

information like the domain, organisation, country name etc. After the Let's Encrypt CA has received the request, it verifies the signature and upon successful validation, it issues a certificate for the requested domain with the public key from the CSR and returns it to the client.

3 REVELIO'S DESIGN

In this section, we give more details about the scope of our work, the motivation behind it, as well as the threat model we adopt and the entities involved in it. Besides that, we are going to describe Revelio's design by outlining our key requirements and elaborating more on the components of its architecture.

3.1 Objective

TEEs provide hardware isolation that can either happen on a process-level boundary or on a VM-level one. On the former one, the operating system is treated as malicious which means that any syscall interrupt coming from the isolated code will immediately exit the TEE context. Consequently, that means that all the applications need retrofitting to work inside the TEEs which hinders the widespread adoption. To that end, the industry is gradually moving towards VM-based TEEs, such as SEV-SNP [2], where the kernel is part of the isolated code/VM and therefore the applications can seamlessly work without any modifications.

With confidential VMs making inroads to the public cloud, we also opt for it as Revelio's foundation, so that we can render our system easily deployable by real-world web-facing services. For these kinds of services, the protection of their data-in-use leveraging hardware-based trusted execution is being decided by the service providers and typically communicated to the end-users to alleviate their security concerns. However, the end-users themselves are largely unaware of the environment such provided services are running in and if the security measures that have been promised are actually in place. By virtue of this inherent lack of control, they have to implicitly place their trust in the service provider for the installed software and the utilised cloud infrastructure without receiving any evidence for it.

Assuming that such evidence, i.e. TEE attestation report, was available to the end-users, they would have to be in a position to verify that an authentic Hardware Root of Trust (HRoT) generated it and that the integrity state of the utilised service matches the expected one. The latter requires the evidence to contain information that reflects the service's state and from the users to have knowledge about what to accept as "correct" and "incorrect" state.

Considering that the whole foundation of the confidential computing paradigm lies in verifying before establishing trust, it is of utmost significance that the end-users have the power to do so. However, currently, they can only take the security assurances from the service providers at face value and have no proof for those. In our work, we address this problem of lack of proof as well as the challenges that come with exposing it to the end-users. More specifically, Revelio

(1) exposes attestation evidence to the end-users in a practical manner,
(2) imprints the cryptographic fingerprint of the utilised service to this evidence,
(3) applies measures to preserve the integrity of the service during runtime,
(4) and enables end-users to verify the provided evidence.

3.2 Threat model

Before describing how Revelio tackles these challenges and fulfils our goals, it's necessary to elaborate more on the threat model we assume during each phase of the VM's lifetime, namely the provisioning, occupancy and decommissioning phase. The main stakeholders involved in those are the:

- Cloud providers; they manage the cloud infrastructure in terms of hardware and the associated software (host OS, hypervisor, firmware), offering VMs to customers as Infrastructure as a Service (IaaS).
- Service providers; they take the role of a cloud customer renting VMs to provide services to end-users. They are in control of the VM's virtual firmware, OS and software.
- End-users; they use the services, as they please, and are largely unaware of technical details of the services' runtime system, like if it is cloud-based or what kind of hardware and software is being used. They use a web browser with Revelio's web extension installed.

Provisioning phase. Before the confidential VM's deployment, its image needs to be built, configured and customised by the service provider so that it has the necessary software packages installed, as well as the network settings and security measures applied. However, the end-users assume that the VM image may contain vulnerabilities to introduce backdoors or other malicious code planted by the service provider so that they can leak their data. Besides the VM image itself, the host platform owned by the cloud provider has to be provisioned as well, so that it has the necessary firmware, kernel and hypervisor to support the TEE. The only component that is considered trusted on the host platform by the service provider and the end-users is the CPU hardware along with the AMD Secure Processor implementation. The virtual firmware, kernel, initrd are assumed to be provided by the service provider to the cloud provider so that they can deploy the confidential VM later with the TEE-enlightened hypervisor. Neither of those are considered trusted by the end-user since they can contain bugs either planted by the cloud provider to enable a privilege escalation from the host VMM or the service provider with the intention of retrieving sensitive data from the service.

Occupancy phase. In line with the confidential computing threat model, we assume that the cloud provider or a host platform intruder has full control over everything except the TEE itself. That basically includes the physical hardware as well as the entire software stack on the host platform, i.e. the OS and the hypervisor, which allows the adversary to perform

a wide range of actions like modification of any file on the system, and monitoring of system activity including network traffic, system logs, user activity etc. We also consider man-in-the-middle attacks where an attacker can intercept network traffic between two parties, spoof or corrupt it, and redirect traffic to a different destination. The cloud provider may also migrate the VM to a different physical host without the knowledge of the service provider exposing it in this way to security risks. In our threat model, we do not target denial-of-service attacks, so a cloud provider's attempt to starve the VM of resources, such as CPU or memory, in an effort to disrupt its operation is out of scope. Ciphertext side-channel attacks on the encrypted VM (by building a dictionary of plaintext-ciphertext pairs) [28] are out of scope. The cloud providers explicitly do not trust the service providers, as they may attempt to get unrestricted access to the cloud resources or tamper with the code and data of other tenants. The end-users assume that the service provider may launch a malicious service in order to retrieve their sensitive data, or/and modify the environment to serve this purpose.

Decommissioning phase. Ensuring the confidentiality of tenants and their secrets also after the node release is essential, so we address any attacks that can be launched by subsequent software running on the node aiming to retrieve any remaining information from the persistent storage.

3.3 Requirements for Revelio

Our main objective is to remove the cloud provider from the trust domain of the end-users that use confidential web-facing services as well as reveal if those have potentially been misconfigured by the service providers, narrowing the trust computing base even further. On that account, Revelio offers end-users cryptographic proof of the state of the utilised service as well as of the confidential VM it is running on top of. To achieve that goal, our system needs to fulfil the following requirements pertaining to functionality (F1-6) as well as deployability (D1-3).

F1: Evidence disclosure Before sending any sensitive information to the service, the end-users need to be able to receive cryptographic evidence about the service's state so they can attest it and assess its trustworthiness.

F2: Trust only by measurement for TCB The evidence offered to the end-users needs to reflect the VM's volatile as well as non-volatile memory (i.e. any block storage device mounted at startup) so that they can have a representative view of the whole TEE that was loaded and not implicitly place their trust on its TCB. This evidence should be cryptographically bound to the Hardware Root of Trust (HRoT).

F3: Confidential VM & service linkage End-users need to make sure that the evidence they receive corresponds to the confidential VM the service is actually running on. Therefore, the confidential service's identity should be cryptographically linked to the hardware-based identity of the VM.

F4: Integrity during runtime Aiming to avert any misconfiguration of the service by a malicious intruder or malevolent service provider, Revelio should render it infeasible for them

to modify the service as well as the runtime environment after booting, namely the network settings, ssh connectivity, approved ports etc.

F5: Reproducibility and deterministic verifiability Build verifiability is an important safety property for software releases [36], as external attesters can audit and verify if software is susceptible to various security problems introduced during the build process (e.g., surveillance malware, compromised cryptographic signatures, supply chain attacks, and untrusted dependencies). By providing a reproducibly built image for the VM, we intend to automatically accommodate a practical and efficient attestation by the interested parties as well.

F6: Persistent state protection Considering that the service's persistent state may contain sensitive users' data, Revelio should aim to protect the storage devices used by the confidential VM against any offline attacks performed by anyone outside of the TEE's context.

D1: Ease of use Revelio needs to enable end-users to perform remote attestation seamlessly as an integral part of the service utilisation. An intuitive way of presenting the attestation results and indicating possible security violations should also be provided by the system.

D2: Flexibility over the verification By using Revelio, end-users should be enabled to perform an extensive verification of the received evidence. However, considering the complexity of this task and the time investment it typically requires even for the technically-skilled users, a flexible approach needs to be provided so that they can either delegate the validation to third parties or perform it themselves.

D3: Scalability Taking into account that Revelio's use case scenario concerns web-facing services that need to be available at all times and serve a significant load of users' requests, Revelio should take measures to support scalability.

3.4 Revelio in a nutshell

To address the requirements set for our system, Revelio has to be involved in the building/provisioning as well as deployment phase of the confidential service which takes place at the service and cloud provider's premises respectively. In the following sections, we present the core parts of our design as well as the requirements each one of them aims to satisfy.

3.4.1 Reproducible build as a basis for practical and efficient remote attestation. Typically, in dynamic systems the system's state frequently changes due to system-centric events such as software package installations, dependencies in libraries, configuration of the build environments, timestamps, file permissions etc. This means that there may be multiple acceptable hashes for the VM depending on the building environment, which makes the reconstruction of the expected hash and consequently the VM's attestation particularly cumbersome in the best case, and in the worst case completely unreliable. By removing any non-determinism in the build process, Revelio guarantees identical built binaries and images for every invocation of the build process, so that the user can reproduce the expected measurement in a systematic and practical way (F5).

3.4.2   Robust configuration to ensure system's integrity during runtime. Considering that we do not assume any runtime monitoring system running on the confidential VM, there is no way to ensure that the VM will stay intact after booting. One of Revelio's requirements (F4) is to take the necessary measures to guarantee integrity during runtime. We address this problem in two ways; first by blocking any inward connections to the confidential VM and second by enforcing a read-only root filesystem and rendering it integrity protected. The first ensures us that no outside attacker or authorised personnel can access the VM after booting and therefore they cannot spawn a malicious process or tamper with the existing ones during runtime. The second one guarantees the integrity of the volume containing the root filesystem (including kernel, services etc.) against corruption by malware or persistent rootkits that hold onto root privileges and compromise devices. Both of those measures are being applied during build time so that the generated VM image has already the correct network configuration and the system integrity is protected.

3.4.3   Expanding coverage of the reported state. Based on the current state of the upstream versions of the hypervisor and virtual firmware that support the confidential VM loading, only the initial state of the VM is subject to the measuring performed by the trusted hardware, which essentially contains only the firmware. To expand what is covered by the reported state in the attestation evidence sent to the users (F2), we deploy the method of measured direct boot (2.1.2). This allows us to imprint in the VM's final measurement, not only the virtual firmware's hash but the ones corresponding to kernel, initrd and kernel's command line respectively. To extend the trust chain even further so that it includes the cryptographic state of the root filesystem, we pass as kernel command line argument the root hash of the block device's integrity metadata. This has already been generated during build time, when the image of the rootfs is being constructed. The code enforcing the integrity protection for the rootfs is part of the initrd and the kernel, which are both measured by the hardware in our system.

3.4.4   Enabling end-users to attest before use. End-users should be able to attest the service before interacting with it and passing their sensitive data, like account credentials, financial details etc. Even if TLS is in place and assumed to provide authentication, integrity and privacy for the data transmitted across the untrusted channel, those are no longer guaranteed if the server endpoint is subverted by a potentially malicious service provider. Therefore, by receiving actual proof that the server is in a good state end-users can benefit from the minimised TCB guaranteed by the confidential computing paradigm. To serve both of the requirements F1 and D1, Revelio offers the attestation evidence via a web extension, so that end-users can perform the attestation using the same web browser interface.

3.4.5   Binding web-service's identity to the TEE. An SSL certificate is essentially the digital identity of the web-facing service end-users interact with. To verify this identity they rely on the CAs that have provided the root and intermediate certificates. This root of trust is independent of the HRoT corresponding to the attestation evidence and therefore the end-users cannot originate the received attestation evidence from the confidential VM that generated them. Revelio addresses this problem (F3) by cryptographically binding the TLS identity of the service to the TEE that it is running on, thereby enabling the end-users to verify that the proof they receive about the service is not only authentic but represents the actual service they interact with.

3.4.6   TLS key sharing to serve scalability. Since our system is meant to be deployed in decentralised secure services for which scalability (D3) is crucial, rate limiting on the SSL certificate creation is a critical problem that we need to take into account. More specifically, certificate authorities like Let's Encrypt, consider a rate limit [24] on the SSL certificates generated in a certain period to ensure fair usage among the users. To address this problem, we assume that all the Revelio-VMs that host instance of a service share the same SSL certificate. The account credentials for the certificate generation reside in a platform belonging to the service provider's infrastructure and they are isolated from the public cloud. This machine will be the one performing the DNS challenges for proving ownership of the service domain and distribute the generated certificate among the nodes after attesting them.

The SSL certificate will be generated based on a key pair of one of the nodes picked by the service provider's machine. This key pair essentially will be the TLS identity corresponding to the service that end-users are going to interact with later on and will be created in the context of one of the confidential VMs the service is running on. The private counterpart of this pair will be later distributed among the rest of the nodes, so that they can serve the user's requests in a secure connection. Before that, the TLS key owner will attest them in a mutual attestation protocol (Section 5) and encrypt its private key with the public one of the other attested node's unique pair.

3.4.7   Tailoring verification of Revelio VMs. In order to make Revelio more deployable in the context of web-facing confidential services that are being used by a wide range of end-users, we have designed it so that the attestation process can be tailored to fit users' trust boundaries (D2). More specifically, after Revelio delivers cryptographic proof to them about the state of the service, they have to compare it against the expected one(s). If they have some technical expertise, they can reconstruct the state of the service on their own premises fetching the corresponding open-source components, compute the final cryptographic measurement and therefore make their own deduction on what state is "good" or not. In case end-users cannot follow this process due to a lack of knowledge for instance, then they can retrieve the "golden" values from a third-party source. This can come either from an auditing company that has been delegated with the task to check the software stack of the confidential VM for bugs and vulnerabilities, or from an on-chain decentralised autonomous

organisation (e.g. Internet Computer's Network Nervous System [37]) where the community votes on the "good" values of the software running on the respective TEEs.

3.4.8 Protecting persistent state. To satisfy requirement F6 and protect any sensitive persistent data between shutdowns of the confidential VM, Revelio leverages the feature of sealing and encrypts the external volumes with the VM's measurement-derived key. In this way, no other VM that has not been in the expected state can decrypt the disk and retrieve any sensitive data belonging to the service, like the private TLS key.

4 USE CASES OF REVELIO

While arbitrary services can be protected by trusted execution, the use of Revelio imposes an additional requirement for transparency. At a high level, a service provider has to make all code and relevant operational configuration data of their built image available to external parties. For VM-based TEEs this goes beyond the actual service code and includes the execution environment itself, like the OS. In the scenario where the end-users' trust boundaries are very narrow, it would be them who will retrieve and verify all of the components comprising the trusted execution context; in a less restricted setting an external authorised party can be given access and perform validation on their behalf.

Although providing access to all of the service's source code and operational configuration data to the public is not a common practice for commodity offerings up until recently, we argue that the shift towards the confidential computing paradigm might render it necessary. In particular, end-users can find Revelio significantly useful in scenarios, where they share sensitive data, such as medical records [10], and in the context of private social exchange. There are also numerous other scenarios where the demand for the service's integrity might be of key interest, like in auction sites, lotteries and any form of e-commerce service. When it comes to the execution environment of the services, like the OS, the firmware etc., the ever-growing trend of open source software [6, 11, 38] already fulfils Revelio's requirement for code availability.

In this section, we discuss two use cases:

- The first example describes Revelio's use to secure a cloud end-to-end encrypted collaboration suite, where data confidentiality is of key interest.
- The second one refers to the need for confidentiality and integrity of the end-user-provided data, shows a greater level of complexity and represents a deployed instance of a protocol translation proxy.

4.1 End-to-end user-encrypted cloud collaboration suite with hardware-based trust

Over the recent years end-to-end encryption has significantly grown in demand as an effective means to prevent the compromise of an agent's private information and is currently used by many messaging protocols, like Matrix [17], where users are ensured that messages can't be spoofed and that only
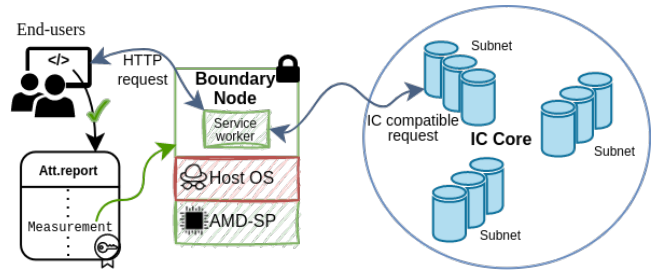

Figure 2: Revelio-protected Boundary Node.

the senders and receivers of them can read the contents. Recent research has shown though that a malicious homeserver can modify the execution environment [3], i.e. deliberately add users to end-to-end encrypted rooms to decrypt future messages sent in the room or add a device in the room they wish to eavesdrop, both of which essentially invalidate the confidentiality guarantees promised by Matrix' threat model.

Collaborative editing platforms, like CryptPad, also use end-to-end encryption and enable users to enforce access control to their data by themselves and not depend on a central authority for that, which usually coincides with the server hosting the content and is a single point of failure. Despite the strong security guarantees that CryptPad offers, it still adopts a more relaxed threat model of an honest but curious cloud server [29] where one still must trust the validity of the Javascript hosted on it or sent to the end-users' browser [13, 14], which can potentially leak sensitive information. We argue that Revelio can be used on that type of applications to fill this security gap and provide the users with the means to verify the trustworthiness of the software installed on the servers as well as preserve their data privacy and integrity.

4.2 Protocol translation proxy: a use case for elevated security

The Internet Computer (IC)[1] [37] is a decentralised platform for the execution of general-purpose decentralized applications in the form of so-called *smart contracts*. A smart contract is a computer program that executes a certain program logic in a decentralized manner. Smart contracts on the IC are called *canisters*; all canisters are hosted on dedicated node machines running the IC protocol. The protocol uses a threshold-signature-based agreement protocol to provide Byzantine fault tolerance. To achieve high scalability the nodes are partitioned in subnets. To securely interact with a canister the IC protocol has to be used, which enables the exchange of threshold-signed messages with the end-user. As expected of a modern Web3 infrastructure, end-users typically access canisters delivering feature-rich web applications via a browser.

While the IC provides the basis for digital democracy and Byzantine fault-tolerance, as well as fits into the conventional Web on a functional level, client-side software and browsers

---

[1]https://internetcomputer.org/

Anna Galanou, Khushboo Bindlish, Luca Preibsch, Yvonne-Anne Pignolet, Christof Fetzer, and Rüdiger Kapitza

have yet to fully adopt the IC protocol. The IC overcomes this issue by utilising the Boundary Nodes that take the role of protocol translation proxies. A Boundary Node is capable of translating an ordinary HTTP request to an IC protocol-compatible message exchange. This can happen either directly by simply receiving the HTTP request of an end-user, transforming it to an IC message and forwarding it to the right canisters or as optionally via a JavaScript-based service worker. This service worker is returned by a Boundary Node on the first request of a user to the IC. Once activated on the browser side, ordinary requests are directly transformed to IC messages inside the browser by the service worker and forwarded to a Boundary Node which in turn delivers them to the IC.

The outlined approach gives any user immediate access to the IC via commodity browsers, however, it also represents a security risk as a malicious Boundary Node could manipulate user requests either directly inside the VM or indirectly by providing a modified service worker. The consequences of such a malicious node could be significant as it compromises the Byzantine fault-tolerance of the IC. Letting Boundary Nodes run inside encrypted VMs facilitates higher security guarantees. By leveraging Revelio, end-users can verify in a practical manner those guarantees, check if the services of the Boundary Node actually run in a hardware-based TEE and more importantly validate the exact software stack as well the configuration used on the node (Fig. 2).

5 IMPLEMENTATION DETAILS OF REVELIO

In this section we will provide implementation details for the main components that Revelio is comprised of, as described in Section 3, and we structure them by the phases of a Revelio VM's lifetime, namely the provisioning of its image, its bootstrapping during the first deployment and its normal operation. For our prototype, we leverage the SEV-SNP hardware for the seamless encryption of the VMs, however, our approach can be based on any VM-based TEE, including Intel's TDX [23] or ARM's Realms [5].

5.1 Image provisioning

Before the VM's deployment in the cloud provider's infrastructure, the service provider needs to build and configure its image accordingly so that it complies with our requirements. This step lays the groundwork for the trust chain establishment and Revelio's bootstrapping later on.

5.1.1 Reproducible build. The Revelio VM's image is essentially an on-disk combination of an SEV-SNP-aware Linux kernel binary, an initrd which contains a root filesystem with all the user-land programs that have been produced in the service provider's CI/CD pipeline as well as their corresponding software dependencies and libraries. Since we assume that the VM is being launched with the method of direct boot, a bootloader is not necessary to exist in the image. To achieve reproducibility, we deploy a deterministic build process which generates repeatedly equivalent images given the same set of source code files, build scripts, and
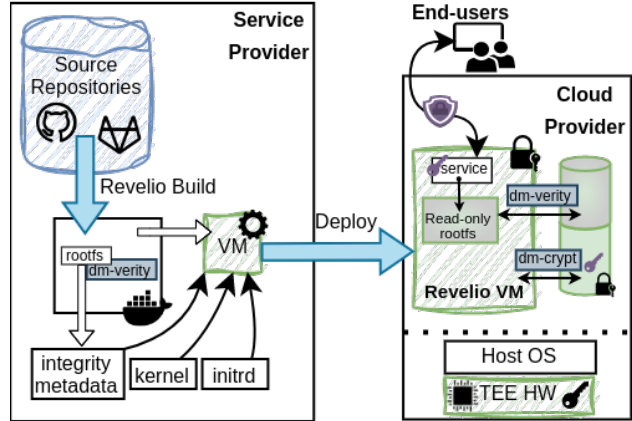

Figure 3: Revelio VM's build and deployment.

build environment leveraging bazel [35] and its hermeticity. Furthermore, the build scripts are modified in order to remediate sources of non-determinism (e.g., timestamps, build paths, file ordering and permissions) by clearing all files that may lead to in-deterministic build (e.g. /var/lib/apt/lists/*, /var/lib/dbus/machine-id/ etc.), squashing all timestamps and specifying a uuid for each partition we create. Besides that, we create two docker images; the first one is meant to pull the needed packages and build the software dependencies while the second is meant to hold the actual binaries that we need which we copy from the former one. In that way, we don't actually pollute the final image with non-deterministic elements.

Besides the non-determinism sources related to the filesystem building, application packages can be another source of such problems since the package versions can change on every invocation of apt-get, dnf etc., thereby leading to different builds. To tackle this problem instead of installing the packages from scratch during every build, we pull a published image instead (i.e., docker image). This image needs to be updated regularly and be publicly available so that it can be retrieved in every build. We also need to ensure that it is trusted and integrity protected, so every time the software dependencies change, the new image is built in a protected environment (i.e. Gitlab runner) and then pushed to the registry. Finally, it should be noted that our system uses as a starting point an official Ubuntu 20.04 LTS Docker image, which is maintained by Canonical and therefore we trust the latter to a certain extent. However the sources and the build processes of such images are public, so they can be subjected to an audit and a review from the community.

5.1.2 Extending the measured envelope. Leveraging direct measured boot[2] [30] allows us to extend the measured envelope from the virtual firmware up to the kernel, initrd and boot arguments. In order to extend this even further

---

[2]https://www.mail-archive.com/qemu-devel@nongnu.org/msg945567.html

and establish a trust chain up to the system's rootfs, we include the root value of the rootfs' integrity metadata to the kernel's command line. The root value along with the rest of metadata of the underlying block device is generated during the build process (Fig. 3), when the rootfs is being constructed, leveraging the dm-verity utility of the Linux kernel. Dm-verity [8] essentially uses a Merkle tree of sha256 hashes computed for the device's blocks and verifies them every time they are being read, ensuring that files have not changed between reboots or during runtime. When dm-verity is being setup for the hardware storage device on which rootfs has been stored, the latter is being remapped through a target driver to produce a new, transparent, virtual storage device. We modify VM's init process so that it mounts the read-only and integrity-protected new virtual storage device corresponding to the rootfs, retrieving the root hash from the kernel command line and the rest of the integrity metadata from the designated partition of our choosing. Since initrd is measured as well in our method, we ensure that if any of those measures are not being applied, it is going to be reflected on the attestation evidence of the VM.

5.1.3 Blocking unauthorized inward connections.
During the build phase, we configure the network services accordingly so that no unauthorised connections are allowed to come through the VM. The network configuration is part of the rootfs and the relevant services are contained in initrd, therefore their integrity is guaranteed with the dm-verity utility and the measurements will be reflected on the attestation report via the direct measured boot.

5.2   Bootstrapping

During the first boot of the Revelio VM, there are certain services that are being triggered pertaining to the encryption and integrity protection of the disk as well as the TLS identity creation.

5.2.1 Disk encryption and integrity protection.
Complying with our threat model necessitates the encryption of the secure-sensitive partitions of our disk and the integrity protection of the ones that do not contain confidential data. Regarding encryption, we leverage the sealing capability of the SEV-SNP hardware and derive a key based on the VM's cryptographic measurement. We choose that policy to ensure that only untampered VMs on the same platform can successfully decrypt the disks and access the data. This is particularly useful if we need to persist data after a shutdown or because of an unplanned power outage. During the first boot, this key is derived on the fly and encrypts the chosen volume (Fig. 3) leveraging Linux kernel's feature, dm-crypt, and cryptsetup [9] to set it up. Regarding integrity protection, as previously mentioned, we modified initramfs to setup dm-verity for the rootfs and boot with the virtual device as the root partition instead of the ordinary one. Before mounting, we leverage veritysetup [9] to verify the integrity of the block device, passing the root hash from the kernel command line and the rest of the integrity metadata that have been

planted on another partition during the build phase. If the verification is successful we can proceed with the booting, otherwise, the VM's launching is terminated.

5.2.2 Unique VM identity creation.
During the first boot, a service is triggered that creates a key pair unique to each VM. This will either be the TLS identity that will correspond to the SSL certificate identifying the web-facing service, or it can be used for secure data exchange between VMs after a mutual attestation has taken place. After the key pair creation, a set of attestation reports are being created. The first one includes as `REPORT_DATA` the hash of the public counterpart of the VM's identity. The second report is related to the SSL certificate issuing, for which a Certificate Signing Request (CSR) is created first for the VM's key pair based on the configuration details on the service's domain. The `REPORT_DATA` contains the hash of the CSR in this case.

5.3   Normal operation

After Revelio's bootstrapping, the VMs on which the web-facing service will run, need to acquire the SSL certificate and the corresponding private key, since they all have to share them to satisfy our scalability requirement. In our implementation, we assume that there is an isolated node running on the service provider's premises that holds the DNS API credentials and is responsible for the SSL certificate issuing, distribution and node attestation. We will refer to this node as SP node in the following sections. For the HTTP server we use nginx and to process the HTTP GET and POST requests as well as trigger the Revelio relevant processes, we use CGI scripts that interact with the server via the FastCGI protocol [7]. To validate the certificate chain while performing remote attestation, we contact AMD Key Distribution Server (KDS) [3] and retrieve the root certificates (for AMD Root Key (ARK) and AMD SEV Key (ASK)) as well as the VCEK certificate after we have specified the Chip ID and TCB version. To validate the measurements of the reports, we compare them against some hard-coded values that have been planted on the VMs at the build time. However, this can be changed and each node can contact a remote Trusted Registry that is maintained on a blockchain infrastructure for instance [37], where the community votes on what is a "good" state or not.

5.3.1 Certificate management.
In order for the SP node to generate an SSL certificate, it needs to acquire first the CSRs from the nodes and pick one as the "leader". Before entrusting one with this role, SP attests the whole set of the nodes by retrieving their report-CSR bundles (Fig. 4) and verifying two things; first that the measurements are as expected and second that the CSR's hash matches the one imprinted on the `REPORT_DATA` of the report. Besides that, during the attestation the node's Chip ID (imprinted in the report) and IP are being checked against a set of approved ones, to avert an impersonator from obtaining the private key of the SSL certificate, even if it presents an authentic

---

[3]Hosted at https://kdsintf.amd.com/

Anna Galanou, Khushboo Bindlish, Luca Preibsch, Yvonne-Anne Pignolet, Christof Fetzer, and Rüdiger Kapitza
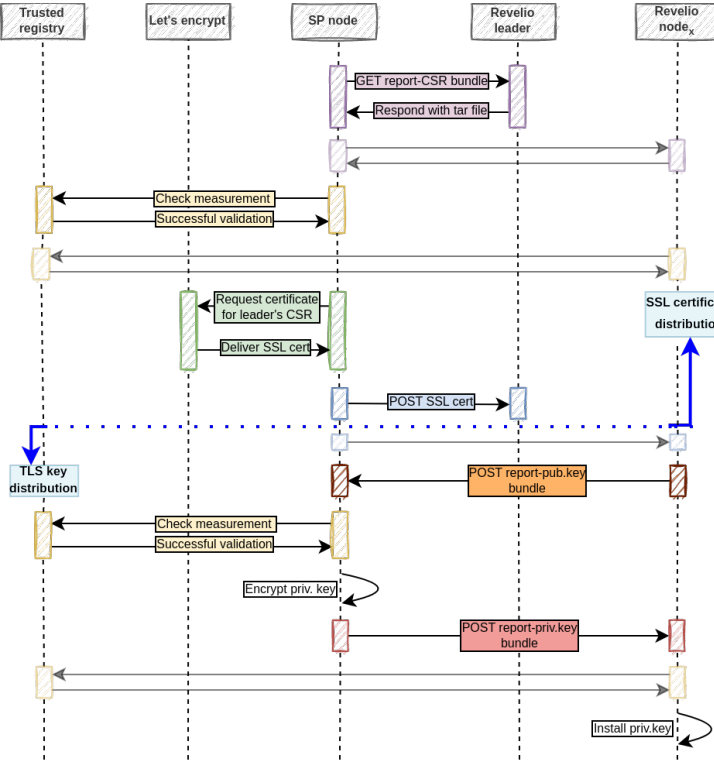

Figure 4: SSL certificate and TLS private key distribution.

and valid attestation report. After the round of attestations has been completed and the certificate has been generated for the leader's CSR, the SP node sends the certificate in a round of HTTP POST requests along with the leader's IP to notify the nodes, who they should contact later to acquire the private key.

For the secure exchange of private key (Fig. 4), each of the service provider's nodes contacts the chosen leader and performs an HTTP POST request with their report-public key bundle, containing a report with the hash of their public key, along with the public key itself. The leader validates the report in a similar fashion as the SP node and then encrypts its private key with each node's public key. Then it begins a round of POST requests sending its own attestation report bundle, containing the encrypted private key. The nodes attest the leader in the same way and after verifying that the SSL certificate corresponds to the received private key, they install both of them in a temporary folder. An incron job [16] is triggered after that, installs the private key in the server-designated folder and restarts the HTTP server which is subsequently ready to serve requests. It is worth noting that the private key is stored in an encrypted partition, so it cannot be leaked at rest.

5.3.2 Remote attestation by end-users. To enable seamless remote attestation of a Revelio VM for end-users, the former must be integrated into the browser. This can be achieved in two ways, either by extending the browser itself or by a

less intrusive way of a web-extension. The former approach would require from major browser vendors to provide the necessary support for remote attestation. For this reason, we designed a web extension that will automatically perform remote attestation for registered sites and enable it to learn about new sites offering Revelio attestation while browsing the web. Therefore, we assume that the validated HTTP server provides an attestation report under a well-known URL (e.g., as in the case of robots.txt [27]) and runs in a Revelio VM.

Register Revelio-conformed websites. We assume two basic approaches of how a website and the associated Revelio VMs can be registered via the Revelio web extension. The first approach requires the end-user to manually register a website that should be validated via a configuration dialogue of the web extension. In the basic case, the end-user needs to provide the domain name of the website and the expected measurement. The measurement has either been computed by the user itself or it has been received via an out-of-band channel. Here, more sophisticated schemes can be assumed but are out of the scope of this work. The second approach is simply to opportunistically learn about the Revelio VMs while browsing the web. To do this, the web extension will detect if a website is hosted by a Revelio VM by checking the well-known URL for the attestation report. If an attestation report can be fetched, the web extension alerts the user who subsequently needs to validate the measurement. The first approach is more secure and should be employed for security-sensitive sites since it does not require an initial untampered contact.

Intercepting requests and performing remote attestation. Assuming a domain has been registered with the web extension for remote attestation, the first access to it for a new browser context is always intercepted by the web extension and the attestation report is fetched from the VM. Next, the report is validated to verify that it originates from a SEV-SNP secured context. After querying AMD KDS for the VCEK key, providing the corresponding Chip ID and TCB version extracted from the attestation report, the web extension validates the certificate chain of VCEK (checking the ASK and ARK certificates) and that the report's signature matches VCEK's public key. Following this, it validates the report's measurement against the locally stored one. If that succeeds, it is confirmed that the Revelio VM is considered trustworthy. As a last step, it has to be validated that the browser has established a secure connection that terminates inside the VM. This is done via querying the browser regarding the public key of the current TLS connection used to connect to the remote site. If the value matches the public key that is part of the attestation report, it is ensured that the endpoint terminates securely in the VM. Should any of the checks fail, this is flagged to the user and they have to make a decision to proceed with or abort the access to the website. After this point, it has to be constantly monitored that the TLS connection is not reset and redirected to a different location. This is possible as the browser only validates

if a connection request is established with a certificate that it trusts. A malicious service provider can, for example, create a new certificate as they control access to DNS and use this new certificate to redirect users away from the secure VM. The web extension prevents this by intercepting all requests to the domain and validates that for each request the connection is based on the public key that was part of the attestation report.

At a technical level, it has to be noted that currently only Mozilla Firefox provides the necessary APIs to run the web extension. The use of the service worker API (Section 4.2) should be avoided for now since the (re-)loading of it can only be partially controlled. In this case, a small extension of the browser-provided APIs could lift this restriction.

6   EVALUATION

In this section, we present a security analysis of Revelio to demonstrate its robustness against attacks assumed in our threat model. Next, we assess the overhead that some of our Revelio-enforced techniques impose in VM's booting, runtime as well as in the interaction with the web-facing service.

6.1   Security Analysis

In order to ascertain the protection of Revelio, we provide a security analysis addressing potential attacks that can be launched by any of the stakeholders assumed in our model.

6.1.1   Loading a modified kernel or initrd. Since the hypervisor used to launch the VMs is under the cloud provider's control in the host platform, it can be instructed to load a malicious guest kernel or a modified initrd which does not enable the integrity protection for the root file system or edit the kernel command-line arguments to pass a different root hash. These attacks can also be performed by the service provider and are averted due to the measured direct boot deployed in our system. If the host uses the wrong kernel, initrd, or command-line, the measurements constructed by the QEMU and then verified by OVMF will not match and the booting will not be successful. If the host replaces the OVMF with a malicious version that does not verify the hashes, then this will be reflected on the measurements taken by the AMD-SP and hence on the attestation report. If the host fills the expected hashes in the table constructed by the QEMU, but passes the wrong kernel/initrd/command-line, then the OVMF will detect that when it measures the individual components and checks them against the stored measurements so the booting will not be successful either.

6.1.2   Tampering with rootfs. If the cloud/service provider modifies the rootfs image (with a malicious service for instance), then the root hash included in the kernel command line arguments will not correspond to it, so the mounting will be unsuccessful since we verify at this stage its integrity. If they try to modify the root hash as well, then this will be measured and reflected on the attestation report, so the attestation of the VM will fail.

Table 1: Revelio imposed delays on first boot

	Latency (ms)		Overhead (%)	
	BN	CP	BN	CP
`dm-crypt` setup	611	481	2.76	4.94
`dm-verity` setup	219	194	0.97	1.94
`dm-verity` verify	4680	3340	25.94	48.61
Identity creation	123	132	0.54	1.31

6.1.3   Modifying the system during runtime. An external attacker may attempt to modify the system during runtime since any changes after boot will not be reflected on the attestation report. In order to do that they would have to remotely access the system which is not possible with our imposed system configuration, or modify the rootfs (see Section 6.1.2). Even if those measures fail, since dm-verity protects the disk at a binary level, even a single bit change anywhere in the disk will cause dm-verity to raise errors, including mounting the disk read-write since that will cause small changes in metadata that get written to disk.

6.1.4   Rollback attacks on the VM image. An ill-intended service provider may launch a VM with an obsolete software stack so that they can exploit an existing bug. The certificate chain, Chip ID and IP validation would succeed, but the verification of the VM's measurement will fail, since we assume that the obsolete cryptographic hashes are being revoked every time there is a newer image rollout to prevent rollback attacks.

6.2   System performance

In this section, we evaluate how much delay Revelio's techniques introduce to the VM's booting during bootstrapping as well as during runtime. Our measurements were taken on a machine equipped with an AMD EPYC 7313 16-core processor and 112 GB RAM. We run Ubuntu 20.04.6 with Linux kernel 5.19.0-rc6 with the SEV-SNP patches installed on the host, and the VM is running Ubuntu 20.04.5, a Linux kernel 5.17.0-rc6 with the SEV-SNP patches installed.

6.3   Booting latency

The services that are relevant to Revelio's implementation, which we evaluate are (see Table 1):

- the encryption service, which retrieves the VM's sealing key derived from its measurement, and then encrypts the chosen volume with cryptsetup,
- the device-mapper service, which manages the integrity protected and readonly virtual volume created with veritysetup,
- the rootfs verification service, which verifies its underlying volume with the provided integrity metadata and root hash,
- and the VM identity creation service, which creates the VM's key pair, CSR, and a pair of reports holding this data.
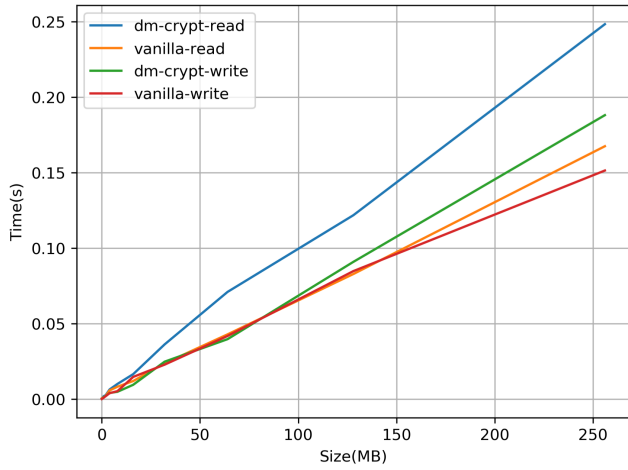
Figure 5: Dm-crypt I/O latency.

Figure 6: Dm-verity read latency.

The encryption service latency is size dependent and the volume it was performed on had a size of 84 MB which was sufficient for both of our use cases. The same applies to the verification of the dm-verity protected rootfs, which in our case is 4 GB. The encryption of the volume along with the VM's identity creation happens only on the first boot. The performance overheads, presented in Table 1, have been calculated against the total time it took for a Revelio-protected Boundary Node (BN) and Revelio-protected Cryptpad Server (CP) to boot and initialise all the system services, which were 22.725 s and 10.211 s respectively. The difference in the booting time stems from the fact that the Boundary Node has a lot more services that are starting up during boot in comparison with the Cryptpad Server that has only the server instance and the Revelio system-related services.

6.3.1 I/O request latency. To further evaluate the overhead that dm-crypt and dm-verity introduce during runtime, we simulate read and write requests in an encrypted volume of 10 GB size and read the files under the read-only integrity protected rootfs of 4 GB respectively. We have configured dm-crypt to use `aes-xts-plain64` for the cipher and `pbkdf2` with 1000 iterations for its derivation. For dm-verity, the chosen hash algorithm is `sha256` with a data and hash block size of 4 kB.

The write and read requests are performed via the `dd` utility with a block size of 4 kB and a total size of up to 256 MB, because for both of our use cases, namely the Boundary Node and the Cryptpad Server, the average file being read or written does not surpass this size. For read requests the minimum overhead introduced by dm-crypt (Fig. 5) is 1.99% and the average is 26.32%, while for write requests those are 0.35% and 12.03% respectively. Considering the evaluation of the integrity-protected volume (Fig. 6), we read the files under the BN's rootfs where the biggest file has a size of 94.8 MB and the read latency presents on average a 9.35× slowdown.
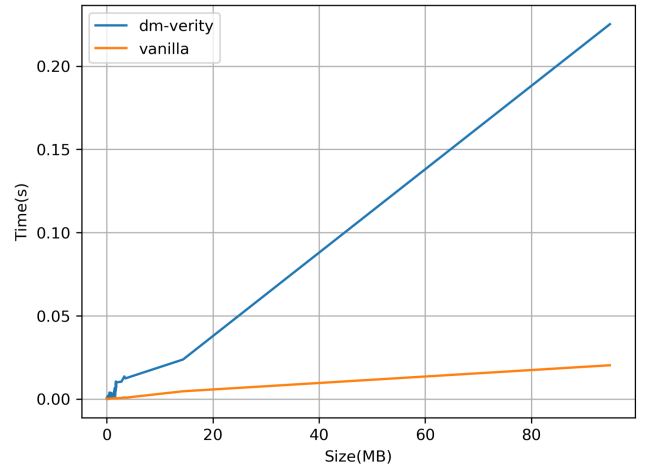
Table 2: SSL certificate generation and distribution

	Latency (ms)
Attestation evidence retrieval	17
Attestation evidence validation	13
SSL certificate generation	2996
SSL certificate distribution	15

6.3.2 SSL certificate operations latency. Regarding the impact of the SSL certificate generation and distribution (see Table 2), we measured the time that takes the AMD-SP to retrieve the attestation evidence from one node – which comprises of the CSR reflecting its key and the attestation report containing the hash of the CSR–, validate the report's signature, certificate chain, measurement and hash, as well as the time to create the SSL certificate with certbot and distribute it to the node. This happens typically once every 90 days when the SSL certificate needs to be renewed and redistributed, so it doesn't affect the runtime performance.

6.4 Client side impact

To determine the impact of performing remote attestation and constantly securing the connection of a web browser we evaluate a mobile user scenario. The client-side consisted of a notebook (Apple M2, 16 GB RAM, macOS Ventura 13.3.1) connected via wireless to the Revelio-protected Boundary Node. An instance of Firefox (13.0.2) equipped with and without our web extension repeatedly accessed a minimal web page (see Table 3) using Selenium (4.12). The base network latency accounted for 5.2 ms and the plain access of the web page with 100.9 ms. In a fresh web session, remotely attesting a Revelio VM takes (including accessing the web page) 778.9 ms on average. Thereby contacting the AMD key server for the VCEK consumes most of the time (427.3 ms). Since the VCEK is the same until the SEV-SNP firmware is updated, it can be cached, and this speeds up the access of

Table 3: Browser-based remote attestation and validation

	Latency (ms)
Network latency	5.2
Plain HTTP GET	100.9
HTTP GET and remote attestation	778.9
HTTP GET and conn. validation	115.0

websites that are frequently visited. Once the remote attestation succeeds, it has to be monitored that the connection is not reset and replaced with a new certificate. This requires for each request to query the browser from the web extension for the connection context. Accessing the test web page while a remote attestation has already been performed requires on average 115.0 ms. If the browser itself would be modified this overhead could be eliminated, as a re-establishment of a connection could simply trigger a re-validation in this situation.

7 RELATED WORK

There has been extensive research regarding remote attestation of TEEs [4, 12, 15, 19, 20, 31, 33, 41] and how it can be used to establish trust in software components and security architectures. Ryoan [22] and TrustJS [18] feature a two-way sandbox, protected inside SGX enclaves, enabling end-users to run their workloads on untrusted systems and attest the execution environment. Due to the construction of the double-sided sandbox, the execution results are only exposed to the end-user. So the data of users stay confidential to the service provider. Revelio presents a more flexible approach for VM-based TEEs giving the means to end-users to validate a utilised service including its configuration to the extent of interest without depending on the cloud or service providers.

Remote attestation via TLS [26] and RATLS [40] aim to seamlessly combine remote attestation with the TLS protocol. Both approaches could be integrated with Revelio. Narayanan et al. [32] showcase an implementation of a vTPM based on which they perform remote attestation of SEV-SNP VMs. This approach could also be applied to Revelio's architecture and enable us to have a runtime monitoring system leveraging the vTPM. Johnson et al. [25] guarantee the confidential execution of containerized workloads by introducing and attesting execution policies which are essentially proof over all the future states of the containers, while Pontes et al. [34] implement a SPIRE plugin to provision verifiable identities to SEV-SNP VMs. Revelio, on the other hand, enables end-users to seamlessly attest a web-facing service, before passing any sensitive data to it, via a web extension that verifies its cryptographic measurement over its loading-time state.

8 CONCLUSION

In this work, we introduced Revelio, a novel security architecture that enables end-users to attest web-based services that run on hardware-protected VM-based TEEs from their browsers and to shield their sensitive data from a malevolent cloud or service provider. With a small performance cost, our system can be applied in a wide range of use case scenarios, enhancing the security guarantees offered by confidential computing and allowing end-users to reap the benefits.

REFERENCES

[1] Josh Aas, Richard Barnes, Benton Case, Zakir Durumeric, Peter Eckersley, Alan Flores-López, J. Alex Halderman, Jacob Hoffman-Andrews, James Kasten, Eric Rescorla, Seth Schoen, and Brad Warren. 2019. Let's Encrypt: An Automated Certificate Authority to Encrypt the Entire Web. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 2473–2487. https://doi.org/10.1145/3319535.3363192

[2] Advanced Micro Devices, Inc. 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.

[3] M. R. Albrecht, S. Celi, B. Dowling, and D. Jones. 2023. Practically-exploitable Cryptographic Vulnerabilities in Matrix. In *2023 2023 IEEE Symposium on Security and Privacy (SP) (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 1419–1436. https://doi.org/10.1109/SP46215.2023.00081

[4] Pedro Antonino, Ante Derek, and Wojciech Aleksander Woloszyn. 2023. Flexible remote attestation of pre-SNP SEV VMs using SGX enclaves. arXiv:2305.09351 [cs.CR]

[5] Arm Ltd. 2023. Introducing Arm Confidential Compute Architecture (Version 1). https://developer.arm.com/documentation/den0125/0100.

[6] Knut Blind, Sivan Pätsch, Sachiko Muto, Mirko Böhm, Torben Schubert, Paula Grzegorzewska, and Andrew Katz. 2021. The impact of Open Source Software and Hardware on technological independence, competitiveness and innovation in the EU economy. https://doi.org/10.2759/430161

[7] Mark R. Brown. 1996. FastCGI: A high-performance gateway interface. In *Fifth International World Wide Web Conference*, Vol. 6.

[8] Milan Broz. 2022. DMVerity. https://gitlab.com/cryptsetup/cryptsetup/-/wikis/DMVerity

[9] Milan Broz. 2023. Cryptsetup and LUKS - open-source disk encryption. https://gitlab.com/cryptsetup/cryptsetup

[10] Bundesamt für Sicherheit in der Informationstechnik 2020. Security requirements for eHealth applications. https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TR03161/TR-03161.pdf

[11] TODO Group (European Chapter). 2020. Why open source software matters to your enterprise. https://project.linuxfoundation.org/hubfs/Reports/Why-open-source-software-matters-to-your-enterprise_090820.pdf?hsLang=en

[12] Guoxing Chen, Yinqian Zhang, and Ten-Hwang Lai. 2019. OPERA: Open Remote Attestation for Intel's Secure Enclaves. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 2317–2331. https://doi.org/10.1145/3319535.3354220

[13] CVE-2017-1000051 2017. Cross-site scripting (XSS) vulnerability in pad export in XWiki labs CryptPad before 1.1.1 allows remote attackers to inject arbitrary web script or HTML via the pad content. National Vulnerability Database. https://nvd.nist.gov/vuln/detail/CVE-2017-1000051

[14] Caleb James Delisle. 2017. Cryptpad Blog:Security growing pains. https://blog.cryptpad.org/2017/03/06/Security-growing-pains/

[15] Karim Eldefrawy, Norrathep Rattanavipanon, and Gene Tsudik. 2017. HYDRA: Hybrid Design for Remote Attestation (Using a Formally Verified Microkernel). In *WiSec '17: Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks* (Boston, Massachusetts) *(WiSec '17)*. Association for Computing Machinery, New York, NY, USA, 99–110. https://doi.org/10.1145/3098243.3098261

[16] Charles Fisher. 2017. Linux filesystem events with inotify. *Linux Journal* 2017, 280 (2017), 2.

[17] The Matrix.org Foundation. 2023. Matrix Specification. https://spec.matrix.org/latest/

[18] David Goltzsche, Colin Wulf, Divya Muthukumaran, Konrad Rieck, Peter Pietzuch, and Rüdiger Kapitza. 2017. TrustJS: Trusted Client-Side Execution of JavaScript. In *Proceedings of the 10th European Workshop on Systems Security* (Belgrade, Serbia) *(EuroSec'17)*. Association for Computing Machinery, New York, NY, USA, Article 7, 6 pages. https://doi.org/10.1145/3065913.3065917

[19] F. Gregor, W. Ozga, S. Vaucher, R. Pires, D. Le Quoc, S. Arnautov, A. Martin, V. Schiavoni, P. Felber, and C. Fetzer. 2020. Trust Management as a Service: Enabling Trusted Execution in the Face of Byzantine Stakeholders. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE Computer Society, Los Alamitos, CA, USA, 502–514. https://doi.org/10.1109/DSN48063.2020.00063

[20] Sarah C. Helble, Ian D. Kretz, Peter A. Loscocco, John D. Ramsdell, Paul D. Rowe, and Perry Alexander. 2021. Flexible Mechanisms for Remote Attestation. *ACM Trans. Priv. Secur.* 24, 4, Article 29 (sep 2021), 23 pages. https://doi.org/10.1145/3470535

[21] Lawrence E. Hughes. 2022. *PKCS #10 Certificate-Signing Request (CSR)*. Apress, Berkeley, CA, 75–91. https://doi.org/10.1007/978-1-4842-7486-6__6

[22] Tyler Hunt, Zhiting Zhu, Yuanzhong Xu, Simon Peter, and Emmett Witchel. 2016. Ryoan: A Distributed Sandbox for Untrusted Computation on Secret Data. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, Savannah, GA, 533–549. https://www.usenix.org/conference/osdi16/technical-sessions/presentation/hunt

[23] Intel Corp. 2021. *Intel Trust Domain Extensions*. Technical Report 1.0. Intel Corp. https://cdrdv2.intel.com/v1/dl/getContent/690419.

[24] Internet Security Research Group (ISRG). 2021. Rate Limits. https://letsencrypt.org/docs/rate-limits

[25] 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]

[26] Thomas Knauth, Michael Steiner, Somnath Chakrabarti, Li Lei, Cedric Xing, and Mona Vij. 2018. Integrating Remote Attestation with Transport Layer Security. https://doi.org/10.48550/ARXIV.1801.05863

[27] Martijn Koster, Gary Illyes, Henner Zeller, and Lizzi Sassman. 2022. Robots Exclusion Protocol. RFC 9309. https://doi.org/10.17487/RFC9309

[28] Mengyuan Li, Yinqian Zhang, Huibo Wang, Kang Li, and Yueqiang Cheng. 2021. CIPHERLEAKS: Breaking Constant-time Cryptography on AMD SEV via the Ciphertext Side Channel. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 717–732. https://www.usenix.org/conference/usenixsecurity21/presentation/li-mengyuan

[29] Aaron MacSween, Caleb James Delisle, Paul Libbrecht, and Yann Flory. 2018. Private Document Editing with Some Trust. In *Proceedings of the ACM Symposium on Document Engineering 2018* (Halifax, NS, Canada) *(DocEng '18)*. Association for Computing Machinery, New York, NY, USA, Article 29, 10 pages. https://doi.org/10.1145/3209280.3209535

[30] Dov Murik and Hubertus Franke. 2021. Securing Linux VM boot with AMD SEV measurement. https://static.sched.com/hosted_files/kvmforum2021/ed/securing-linux-vm-boot-with-amd-sev-measurement.pdf.

[31] Jämes Ménétrey, Marcelo Pasin, Pascal Felber, and Valerio Schiavoni. 2022. WaTZ: A Trusted WebAssembly Runtime Environment with Remote Attestation for TrustZone. In *2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS)*. 1177–1189. https://doi.org/10.1109/ICDCS54860.2022.00116

[32] Vikram Narayanan, Claudio Carvalho, Angelo Ruocco, Gheorghe Almási, James Bottomley, Mengmei Ye, Tobin Feldman-Fitzthum, Daniele Buono, Hubertus Franke, and Anton Burtsev. 2023. Remote attestation of SEV-SNP confidential VMs using e-vTPMs. arXiv:2303.16463 [cs.CR]

[33] Simon Ott, Monika Kamhuber, Joana Pecholt, and Sascha Wessel. 2023. Universal Remote Attestation for Cloud and Edge Platforms. In *Proceedings of the 18th International Conference on Availability, Reliability and Security* (Benevento, Italy) *(ARES '23)*. Association for Computing Machinery, New York, NY, USA, Article 12, 11 pages. https://doi.org/10.1145/3600160.3600171

[34] Davi Pontes, Fernando Silva, Eduardo Falcão, and Andrey Brito. 2023. Attesting AMD SEV-SNP Virtual Machines with SPIRE. In *Proceedings of the 12th Latin-American Symposium on Dependable and Secure Computing* (La Paz, Bolivia) *(LADC '23)*. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/3615366.3615419

[35] Bazel Project. 2023. Hermeticity. https://bazel.build/basics/hermeticity

[36] Yong Shi, Mingzhi Wen, Filipe R. Cogo, Boyuan Chen, and Zhen Ming Jiang. 2022. An Experience Report on Producing Verifiable Builds for Large-Scale Commercial Systems. *IEEE Transactions on Software Engineering* 48, 9 (2022), 3361–3377. https://doi.org/10.1109/TSE.2021.3092692

[37] DFINITY Team. 2022. The Internet Computer for Geeks. https://internetcomputer.org/whitepaper.pdf

[38] Linux Foundation Research Team. 2022. Addressing Cybersecurity Challenges in open source Software. https://8112310.fs1.hubspotusercontent-na1.net/hubfs/8112310/LF%20Research/Addressing%20Cybersecurity%20Challenges%20in%20Open%20Source%20Software%20-%20Report.pdf

[39] Christian Tiefenau, Emanuel von Zezschwitz, Maximilian Häring, Katharina Krombholz, and Matthew Smith. 2019. A Usability Evaluation of Let's Encrypt and Certbot: Usable Security Done Right. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 1971–1988. https://doi.org/10.1145/3319535.3363220

[40] Robert Walther, Carsten Weinhold, and Michael Roitzsch. 2022. RATLS: Integrating Transport Layer Security with Remote Attestation. In *Applied Cryptography and Network Security Workshops: ACNS 2022 Satellite Workshops, AIBlock, AIHWS, AIoTS, CIMSS, Cloud S&P, SCI, SecMT, SiMLA, Rome, Italy, June 20–23, 2022, Proceedings* (Rome, Italy). Springer-Verlag, Berlin, Heidelberg, 361–379. https://doi.org/10.1007/978-3-031-16815-4__20

[41] Sachiko Yoshihama, Tim Ebringer, Megumi Nakamura, Seiji Munetoh, and Hiroshi Maruyama. 2005. WS-Attestation: Efficient and Fine-Grained Remote Attestation on Web Services. In *Proceedings of the IEEE International Conference on Web Services (ICWS '05)*. 743–750. https://doi.org/10.1109/ICWS.2005.136