

T-KVM: A Trusted architecture for KVM ARM v7 and v8 Virtual Machines

Securing Virtual Machines by means of KVM, TrustZone, TEE and SELinux

Michele Paolino, Alvise Rigo, Alexander Spyridakis, Jérémy Fanguède, Petar Lalov and Daniel Raho

Virtual Open Systems

Grenoble, France

Email: {m.paolino, a.rigo, a.spyridakis, j.fanguede, p.lalov, s.raho} @virtualopensystems.com

Abstract—The first market release of Advanced RISC Machines (ARM) v8 System on Chips (SoCs) has created big expectations from smart devices, servers and network equipment vendors, who see compelling advantages in integrating them into their systems. As a consequence software stack deployments for ARMv8 platforms translate market requirements to support OpenStack, Network Functions Virtualization (NFV), Mobile Edge Computing (MEC), In-Vehicle Infotainment (IVI) automotive functions. At the same time, ARMv8 will empower Internet of Things (IoT), Cyber Physical Systems (CPS) and user convergence devices. In this context, virtualization is a key feature to enable the cloud delivery model, to implement multitenancy, to isolate different execution environments and to improve hardware/software standardization and consolidation. Since guaranteeing a strict ownership of both the data and the code executed in Virtual Machines (VMs), which belong to governments, companies, telecom operators and private users, counts more than ever, the security of the hypervisor and its guests has become dramatically important. In this paper, Trusted Kernel-based Virtual Machine (T-KVM), a novel security architecture for the KVM-on-ARM hypervisor, is proposed to satisfy the current market trend. T-KVM integrates software and hardware components to secure guest Operating Systems (OSes) and enable Trusted Computing in ARM virtual machines. The proposed architecture combines four isolation layers: ARM Virtualization and Security Extensions (also known as ARM VE and TrustZone), GlobalPlatform Trusted Execution Environment (TEE) APIs and SELinux Mandatory Access Control (MAC) security policy. The T-KVM architecture can be implemented on platforms based on ARM v7 and v8 architectures, without requiring additional custom hardware extensions, since, starting from Cortex-A15 (ARM v7 architecture) released in 2012, both the ARM VE and TrustZone are made available. In this paper the T-KVM architecture is described in details, as well as its key implementation challenges and system security considerations. Lastly, a performance evaluation of the proposed solution is presented.

Keywords—Trusted KVM, KVM Security, ARMv8 Trusted Computing, KVM TrustZone, ARM Virtualization, SELinux, TEE, ARM VMs Security.

I. INTRODUCTION

The use of virtualization in ARM platforms is rapidly increasing due to the deployment of SoCs based on this architecture in different environments such as: servers, Cloud and High Performance Computing (HPC), NFV, MEC, IoT, CPS, smart devices, etc. This technology enables multiple OSes to run unmodified on the same hardware, thus sharing system's resources such as memory, CPUs, disks and other devices. These resources are frequently target of specific virtualized environment attacks (e.g., CPU cache [1], memory bus [2] and VM's devices [3] [4]). For this reason, the security of the virtualized systems is critical. Historically, isolation has been used

to enhance the security of these systems [5], because it reduces the propagation risks in compromised environments [6].

The aim of this paper is to propose T-KVM, a novel security architecture for virtualized systems based on four isolation layers: KVM, ARM TrustZone, GlobalPlatform TEE API and SELinux. The former is considered the most popular hypervisor deployed in OpenStack [7], which is a key solution for Cloud, NFV and HPC computing. KVM for ARM is part of the Linux kernel starting from the version 3.9; it is the Linux component that, exploiting the ARM Virtualization Extension, allows to create a fully-featured virtualization environment providing hardware isolation for CPU, memory, interrupts and timers [8]. TrustZone is an hardware security extension for ARM processors and Advanced Microcontroller Bus Architecture (AMBA) devices [9] designed to drastically improve security inside the ARM ecosystem. The extension starts from the assumption that a system, in order to deliver secure services, has to decouple the resources used for general purpose applications from those that handle security assets. To this end, TrustZone creates two hardware isolated partitions in the system: the Secure and the Non Secure World. While the Non Secure World runs a standard OS with optionally a hypervisor, the Secure World contains, handles and protects all the sensitive data of the system. These two worlds are linked together through the GlobalPlatform TEE API [10] [11], a set of specifications for a secure Remote Procedure Call (RPC) mechanism between the trusted and non-trusted compartment of the system. At the time of writing, these specifications don't support virtualization, preventing the use of TPM services inside virtual machines. This work addresses this limitation proposing a design of a set of virtualization extensions to enable the guest operating systems to make use of TPM services provided by the TrustZone Secure World. The latter T-KVM isolation layer is Security-Enhanced Linux (SELinux), a Mandatory Access Control (MAC) solution which brings type enforcement, role-based access control and Multi-Level Security (MLS) to the Linux kernel [12]. By means of these, SELinux confines processes in security domains, where the interaction with other processes and files is permitted only if there is a specific SELinux policy rule which allows it.

In T-KVM, the above technologies are combined and adapted to work together, providing a high security level to the guest applications, without the need of specific hardware or software. As a matter of fact, the proposed architecture relies on open source (KVM and SELinux) components, public specifications (GlobalPlatform TEE Internal and Client APIs) and available hardware features (ARM TrustZone and VE). For these reasons, T-KVM can be easily ported to currently available ARM platforms and Cloud Infrastructure systems

such as OpenStack.

The remaining part of this paper is organized as follows: Section II provides more details about the main security components of the proposed architecture. Section III contains details about the T-KVM architecture, its implementation and security considerations while Section IV present a performance analysis of the overhead introduced by the proposed solution. The related work is presented in Section V and Section VI concludes the paper.

II. THE SECURITY COMPONENTS

In this section, the isolation layers which characterize the T-KVM architecture are described.

A. KVM hypervisor

A hypervisor is a software layer which is able to create virtual instances of hardware resources such as CPUs, memory, devices, etc. in order to enable the execution of multiple operating systems on the same hardware. Different implementation approaches lead to different hypervisor types: a type 1 hypervisor, is a bare metal hypervisor which runs directly on the hardware (XEN or VMWare ESX). A type 2 hypervisor is, on the other hand, a hypervisor which runs inside an operating system (Oracle VirtualBox or VMWare Workstation) at the application layer. Usually, the latter is used in less critical applications [13] because of its dependency from the underlying operating system.

KVM is a hypervisor included in the Linux kernel and available for ARM, x86 and s390 architectures. It is neither a type 2 hypervisor because it does not run as a normal program inside Linux, nor is a typical type 1 hypervisor, because it relies on the Linux kernel infrastructure to run. KVM exploits the CPU Virtualization Extensions to execute guest's instructions directly on the host processor and to provide VMs with an execution environment almost identical to the real hardware. Each guest is run in a different instance of this execution environment, thus isolating the guest operating system. For this reason, this isolation has been used for security purposes [14] [15] [16] [17] in many scientific works. In the ARM architecture, the KVM isolation involves CPU, Memory, Interrupts and timers [8].

B. TrustZone

ARM TrustZone is a set of hardware security extensions for ARM processors and AMBA devices. With TrustZone, the hardware platform is split in two parts, the Secure and the Non Secure Worlds. In order to isolate these two compartments, TrustZone requires: CPU with ARM Security Extensions (SE) along with TrustZone compliant MMU, AMBA system bus, interrupt and cache controllers. Hence the isolation provided by TrustZone includes CPU, AMBA devices, interrupts, memory and caches.

The Secure World is considered trusted, and is responsible for the boot and the configuration of the entire system. In fact, the CPU has banked registers for each World, and security specific configurations can be performed in Secure World mode only. This compartment contains the root of trust of the system and protects sensitive data. The access to AMBA peripherals such as fingerprint readers, cryptographic engines, etc. can be restricted only to the Secure World, thus protecting security devices.

On the other hand, the Non Secure World is intended to be the user's World. In this untrusted compartment, a standard operating system (i.e., Android or Linux) is run. Secure operations such as the access to a secret or the execution of a security algorithm are provided to the user's application running in this compartment by the services run in the Secure World.

These two compartments interact with each other through a specific CPU mode, namely the Monitor Mode. It typically runs a secure context switch routine and is capable of routing interrupts, depending on the configuration, either to the Secure or Non Secure World.

Moreover, the use of the ARM VE Extensions, and consequently of KVM, is possible only in the Non Secure World.

ARM TrustZone is compliant with the GlobalPlatform TEE System Architecture specification [18], which defines the attributes that the hardware must have to properly execute a TEE.

C. GlobalPlatform TEE

GlobalPlatform defines the TEE as an execution environment, which provides security features such as isolated execution, integrity of Trusted Applications (i.e., applications run in the TEE) along with confidentiality of their assets [18]. This isolation protects Trusted Applications (TA) and their data from the Rich Execution Environment (REE), the environment where a standard operating system such as Linux or Android is run. Figure 1 depicts the standard architecture of a GlobalPlatform TEE compliant system.

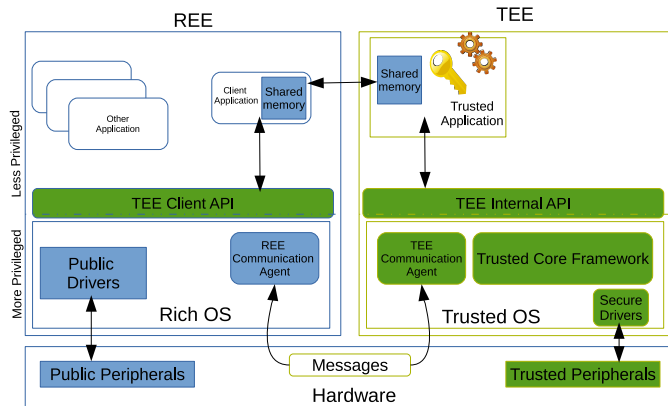


Figure 1. Standard TEE architecture

In order to isolate the TEE from the REE, GlobalPlatform provides a set of specifications which include the following software components [10] [11]:

- The Trusted Core Framework is a common abstraction layer which provides to the TEE Internal API OS-like functions, such as memory management, entry points for TAs, panic and cancellation handling, TA properties access, etc.
- The TEE Client API is an Inter Process Communication (IPC) API that deals with the communication between the REE and the TEE. It allows the applications in the REE (Client Applications or CAs) to leverage the services offered by the TEE.

- The (TEE and REE) communication agents provide support for messaging between the CAs and the TEE. They interact with the Monitor mode to request a context switch between the two Worlds.
- The TEE Internal API allows the TAs to leverage the services offered by the TEE through the following APIs: Trusted Storage for Data and Keys, TEE Cryptographic Operations, Time, and TEE Arithmetical.

Lastly, it is worth to mention that the deployment of a GlobalPlatform compliant solution enables the use of existing TA and CA applications. This is a very important factor, especially in an environment such as the embedded trusted computing, where by tradition the security solutions were developed each time from scratch to address new device families.

D. SELinux

SELinux is a software implementation of the MAC security policy available in the Linux kernel as Linux Security Module (LSM). The key feature of MAC is that the access control decisions are not at discretion of individual users, root included [12]. Thus, once the system security policies have been defined and loaded at boot time in the kernel, they can not be modified. In this way, the subject (e.g., a process) access to objects (e.g., file, socket, etc.) is enforced in the system.

The very same concept can be applied to virtual machines using sVirt, which is a feature of the libvirt library. sVirt installs a set of virtualization specific security policies and automatically tags VMs and their resources in order to isolate guest systems. This isolation prevent any access to VM's resources (disk/kernel files, shared memory, etc.) from external subjects (other VMs, the root user, etc.).

For this and for performance reasons [19], the use of SELinux in virtualized systems is encouraged.

III. THE TRUSTED HYPERVISOR: T-KVM

T-KVM is a secure hypervisor architecture based on KVM, which combines a Trusted Computing solution such as TrustZone with GlobalPlatform TEE and SELinux. In Figure 2, all the components described in Section II are shown together, composing the T-KVM architecture.

In T-KVM, the GlobalPlatform TEE and REE are respectively implemented inside the TrustZone Secure and Non Secure Worlds. For this reason in the remaining part of this paper, Secure World/TEE and Non Secure World/REE are used as synonyms. These two hardware-isolated environments are linked together with a virtualization-enabled implementation of the GlobalPlatform TEE specifications. The virtualization provided by KVM further isolates the user's applications, enabling the use of different operating systems. This eases multitenancy in server and Cloud environments, and enables Bring Your Own Device (BYOD) [20] paradigm in smart devices. In addition, SELinux isolates in software the virtual machines, protecting guests and their resources from the other virtual machines and the host itself (e.g., malicious cloud administrators, host privilege escalation exploits, etc.). Lastly libvirt, the main virtualization API used by OpenStack to interact with KVM, takes automatically care of the policy configuration and the tag assignment through its component sVirt.




Figure 2. T-KVM architecture, which includes KVM, TrustZone, SELinux and virtualized TEE

A. Implementation details

The following part of this manuscript lists the T-KVM implementation challenges and proposes viable solutions.

1) *Trusted boot*: The first step of the system's chain of trust is performed during the boot procedure. In fact when the machine boots, the security configuration of the system is not yet in place, and as a result the system is vulnerable to attacks which target to replace the boot procedure.

In order to minimize this risk, the T-KVM's first stage bootloader is a tiny program stored in a on-chip ROM along with the public key needed for the attestation of the second stage bootloader. Since the first stage bootloader is stored in a read-only memory, it can not be updated and it's therefore critical for the security of the system. Soon after the initialization of the key system components, it checks the integrity and boots the second stage bootloader which is located in an external non-volatile memory (e.g., flash memory). The second stage bootloader then loads the microkernel binary in the system's Secure World memory and boots it.

The third stage bootloader of the T-KVM Trusted boot mechanism is a Trusted Application inside the TEE. In fact, when the Secure World OS is up and running, a specific TA checks the integrity of the Non Secure World OS binary (i.e the Linux kernel) and its bootloader (fourth stage). If this last security check is successful, the fourth stage bootloader runs the Non Secure OS and the system can be considered running. On the other hand, if only one of these checks fails, the boot process will be stopped and the machine will enter in a secure and not operational state. Figure 3 shows the T-KVM Trusted boot chain of Trust.

The Trusted boot process is the key element for the attestation of the user space applications because it ensures the integrity of the chain of trust. T-KVM runs in the TEE an attestation service, which is able to check at any moment the integrity of its key components i.e., Quick Emulator (QEMU) [21], libvirt, the VMs and their resources, etc. libvirt in particular, is extended to attest the VMs identity and integrity at each boot and in an event-driven manner, assuring to users and cloud administrators/providers the authenticity

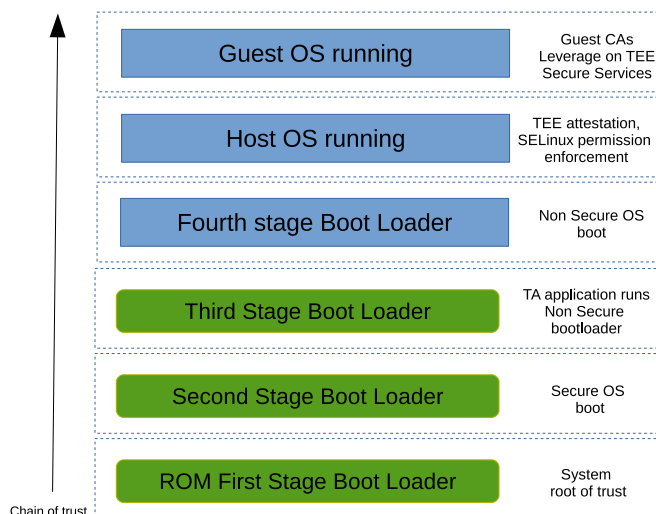


Figure 3. T-KVM Trusted boot procedure

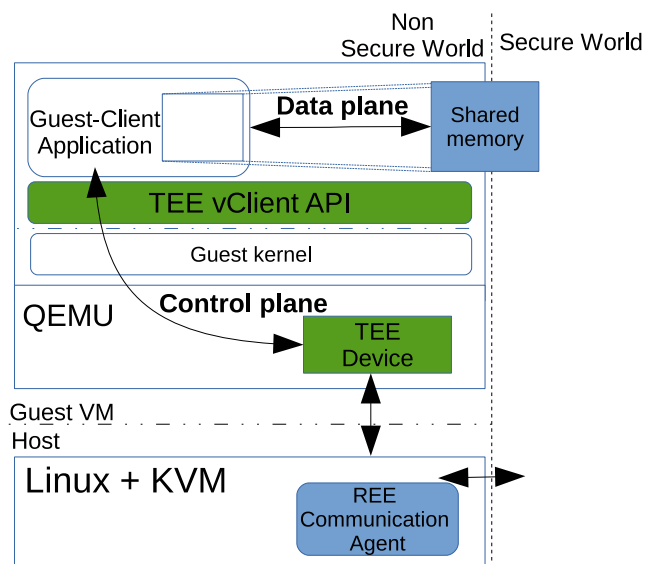


Figure 4. TEE support for Virtual Machines in T-KVM

of the workloads run on the hardware. The security assets (fingerprints, keys, etc.) of these binaries are stored in the Secure World and by consequence can be updated frequently.

2) *GlobalPlatform TEE support for virtualization:* The main novelty introduced by T-KVM is the support for Trusted Computing inside the Virtual Machines. The virtualization of the TEE functions is of utmost importance for the T-KVM architecture, because it links together the applications run in the VMs with the secure services available in the TrustZone Secure World. At the time of writing, the use of the TEE functions in guest operating systems is not included in the GlobalPlatform API Specification. To enable this feature, T-KVM virtualizes the GlobalPlatform TEE APIs, executing the TEE Client API directly in the Guest Operating System. In order to be as much as possible compliant with the GlobalPlatform Specification and to be able to run CAs also at the host level, T-KVM TEE Client API is the only virtualization aware component.

This awareness needs support from the hypervisor infrastructure. For this reason, as depicted in Figure 4, a specific QEMU device is used to implement the TEE control plane and set up its data plane. All the requests (e.g., initialization/close session, invoke command, etc.) and notification of response are sent to the TEE Device, which delivers them either to the TAs or to the CAs running on the guest OS. To provide good data throughput and latency performance, the data plane is based on a shared memory mechanism. Thus when a response notification arrives from the TrustZone Secure World, the TEE device notifies with an interrupt its driver, which forwards the related information to the Guest-Client Application. The Guest-CA is now able to read the data from the shared memory, without involving the TEE device in the data transfer.

3) *Shared memory:* T-KVM needs a zero copy shared memory mechanism to share data between the two TrustZone Worlds and between the virtual machine and the host. The latter in particular is very important in systems where VMs need to communicate with each other frequently e.g., NFV, HPC, MEC, etc. Host-guest shared memory mechanisms which provide high performance and low latency already exist for the KVM hypervisor [22] [23]. What these mechanisms actually lack is the support for TrustZone.

By design, the TrustZone Secure World is able to access the full Non Secure World address space. For this reason, the Trusted Applications are able to read/write the content of the VMs shared memory unless they know the address where the shared memory area begins. In order to pass this information, the TEE device control plane extends the T-KVM shared memory mechanism, enabling it to send the shared memory address to the Secure World applications. This mechanism needs to be secured, especially in the Non Secure World, to prevent attacks and information leakage. T-KVM relies on SELinux to define specific access rules for shared memory and to enforce the shared memory access only to the interested parties.

The encryption of the shared memory area is mindfully not considered because, unless hardware accelerators are present in the platform, there would be a performance loss.

4) *Secure World:* One of the most important parts of the T-KVM architecture is the software running in the TrustZone Secure World. The operating system running in the Secure World should be fast, secure, ideally real-time and free of programming errors (implementation correctness).

The T-KVM architecture empowers the Secure World environment with a microkernel. The primary motivation behind microkernels is the small code footprint, which lead to a smaller attack surface and an easier process of formal verification of the code. In this context, a good candidate is for example seL4, an opensource third-generation microkernel based on L4 and formally verified for functional correctness [24].

The microkernel will run in its userspace the implementation of the GlobalPlatform APIs, the secure device drivers and the TAs. In order to do this, the Secure World OS does not use the main platform storage device to store files and the security assets of the system. An external, non-volatile memory configured by the Secure World as not accessible by the Non Secure World, is used to this purpose.

Finally, a possible alternative to microkernels is a secure

library running in the Secure World such as OPTEE [25]. Despite this solution has a code footprint even smaller than a microkernel, T-KVM uses a microkernel because of its real time features and a higher flexibility for TA developers.

B. Security considerations

Virtual Machines are widely used because of their flexibility and capability to run any operating system. Nonetheless, in order to achieve a higher security level of the system, it is suggested to run single-application operating systems in the T-KVM virtual machines. An example of such an operating system is OSv [26], an opensource solution which is going to be ported to the ARM v8 architecture. For this reason, this work does not discuss the security of the applications inside the virtual machines.

The threat model considered in this paper allows the attacker to completely control one or more virtual machines, both at user and kernel space level. In addition, the cloud administrator, who is permitted to remotely control the virtualized system, is considered as a potential attacker for the identity and integrity of the data.

The security of T-KVM is based on two main assumptions: the attacker does not have physical access to the virtualized system and the first stage bootloader is flawless (thus, the chain of trust is not compromised).

T-KVM has been designed to be compliant with additional hardware accelerators and security modules. For example, SecBus [27] can be used to protect the system against physical attacks on the memory components (e.g., cold boot), Direct Memory Access (DMA) attacks and on-board probing of the external memory bus. Solutions like NoC Firewall [28], can enhance the compartment isolation granularity at VM level, protecting the system against logical attacks (virus, Trojans) or security vulnerabilities, e.g., corrupt DMA engines.

IV. EXPERIMENTAL RESULTS

The different isolation layers which compose T-KVM provide high security, but at a cost of additional overhead. As a matter of fact, a request for a security service from a virtual machine has to pass through the TEE, the host system and SELinux to arrive in the TrustZone Secure World.

For the T-KVM performance analysis of this paper, we focused on the hardware isolation provided by the hypervisor and TrustZone, as the SELinux performance has been measured by the authors in the past [19], and the TEE overhead will be detailed in future works.

For this reason, following the path of a Secure Monitor Call (SMC) from the guest to the Secure World (and then back in the guest) we measured the overhead introduced by T-KVM.

SMC has been added to the ARM instruction set by the ARM Security Extension and it is used to request the execution of a software routine in the TrustZone Secure World passing through the TrustZone Monitor. In the standard KVM implementation, when the SMC instruction is run by a guest OS, its execution is trapped by KVM, which injects an undefined instruction in the guest, forcing it to handle this accordingly. In T-KVM instead, when such instruction is run, the hypervisor traps the guest SMC execution, modifies its arguments and forwards them to the TrustZone Secure World.

In this scenario, two SMC context switches are involved: firstly from Guest to Host, then from Non Secure to Secure World Mode. The overhead assessment of these two context switch operations is the target of the following analysis.

For the first measurement, we implement a bare metal binary blob for KVM which initializes the Performance Monitoring Unit (PMU) and executes the SMC instruction in the guest. The SMC is then trapped by KVM, which has been modified to immediately return the control to the VM. As soon as the program flow returns back to the guest, it checks the PMU cycle counter status and calculates the overhead. In this way, we are able to measure the overhead of a round-trip context switch between the Guest and the Host when an SMC call is executed.

On the other hand, for the measurements of the context switch overhead between the Non Secure and the Secure Worlds, the PMU cycle counter is set by a Linux kernel module in the host, which executes soon after the world switch request (i.e., the SMC instruction). This provokes an additional switch to the TrustZone Monitor, which saves the Normal World registers, loads the Secure World status and finally jumps to the Secure World. In order to measure the T-KVM context switch cost and not add further overhead, the Secure World immediately runs the SMC call, without executing any meaningful operation. When this instruction is executed in the Secure World, it provokes again a switch to the Monitor Mode, which will now save the secure world context, restore the non secure context, and jump back to the Non Secure World. As soon as the context switch is completed, the Linux kernel reads the PMU cycle counter state and computes the overhead.

The results of both of the above measurements are defined minimal because they are not considering any additional work performed at the destination where they are trapping to. However, in a real scenario a trap is followed by the execution of emulation code for a Guest-Host trap, or some secure service for a Non Secure-Secure trap.

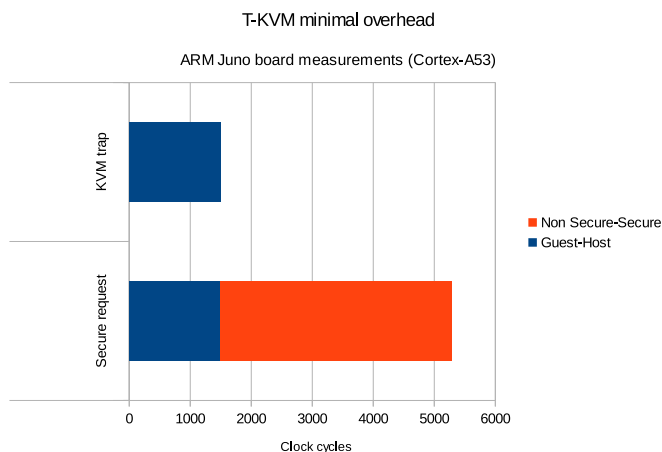


Figure 5. T-KVM overhead measurements

The performance overhead analysis has been performed on an ARMv8 Juno board, which is equipped with two Cortex-A57 and four Cortex-A53 in big.LITTLE configuration. Each test has been repeated five hundred times, running Linux

version 3.17 in both the host and the guest environments. ARM Trusted Firmware [29] has been used as firmware infrastructure. All the workloads for the measurements have been executed on the Cortex-A53, which is the default CPU that the platform uses to execute the Secure World OS.

Figure 5 shows that for a Secure request, which is the sum of the the two measurements described before in this section, the cost is in average 5200 clock cycles. This value represents the minimum overhead that a guest system has to pay to request a secure service to the T-KVM Secure World (the Guest-Host minimum overhead is about 1400, while the Non Secure-Secure is about 3700). This has been compared with the minimum overhead that KVM spends to trap the SMC instruction and perform a context switch between the guest and the host, which is what has been described above in the Guest-Host context switch measurement description. This result (in average about 1400 clock cycles) has been measured with the SMC instruction, but it is valid for all the instructions trapped in KVM, as we did not add any specific code to trap the SMC instruction to the standard KVM implementation.

Finally, it is important to notice that the Non Secure-Secure context switch is 2.5 times slower than the Guest-Host. The main reason for this behaviour is in the number of instructions needed to complete the two operations. In particular, the number of registers that the system has to save and restore for the Guest-Host context switch is significantly lower.

V. RELATED WORK

The hypervisor security is a controversial topic in literature. As a matter of fact, solutions like NoHype [30] [31] propose to secure the virtual machines removing the hypervisor, while others use the hypervisor isolation for security applications [15] [17]. In other scientific works, when compared with TrustZone as a security solution, the hypervisor proves a better flexibility e.g., the Secure World is not able to interpose on and monitor all the important events in the Non-Secure World [16]. The proposed architecture considers the hypervisor as an additional isolation layer, while protecting the security assets through the ARM Security Extensions (TrustZone). T-KVM, combining both solutions and relying on the attestation enabled by the secure boot's chain of trust, is able to provide monitoring features and high security.

In fact, attestation and integrity checks are of paramount importance for the security systems because they allow system designers and administrators to consider a software component as trusted. SecVisor [32], HyperSentry [33] and SPROBES [34] propose different solutions designed for this purpose: the first checks the integrity of commodity OS kernels running the attestation code in hypervisor mode, thus not addressing virtualization. The second enables integrity measurement of a running hypervisor (i.e., XEN) through Intel TxT, hence targeting the x86 architecture. The latter uses TrustZone to enforce kernel code integrity, but without mentioning the attestation challenges in virtualized systems.

These systems are explicitly addressed by solutions such as vTPM [35] or sHype [36], both focusing their efforts on Intel architectures. vTPM proposes a mechanism for the virtualization of TPM functions which dedicates a VM to route and manage the TPM requests, while sHype integrates the MAC security policy directly inside a typical type 1 hypervisor (i.e., XEN).

On the other hand, the solutions proposed by Narari [37] and Lengyel [38] are designed for ARM devices with virtualization extensions. The first proposes a security architecture with TrustZone, SELinux and virtualization, targeting resource constrained devices. The second combines a hypervisor (XEN) and MAC Security policies (XEN Security Modules), targeting high isolation between VMs but without mentioning TPM access for guest OSes. Both proposals lack a solution to standardize the access to TPM functions such as the TEE.

VI. CONCLUSION AND OUTLOOK

This paper proposes T-KVM, a new security architecture for ARM v7 and v8 virtualized systems, providing architecture details and a preliminary performance analysis. T-KVM's architecture offers strong isolation for guest applications by means of the KVM hypervisor, ARM TrustZone, SELinux and a virtualization enabled implementation of the GlobalPlatform TEE API. The last component design is the main contribution of this paper, as it enables the support of Trusted Computing features in ARM based virtual machines through the use of a QEMU device and a shared memory.

The benefits of the proposed solution are its flexibility and compatibility with the existing Cloud and smart devices architectures. In fact, since T-KVM adapts and combines existing opensource components which are already part of virtualized systems and OpenStack Cloud Computing infrastructure (i.e., KVM, libvirt, qemu, etc.), the support of T-KVM in these environments is straightforward to implement. Moreover, the possibility to install a real-time operating system in the TrustZone Secure World adds the support for real-time applications, which is interesting for automotive, avionics and networking applications. Lastly, monitoring and (remote) attestation are eased by the combination of a standard hypervisor with TrustZone.

On the other hand, the lack of the ARM VE support in the Secure World does not allow the hardware assisted virtualization of the TEE. Nonetheless, it is possible to functionally implement multiple TEE (or vTPMs) using paravirtualization or virtualization at the application layer.

As for the measurements and the analysis of the overhead introduced by the proposed solution, we can claim that the performance cost of T-KVM is acceptable. In fact, even if the secure request overhead is significantly higher than a trap in the KVM hypervisor, the execution frequency of the former is expected to be lower than the latter. In fact a service request is issued by a T-KVM guest only to control the communication between the guest and the Secure World, as all the data exchanges will be performed through shared memory.

Finally, the future work includes the implementation of a complete T-KVM ARMv8 prototype. Of interest is also the support and integration of T-KVM in OpenStack, which would enable a complete new set of Cloud features based on Trusted Computing such as location aware scheduling of new instances, trusted multitenancy, etc.

ACKNOWLEDGMENT

This research work has been supported by the FP7 TRESCCA project under the grant number 318036.

REFERENCES

- [1] M. Wei, B. Heinz, and F. Stumpf, "A cache timing attack on aes in virtualization environments," in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, A. Keromytis, Ed. Springer Berlin Heidelberg, 2012, vol. 7397, pp. 314–328. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-32946-3_23
- [2] B. Saltaformaggio, D. Xu, and X. Zhang, "Busmonitor: A hypervisor-based solution for memory bus covert channels," 2013.
- [3] G. Pék, A. Lanzi, A. Srivastava, D. Balzarotti, A. Francillon, and C. Neumann, "On the feasibility of software attacks on commodity virtual machine monitors via direct device assignment," in *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security*, ser. ASIA CCS '14. New York, NY, USA: ACM, 2014, pp. 305–316. [Online]. Available: <http://doi.acm.org/10.1145/2590296.2590299>
- [4] N. Elhage, "Virtunoid: A kvm guest- ζ host privilege escalation exploit," 2011.
- [5] S. E. Madnick and J. J. Donovan, "Application and analysis of the virtual machine approach to information system security and isolation," in *Proceedings of the Workshop on Virtual Computer Systems*. New York, NY, USA: ACM, 1973, pp. 210–224. [Online]. Available: <http://doi.acm.org/10.1145/800122.803961>
- [6] M. Paolino, "Isolating arm platforms: towards secure virtualized embedded systems," May 2014, last visited March 2015. [Online]. Available: http://www.cspforum.eu/uploads/Csp2014Presentations/Track_8/Strong%20Isolation%20in%20ARM%20Platforms.pdf
- [7] G. Chen, "Kvm open source virtualization for the enterprise and openstack clouds," IDC, 2014.
- [8] C. Dall and J. Nieh, "Kvm/arm: Experiences building the linux arm hypervisor," 2013.
- [9] T. Alves and D. Felton, "Trustzone: Integrated hardware and software security," ARM white paper, vol. 3, no. 4, 2004, pp. 18–24.
- [10] GlobalPlatform, "Tee client api specification, public release v1.0," 2010.
- [11] —, "Tee internal api specification, public release v1.0," 2011.
- [12] X. Xu, C. Xiao, C. Gao, and G. Tian, "A study on confidentiality and integrity protection of selinux," in *Networking and Information Technology (ICNIT)*, 2010 International Conference on, June 2010, pp. 269–273.
- [13] A. Jasti, P. Shah, R. Nagaraj, and R. Pendse, "Security in multi-tenancy cloud," in *Security Technology (ICCST)*, 2010 IEEE International Carnahan Conference on, Oct 2010, pp. 35–41.
- [14] A. Vahidi and C. Jämthagen, "Secure rpc in embedded systems: Evaluation of some globalplatform implementation alternatives," in *Proceedings of the Workshop on Embedded Systems Security*, ser. WESS '13. New York, NY, USA: ACM, 2013, pp. 4:1–4:7. [Online]. Available: <http://doi.acm.org/10.1145/2527317.2527321>
- [15] J.-E. Ekberg, K. Kostiaainen, and N. Asokan, "The untapped potential of trusted execution environments on mobile devices," *Security Privacy, IEEE*, vol. 12, no. 4, July 2014, pp. 29–37.
- [16] C. Gehrman, H. Douglas, and D. Nilsson, "Are there good reasons for protecting mobile phones with hypervisors?" in *Consumer Communications and Networking Conference (CCNC)*, 2011 IEEE, Jan 2011, pp. 906–911.
- [17] F. Liu, L. Ren, and H. Bai, "Secure-turtles: Building a secure execution environment for guest vms on turtles system," *Journal of Computers*, vol. 9, no. 3, 2014, pp. 741–749.
- [18] GlobalPlatform, "Tee system architecture v1.0," 2011.
- [19] M. Paolino, M. M. Hamayun, and D. Raho, "A performance analysis of arm virtual machines secured using selinux," in *Cyber Security and Privacy*, ser. Communications in Computer and Information Science, F. Cleary and M. Felici, Eds. Springer International Publishing, 2014, vol. 470, pp. 28–36. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-12574-9_3
- [20] K. W. Miller, J. Voas, and G. F. Hurlburt, "Byod: Security and privacy considerations," *IT Professional*, vol. 14, no. 5, 2012, pp. 53–55.
- [21] F. Bellard, "Qemu, a fast and portable dynamic translator," in *Proceedings of the Annual Conference on USENIX Annual Technical Conference*, ser. ATEC '05. Berkeley, CA, USA: USENIX Association, 2005, pp. 41–46.
- [22] C. Macdonell, X. Ke, A. W. Gordon, and P. Lu, "Low-latency, high-bandwidth use cases for nahanni/ivshmem," 2011.
- [23] M. Paolino, "A shared memory zero-copy mechanism for arm vms:vosyshmem," last visited March 2015. [Online]. Available: <http://www.virtualopensystems.com/en/products/vosyshmem-zero-copy/>
- [24] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood, "sel4: Formal verification of an os kernel," in *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles*, ser. SOSP '09. New York, NY, USA: ACM, 2009, pp. 207–220. [Online]. Available: <http://doi.acm.org/10.1145/1629575.1629596>
- [25] "Optee," last visited March 2015. [Online]. Available: <https://github.com/OP-TEE>
- [26] "Osv," last visited March 2015. [Online]. Available: <http://www.osv.io>
- [27] J. Brunel, R. Pacalet, S. Ouaraab, and G. Duc, "Secbus, a software/hardware architecture for securing external memories," in *Mobile Cloud Computing, Services, and Engineering (MobileCloud)*, 2014 2nd IEEE International Conference on, April 2014, pp. 277–282.
- [28] M. D. Grammatikakis, "System-level modeling of a noc firewall on spidergon stnoc," May 2014, last visited March 2015. [Online]. Available: http://www.cspforum.eu/uploads/Csp2014Presentations/Track_8/System-Level%20Modeling%20of%20a%20NoC%20Firewall%20on%20Spidergon%20STNoC.pdf
- [29] "Arm trusted firmware," last visited March 2015. [Online]. Available: <https://github.com/ARM-software/arm-trusted-firmware>
- [30] E. Keller, J. Szefer, J. Rexford, and R. B. Lee, "Nohype: Virtualized cloud infrastructure without the virtualization," *SIGARCH Comput. Archit. News*, vol. 38, no. 3, Jun. 2010, pp. 350–361. [Online]. Available: <http://doi.acm.org/10.1145/1816038.1816010>
- [31] J. Szefer, E. Keller, R. B. Lee, and J. Rexford, "Eliminating the hypervisor attack surface for a more secure cloud," in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 401–412. [Online]. Available: <http://doi.acm.org/10.1145/2046707.2046754>
- [32] A. Seshadri, M. Luk, N. Qu, and A. Perrig, "Secvisor: A tiny hypervisor to provide lifetime kernel code integrity for commodity oses," *SIGOPS Oper. Syst. Rev.*, vol. 41, no. 6, Oct. 2007, pp. 335–350. [Online]. Available: <http://doi.acm.org/10.1145/1323293.1294294>
- [33] A. M. Azab, P. Ning, Z. Wang, X. Jiang, X. Zhang, and N. C. Skalsky, "Hypersentry: Enabling stealthy in-context measurement of hypervisor integrity," in *Proceedings of the 17th ACM Conference on Computer and Communications Security*, ser. CCS '10. New York, NY, USA: ACM, 2010, pp. 38–49. [Online]. Available: <http://doi.acm.org/10.1145/1866307.1866313>
- [34] X. Ge, H. Vijayakumar, and T. Jaeger, "Sprobes: Enforcing kernel code integrity on the trustzone architecture," *CoRR*, vol. abs/1410.7747, 2014. [Online]. Available: <http://arxiv.org/abs/1410.7747>
- [35] R. Perez and et al., "vtpm: virtualizing the trusted platform module," in *Proc. 15th Conf. on USENIX Security Symposium*, 2006, pp. 305–320.
- [36] R. e. a. Sailer, "shype: Secure hypervisor approach to trusted virtualized systems," 2005.
- [37] H. Nahari, "Trusted secure embedded linux," in *Proceedings of 2007 Linux Symposium*. Citeseer, 2007, pp. 79–85.
- [38] T. K. Lengyel, T. Kittel, J. Pfohl, and C. Eckert, "Multi-tiered security architecture for arm via the virtualization and security extensions," 2014, pp. 308–312. [Online]. Available: <http://dx.doi.org/10.1109/DEXA.2014.68>