



Femto-Containers: DevOps on Microcontrollers with Lightweight Virtualization & Isolation for IoT Software Modules

Koen Zandberg, Emmanuel Baccelli

► To cite this version:

Koen Zandberg, Emmanuel Baccelli. Femto-Containers: DevOps on Microcontrollers with Lightweight Virtualization & Isolation for IoT Software Modules. 2021. hal-03263164v2

HAL Id: hal-03263164

<https://inria.hal.science/hal-03263164v2>

Preprint submitted on 3 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Femto-Containers: DevOps on Microcontrollers with Lightweight Virtualization & Isolation for IoT Software Modules

Koen Zandberg
Inria, France

Emmanuel Baccelli
Inria, France
Freie Universität Berlin, Germany

ABSTRACT

Development, deployment and maintenance of networked software has been revolutionized by DevOps, which have become essential to boost system software quality and to enable agile evolution. Meanwhile the Internet of Things (IoT) connects more and more devices which are not covered by DevOps tools: low-power, microcontroller-based devices. In this paper, we contribute to bridge this gap by designing Femto-Containers, a new architecture which enables containerization, virtualization and secure deployment of software modules embedded on microcontrollers over low-power networks. As proof-of-concept, we implemented and evaluated Femto-Containers on popular microcontroller architectures (Arm Cortex-M, ESP32 and RISC-V), using eBPF virtualization, and RIOT, a common operating system in this space. We show that Femto-Containers can virtualize and isolate multiple software modules, executed concurrently, with very small memory footprint overhead (below 10%) and very small startup time (tens of microseconds) compared to native code execution. We show that Femto-Containers can satisfy the constraints of both low-level debug logic inserted in a hot code path, and high-level business logic coded in a variety of common programming languages. Compared to prior work, Femto-Containers thus offer an attractive trade-off in terms of memory footprint, energy consumption, agility and security.

1 INTRODUCTION

An estimated 250 billion microcontrollers are in use today [1]. An increasing percentage of these microcontrollers are networked and take part in distributed cyber-physical systems and the Internet of Things (IoT) we increasingly depend upon.

With the availability of low-power operating systems [2] and network stacks (e.g. 6LoWPAN), low-power IoT software has made giant leaps forward ; but fundamental gaps remain compared to current practices for networked software. In fact, current state-of-the-art for managing, programming, and maintaining fleets of low-power IoT devices resembles more PC system software workflow from the 90s than today's common software practices: simplistic application programming interfaces (APIs) offer basic performance and connectivity, but no additional comfort.

However, since the 90s, networked software was revolutionized many times over. Networked software has entered the age of agility. DevOps [3] drastically shortened software development/deployment life cycles to provide continuous delivery of higher software quality. Additional layers providing

cybersecurity, flexibility and scalability thus became crucial : ubiquitous script programming (e.g., Python, Javascript), light-weight software containerization (e.g., Docker), hypervisors and software virtualization, deployment and management tools for swarms of virtualized software instances (e.g., Kubernetes or AWS), and frameworks for decentralizing system software updates, development and maintenance on platforms such as Linux, Android, iOS, Windows etc.

In such a context, low-power IoT devices are the new 'weakest link' within distributed cyber-physical systems upon which relies IoT and related services. Indeed, state-of-the-art DevOps mechanisms are not commonly applicable on low-power devices : they are either not applicable on microcontrollers (e.g., Docker), too prohibitive in terms of hosting engine memory resource requirements (e.g., OS virtualization or standard Java virtual machines), or restricted to very specific use cases (e.g., JavaCard). This lackluster create bottlenecks which severely impact both flexibility and cybersecurity in IoT.

A key question emerges: can we provide new concepts adequate for software containerization and rapid deployment on swarms of IoT devices, combining agility, low-power consumption and cybersecurity? The goal we pursue in this paper is to explore practical solutions for software virtualization, containerization and deployment applicable to fleets of low-power, connected, microcontroller-based IoT devices.

Contributions – In this paper, the work we present mainly consists in the following:

- we survey existing techniques for process isolation & virtualization for microcontrollers;
- we design Femto-Containers, a novel secure DevOps architecture for constrained IoT devices;
- we implement Femto-Containers based on eBPF virtualization, and small containers hosted in a common low-power IoT operating system (RIOT);
- we evaluate the performance of Femto-Containers in a variety of use cases, on popular microcontroller architectures (Arm Cortex-M, ESP32 and RISC-V);
- we compare Femto-Containers with prior work based on WebAssembly, JavaScript (RIOTjs) and Python (microPython). We show Femto-Containers offer an attractive trade-off in terms of memory footprint, energy consumption, and security.

2 SCENARIOS & THREAT MODEL

As low-power embedded IoT software complexifies on various devices, it becomes necessary for security (and sometimes also privacy) reasons to delegate maintenance and updates of

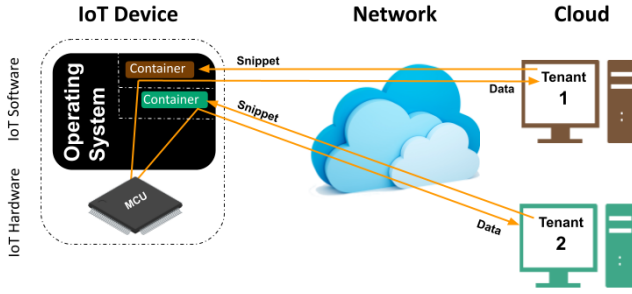


Figure 1: DevOps use-cases on IoT microcontrollers.

different parts of the embedded software to distinct entities with limited mutual trust (as described in [4] for instance). Furthermore, enabling on-the-fly and safe remote instrumentation of already-deployed software is essential to maintain cyberphysical systems involving low-power IoT devices. We thus consider different categories of use-cases, depicted in Figure 1:

- (1) Use-case 1: Hosting and isolating some high-level business logic, updatable on-demand remotely over the low-power network. The execution of this type of logic is typically periodic in nature, and has loose (non-real-time) timing requirements.
- (2) Use-case 2: Hosting and mutually isolating several virtual machines, managed by several different tenants.
- (3) Use-case 3: Hosting and isolating some debug/monitoring code applications at low-level, inserted and removed on-demand, remotely, over the network. Comparatively, this type of logic is short-lived and exhibits stricter timing requirements.

Threat Model. We consider both malicious tenants which can deploy malicious containers and malicious clients which can maliciously interact with deployed containers.

Malicious Tenant: While a tenant has to work within the permissions granted by the host system, it can make free use of the granted resources. A tenant could provide a vulnerable or malicious application code for execution inside a container. To protect against this, the host system must ensure that hosted applications are constrained in the resources they can access and have a fair share of the processing time and network bandwidth.

Malicious Client: The malicious client makes use of a vulnerable tenant application. The client can send requests to networked applications, including arbitrary packets. Assuming a vulnerable application, the client can access any resource accessible by the tenant application. While the protection of already vulnerable applications is considered out of scope, the hosting engine must isolate other tenants' memory access from a compromised application.

3 ARCHITECTURAL DESIGN

The standard system architecture enabling modern DevOps at scale in the cloud is depicted in Figure 2.2. A provider-controlled operating system hosts one or more (typically many) virtual machines sharing the same hardware resources. Each virtual machine (VM) virtualizes an OS maintained by a third-party. Each VM can host one or more (typically many) containers. This architecture is favored because it naturally provides the below properties, which are crucial in terms of code mobility and cyber-security:

- (1) **Hardware abstraction:** the VM abstracts the hardware on top of which software is running, facilitating code portability across different hardware architectures and configurations.
- (2) **OS abstraction:** the container further enhances code mobility by offering standardized access to OS services, which facilitates code portability across different OS.
- (3) **Isolation:** the VM provides a natural defence line. If an attacker breaks out of a container, it only compromises the one VM it is in, other VMs and the host system are safe.

However, the standard "cloud-native" DevOps architecture depicted Figure 2.2 faces issues on low-power IoT devices with constrained resources:

- Full OS virtualization leads to a prohibitive toll on resources and/or execution speed;
- Standard containers solutions are not applicable on constrained IoT devices.

For these reasons, we introduce Femto-Containers, an alternative architecture targeting constrained IoT devices, as described in the following.

Femto-Container Architecture

Instead of using traditional containers inside large VMs which virtualize a full-blown OS, we host smaller VMs inside simplified containers, running on top of a real-time operating system (RTOS), as depicted in Figure 2.1. By flipping around virtualization and containerization, we can retain the crucial properties w.r.t. code mobility and cyber-security (as we still combine isolation, hardware/OS abstraction) but we are able to drastically reduce the scope of virtualization and its cost on constrained IoT devices.

Compared to the large scale cloud architecture, the Femto-Container architecture trades features for a more lightweight approach suitable for small embedded devices. The Femto-Container architecture relies on a number of RTOS-provided features and a set of assumptions, listed below.

Use of an RTOS with Multi-Threading. It is assumed that the RTOS supports real-time multi-threading with a scheduler. Each Femto-Container runs in a separate thread. Well-known operating systems in this space can provide for that, such as RIOT [5] or FreeRTOS [6] and others. These can run on the bulk of commodity microcontroller hardware available. Note that RTOS facilities for scheduling enable simple controlling

of how Femto-Containers interfere with other tasks in the embedded system.

Use of Simple Containerization. A slim environment around the virtual machine (VM) exposes RTOS facilities to the VM. The container sandboxing a VM allows this VM to be independent of the underlying operating system, and provide the facilities as a generic interface to the VM. Simple contracts between container and RTOS can be used to define and limit the privileges of a container regarding its access to OS facilities. Note that such limitations must be enforced at run-time to safely allow 3rd party module reprogramming.

Use of Light-weight Virtualization. The virtual machine provides hardware agnosticism, and should therefore not rely on any specific hardware features or peripherals. This allows for running identical application code on heterogeneous hardware platforms. The virtual machine must have a low memory footprint, both in Flash and in RAM, per VM. This allows to run multiple VMs in parallel on the device. Note that, since we aim to virtualize less functionalities, the VM can in fact implement a reduced feature set. For instance, virtualized peripherals such as an interrupt controller are not required, and we give up the possibility of virtualizing a full OS.

Isolation & Sandboxing through Virtualization. The OS and Femto-Containers must be mutually protected from malicious code, as described in section 2. This implies in particular that code running in the VM must not be able to access memory regions outside of what is allowed. Here again, simple contracts can be used to define and limit memory and peripheral access of the code running in the Femto-Container.

Event-based Launchpad Execution Model. Femto-containers are executed on-demand, when an event in the RTOS context calls for it. Femto-container applications are rather short-lived and have a finite execution constraint. This execution model fits well with the characteristics of most low-power IoT software. To simplify containerization and security-by-design, we mandate that Femto-Containers can only be attached and launch from predetermined launch pads, which are sprinkled throughout the RTOS firmware. Where applicable however, the result from the Femto-Container execution can modify the control flow in the firmware as defined in the launch pad.

In the following, with a view to implement the Femto-Container architecture, we explore building blocks alternatives.

4 SOFTWARE-BASED TECHNIQUES FOR PROCESS ISOLATION & VIRTUALIZATION

As seen in Section section 3, process isolation and virtualisation is a critical building block on which the femto-container architecture depends. Different approaches are possible in order to virtualize and/or isolate software modules running on a microcontroller.

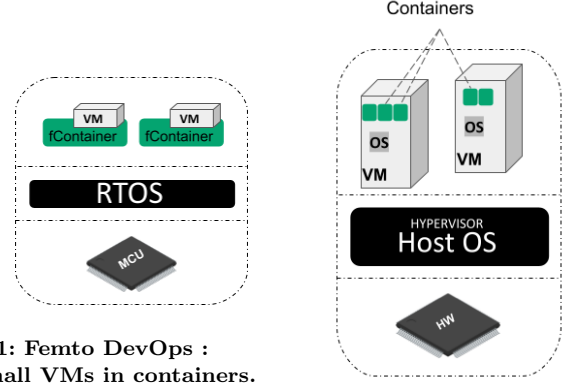


fig1: Femto DevOps : Small VMs in containers.

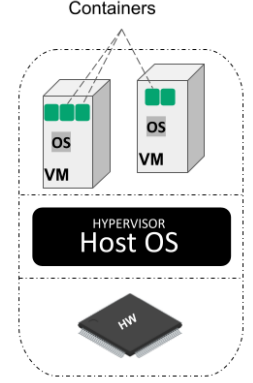


fig2: Cloud-native DevOps : Containers in large VMs.

Figure 2: DevOps architecture: Femto *vs* Cloud.

Software-based Online Approach

One type of approach to isolate software modules is to modify the embedded hardware architecture. Prominent examples of this trend include TrustZone modifying the Arm Cortex-M architecture [7], Sanctum modifying the RISC-V architecture [8], or Sancus complementing the MSP430 architecture [9, 10]. However, such *hardware-based* approaches are by nature specific to each hardware architecture. To retain more general applicability, including on legacy IoT hardware, in this paper we aim instead for a *software-based* approach which does not require specific hardware-based memory protection mechanism.

Taking a different angle, software security guarantees (such as process isolation, or functional correctness) can be determined with *offline* techniques such as formal verification, which can prove properties on software that is known a priori, before it is actually deployed and running. In this paper, we do not pursue this type of approach. We instead focus on *online* techniques, which enforce checks and guarantees on-the-fly, on previously unknown code which is deployed (and runs tentatively). The reason for this choice is mainly that our use cases include scenarios where distinct entities indepently update different software modules running on the same IoT device.

Nevertheless, note that our approach does not preclude the complementary use of hardware-based mechanisms and/or offline mechanisms.

4.1 Survey of Software-based Techniques

Different categories of light-weight, software-based techniques for process isolation and virtualization have been developed in prior work.

Virtual machines. One category of techniques consists in small virtual machine, used to host and isolate a process from other processes running on the microcontroller. One example is WebAssembly (Wasm [11]), a VM specification

with a stack-based architecture, designed for process isolation in Web browsers, which has recently been ported to microcontrollers [12]. Another recent example is Velox [13] a VM able to host and isolate high-level functional programming logic on microcontrollers. On the other hand, Darjeeling [14] is a subset of the Java VM, modified to use a 16 bit architecture, designed for 8- and 16-bit microcontrollers. In fact, beyond the low-power IoT domain, small Java VMs have also been used in other contexts for a long time. For instance JavaCard [15] uses a small Java VM running on smart cards.

A different VM approach comes from the Linux ecosystem, based on eBPF (extended Berkeley Packet Filter [16, 17]), which enables small VM hosting and isolating, for debug and inspection code inserted in the Linux kernel at run-time. Recently, a preliminary prototype adapting an eBPF virtual machine hosted on a low-power microcontroller was developed in [18].

Scripted logic containers. Yet another type of approach uses scripted logic interpreters to virtualize and/or isolate some processes. For instance, MicroPython [19] is a very popular scripted logic interpreter used on microcontrollers, offering partial Python scripting support. Another popular scripted logic interpreter is JerryScript, which offers full ECMA5.1 scripting support. Prior work such as RIOTjs [20] provides a small JavaScript run-time container, which can host (updatable) business logic interpreted on-board a microcontroller, using JerryScript glued atop a real-time OS (RIOT). However, neither MicroPython nor RIOTjs/JerryScript provide specific container isolation guarantees. Complementary mechanisms can however guarantee mutual isolation between scripts. For instance, the SecureJS [21], provides a Javascript-to-Javascript compiler (used on top of JerryScript). However SecureJS does not target low-power microcontrollers specifically.

OS-level mechanisms. Yet another category of solution uses OS-level mechanisms for process isolation. For instance, Tock [22] is an OS written in the Rust programming language, which offers strong isolation between its kernel and application logic processes. However, Tock requires that microcontroller hardware provides a memory protection unit (MPU). More distant related work can also be found in the domain of network function virtualization. Compact kernels such as EdgeOS [23] provide light-weight instance spinning and isolated execution mechanisms. However such kernels are designed for high-throughput middleboxes (which are Linux-capable) instead of low-throughput, low-power microcontrollers.

4.2 Candidate Techniques pre-Selection

We next aim to provide a reality check gauging the potential of the different categories of approaches. For this, we pre-select and compare a representative subset of the existing solutions we identified.

In the *virtual machines* category, we selected WebAssembly [12] and rBPF [18]. Our reasoning motivating our choices here is that WebAssembly is a well-known solution for strong,

generic software module isolation, while rBPF promises very small memory footprint according to preliminary prior work [18].

In the *scripted logic containers* category, we selected MicroPython [19] and RIOTjs [20], which are good representatives of prominent high-level scripting languages on microcontrollers: Python and JavaScript, respectively. We also remark that the performance of RIOTjs is an upper bound for that of SecureJS (since in essence SecureJS adds a layer on top of RIOTjs).

We however chose not to further pursue techniques using OS-specific mechanisms, because we aim to retain generic applicability to multiple OS in this space. We now overview the essential aspects of each preselected candidate technique, before proceeding to comparative benchmarks in the next section.

4.2.1 Web Assembly.

WebAssembly (Wasm [11]) is standardized by the World Wide Web Consortium (W3C). Initially aimed at portable web applications, Wasm has been adapted to microcontrollers.

Architecture. Wasm is a virtual instruction set architecture (ISA) with flexible instruction size. This ISA allows for small binary size – decreasing the time needed to transport logic over the network, and necessary memory footprint on the IoT device. The WebAssembly VM is stack-based, using both a stack and a flat heap for memory storage. While heap and stack sizes are flexible, Wasm specifications mandate memory allocations in chunks of 64 KiB (pages).

DevOps Toolchain. Wasm uses the LLVM compiler: Wasm applications code can be written in any language supported by LLVM such as C, C++, Rust, TinyGo, or D, among others. The Wasm code development and execution workflow is shown in Figure 3. Note that for C and C++, the WebAssembly binaries are created using the emcc toolchain, which combines the EmSDK with LLVM. Furthermore, a POSIX-like interface is specified for host OS access, called WASI [24]. WASI standardizes access to operating system facilities such as files, network sockets, clocks and random number etc.

Interpreter. Once the Wasm binary is compiled with LLVM, the resulting bytecode can be transferred to the IoT device, on which it is interpreted and executed, as shown in Figure 3. Several interpreters have been developed. In this paper we use the WASM3 [12] interpreter. WASM3 is based on a two-stage approach: in a first phase, the loaded application is transpiled to an executable, then in a second phase, it is executed in the interpreter.

Security & Isolation. The sandbox provided by Wasm offers strong security guarantees on memory access. The memory space accessible by the virtual machine is a virtual space mapping different real memory regions. This prevents hazardous access to the host memory.

4.2.2 eBPF.

Extended Berkeley Packet Filter (eBPF [17]) is a small in-kernel VM stemming from the Linux ecosystem, compatible

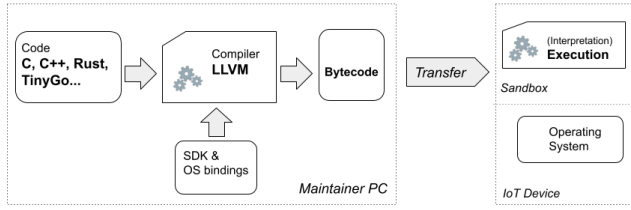


Figure 3: VM toolchain.

with Unix-like operating systems. eBPF provides a tiny facility able to run custom VM code, inside the kernel, hooking into various subsystems.

Architecture. eBPF is 64-bit register-based VM, using fixed-size 64bit instructions and a small ISA. eBPF uses a fixed-sized stack (512B) and no heap, which limits VM memory overhead in RAM. As a replacement for a heap, a key-value store is used for storage between invocations. Very recently, the eBPF instruction set has been ported to microcontrollers [18].

DevOps Toolchain. The toolchain and workflow with eBPF is akin to that of Wasm shown in Figure 3. In particular, eBPF also uses the LLVM compiler to produce bytecode, and supports WM logic written in any language supported by LLVM (C, C++, Rust, TinyGo, D...).

Interpreter. Contrary to Wasm, the bytecode produce by LLVM does not need a preliminary phase to transpile, and can directly be executed by the eBPF interpreter, on the IoT device. Several interpreters have been developed in prior work. In this paper, we specifically consider the rBPF interpreter [18].

Security & Isolation. The sandbox provided by rBPF offers security guarantees on memory access and code execution. All memory access, including to the stack, happen via register load and store instructions and are checked against simple memory access controls. Furthermore, there are limitations on branch and jump instruction targets. The application has no access to the program counter via registers or instructions and a jump is always direct, and relative to current program counter. These characteristics facilitate implementations of the necessary checks at runtime to limit access and execution and thus eliminate this attack surface. rBPF should not be vulnerable against hazardous memory access and code execution. However according to the authors there are no limits on the execution time granted to the application.

4.2.3 RIOTjs.

RIOTjs is a Javascript execution environment, integrated in the RIOT operating system. It executes high-level logic snippets written in the Javascript language, loadable at runtime over the low-power network, on the IoT device.

Architecture. The architecture is based on a two step approach where the Javascript code loaded in the runtime container is first compiled into a compact bytecode format and then interpreted – all of which happens directly on the

IoT device. The bytecode is a CISC-like instruction set with main focus on a compact representation. To achieve this, the instruction set uses single instructions to cover multiple atomic tasks.

Interpreter & DevOps Toolchain. Using a small interpreter provided by the JerryScript engine [25], RIOTjs provides a relatively lightweight VM. JerryScript performs a pre-execution phase which parses the Javascript code. The parser itself is implemented as a recursive descent parser to convert the javascript source into bytecode, without requiring an abstract syntax tree. Thanks to this parsing phase, the toolchain is simplified: the container developer only needs a text editor (which is a major advantage of scripted logic approaches).

Security & Isolation. With RIOTjs, compiled bytecode is executed within a virtual machine. However, RIOTjs is not specifically designed for JavaScript runtime container security and isolation. The hardware-specific mechanisms (e.g. hardware memory protection) or additional layers and complementary mechanisms in software (e.g. SecureJS [21] already mentioned earlier) would be necessary to provide strong security and isolation. guarantees.

An alternative which one can envision is to offload parsing and bytecode generation to the maintainer PC. However, in this case, an additional layer providing mechanisms that check bytecode correctness would be required on the IoT device.

Hence, the benchmarks results we provide in the next section are to be considered as an upper bound on what can be achieved with this type of approach.

4.2.4 MicroPython.

MicroPython [19] is akin to the JerryScript engine, but for Python code. Bare-metal operation of MicroPython is possible on some IoT hardware such as the pyboard. Variants of MicroPython such as CircuitPython [26] can run bare-metal on more hardware. Integration in operating systems is also available (such as in [27], which is similar to RIOTjs, but for MicroPython).

Architecture. MicroPython is based on a two stage approach. In a first phase, based on an abstract syntax tree, a parser/lexer compiles Python code to native bytecode. In a second phase, an interpreter executes the bytecode. Both stages can be performed directly on the IoT device. Within the Python logic hosted in the container, memory management is abstracted away for the developer: it automated with a heap and garbage collector.

Interpreter & DevOps Toolchain. The bytecode interpreter machine implements a stack-based architecture. Thanks to the parsing phase which can be performed by MicroPython, a minimal toolchain is possible: the container developer only requires a text editor.

Security & Isolation. MicroPython is not specifically designed for Python runtime container security and isolation. Additional layers and complementary mechanisms would be necessary to provide strong security and isolation guarantees.

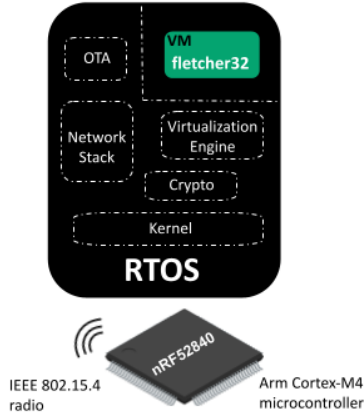


Figure 4: Benchmark setup. VM executing Fletcher32 logic, hosted in a typical RTOS configuration, on an nRF52840 microcontroller.

Hence, the benchmarks results we provide in the next section are to be considered as an upper bound on what can be achieved with this type of approach.

5 BENCHMARKS OF PROCESS VIRTUALIZATION & ISOLATION TECHNIQUES

In this section, to get an idea of what to expects in terms of ballpark performance, we carry out benchmarks comparing the candidate virtualization and isolation techniques we pre-selected. Based on these results, we aim to discuss and select an approach upon which to base our Femto-Containers design.

5.1 Hardware & Software Setup

The hardware and software setup is depicted in Fig. 4. For each virtualization candidate, the virtual machine is loaded with an application performing a Fletcher32 checksum on a 360 B input string. All benchmarks are run COTS hardware: a Nordic nrf52840dk development kit, which is based on a common ARM Cortex-M4 processor running at 64 MHz. The operating system hosting the VMs is RIOT [5]. As base, we take RIOT Release 2021.04, configured to be IoT-ready, providing standard low-power networking connectivity, leveraging the board’s IEEE 802.15.4 radio chip and a resource-efficient IPv6-compliant stack (6LoWPAN, UDP, CoAP and SUIT).

5.2 Results & Preliminary Analysis

Our benchmarks results comparing different scripting and virtualization techniques are shown in Table 1 and Table 2. The results highlight how much the footprint of hosting logic in a VM can vary, depending on the virtualization technique being used.

Looking at size. While the size of applications are roughly comparable accross virtualization techniques (see Table 2)

	ROM size	RAM size
WASM3 Interpreter	64 KiB	85 KiB
rBPF Interpreter	4.4 KiB	0.6 KiB
RIOTjs	121 KiB	18 KiB
MicroPython	101 KiB	8.2 KiB
Host OS (without VM)	52.5 KiB	16.3 KiB

Table 1: Memory requirements for VM interpreters.

	code size	startup time	run time
Native C	74 B	–	27 μ s
WASM3	322 B	17 096 μ s	980 μ s
rBPF	456 B	1 μ s	2133 μ s
RIOTjs	593 B	5589 μ s	14 726 μ s
MicroPython	497 B	21 907 μ s	16 325 μ s

Table 2: Size and performance of fletcher32 logic hosted in different VMs.

the memory required on the IoT device differs wildly. In particular, techniques based on script interpreters (RIOTjs and MicroPython) require the biggest dedicated ROM memory budget, above 100 KiB.

For comparison, the biggest ROM budget we measured requires 27 times more memory than the smallest budget. Similarly, RAM requirements vary a lot. Note that we could not determine with absolute precision the lower bound for script interpreters techniques (due to some flexibility given at compile time to set heap size in RAM). Nevertheless, our experiments show that the biggest RAM budget requires 140 times more RAM than the smallest budget. We remark that, as noted in prior work [18] the minimum required page size of 64 KiB to comply with the WebAssembly specification explains why WASM performs poorly in terms of RAM. One can envision enhancements where this requirement is relaxed. However the RAM budget would still be well above an order of magnitude more than the lowest RAM budget we measured (rBPF).

Last but not least, let’s give some perspective by comparison with a typical memory budget for the *whole* software embedded on the IoT device. As a reminder, in the class of devices we consider, a microcontroller memory capacity of 64kB in RAM and 256kB in Flash (ROM) is not uncommon. A typical OS footprint for this type of device is shown in the last row of Table 1. For such targets, according to our measurements, adding a VM can either incur a tremendous increase in total memory requirements (200% more ROM with MicroPython) or a negligible impact (8% more ROM with rBPF) as visualized in Figure 5.

Looking at speed. To no surprise, beyond size overhead, virtualization also has a cost in terms of execution speed. But here again, performance varies wildly depending on the virtualization technique. On one hand, solutions such as MicroPython and RIOTjs directly interpret the code snippet and execute it. On the other hand, solutions such as rBPF

and WASM3 require a compilation step in between to convert from human readable code to machine readable.

Our measurements show that script interpreters incur an enormous penalty in execution speed. Compared to native code execution, script interpreters are a whopping 600 times slower. Compared to the same base (native execution) WASM is only 37 times slower, and rBPF 77 times slower.

One last aspect to consider is the startup time dedicated to preliminary pre-processing when loading new VM logic, before it can be executed (including steps such as code parsing and intermediate translation, various pre-flight checks etc.). Depending on the virtualization technique, this startup time varies almost 1000 fold – from a few microseconds compared to a few milliseconds.

5.3 Discussion

We now aim to reason a choice for an approach to design tiny VMs which efficiently address our target use cases (described in Section 2) which involve hosting and mutually isolating multiple VMs which may contain either high-level business logic, or low-level debug/monitoring code snippets.

Considering architecture & security. There are notable architectural differences amongst the solutions we pre-selected and looked at in this section. For instance, WASM, MicroPython and RIOTjs each require some form of heap on which to allocate application variables. On the other hand, rBPF does not require a heap. With a view to accommodating several VMs concurrently, a heap-based architecture presents on the one hand some potential advantages in terms of memory (pooling) efficiency, but on the other hand some potential drawbacks in terms of security (mutual isolation of the VMs’ memory).

From another angle: security guarantees call for a formally verified implementation of the hosting engine, down the road. A typical approximation is: less LoC (lines of code) means less effort produce a verified implementation. For instance, the rBPF implementation is 1.5k LoC, while the WASM3 implementation is 10k LoC. The other implementations we considered in our pre-selection (RIOTjs and MicroPython) encompass significantly more LoC.

Considering performance. Our benchmarks indicate that both from a memory overhead and from a startup time standpoint an eBPF-based approach is the most attractive, by far. From an execution time point of view however, a WebAssembly approach does offer faster execution times than an eBPF-based approach. We nevertheless deem safe to consider that, for the use cases we target, this difference is negligible. These results both extend and confirm independent results presented in prior work [18]. All in all, both due to much larger memory footprint and enormous execution time penalties, Python and JavaScript approaches could not be considered beyond rapid prototyping – in particular when considering one of our use cases: a virtual machine hosting debug applications in a hot code path.

Ahead-of-Time vs Just-in-Time. One approach to speed up embedded execution time is to perform a translation into device-native code. One way to offload the device is to use more Ahead-of-Time (AoT) compilation/interpretation, and less Just-in-Time (JiT) processing on-device. However, using AoT pre-compiled code can both complicate run-time security checks on-board the IoT device, and reduce the portability of the code deployed on the device. For these reasons, in this paper, we consider primarily JiT.

Intermediate conclusion. Based on our discussion and our benchmarks, we derive the preliminary conclusion that an eBPF-based architecture is a promising approach to design efficient and secure tiny concurrent containers to host and execute logic on a microcontroller. In the following, we thus design, implement and experimentally evaluate a novel virtualization and isolation mechanism for software modules, based on eBPF.

6 FEMTO-CONTAINER IMPLEMENTATION

As proof of concept, we implemented the femto-container architecture, with containers hosted in the operating system RIOT and virtualization using an instruction set compatible with the eBPF instruction set. This implementation is open source (published in [28]). We detail below its main characteristics.

Use of RIOT Multi-Threading

Each Femto-Container application instance running on the RTOS is scheduled as a regular thread in RIOT. The native OS thread scheduling mechanism can thus simply execute concurrently and share resources amongst multiple femto-containers and other tasks, spread over different threads. An overview of how Femto-Containers integrates in the operating system is shown in Figure 6. A femto-container instance requires minimal RAM: a small stack and the register set, but no heap. The host RTOS bears thus a very small overhead per femto-container instance.

The hardware and peripherals available on the device are not accessible by the Femto-Container instances. All interaction with hardware peripherals passes through the host RTOS via the system call interface. As the Femto-Container virtual machine does not virtualize its own set of peripherals, no interrupts or pseudo-hardware is available to the Femto-Container application. This also removes the option to interrupt the application flow inside a Femto-Container.

Basic Containerization

Simplified containers provides a uniform environment around the VM, independent of the operating system (RIOT in our implementation). Access from the femto-container to the required OS facilities is allowed through system calls to services provided by RIOT. These system calls can be used by the loaded applications via the eBPF native `call` instruction. Furthermore, the OS can share specific memory regions with

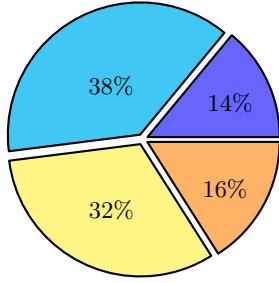


fig1: RIOT without hosting engine(53kBytes).

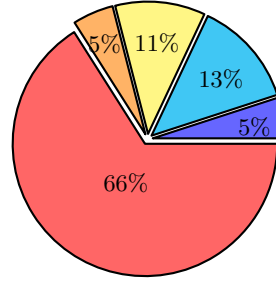


fig2: RIOT with MicroPython VM (154kBytes).

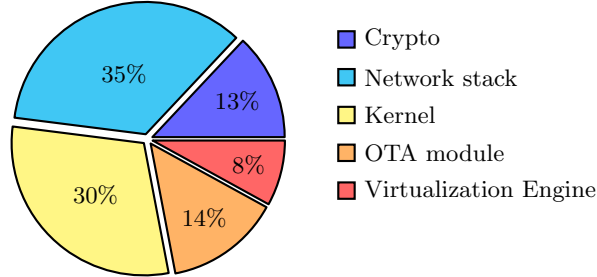


fig3: RIOT with rBPF virtual machine (57kBytes).

Figure 5: Flash memory distribution. RIOT with 6LoWPAN, CoAP, SUIT-compliant OTA and different application hosting engines.

the container.

Key-value store. In lieu of a file system, applications hosted in femto-containers can load and store simple values, by a numerical key reference, in a key-value store. This provides a mechanism for persistent storage, between application invocations. Interaction with this key-value store is implemented via a set of system calls, keeping it independent of the instruction set. By default, two key-value stores are provided by the OS. The first key-value store is local to the application, for values that are private to the VM accommodated in the container. The second key-value store is global, and can be accessed by all applications, used to communicate values between applications. An optional third intermediate-level of key-value store is possible to facilitate sharing data across a set of VMs from the same tenant, while isolating this set of VMs from other tenants' VMs.

Light-weight Virtualisation

Application code is virtualized using the eBPF instruction set and the rBPF interpreter.

Register-based VM. The virtual machine operates on eleven registers of 64 bits wide. The last register (**r10**) is a read-only pointer to the beginning of a 512B stack provided by the femto-container hosting engine. Interaction with the stack happens via load and store instructions. Instructions are divided into an 8 bit opcode, two 4 bit registers: source and destination, an 16 bit offset field and an 32 bit immediate value. Position-independent code is achieved by using the reference in **r10** and the offset field in the instructions.

Jumtable & Interpreter. The interpreter parses instructions and executes them operating on the registers and stack. The machine itself is implemented as a computed jumtable, with the instruction opcodes as keys. During execution, the hosting engine iterates over the instruction opcodes in the application, and jumps directly to the instruction-specific code. This design keeps the interpreter itself small and fast.

Isolation & Sandboxing

To control the capabilities of Femto-Containers, and to protect the OS from memory access by malicious applications, a simple but effective memory protection system is used. By default each virtual machine instance only has access to its VM-specific registers and its stack.

Memory access checks at runtime. Whitelist can be configured (attached in the hosting engine) to explicitly allow a VM instance access to other memory regions. These memory regions can have individual flags for allowing read/write access. For example, a firewall-type trigger can grant read-only access to the network packet, allowing the virtual machine to inspect the packet, but not to modify it. As the memory instructions allow for calculated addresses based on register values, memory accesses are checked at runtime against the resulting address, as show in Figure 7. Illegal access aborts execution.

Pre-flight instruction checks. A femto-container verifies the application before it is executed for the first time. These checks includes checks on the instruction fields. For

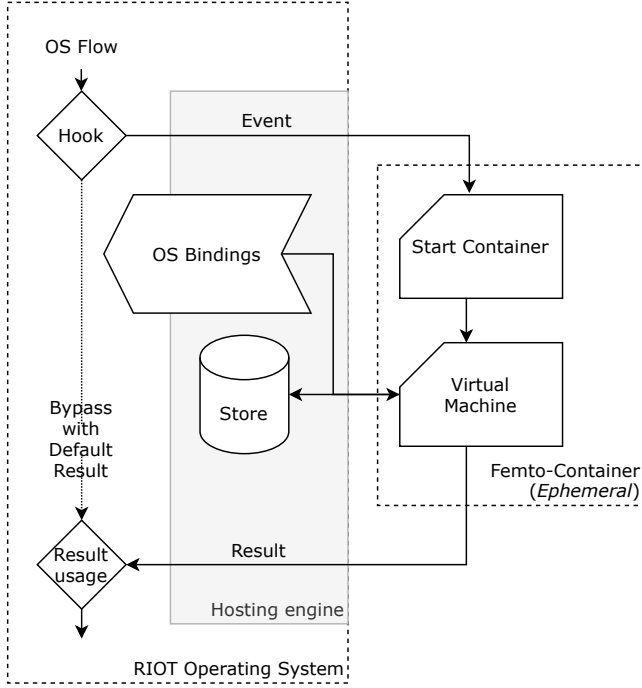


Figure 6: Femto-Container RTOS integration.

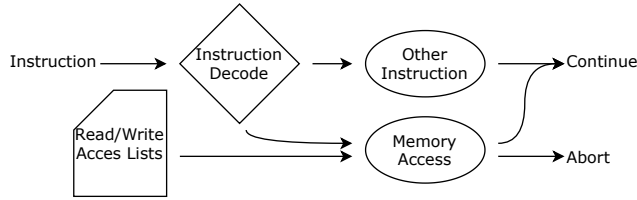


Figure 7: Interaction between memory instructions and the access lists.

example, as there are only 11 registers, but space in the instruction for 16 registers, the register fields must be checked for out-of-bounds values. A special case here is register `r10` which is read-only, and thus is not allowed in the destination field of the instructions.

The jump instructions are also checked to ensure that the destination of the jump is within the address space of the application code. As calculated jump destinations are not supported in the instruction set, the jump targets are known before executions and are checked during the pre-flight checks. During the execution of the application, the jump destinations no longer have to be verified and can be accepted as valid destinations.

Finite execution is also enforced, by limiting both the total number of instructions N_i , and the number of branch instructions N_b that are allowed. In practice, this limits the total number of instructions executed to: $N_i \times N_b$.

Hooks & Event-based Execution

The Femto-container hosting engine instantiates and runs containers as triggered by events within the RTOS. Such events can be a network packet reception, sensor reading input or an operating system scheduling events for instance. Business logic applications can be implemented either by directly responding to sensor input or by attaching to a timer-based hook to fire periodically.

Simple hooks are pre-compiled into the RTOS firmware, providing a pre-determined set of pads from which femto-containers can be attached and launched.

```

sched_ctx_t context = {
    .previous = active_thread,
    .next = next_thread,
};

int64_t result;

fc_hook_execute(BPF_HOOK_SCHED, &context,
               sizeof(context), &result);

```

Listing 1: Example hook implementation.

An example of a hook integrated in the firmware is shown in Listing 1. The firmware has to set up the context struct for the Femto-Containers after which it can call the hosting engine to execute the containers associated with the hook.

Low-power Secure DevOps

Launching a new femto-container or modifying an existing femto container can be done without modifying the RTOS firmware. However, updating the hooks themselves requires a firmware update. In our implementation, both types of updates use CoAP network transfer and software update metadata defined by SUIT (CBOR, COSE) to secure updates end-to-end over network paths including low-power wireless segments. Updating a Femto-Container application attached to a hook is done via a SUIT manifest. The exact hook to attach the new Femto-Container to is done by specifying the hook as unique identifier (UUID) as storage location in the SUIT manifest. A rapid develop-and-deploy cycle only requires a new SUIT manifest with the storage location specified every update. Sending this manifest to the device triggers the update of the hook after the new Femto-Container application is downloaded to the device and stored in the RAM. Multiple Femto-Containers can be deployed on a single hook, where it depends on the hook implementation how conflicting return values from the different Femto-Containers are used. For example, a timer-based hook for periodic container execution support attaching multiple Femto-Containers from different tenants. This hook then periodically executes all containers attached without using any of the return values.

7 USE-CASE PROTOTYPING WITH FEMTO-CONTAINERS

In this section, we use Femto-Containers to prototype the implementation of several use cases involving one or more applications, hosted concurrently on a microcontroller, matching

targets we identified initially (in section 2). In the prototype implementation we show below, we used C to code logic hosted in Femto-Containers. However, any other language compiled with LLVM could be used instead (C++, Rust, TinyGo, D...).

7.1 Kernel Debug Code Example

The first prototype consists in a single application, which intervenes on a hot code path: it is invoked by the scheduler of the OS, to keep an updated count of threads' activations. The logic hosted in the Femto-Container is shown in Listing 2. A small struct is passed as context, which contains the previous running thread ID and the next running thread ID. The application maintains a value for every thread, incrementing it every time the thread is scheduled. External code can request these counters and provide debug feedback to the developer.

```
#include <stdint.h>
#include "bpf/bpfapi/helpers.h"

#define THREAD_START_KEY 0x0

typedef struct {
    uint64_t previous; /* previous thread */
    uint64_t next;    /* next thread */
} sched_ctx_t;

int pid_log(sched_ctx_t *ctx)
{
    /* Zero pid means no next thread */
    if (ctx->next != 0) {
        uint32_t counter;
        uint32_t thread_key = THREAD_START_KEY +
            ctx->next;
        bpf_fetch_global(thread_key,
            &counter);
        counter++;
        bpf_store_global(thread_key,
            counter);
    }
    return 0;
}
```

Listing 2: Thread counter code.

7.2 Networked Sensor Code Example

For the second prototype we add two Femto-Containers from another tenant to the setup of the first prototype. Interaction between these two additional containers is achieved via a separate key-value store, as depicted in Figure 8. The logic hosted in the first Femto-Container, periodically triggered by the timer event, reads, processes and stores a sensor value. The code for this logic is shown in [29]. The second container's logic is triggered upon receiving a network packet (CoAP request), and returns the stored sensor value back to the requestor. The code for this logic is shown in [30].

In this toy example, the sensor value processing is a simple moving average, but more complex post-processing is possible instead, such as differential privacy or some federated learning logic, for instance. This example sketches both how multiple tenants can be accommodated, and how separating the concerns between different containers is possible (between

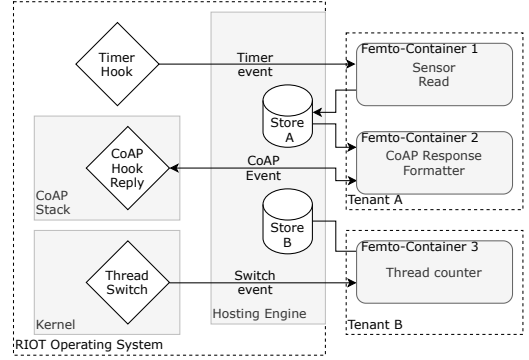


Figure 8: Event and value flow when hosting multiple containers for different tenants.

sensor value reading/processing on the one hand, and on the other hand the communication between the device and a remote requestor).

8 PERFORMANCE EVALUATION

In this section we evaluate the performance of Femto-Containers on low-power IoT hardware. We use RIOT Release 2021.04 as a base for our benchmarks.

8.1 Hardware Testbed Setup

We carry out our measurements on popular, commercial, off-the-shelf IoT hardware, representative of the landscape of modern 32-bit microcontroller architecture that are available: Arm Cortex-M, ESP32, and RISC-V. More precisely, we build and run the code on the following boards:

- a Nordic nRF52840 Development Kit, using an Arm Cortex-M4 microcontroller with 256 KiB RAM, 1 MiB Flash, and a 2.4 GHz radio transceiver (BLE/802.15.4)
- a WROOM-32 board, using an ESP32 module which provides two low-power Xtensa® 32-bit LX6 microprocessors with integrated Wi-Fi and Bluetooth, 520 KiB RAM, 448 KiB ROM and 16 kB RTC SRAM.
- a Sipeed Longan Nano GD32VF103CBT6 Development Board, which provides a RISC-V 32-bit microcontroller with 32 KiB RAM and 128 KiB Flash.

An open-access testbed such as IoT-Lab [31] also provides some of this hardware, for reproducibility.

8.2 Femto-Container Engine Code Analysis

The Femto-Containers hosting engine code size is small: 1874 lines of code in total. This includes bindings to the operating system facilities. Compared to the rBPF hosting engine for example (1615 lines of code) this represents a relatively modest increase (15%) which remains in the same ballpark.

The in-memory structures required to run Femto-Containers are also small. There are two important structures used to manage Femto-Containers. The first structure contains the full state of the virtual machine and any flags required to

manage the VM. This structure requires 664 B in total and includes the stack for the VM instance. The second structure is a small 16 B structure used to whitelist different memory regions for additional access for the VM. The virtual machine state structure already includes such a memory whitelist structure to grant access to the stack space of the VM.

8.3 Experiments with a Single Container

In this section, we evaluate the footprint and the speed of execution with a Femto-Container, on various 32-bit microcontrollers. Again, we compare to the performance of using a basic rBPF virtual machine. Femto-Containers proceeds to security checks on the application bytecode, prior to actually launching the VM. First, this verification stage checks that the registers in all instructions are within the bounds of the eleven available registers, where the source address must be one of these registers and the destination must be one of register `r0` to `r9`, as register `r10` is read-only per specification. Second, the verification checks the destination of the branch-type instructions. Compared to rBPF which performs more limited checks, for instance, the Femto-Container engine increases security measures. We evaluate the impact of these pre-flight checks on speed in Table 3. While still much faster than alternative virtualization techniques (recall Table 2), Femto-Containers startup time is significantly longer than those in rBPF, across the board. Note however that pre-flight checks are run only the first time the VM runs, and are skipped onwards, after the second time the application is run. The increased security measures (in the pre-flight checks) also show when looking at the memory footprint of the Femto-Containers, reported in Table 4. Compared to rBPF the required ROM for the hosting engine increased slightly (less than 10%). The required RAM per virtual machine also increased slightly because of the additional security features. In Table 5, we also measure the RAM memory required to run the kernel debug application example (described in subsection 7.1) which is also small, at 700 Bytes. Last but not least, we observe in Table 3 that the execution time with a Femto-Container is slightly faster than execution time with rBPF virtual machine.

	rBPF		Femto-Containers	
	Install	Run	Install	Run
Cortex-M4	1 μ s	2133 μ s	28 μ s	2061 μ s
ESP32	15 μ s	1817 μ s	66 μ s	1770 μ s
RISC-V	1 μ s	1248 μ s	17 μ s	1238 μ s

Table 3: Speed of a Femto-Container (hosting fletcher32 logic).

8.4 Experiments with Multiple Containers

We now measure in Table 5 the memory required to concurrently host multiple containers from multiple tenants on the same microcontroller, from the examples we described in section 7. In general, each Femto-Container needs memory to

	ROM size	RAM size
Femto-Container	4742 B	664 B
rBPF Interpreter	4440 B	660 B

Table 4: Memory footprint of a Femto-Container hosting fletcher32 logic (on Arm Cortex-M4).

- store the application bytecode
- handle the virtual machine state and stack

This minimal default memory footprint used by a Femto-Container amounts to 664 B, which is for storing the VM stack, housekeeping structs, information about memory regions etc. Note that on top of this basic footprint, a small additional overhead is necessary to code memory permissions. For example, the CoAP handler container (see Figure 8) requires additional read/write permissions to two memory regions to handle the CoAP packet, which increases the overhead by 16 B per region. Furthermore, the key-value stores are also in RAM. In this case the total RAM used by the key value stores (and housekeeping) for different tenants was 340 B. Hence, the required RAM memory we measured so as to run the example with 3 containers and 2 tenants is 3.2 KiB.

Beyond these examples, if we consider more containers hosting larger applications (e.g., ≈ 2000 Bytes) an Arm Cortex-M4 microcontroller with 256 KiB RAM, the density of containers achievable would be of ≈ 100 instances, next to running the OS. Femto-Containers thus allow an almost arbitrarily high density of VMs, even on small microcontrollers.

	Bytecode	Container RAM	Total RAM
Thread Counter	104 B	664 B	768 B
Sensor Reader	496 B	664 B	1160 B
CoAP Handler	264 B	696 B	960 B

Table 5: RAM required to host multiple concurrent Femto-Containers applications.

8.5 Overhead Added by Hooks

One key question is how performance is affected by preprovisioning launchpads (hooks) in the RTOS firmware. We measure in Table 6 the overhead caused by adding a hook to the RTOS workflow. This overhead amounts to ≈ 100 clock ticks on all the hardware we tested. Compared to the number of cycles needed for an average task in the operating system, this impact is low. Furthermore, this overhead is less than 10% of the number of cycles needed to execute the logic hosted in a Femto-Container. From this observation, we can conclude that, even if this hook is on a very hot code path (as for the Thread Counter example) the performance loss is tolerable. Conversely, the perspective of adding many hooks sprinkled in many places in the RTOS firmware is realistic without incurring significant performance loss.

	Empty Hook	Hook with Application
Cortex-M4	109	1750
ESP32	83	1163
RISC-V	106	754

Table 6: Hook overhead in clock ticks for the thread switch example

9 DISCUSSION

Virtualization vs Power-Efficiency. Inherently, virtualization causes some execution overhead, due to interpretation of the code. Thus Femto-Containers increase power consumption for functionality executed within the VM, compared to native code execution. However, this drawback is mitigated by several other factors. First, the absolute power consumption overhead may be negligible, e.g. if the hosted logic is not performing long-lasting, heavy-duty tasks. Second, network transfer costs, power consumption and downtime are saved if software updates modify a Femto-Container instead of the full firmware.

Controlling Tenant Privileges. Controlling and granting access to specific RTOS resources to different containers or tenants is a complex challenge. Our design includes a basic permission system based on preprovisioned hooks, system calls, and simple contracts between the hosting engine (on behalf of the OS) and a given container. Basically: the OS restricts the set of privileges that can be granted, the container specifies the set of privileges it requires, and the hosting engine grants the intersection of these sets. One limitation of our current simplified design is that there is only one fixed set of privileges possible per hook. In case 2 tenants have different privileges, a second hook must be made available. Additional mechanisms would be required to overcome this limitation and/or to enable dynamic privilege levels.

Install Time vs Execution Time. As mentioned before, one limitation due to virtualization is the inherent slump in execution speed, compared to native code execution. One way to remove this overhead is to transpile the portable eBPF bytecode into native instruction code. This could be done in a single pass to convert the whole application into native instructions in an installation step. This can result into a speed-up at the cost of extra install-time overhead. To avoid the issues described before on complicating the run-time security checks, this compilation into native code has to be done at run-time by the device deploying the code.

Fixed- vs Variable-length Instructions. Originally, eBPF scripts are optimized for fast execution on 64-bit platforms. Compared to other virtual machines such as Wasm, the resulting bytecode is relatively large. In fact, most of the instructions have bit fields that are fixed at zero. A possible way to reduce the size of these scripts is to compress the instructions into a variable size instruction set, removing these fields from the instructions where possible. This would create a variable

length instruction set based on the eBPF set. For example the immediate field is not used with half of the instructions and would reduce the instructions to 32 bits in size when removed.

Formal Verification Perspectives. With such a small size in terms of LoC, there is a clear opportunity to go further and attempt at producing a formally verified implementation of the Femto-Containers hosting engine, providing proof that the memory of the host system and of other tenants are protected against malicious tenants and clients. So far, we have been using fuzzing tools on the application loaded into the virtual machine in an attempt to find bugs and vulnerabilities in the hosting engine. However, this does not ensure that all opcodes are executed according to their specification, it only checks whether instructions are able to crash the running system.

rBPF Logic & Execution Limitations. Inherent limitations due to the eBPF instruction set, combined with the absence of virtualized hardware, restrict what logic can be deployed in Femto-Containers currently. Femto-Containers are designed to host logic that is rather script-like, short-lived, and not computation-intensive. On the one hand, such characteristics increase security-by-design. On the other hand they reduce flexibility. For instance, asynchronous operation is not supported: there is no option to interrupt the control flow inside a Femto-Container from outside the virtual machine. Another limitation is the fixed, small size of the stack (512 Bytes) dictated by the eBPF specification. More memory-consuming tasks would need special handling to provide additional memory. Allowing the application to request more stack from the RTOS, for example via the contracts, could solve part of this issue. More computation- and memory-intensive tasks could also make use of additional system calls provided by the RTOS, which could execute generic primitives at native speed.

10 RELATED WORK

We have already provided a survey of related work in section 4. To the best of our knowledge, the closest related work is rBPF [18]. Compared to rBPF:

- Femto-Containers are applicable not only to single container use-cases, but also to use cases hosting multiple containers and distinct tenants, concurrently, on the same microcontroller;
- Femto-Containers provide additional containerization, security and isolation mechanisms;
- Femto-Containers improve performance in the single-container case.

Furthermore, on the experimental side, we compare the performance of more diverse containers techniques, on a wider variety of low-power IoT hardware architectures.

11 CONCLUSION

In this paper we have introduced Femto-Containers, a new architecture we designed to enable modern DevOps on fleets

of heterogeneous low-power IoT hardware. Using Femto-Containers, authorized maintainers of IoT device software can manage (via the network) mutually isolated software modules embedded on the same microcontroller-based device. We implemented a Femto-Container hosting engine on a common low-power IoT operating system, porting the eBPF instruction set to RIOT. We demonstrated experimentally its performance, without requiring any specific hardware-based memory isolation mechanism, on the most common 32-bit microcontroller architectures (including Arm Cortex-M, RISC-V, ESP32). While requiring negligible Flash and RAM memory overhead (less than 10%), Femto-Containers improve state-of-the-art containerization, virtualization and eBPF use on IoT microcontrollers, by increasing security, isolation and execution speed. In effect, Femto-Containers enables hosting (tens of) applications executing concurrently, and multiple tenants, on a single low-power IoT device.

REFERENCES

- [1] Huston Collins. Why TinyML is a giant opportunity. <https://venturebeat.com/2020/01/11/why-tinyml-is-a-giant-opportunity/>.
- [2] Oliver Hahm, Emmanuel Baccelli, Hauke Petersen, and Nicolas Tsiftes. Operating Systems for Low-end Devices in the Internet of Things: a Survey. *IEEE Internet of Things Journal*, 3(5):720–734, 2015.
- [3] Len Bass, Ingo Weber, and Liming Zhu. *DevOps: A software architect's perspective*. Addison-Wesley Professional, 2015.
- [4] Ian Thomas, Shinji Kikuchi, Emmanuel Baccelli, Kaspar Schleiser, Joerg Doerr, and Andreas Morgenstern. Design and Implementation of a Platform for Hyperconnected Cyber Physical Systems. volume 3, pages 69–81. Elsevier, 2018.
- [5] Emmanuel Baccelli et al. RIOT: an Open Source Operating System for Low-end Embedded Devices in the IoT. *IEEE Internet of Things Journal*, 2018.
- [6] R. Barry. FreeRTOS, a FREE open source RTOS for small embedded real time systems. <http://www.freertos.org>.
- [7] Sandro Pinto and Nuno Santos. Demystifying Arm TrustZone: A Comprehensive Survey. *ACM Computing Surveys (CSUR)*, 51(6):1–36, 2019.
- [8] Victor Costan, Ilia Lebedev, and Srinivas Devadas. Sanctum: Minimal hardware extensions for strong software isolation. In *25th USENIX Security Symposium*.
- [9] Fritz Alder, Jo Van Bulck, Frank Piessens, and Jan Tobias Mühlberg. Aion: Enabling open systems through strong availability guarantees for enclaves, 2021.
- [10] Job Noorman et al. Sancus 2.0: A low-cost security architecture for iot devices. *ACM TOPS*, 2017.
- [11] Andreas Haas et al. Bringing the web up to speed with WebAssembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 185–200, 2017.
- [12] Volodymyr Shymanskyy. WASM3: A high Performance WebAssembly Interpreter Written in C. <https://github.com/wasm3/wasm3>.
- [13] Nicolas Tsiftes and Thiemo Voigt. Velox vm: A safe execution environment for resource-constrained iot applications. *Journal of Network and Computer Applications*, 118:61–73, 2018.
- [14] Niels Brouwers et al. Darjeeling, a feature-rich vm for the resource poor. In *ACM SenSys*, 2009.
- [15] Oracle. Java Card 3.1. <https://www.oracle.com/java/technologies/java-card-tech.html>, 2019.
- [16] Steven McCanne and Van Jacobson. The BSD Packet Filter: A New Architecture for User-level Packet Capture. In *USENIX*, volume 46, 1993.
- [17] Matt Fleming. A Thorough Introduction to eBPF. *Linux Weekly News*, 2017.
- [18] Koen Zandberg and Emmanuel Baccelli. Minimal virtual machines on iot microcontrollers: The case of berkeley packet filters with rbpf. In *2020 9th IFIP International Conference on Performance Evaluation and Modeling in Wireless Networks (PEMWN)*, pages 1–6. IEEE, 2020.
- [19] MicroPython. <https://micropython.org/>.
- [20] Emmanuel Baccelli et al. Scripting Over-The-Air: Towards Containers on Low-end Devices in the Internet of Things. In *IEEE PerCom*, March 2018.
- [21] Yoonseok Ko, Tamara Rezk, and Manuel Serrano. Securejs compiler: Portable memory isolation in javascript. In *SAC 2021-The 36th ACM/SIGAPP Symposium On Applied Computing*, 2021.
- [22] Amit Levy et al. Multiprogramming a 64kb computer safely and efficiently. In *ACM SOSP*, 2017.
- [23] Yuxin Ren, Guyue Liu, Vlad Nitu, Wenyuan Shao, Riley Kennedy, Gabriel Parmer, Timothy Wood, and Alain Tchana. Fine-grained isolation for scalable, dynamic, multi-tenant edge clouds. In *USENIX*, pages 927–942, 2020.
- [24] W3C. WASI: libc Implementation for WebAssembly. <https://github.com/WebAssembly/wasi-libc>.
- [25] JerryScript: JavaScript engine for the Internet of Things. <https://github.com/jerryscript-project/jerryscript>.
- [26] CircuitPython. <https://circuitpython.org/>.
- [27] MicroPython in RIOT. http://doc.riot-os.org/group__pkg__micropython.html.
- [28] Femto-Containers RIOT Implementation. https://github.com/future-proof-iot/Femto-Container_tutorials.
- [29] Femto-container sensor readout application. https://github.com/future-proof-iot/Femto-Container_tutorials/blob/main/snippets/sensor_process.c.
- [30] Femto-Container CoAP sensor value handler. https://github.com/future-proof-iot/Femto-Container_tutorials/blob/main/snippets/counter_fetch_gcoap.c.
- [31] Cedric Adjih et al. FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed. In *IEEE WF-IoT*, 2015.