Virtual Environment and Tool-boxing for Trustworthy Development of RISC-V based Cloud Services

Vitamin-V: pushing for a complete and secure hardware/software RISC-V stack for IoT/Cloud

Ramon Canal – 9/6/2023
RISER workshop @RISC-V Summit Europe 2023

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Horizon programme
HORIZON-CL4-2022-DIGITAL-EMERGING-01-26
Start date: January 1st, 2023
End date: December 31st, 2025

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Vitamin-V aims to develop further RISC-V …

1) open-source virtual environments,
2) software validation suites, and
3) cloud hardware-software stacks.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Objectives (1)

Develop a complete **virtual execution environment** to favor and support software development for future cloud architectures

- Three virtual environments: **QEMU**, **gem5** and **Cloud-FPGA** prototype (Atrevido core).
- Extensions include:
  - **Virtualization**, **cryptography, bit manipulation** and **vectorization** ISA extensions
  - **Memory and communication interfaces**
  - **Extensive hardware monitoring**
  - Accelerated **multi-node simulation/emulation/execution**

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Development platforms

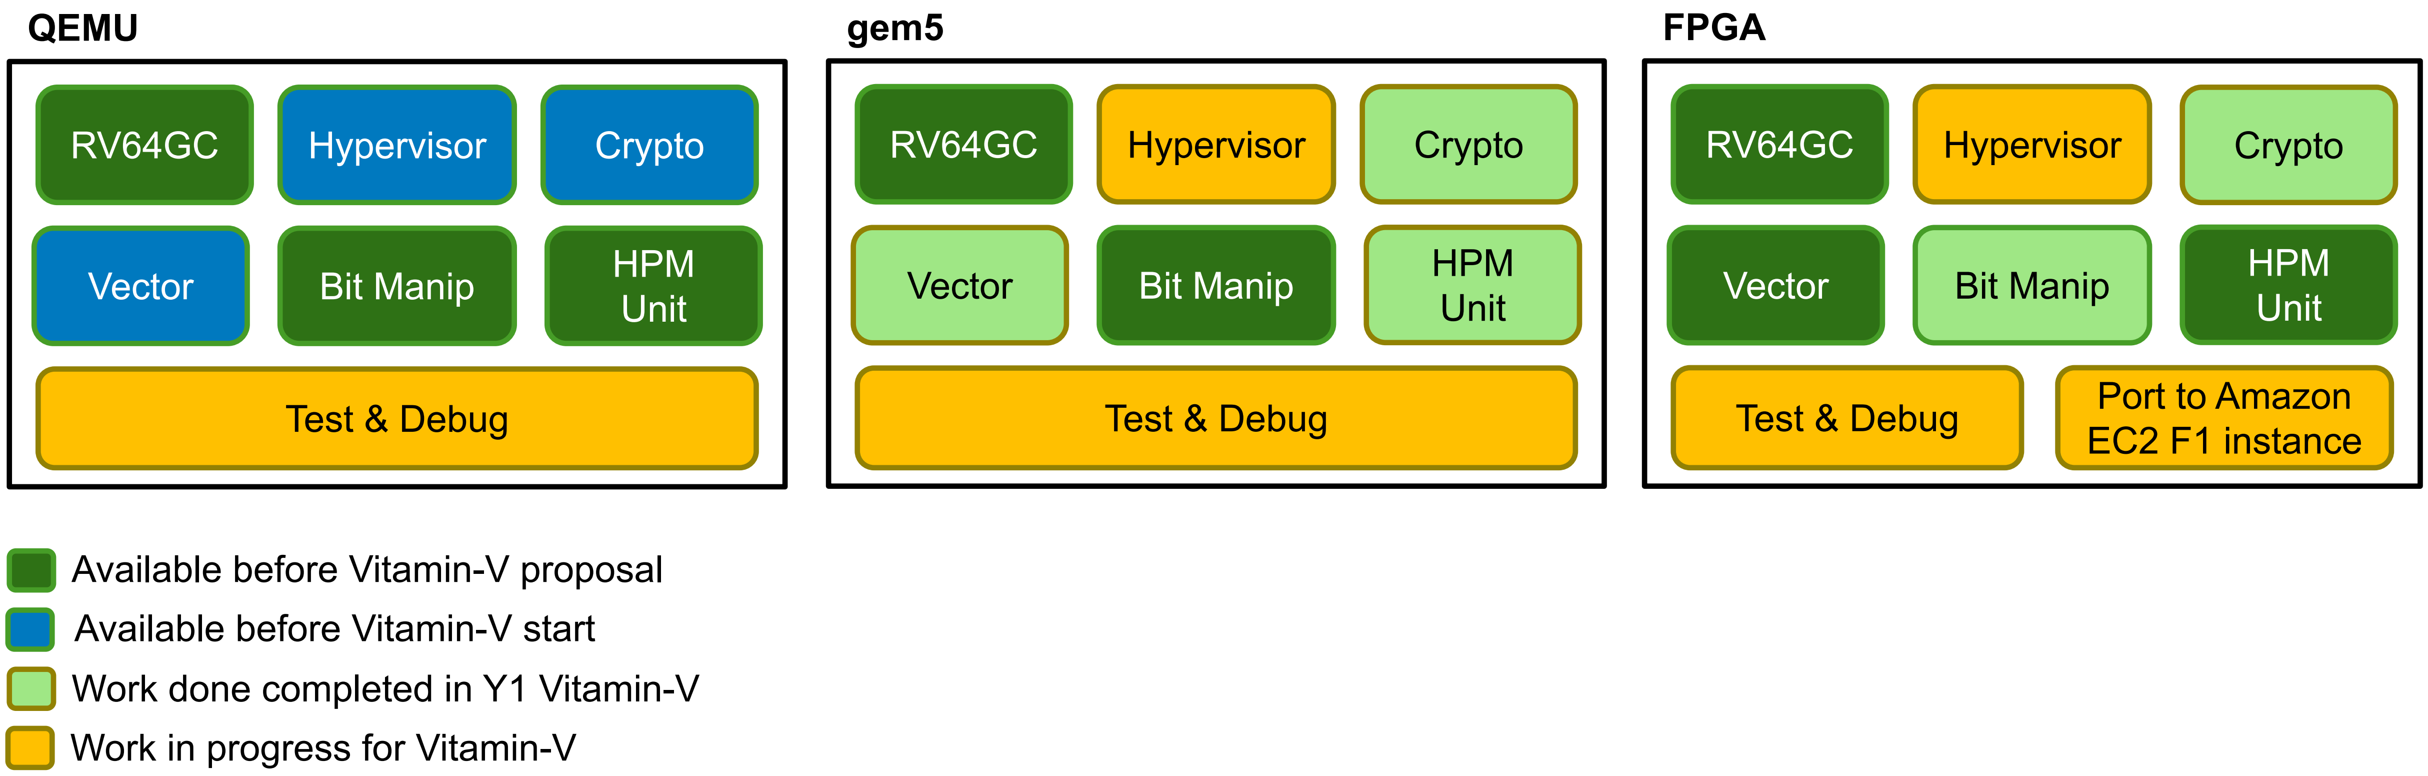

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Objectives (2)

Deliver a **toolset** to support software developers and allow **development**, **validation, verification, and testing of the trustworthiness** of software layers over RISC-V architectures.

- Extensions of **compiler (LLVM-based) and toolchain** support for the ISA extensions.
- Novel AI approaches for the validation, verification and test of **software trustworthiness**.
  - Software bugs (e.g. memory leaks), malicious code, …
  - Based on hardware monitors and source-code analysis.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

T2.2 Compiler and toolchain – Context

- Creation of toolchains that enable work on the RISC-V virtual execution environments
  - Focus on Rust, Go, C/C++
- Two items to address

What is needed to create the toolchains themselves

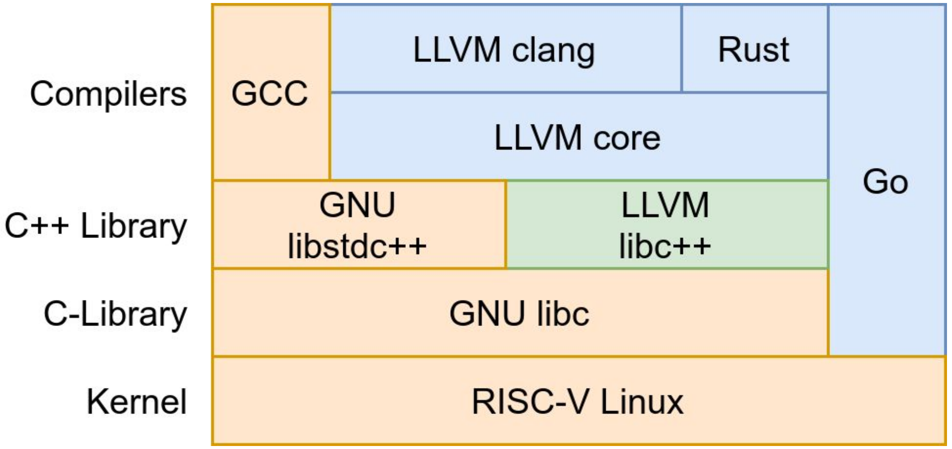

What is needed to run applications created using the toolchains

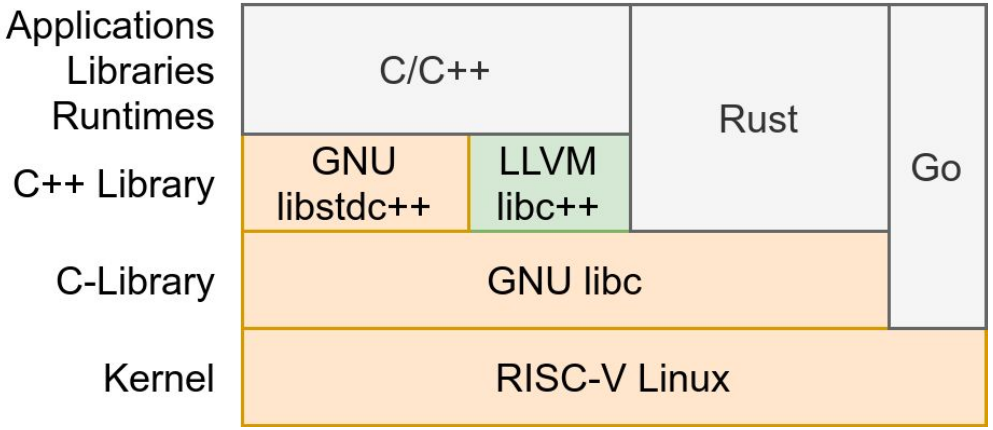

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

HPM and and AI-based ANOMALY detection

- Initial analysis on x86 architectures to use available hardware counters
- Used available Meltdown, Spectre and ZombieLoad attacks
- Used Youtube (on Firefox) and Bitcount and BasicMath (from MiBench)
- **Attacks clearly distinguishable.**

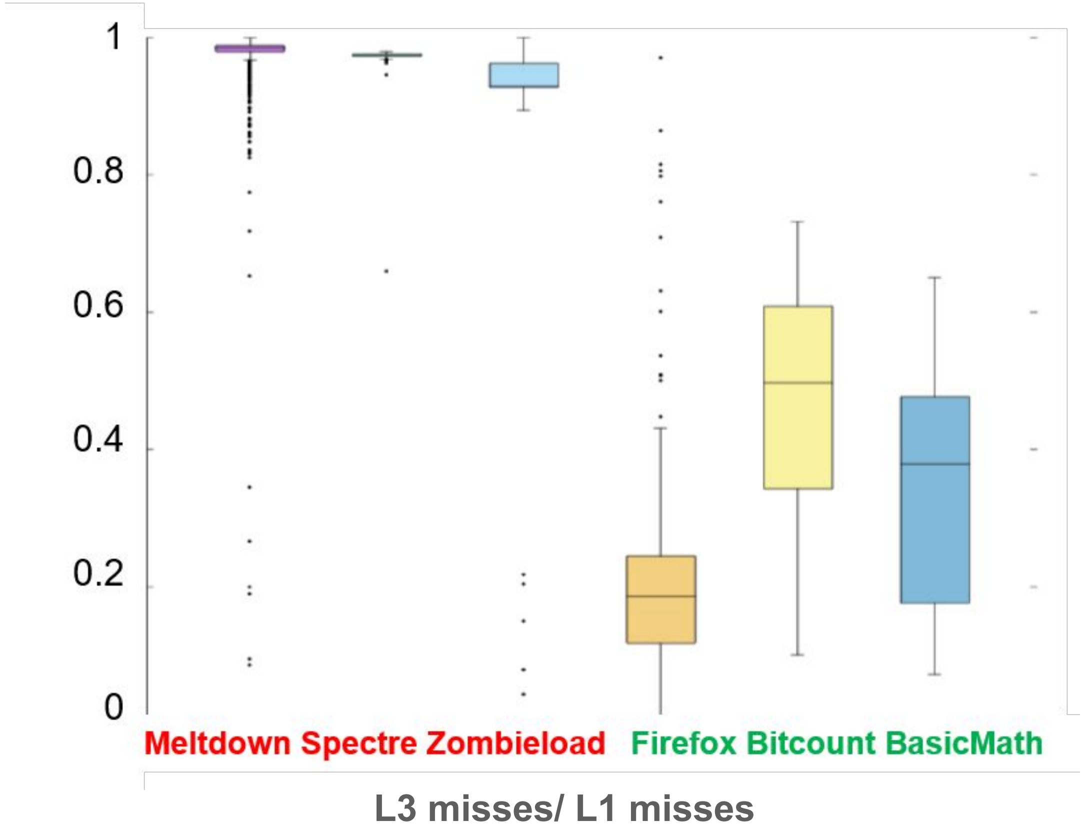

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Objectives (3)

Develop a complete **RISC-V OSS stack for Cloud services** with iso-performance to the Cloud-dominant x86 for three cloud setups

- *Classical approach*: Entire virtualized system (VMs) managed by OpenStack
- *Modern approach*: Lightweight virtualization (Containers) managed by Kubernetes
- *Serverless approach*: FaaS workloads managed by Kata + Kubernetes, RustVMM

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Objectives (3)

Produce **fundamental software ports & adaptions**, building the full OSS Cloud Stack

- **VMMs and container** suites
  - I.e. VOSySmonitor, KVM, QEMU, Docker, RustVMM and Kata containers
- **Cloud management** middleware
  - I.e., OpenStack and Kubernetes, together with their software and libraries dependencies, i.e. JVM and Python.
- **AI & BigData** applications
  - I.e., Apache Spark and Google Tensorflow.

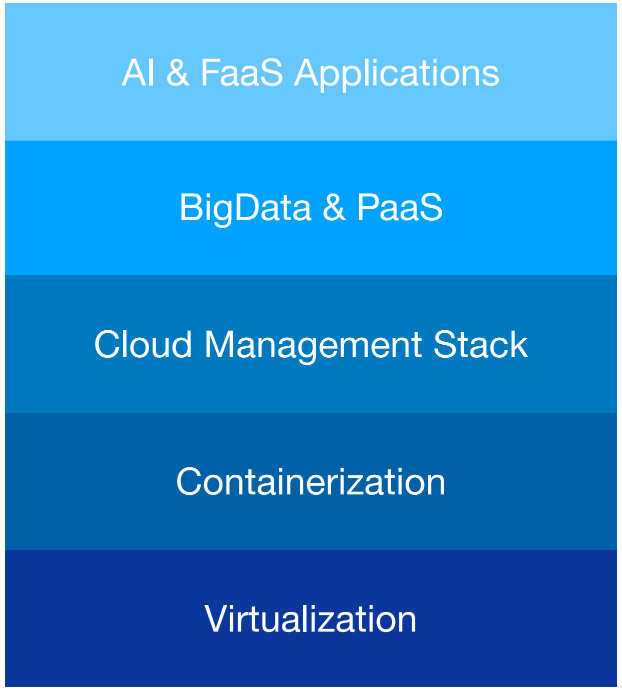

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Objectives (4)

Demonstrate AI, Big-Data, and Serverless applications running on top of classic, modern, and serverless cloud environments

- Demonstrate *standard* AI+ BigData representative applications
  - AI: VGG / ResNet / etc… SoA NNs in use by industry/science
  - Big-Data: TPC benchmarking suites (i.e. TPC-DS)
- Demonstrate Serverless and FaaS architectures
  - Srvless / FaaS: I.e., FunctionBench and ServerlessBench

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Example: Focusing on Data Analytics

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Example: Focusing on Data Analytics

- AI technologies (i.e., TensorFlow)
  - Currently there are no [functional] implementations for RISC-V. Only TensorFlow lite (inference).
  - XNNPACK is the new library for mathematical optimizations. It does have a limited RISC-V version ported for half-precision floating operations.
    - Missing full-precision (FP32).
    - Not all neural network operations are ported!
    - Performance compared to non-optimized libraries is the same. Room for improvement.

- Objectives
  - Port entire TF stack to RISC-V + target XNNPACK operations related to mostly used NNs and optimize
    - XNNPACK is an open-source library. Contributes to the community.
    - Porting FP32 operations can be a big contribution.
    - Target: 1x performance compared to x86.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Example: Focusing on Data Analytics

- Big-Data technologies (i.e., Spark)
    - Relies entirely on Java. JVM can run on RISC-V.
    - Initial results show slow performance compared to x86.
    - ML-related libraries can run on BLAS libraries → leverage RISC-V BLAS-like libraries
- Objectives
    - Implement target mathematical operations to optimize performance
    - To compile it using LLVM and optimize compilation for RV.
        - Target: 1x performance compared to x86.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

VITAMIN-V

Virtual Environment and Tool-boxing for Trustworthy Development of RISC-V based Cloud Services

https://vitamin-v.upc.edu
@VitaminVProject

Thank you for your attention!!

Stay in touch for up-to-date releases!!

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the HaDEA. Neither the European Union nor the granting authority can be held responsible for them. Project number: 101093062. https://vitamin-v.upc.edu