

VITAMIN-V: Virtual Environment and Tool-Boxing for Trustworthy Development of RISC-V Based Cloud Services

Original

VITAMIN-V: Virtual Environment and Tool-Boxing for Trustworthy Development of RISC-V Based Cloud Services / Canal, Ramon; Chenet, Cristiano; Arelakis, Angelos; Arnau, José-Maria; Berral, Josep Ll.; Call, Aaron; Di Carlo, Stefano; Costa, Juan José; Gizopoulos, Dimitris; Karakostas, Vasileios; Lubrano, Francesco; Nikas, Konstantinos; Nikolakopoulos, Yiannis; Otero, Beatriz; Papadimitriou, George; Papaefstathiou, Ioannis; Pnevmatikatos, Dionisios; Raho, Daniel; Rigo, Alvisé; Rodríguez, Eva; Savino, Alessandro; Scionti, Alberto; Tampouratzis, Nikolaos; Torregrosa, Alex. - STAMPA. - (2023), pp. 302-308. (26th Euromicro Conference on Digital System Design (DSD) Golem (Albania) 06-08 September 2023) [10.1109/dsd60849.2023.00050]
This version is available at: 11583/2987223 since: 2024-03-22T07:05:15Z

Publisher:

IEEE Computer Society

Published

DOI:10.1109/dsd60849.2023.00050

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

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

Ramon Canal¹, Cristiano Chenet², Angelos Arelakis¹⁰, José-Maria Arnau⁸, Josep Ll. Berral^{1,6}, Aaron Call⁶, Stefano Di Carlo², Juan José Costa¹, Dimitris Gizopoulos³, Vasileios Karakostas³, Francesco Lubrano⁵, Konstantinos Nikas⁷, Yiannis Nikolakopoulos¹⁰, Beatriz Otero¹, George Papadimitriou³, Ioannis Papaefstathiou⁴, Dionisios Pnevmatikatos⁷, Daniel Raho⁹, Alvise Rigo⁹, Eva Rodríguez¹, Alessandro Savino², Alberto Scionti⁵, Nikolaos Tampouratzis⁴, Alex Torregrosa⁸

¹Universitat Politècnica de Catalunya, Spain

²Politecnico di Torino, Italy;

³University of Athens, Greece;

⁴Exapsys, Greece;

⁵LINKS Foundation, Italy;

⁶Barcelona Supercomputing Center, Spain;

⁷Institute of Communication & Computer Systems, Greece;

⁸Semidynamics, Spain;

⁹Virtual Open Systems, France;

¹⁰ZeroPoint Technologies AB, Sweden;

Contact email: rcanal@ac.upc.edu

Abstract—VITAMIN-V is a 2023-2025 Horizon Europe project that aims to develop a complete RISC-V open-source software stack for cloud services with comparable performance to the cloud-dominant x86 counterpart and a powerful virtual execution environment for software development, validation, verification, and testing that considers the relevant RISC-V ISA extensions for cloud deployment. VITAMIN-V will specifically support the RISC-V extensions for virtualization, cryptography, and vectorization in three virtual environments: QEMU, gem5, and cloud FPGA prototype platforms. The project will focus on European Processor Initiative (EPI) based RISC-V designs and accelerators. VITAMIN-V will also support the ISA extensions by adding the compiler and toolchain support. Furthermore, it will develop novel software validation, verification, and testing approaches to ensure software trustworthiness. To enable the execution of complete cloud stacks, VITAMIN-V will port all necessary machine-dependent modules in relevant open-source cloud software distributions, focusing on three cloud setups. Finally, VITAMIN-V will demonstrate and benchmark these three cloud setups using relevant AI, big-data, and serverless applications. VITAMIN-V aims to match the software performance of its x86 equivalent while contributing to RISC-V open-source virtual environments, software validation, and cloud software suites.

Index Terms—Cloud computing, RISC-V, virtual environments, trustworthy applications

Funded by the European Union. Views and opinions expressed are, however, those of the authors 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

I. INTRODUCTION

Cloud services are a rapidly growing sector in the technology industry, allowing businesses and individuals to store, process, and access data and applications remotely over the Internet [1]. The European Union (EU) is working towards achieving technological sovereignty in crucial strategic areas, including integrated circuits (ICs). As ICs are essential components of modern computing systems, including those used for cloud computing, they play a critical role in the cloud value chain [2]. More specifically, the cloud value chain refers to the processes of delivering cloud computing services, from designing and manufacturing hardware components to developing and running software applications. ICs provide cloud computing systems' processing power and storage capacity in this value chain. Currently, most ICs used in cloud computing systems are manufactured outside Europe, primarily in Asia and the United States. This means that Europe is heavily dependent on imports of these critical components, which could potentially pose a threat to its technological sovereignty and national security. By controlling ICs for the cloud value chain, Europe can ensure it has access to the necessary technologies and expertise to develop and deploy cloud computing systems independently. This would reduce its reliance on imports and enhance its ability to compete globally in the digital economy [3].

In recent years, the RISC-V Instruction Set Architecture

(ISA) has emerged as a promising platform for developing cloud services due to its open-source nature, scalability, and energy efficiency. In addition, there is a strong effort from the European governing bodies to invest in this ISA [4]. RISC-V is a free and open-source ISA, allowing anyone to design and implement processors based on the RISC-V specification without paying licensing fees. This has led to a growing ecosystem of RISC-V processors, tools, and software that drives innovation and reduces costs for cloud service providers.

One of the critical benefits of RISC-V for cloud services is its scalability. RISC-V processors can be designed to support a wide range of performance and power requirements, from tiny microcontrollers to high-performance data center servers. This makes RISC-V well-suited for cloud services, whose computing needs may vary significantly (i.e., scale-up and scale-out workloads) and require computing platforms with different resources for optimal performance and energy efficiency.

However, some challenges must be addressed before RISC-V can be widely adopted for cloud applications. Although RISC-V has gained significant momentum in recent years, the ecosystem around RISC-V processors is still developing. This includes hardware and software tools and the number of vendors and support services available. Also, hardware availability is still limited, especially in the high-performance range. Available RISC-V cores may not match the performance of more established architectures like x86 or ARM in specific applications. This can be a barrier to adoption in performance-sensitive cloud applications. In parallel to these problems, compatibility is a primary issue. Many cloud applications are designed to run on x86 or ARM architectures, and their porting to a new architecture is not straightforward.

This paper introduces VITAMIN-V, a research project funded under the Horizon Europe framework program to develop a complete open-source software stack for cloud services using the RISC-V architecture. The goal is to achieve comparable performance to the cloud-dominant x86 counterpart and provide a robust virtual execution environment for software development, validation, verification, and testing. This project will consider the relevant RISC-V ISA extensions for cloud deployment, including virtualization, cryptography, and vectorization. Commercial cloud systems currently use hardware features unavailable in RISC-V virtual environments, making it challenging to start working on software porting before actual hardware is available. VITAMIN-V will support these features in three virtual environments: QEMU [5], gem5 [6], [7], and cloud-FPGA prototype platforms. The project will focus on European Processor Initiative (EPI) based RISC-V designs for the main CPUs and cloud-important accelerators for memory compression [8], [9]. VITAMIN-V will also support the ISA extensions by adding the compiler (LLVM [10]) and toolchain support. Furthermore, it will develop novel software validation, verification, and testing approaches to ensure software trustworthiness. To enable the execution of complete cloud stacks on the VRISC-V virtual execution environment, VITAMIN-V will port all necessary machine-

dependent modules in relevant open-source cloud software distributions, focusing on three cloud setups. The targeted modules include support for running entire Virtual Machines (VMs), containers, lightweight VMs, safety-security trusted execution environments, cloud management software, and AI and Big Data libraries. Finally, VITAMIN-V will demonstrate and benchmark these three cloud setups using relevant AI, big data, and serverless applications.

The paper is organized as follows: Section II overviews the main ideas and methodologies introduced by VITAMIN-V while Sections III, IV, and V detail our plan to build a complete cloud environment from hardware emulation to complex benchmarking considering different cloud deployment environments. Finally, VI summarizes the paper’s main contributions.

II. CONCEPT AND METHODOLOGIES

The objective of VITAMIN-V is to support the creation of a complete RISC-V cloud software stack that can compete in performance with the dominant x86-based cloud systems, as shown in Figure 1. However, a significant challenge in achieving this goal is the lack of RISC-V hardware cores, which makes it difficult to port and evaluate advanced cloud setups and software stacks. Commercial cloud systems use hardware features unavailable in RISC-V virtual environments and commercial hardware cores, such as virtualization, cryptography, and vector extensions.

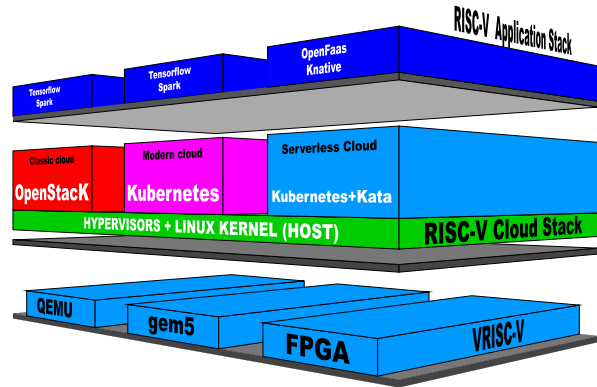


Fig. 1. VITAMIN-V concept and architecture.

To overcome this challenge, one of the first products VITAMIN-V aims to build is a cutting-edge virtual execution environment called VRISC-V, providing three key emulation/simulation technologies: (i) functional emulation (QEMU [5]) for fast software prototyping, (ii) cycle-accurate simulation (gem5 [7], [6]) for detailed architectural performance analysis, and (iii) an FPGA-based hardware system node capable of running on AWS EC2 F1 FPGAs [11] (and thus scalable to hundreds of nodes) for large cloud systems emulation. These technologies provide unique features crucial during software development, validation, verification, and testing. Apart from the RISC-V main CPUs, VRISC-V will

also support accelerators developed by EPI, particularly for memory compression, which are suitable for cloud setups.

To enable rapid software development in new RISC-V hardware setups, the project will provide a mature compiler toolchain based on the Low-Level Virtual Machine (LLVM) compiler [10] to handle the complete RISC-V ISA and its extensions. Additionally, VITAMIN-V will develop a Validation, Verification, and Testing (VVT) toolset to help identify software bugs and illegal or malicious code sequences, thus enhancing the trustworthiness of software layers in future RISC-V systems. VVT tools will use the VRISC-V platform to facilitate software porting and prototyping before mature hardware is available. This is crucial for migrating to new RISC-V servers in the cloud.

To enable the execution of complete cloud stacks on the VRISC-V virtual execution environment, VITAMIN-V will port all necessary machine-dependent modules in relevant open-source cloud software distributions. The targeted modules include support for running entire Virtual Machines (VMs), containers, and lightweight VMs (KVM [12], QEMU [5], Docker [13], RustVMM [14]), safety-security trusted execution environments (VOSySMonitRV [15]), cloud management software (OpenStack [16], Kubernetes [17], Kata Containers [18]), and Artificial Intelligence (AI) and Big Data libraries (Tensorflow [19], Spark [20]). The project will address the three cloud stacks described in Figure 2: classic, modern, and serverless, each targeting a specific cloud setup.

VITAMIN-V will demonstrate and benchmark the three cloud setups using relevant AI, big data, and serverless applications. The goal is to achieve performance that matches its x86 equivalent, currently the dominant ISA in cloud servers.

III. VITAMIN-V

VITAMIN-V will develop VRISC-V, a complete, multi-layered, high-performance virtual execution environment. VRISC-V will enable and support the software development for future cloud architectures powered by EPI RISC-V based cores and systems. VRISC-V is based on three cutting-edge technologies: (i) functional emulation (QEMU [5]), (ii) cycle-accurate microarchitectural simulation (gem5 [7]), and (iii) FPGA hardware prototyping.

A. Functional Emulation with QEMU

RISC-V based systems and platforms are still under development and not widely available for experimentation. Hence, using RISC-V based emulation platforms is mandatory to develop the necessary system and application software support for cloud computing in parallel with developing real processors, boards, and systems.

VITAMIN-V focuses on QEMU [5], the most widespread and mature emulator available for several commercial ISAs. QEMU is a virtual emulation platform that reaches high functional emulation throughput (typically less than one order of magnitude slower than native execution). This high emulation speed allows QEMU to be used for software development, and

thus, QEMU becomes a crucial component for porting cloud software stacks to RISC-V based systems.

VITAMIN-V will extend QEMU to support open-source system emulation with the RISC-V ISA extensions crucial for deploying RISC-V hardware in the cloud. VITAMIN-V will also extend QEMU with a distinctive, functionally accurate EPI SGA2 feature that targets memory compression. This extension will involve custom instructions, memory compression and decompression algorithms, and a hardware accelerator model responsible for the memory compaction, bookkeeping, and management of the compressed memory. With these extensions, QEMU can emulate a full cloud stack, enable software porting, and provide a software verification, validation, and test framework (see Section IV).

B. Microarchitectural cycle-accurate simulation with gem5

QEMU emulator is designed for speed and code compatibility and, thus, does not provide cycle-accurate performance evaluation. However, accurate performance evaluation is mandatory when porting virtualization technologies to an emerging ISA, and different microarchitectures can be employed. Without abundant RISC-V hardware, the importance of fast (multiple hundreds of thousands of simulated instructions per second) cycle-accurate performance simulation is dramatic.

VITAMIN-V focuses on gem5 [7], [6], the state-of-the-art, open-source, full-system simulator that can provide accurate performance evaluation results and supports many commercial CPUs ISAs. Gem5 flexibly supports configuring a rich set of microarchitectural parameters (number, size, organization of cores, pipelines, caches, buffers, queues, speculation structures, etc.). VITAMIN-V will extend gem5 to provide a complete open-source detailed system simulation framework for accelerated evaluation of RISC-V software deployment in multi-node cloud configurations along the following five axes.

1) *Full-system support for RISC-V*: The gem5 simulator can operate with the functional (atomic) and the cycle-accurate (detailed) microarchitecture detail. Moreover, it can work by performing syscall emulation (SE) or full system (FS) simulation. Gem5's FS mode is of interest to VITAMIN-V for porting and evaluating cloud services on top of RISC-V CPUs. VITAMIN-V will extend the FS simulation support of gem5, enabling a broad design space exploration of alternative RISC-V CPU microarchitectures for the entire software stack.

2) *Support for the GHKBV ISA extensions*: VITAMIN-V will implement the H RISC-V ISA hypervisor extension that targets improved virtualization performance. VITAMIN-V will also implement in gem5 the hardware side components for complete system modeling and accurate simulation, e.g., two-stage address translation, trap and interrupt handling, and system/platform level support. Gem5 can execute hypervisor modules (KVM Linux) and manage virtual machines with these modifications. VITAMIN-V will also extend gem5 to implement the K, B, and V RISC-V ISA extensions to support the targeted cloud setups and applications.

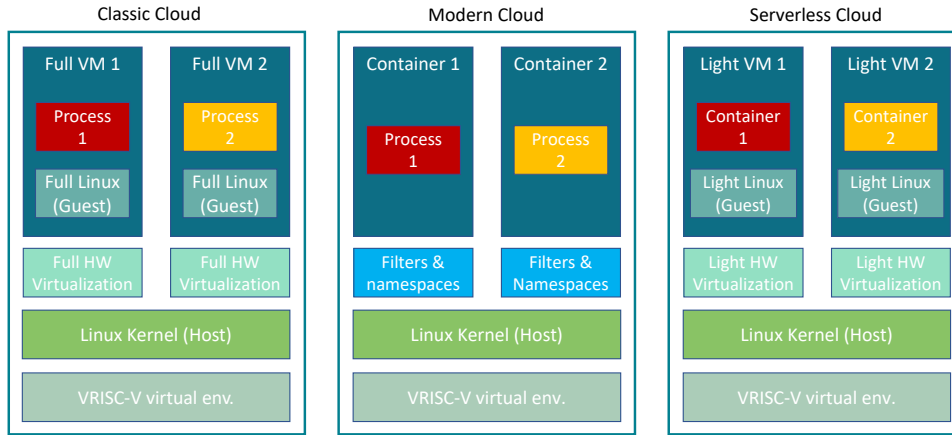


Fig. 2. VITAMIN-V cloud scenarios.

3) *Support for the Hardware Performance Monitoring extensions:* VITAMIN-V will flexibly implement the latest definition of the RISC-V Hardware Performance Monitor (HPM) extensions in gem5. The implementation of HPM will focus on several architectural and microarchitectural events (that will be made accessible to software) and hardware assertions in gem5, a valuable debugging tool for performance, design, and software bugs.

4) *Support for memory compression:* VITAMIN-V aims to extend gem5 to simulate in detail the hardware support for memory compression in the context of the EPI SGA2 feature that targets memory compression. The gem5 platform will also enable the validation and assessment of the system-software stack of the memory compression IP.

5) *Support for multi-node simulation:* Finally, VITAMIN-V will integrate the gem5 simulation engine with flexible network simulation infrastructure [21], [22], enabling cloud-level large-scale (multi-node) simulations at cycle accuracy.

C. FPGA Prototyping

Cycle-accurate performance evaluation alone is insufficient when scaling up the system’s complexity to complex cloud stacks deployed over hundreds of computing nodes. FPGAs have been extensively used to accelerate the system’s performance evaluation [23]. However, despite several FPGA-ready RISC-V cores being under development, no known public design of a complete EPI-compliant RISC-V system able to run a full cloud software stack is available for software developers. This creates a bottleneck in the widespread adoption of the RISC-V architecture in the cloud domain.

VITAMIN-V focuses on using cloud FPGA nodes to cover this gap. Cloud provider FPGA instances, such as Amazon EC2 F1 [11], provide high-speed FPGAs to deliver custom hardware designs over several hundreds of nodes.

VITAMIN-V will provide an FPGA-based prototype of the Semidynamics’ Atrevido RISC-V RV64GCV [24]. Atrevido is an out-of-order RISC-V 64GCV core that can boot Linux and has been synthesized in several FPGAs for testing. This

prototype core will be extended to run a complete software stack (from firmware through Linux up to the middleware) at approximately 50MHz speed with 100% cycle accuracy. The core will also be extended to implement the H, K, B, and V ISA extensions that are the basis for deploying VMMs and executing cloud workloads. Furthermore, the core will be enhanced with a RISC-V HPM infrastructure capable of producing detailed, cycle-by-cycle events related to the executed instructions. The (encrypted) core and supporting system software will be containerized to be deployed on cloud FPGA instances, such as Amazon EC2 F1, thus enabling the creation of a fully connected cloud node prototype. The container will be publicly released to foster the RISC-V ecosystem.

IV. SOFTWARE VALIDATION, VERIFICATION, AND TEST (VVT)

VITAMIN-V aims to support rapid software development in new RISC-V hardware setups by providing a mature compiler toolchain based on LLVM [10], which can handle the complete RISC-V ISA and its extensions. The project will also develop a validation, verification, and testing toolset for developers to identify software bugs and illegal or malicious code sequences. These tools will use the VRISC-V platform to foster software porting and prototyping before mature hardware is released.

A. Toolchain

A toolchain is required to compile the source code for the target RISC-V architecture to deploy state-of-the-art cloud stacks on RISC-V computing nodes. VITAMIN-V focuses on LLVM [10], which supports various architectures and can be easily modified and extended. However, the full RISC-V ISA provides a large set of extensions, which need support from the compiler and toolchain. The maturity of RISC-V porting of LLVM is currently limited and does not allow applications to exploit the full range of features of this hardware architecture.

VITAMIN-V will deliver a complete and mature RISC-V build toolchain based on LLVM, supporting the required ISA extensions to develop and deploy applications ready to be executed on the VRISC-V platform. This porting activity

will involve all related technologies needed to fully integrate LLVM with high-level languages required to deploy relevant cloud software distributions (i.e., Golang [25] for Kubernetes, Rust for rust-vmm, Java Virtual Machine, and Python3 for Apache Spark and TensorFlow). As a result, cloud providers can deploy such stacks in new computing nodes equipped with RISC-V processors and deploy their custom applications through the provided development toolchain.

B. VVT Support

In the era of cloud computing, transferring software workloads between different computing platforms can pose challenges in terms of performance and trustworthiness [26], [27]. Several factors, such as poor software implementation, inefficient data structures, and limited caching, can decrease performance when moving from one architecture to another. Additionally, the software can be vulnerable to bugs or malicious code at any point in its lifecycle. To optimize and ensure the trustworthiness of an application, it is crucial to monitor its execution, identify performance bottlenecks, and customize the software to fit the underlying hardware.

However, this task cannot be achieved solely through simple performance metrics like execution time or clock cycles. Several factors must be considered when mapping software to a modern computing system, such as in-order vs. out-of-order execution engines, pipeline stages, execution ports with corresponding latencies, re-order buffers, load/store queues, and cache organization. As a result, capturing and analyzing detailed performance metrics is becoming increasingly necessary to enable in-depth architecture modeling and optimization procedures [28].

VITAMIN-V will develop full support for the latest specifications of the RISC-V HPM, starting from hardware emulation in VRISC-V up to the hypervisor and integration in the operating system through consolidated open-source libraries (i.e., Linux perf and PAPI). The RISC-V HPM unit is significantly less complex than the x86 counterpart, and the relevant performance analysis tools are not comparable to ARM's and x86's monitoring solutions. Nevertheless, the RISC-V HPM specification is a flexible generic performance monitoring solution, and being open-source allows any degree of implementation freedom. VITAMIN-V uses the RISC-V HPM to provide a powerful infrastructure for software VVT.

Software bugs may arise in various development life cycle stages, and simple debugging tools cannot guarantee software trustworthiness in complex cloud projects. AI techniques supported by information collected from a powerful RISC-V HPM are promising solutions to help designers in this challenging task.

Assessing the reliability of software before execution is crucial. To achieve this, VITAMIN-V is considering implementing a machine learning (ML) tool to analyze the static content of executable files and distinguish between benign and malicious code. Previous research has shown that static analysis of executable files can detect software bugs and security threats [29], [30]. VITAMIN-V intends to incorporate

deep learning (DL) techniques to identify complex patterns from labeled data. However, in cases where the dataset is insufficient or zero-day attack detection is needed, transfer learning (TL) can improve detection accuracy. TL involves utilizing a pre-trained model for a different task [31].

Even if an application passes static verification, it can still be vulnerable to corruption during runtime, making dynamic monitoring necessary [32]. Environmental conditions, design flaws, aging, intentional attacks, and manufacturing imperfections can lead to abnormal processor behavior and jeopardize modern computing systems' safety, security, and reliability. Identifying data patterns that deviate from expected behavior is crucial to detect such anomalies. VITAMIN-V plans to use the RISC-V HPM unit to accomplish this. VITAMIN-V intends to develop AI-powered techniques to distinguish between benign and malicious applications using hardware counters, creating an anomaly detection system to identify data patterns that deviate from expected behavior [33], [34], [35].

V. SOFTWARE PORTS AND DEMONSTRATION

VITAMIN-V will port all necessary machine-dependent modules in relevant open-source cloud software distributions to enable the execution of complete cloud stacks on the VRISC-V virtual execution environment. Targeted modules include support for running entire Virtual Machines (VMs), containers (Docker), and lightweight VMs (KVM, QEMU, RustVMM), safety-security trusted execution environments (VOSySMonitoRV), cloud management software (OpenStack, Kubernetes, Kata Containers), and AI and Big Data libraries (Tensorflow, Spark). VITAMIN-V will address a classic cloud stack that targets the execution of entire VMs managed by OpenStack, a modern cloud setup that targets entire VMs and containers managed by Kubernetes, and a serverless cloud stack that targets the execution of lightweight VMs managed by Kubernetes with Kata Containers.

A. Cloud environments ports

VITAMIN-V will provide virtualization and containerization technologies to the VRISC-V framework. KVM [12], the Linux in-kernel hypervisor, will be extensively tested against typical cloud use cases and optimized, where necessary, to sustain intensive workloads spread among multiple guests. Moreover, VITAMIN-V will extend KVM to utilize the aforementioned memory compression features in VRISC-V. The scope will be to utilize KVM's memory management knowledge for targeted training of compression algorithms and provide mechanisms so that the guest virtual machines can use the memory space gained thanks to compression. This will allow the exploitation of virtualization solutions like QEMU [5] and rust-vmm [14]; the former as a complete system emulator capable of instantiating features-rich VMs, the latter as a powerful set of building blocks that are written in Rust and are designed to implement flexible virtualization solutions for fast start-up times and reduced memory footprint. Once the rust-vmm port to RISC-V is complete, Firecracker [36] and Cloud Hypervisor [37] that are built on

top of it, will become available, constituting precious assets to tackle serverless computing.

Docker [13] and related components will also be part of the VRISC-V framework, allowing developers to leverage containers as a more lightweight virtualization solution capable of fast deployment and quick prototyping.

In addition, VITAMIN-V will consider Kata containers [18] to provide the flexibility of containers and the strong isolation of virtual machines. This technology will also be adapted and optimized to run on RISC-V, targeting the integration with Kubernetes [17] to enable the deployment of scalable distributed systems. In the context of VITAMIN-V, less conventional forms of virtualization will also be explored. This is the case of VOSySmonitorRV [15], a proprietary low-level system partitioner capable of defining multiple hardware partitions where to deploy bare-metal operating systems. VITAMIN-V is committed to making publicly available all the contributions to the aforementioned open-source projects and engaging the relevant communities before submitting patches or RFCs.

VITAMIN-V will demonstrate the capability of running modern cloud environments by porting the Kubernetes [17] framework on top of the developed simulation/emulation solutions. Kubernetes gained popularity as it allows orchestrating and managing cloud applications in the form of lightweight containerized (micro-)services. Kubernetes abstracts the underlying infrastructure, automatically manages computing, storage, and networking resources, and automates several management tasks ensuring deployed applications' integrity, resiliency, and scalability. This porting includes adapting all technical elements related to supporting the execution of containers and the compilation of the framework itself (e.g., Golang [25], OS libraries, etc.). Porting cloud management applications on RISC-V will promote adopting such open-source technology in data centers and modern edge environments.

Moreover, VITAMIN-V will port OpenStack [16], enabling the deployment of docker containers through the previously ported Kubernetes framework. OpenStack is one of the world's most widely deployed open-source cloud software. The porting includes the adaptation to integrate the ported hypervisor into RISC-V, the host and guest OS, the runtime libraries, runtime platforms, and benchmarking applications as representative client processes running on a cloud system.

B. Big Data and Data Analytics environments

VITAMIN-V will also provide Big Data and Data Analytics runtimes, whose main challenge is processing data as fast as possible and holding vast amounts of data not fitting into a single node's memory. Thus, workloads in such environments require many compute nodes, and coordination between them is necessary. However, current RISC-V hardware implementations lack some extensions, such as vector extensions, critical for performance gains in both environments. Therefore, achieving the same performance on RISC-V on existing hardware as on x86 or AMD architectures is impossible.

The main runtimes of such environments are TensorFlow [19] and Spark [20]. The first provides a suitable

framework for neural networks, while the second provides support for handling and computing large amounts of data across hundreds or thousands of nodes. Regarding TensorFlow, the runtime has only been partially ported to RISC-V; hence, some features are missing.

Currently, there are only ports of TensorFlow Lite [38] to RISC-V. However, TensorFlow Lite only supports inference. Thus, it is a good use case for edge-computing applications where the edge node applies inference based on a previously trained model. For instance, this model could have been trained on a cloud server based on an x86 architecture. Since VITAMIN-V will also enable RISC-V on cloud environments, it will port the full TensorFlow stack, including the training stage. It will make the necessary adjustments to compile the TensorFlow runtime and train relevant neural network models, such as VGG19 [39], ResNet50 [40], and MobileNet [41], representing the most real-world use cases.

Spark is a big data runtime that has become a de facto standard for this kind of applications. VITAMIN-V will adapt the Spark runtime to be executed in its VRISC-V environment and exploit the memory optimization features developed within the project. The project intends to use the TPC-H [42] benchmark. The TPC-H is a decision support benchmark. It consists of a suite of business-oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance. This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions.

Moreover, VITAMIN-V will also use Variant-Interaction Analytics (VIA) workload. VIA is a genomics workload requiring computing as many variants of the human genome in the shortest time possible. Thus, it requires processing a large amount of data (a single genome might occupy up to 1TB of data) in the lesser time possible. VITAMIN-V has already contributed with an open-source data repository [43] comprising VIA workload results on the HiFive [44] Unmatched RISC-V development boards and an under revision paper.

VI. CONCLUSIONS

The computing industry is on the brink of a new era of open and collaborative computing, which could significantly impact how we approach technology development and innovation. One of the key technologies driving this revolution is RISC-V, an open-source instruction set architecture (ISA) that is gaining popularity in various applications. By integrating RISC-V into cloud computing, we can revolutionize the industry with unparalleled flexibility, efficiency, and scalability.

VITAMIN-V is an ambitious project that aims to play a crucial role in this challenging revolution. VITAMIN-V seeks to showcase the benefits of RISC-V in cloud computing and enable service providers, system architects, and application developers through its focus on benchmarking, porting, and adapting benchmarking tools and applications to the RISC-V architecture to deploy and run different benchmarks on data

analytics and serverless platforms. Interested readers can follow VITAMIN-V updates through the projects' website: <https://www.vitamin-v.eu/>.

REFERENCES

- [1] H. Taherdoost, "An overview of trends in information systems: Emerging technologies that transform the information technology industry," *Cloud Computing and Data Science*, pp. 1–16, 2023.
- [2] R. Huggins, A. Johnston, M. Munday, and C. Xu, "Competition, open innovation, and growth challenges in the semiconductor industry: the case of Europe's clusters," *Science and Public Policy*, 03 2023, scad005. [Online]. Available: <https://doi.org/10.1093/scipol/scad005>
- [3] J. Bish, S. Duncan, K. Ramachandran, P. Lee, and S. Beerlage, "A new dawn for european chips," [Online] <https://www2.deloitte.com/us/en/insights/industry/technology/semiconductor-chip-shortage-supply-chain.html>, April 2023.
- [4] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanović, "The risc-v instruction set manual, volume i: User-level isa, version 2.0," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2014-54, May 2014. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-54.html>
- [5] "QEMU, A generic and open source machine emulator and virtualizer," [Online] <https://www.qemu.org/>, accessed: 2023-03-17.
- [6] J. Lowe-Power, A. M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, B. Beckmann, S. Bharadwaj, G. Black, G. Bloom, B. R. Bruce, D. R. Carvalho, J. Castrillon, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, C. Escuin, M. Fariborz, A. Farmahini-Farahani, P. Fotouhi, R. Gambord, J. Gandhi, D. Gope, T. Grass, A. Gutierrez, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. A. R. Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T. M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, M. Moreto, T. Mück, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L. E. Olson, M. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M. D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, W. Wang, Z. Wang, N. Wehn, C. Weis, D. A. Wood, H. Yoon, and É. F. Zulian, "The gem5 simulator: Version 20.0+," 2020.
- [7] "gem5," [Online] <https://github.com/gem5/gem5>, accessed: 2023-03-17.
- [8] "European Processor Initiative," [Online] <https://www.european-processor-initiative.eu/>, accessed: 2023-03-17.
- [9] M. Kovač, "European processor initiative: The industrial cornerstone of eurohpc for exascale era," in *Proceedings of the 16th ACM International Conference on Computing Frontiers*, 2019, p. 319.
- [10] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis and transformation," San Jose, CA, USA, Mar 2004, pp. 75–88.
- [11] "Amazon F1," [Online] <https://aws.amazon.com/ec2/instance-types/f1/>, accessed: 2023-03-17.
- [12] "Kernel virtual machine," [Online] https://www.linux-kvm.org/page/Main_Page, accessed: 2023-05-06.
- [13] "Docker, accelerated, containerized application development," [Online] <https://www.docker.com/>, accessed: 2023-05-06.
- [14] "rust-vm," [Online] <https://github.com/rust-vm>, accessed: 2023-05-06.
- [15] F. Caforio, P. Iannicelli, M. Paolino, and D. Raho, "VOSySmonitoRV: a mixed-criticality solution on linux-capable RISC-v platforms," in *2021 10th Mediterranean Conference on Embedded Computing (MECO)*. IEEE, jun 2021. [Online]. Available: <https://doi.org/10.1109/2Fmec052532.2021.9460246>
- [16] "Openstack: Open source cloud computing infrastructure," [Online] <https://www.openstack.org/>, accessed: 2023-05-06.
- [17] "Kubernetes," [Online] <https://kubernetes.io/>, accessed: 2023-05-06.
- [18] "Kata containers - open source container runtime software," [Online] <https://katacontainers.io>, accessed: 2023-05-06.
- [19] "Tensorflow," [Online] <https://www.tensorflow.org/>, accessed: 2023-05-06.
- [20] "Apache spark - unified engine for large-scale data analytics," [Online] <https://spark.apache.org/>, accessed: 2023-05-06.
- [21] N. Tampouratzis and I. Papaefstathiou, "A novel, simulator for heterogeneous cloud systems that incorporate custom hardware accelerators," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 4, no. 4, pp. 565–576, 2018.
- [22] A. Brokalakis, N. Tampouratzis, A. Nikitakis, I. Papaefstathiou, S. Andrianakis, D. Pau, E. Plebani, M. Paracchini, M. Marcon, I. Sourdis, P. R. Geethakumari, M. C. Palacios, M. A. Anton, and A. Szasz, "Cossim: An open-source integrated solution to address the simulator gap for systems of systems," in *2018 21st Euromicro Conference on Digital System Design (DSD)*, 2018, pp. 115–120.
- [23] D. Chiou, D. Sunwoo, J. Kim, N. A. Patil, W. Reinhart, D. E. Johnson, J. Keefe, and H. Angepat, "Fpga-accelerated simulation technologies (fast): Fast, full-system, cycle-accurate simulators," in *Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO 40. USA: IEEE Computer Society, 2007, pp. 249–261.
- [24] "SemiDynamics High Bandwidth Vector-capable RISC-V Cores," [Online] <https://semidynamics.com/uploads/semidynamics/pdf/2021.03.31-RISCV-WEEK.pdf>, accessed: 2023-04-06.
- [25] "The go programming language," [Online] <https://go.dev/>, accessed: 2023-05-06.
- [26] N. Tampouratzis and I. Papaefstathiou, "A novel, simulator for heterogeneous cloud systems that incorporate custom hardware accelerators," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 4, no. 4, pp. 565–576, 2018.
- [27] M. Alonso, D. Andreu, R. Canal, S. D. Carlo, C. Chenet, J. Costa, A. Girones, D. Gizopoulos, V. Karakostas, B. Otero, G. Papadimitriou, E. Rodriguez, and A. Savino, "Validation, verification, and testing (vvtt) of future risc-v powered cloud infrastructures: the vitamin-v horizon europe project perspective," in *Proceedings of the IEEE European Test Symposium*, 2023.
- [28] D. Marques, A. Ilic, Z. A. Matveev, and L. Sousa, "Application-driven cache-aware offline model," *Future Generation Computer Systems*, vol. 107, pp. 257–273, 2020.
- [29] Y. Ding, W. Dai, S. Yan, and Y. Zhang, "Control flow-based opcode behavior analysis for malware detection," *Computers & Security*, vol. 44, pp. 65–74, 2014.
- [30] H. HaddadPajouh, A. Dehghantanha, R. Khayami, and K.-K. R. Choo, "A deep recurrent neural network based approach for internet of things malware threat hunting," *Future Generation Computer Systems*, vol. 85, pp. 88–96, 2018.
- [31] E. Rodríguez, P. Valls, B. Otero, J. J. Costa, J. Verdú, M. A. Pajuelo, and R. Canal, "Transfer-learning-based intrusion detection framework in iot networks," vol. 22, no. 15. MDPI, 2022, p. 5621.
- [32] C. P. Chenet, A. Savino, and S. D. Carlo, "A survey of hardware-based malware detection approach," [Online] <https://arxiv.org/abs/2303.12525>, 2023.
- [33] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, jul 2009. [Online]. Available: <https://doi.org/10.1145/1541880.1541882>
- [34] S. Dutto, A. Savino, and S. Di Carlo, "Exploring deep learning for in-field fault detection in microprocessors," in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2021, pp. 1456–1459.
- [35] D. Kasap, A. Carpegna, A. Savino, and S. Di Carlo, "Micro-architectural features as soft-error induced fault executions markers in embedded safety-critical systems: a preliminary study," in *Proceedings of the 28th IEEE European Test Symposium (ETS)*, May 2023.
- [36] "Firecracker," [Online] <https://firecracker-microvm.github.io/>, accessed: 2023-05-06.
- [37] "Cloud hypervisor," [Online] <https://www.cloudhypervisor.org/>, accessed: 2023-05-06.
- [38] "Tensorflow lite - ml for mobile and edge devices," [Online] <https://www.tensorflow.org/lite>, accessed: 2023-05-06.
- [39] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014. [Online]. Available: <https://arxiv.org/abs/1409.1556>
- [40] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015. [Online]. Available: <https://arxiv.org/abs/1512.03385>
- [41] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," 2017. [Online]. Available: <https://arxiv.org/abs/1704.04861>
- [42] "Transaction processing performance council," [Online] <https://www.tpc.org/>, accessed: 2023-03-16.
- [43] "Genomics risc-v open-data repository," [Online] https://github.com/MortI2C/genomics_riscv_openrepo/, accessed: 2023-03-23.
- [44] "Hifive unmatched risc-v board," [Online] <https://www.sifive.com/boards/hifive-unmatched>, accessed: 2023-03-23.