

Open Challenges for a Production-ready Cloud Environment on top of RISC-V hardware

Guillem Senabre¹ *, Aaron Call¹, Ramon Nou¹

¹Barcelona Supercomputing Center

Abstract

As part of the Vitamin-V European project, we have built a prototype of a RISC-V cluster managed by OpenStack, with the goal of realizing a functional RISC-V cloud ecosystem. In this poster we explain the hardware and software challenges encountered while porting some elements of OpenStack. We also discuss the current performance gaps that challenge a performance-ready cloud environment over such new ISA, an essential element to fulfill in order to achieve european technological sovereignty.

Introduction and motivation

The RISC-V Instruction Set Architecture (ISA) is at the core of the European Union’s technological sovereignty plans. One key initiative in this direction is the European Processor Initiative (EPI) [1], which aims to produce processors ready for mass production based on the RISC-V open-source ISA. These processors are intended to be used in various applications, including cloud computing and data centers.

Large-scale supercomputers and datacenters become essential in many scientific applications requiring big-data processing: discovering relevant clinical, social, economic, or environmental indicators at any scale (e.g. 1 PetaByte of genomic datasets [2]). However, most current computing architectures are proprietary and closed-source technologies such as x86 and ARM, which creates concerns about the reliability of privacy and security.

The Vitamin-V[3] project emerged as an effort to build European data centers based on the EPI processor as a reliable alternative to traditional proprietary systems like x86 and ARM. As a part of the Vitamin-V project, we have devoted the last year to developing a functional OpenStack cluster utilizing real hardware instead of emulators. To this end, we are using a hardware platform based on RISC-V development boards. In this poster, we show how we are using Spieed Lichee PI 4A [4] as a demonstrator in a RISC-V-based cloud environment.

In this poster, we explain some open challenges in RISC-V that need to be addressed to compete with clouds built on top of commodity hardware, as well as the challenges that we have had to overcome in order to build a functional cloud stack on top of RISC-V, and particularly a cloud stack based on the OpenStack framework.

Ecosystem maturity

Hardware

When transitioning from QEMU-based emulation to physical hardware, the limited availability of commercial RISC-V boards becomes evident. The available options are either relatively expensive or suffer from significant performance limitations. Notably, only one commercially available HPC board [5] supports vector extensions (RVV1.0). However, its release was delayed by six months, and it sold out within three months despite its 600 € price tag.

Board Name	Price (USD)	Availability
HiFive Pro P550	400-500	Limited stock
HiFive Unmatched	295	Available
Milk-V Jupiter	599	Sold out
VisionFive 2	100	Available
LicheePi 4A	100	Available

Table 1: High-Performance RISC-V boards, their prices, and availability.

Moreover, the boards have development purposes in mind, not business-ready use cases. So even in those cases where the boards are based on top of powerful chips, there is a significant performance gap. The reasons for this are mostly because RISC-V ISA is yet being properly defined in numerous extensions and many designers have not yet implemented those in real chips. For instance, RVV1.0 extension was ratified just in 2024, and only one implementation is being commercialized. On the other hand, hypervisor extensions are defined but not built on any chips. Other extensions face similar issues.

As the hardware platform we use a RISC-V development board by Sipeed, the Lichee PI 4A owing to their balanced pricing and capabilities. More specifically, the utilized development platform provides a

*Corresponding author: guillem.senabre@bsc.es

TH1520 RISC-V CPU (4 Threads), 16GB of RAM, 128GB storage. In addition it also provides a dual Gigabit Ethernet a feature particularly interesting when building an OpenStack cluster.

We currently have a controller node, which acts like a bastion, four compute nodes, and a storage node connected to a 2TB USB SSD as shown in figure 1.

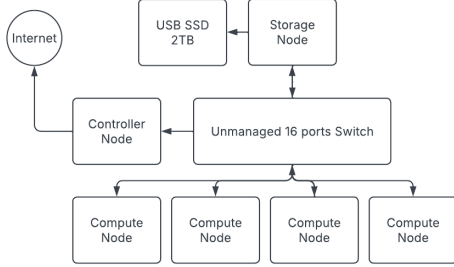


Figure 1: Openstack cluster using LicheePi4a boards.

We compared this platform against a cluster composed of 8 nodes based on Intel Xeon Silver 4114 operating at 2.2GHz, with 126GB DDR5 of RAM, interconnected via a 10GBps ethernet network plus a node dedicated to storage with 10TB of capacity. The results of this performance comparison are shown in table 2. While it is obvious that a performance gap should exist due to the difference in maturity and performance of both platforms, the gap is way bigger than one could think, in some cases, RISC-V being 10 times slower. These performance issues make it currently unfeasible to enable production-ready cloud environments such as OpenStack, as performance, scalability, and network performance are critical elements.

Test	RISC-V	x86
Coremark (iter./sec)	8500	22730
Mem. Latency 16KB (ns)	116	1497
Disk I/O SeqW. (MB/s)	204	1451
Disk I/O RandW.(MB/s)	28.8	1316
Network Throughput (MB/s)	118	1200
Power Consumption (W)	11	270-350

Table 2: Performance Comparison between LicheePi 4A (RISC-V) and Intel Xeon Silver 4112 @ 2GHz (x86) machine.

Software

Additionally, the software itself also poses challenges. When enabling OpenStack one realizes that it consists of hundreds of Python libraries and depends on precise package versions. Keeping these packages up-to-date is crucial for stability and compatibility.

The default Linux distribution for the LicheePi 4A is a Debian-based variant provided by Sipeed, which

relies on custom repositories. However, these repositories lacked many essential Debian packages, particularly development tools. Many of those packages are, moreover, not yet ported, so we either had to port them ourselves or wait until the community did. As the project progressed, Debian’s official repositories expanded their RISC-V package coverage to approximately 95%, significantly improving package availability. Consequently, we transitioned entirely to Debian’s repositories to ensure access to the latest software versions.

During OpenStack deployment, we encountered multiple issues related to unported libraries and abandoned dependencies still referenced by other packages. One notable case involved bcrypt and passlib, both of which are used by the OpenStack Identity service (Keystone) for password hashing.

Moreover, the LicheePi 4A was shipped with kernel 5.10.113, and obtaining an updated image with kernel 6.6.48 required over a month of discussions with the vendor. This slow update cycle highlighted significant maintenance challenges for long-term development and deployment.

Recently, Ubuntu updated its repositories with stable OpenStack packages. Once we discovered this, we started testing a minimal Devstack deployment on a RISC-V Ubuntu image. A patch is still needed to ensure that Nova compute correctly interacts with lower-level libraries like LibVirt or QEMU to determine supported architectures. Although further testing is required, no package mismatches were found, and OpenStack services were installed successfully. After applying workarounds to address issues during the instance boot stages, we now have a functional OpenStack cluster.

References

- [1] Mario Kovač. “European Processor Initiative: The Industrial Cornerstone of EuroHPC for Exascale Era”. In: *Proceedings of the 16th ACM International Conference on Computing Frontiers*. CF ’19. Alghero, Italy: Association for Computing Machinery, 2019, p. 319. ISBN: 9781450366854. DOI: 10.1145/3310273.3323432. URL: <https://doi.org/10.1145/3310273.3323432>.
- [2] Michael Eisenstein. “Big data: The power of petabytes”. In: *Nature* 527.7576 (Nov. 2015), S2–S4. ISSN: 1476-4687. DOI: 10.1038/527S2a. URL: <https://doi.org/10.1038/527S2a>.
- [3] Ramon Canal et al. “VITAMIN-V: Virtual Environment and Tool-Boxing for Trustworthy Development of RISC-V Based Cloud Services”. In: *2023 26th Euromicro Conference on Digital System Design (DSD)*. Los Alamitos, CA, USA: IEEE Computer Society, Sept. 2023, pp. 302–308. DOI: 10.1109/DSD60849.2023.00050.
- [4] SiPEED. *Sipeed Liche Pi4A*. <https://sipeed.com/licheepi4a>. Accessed: 2024-01-25.
- [5] Milk-V. *milk-v jupiter board*. Accessed: 6 Feb. 2025. 2024. URL: <https://milkv.io/jupiter>.