



So... how do you start?

# What we'll cover

The approaches we've taken to fix this.  
What we learned. What works. What doesn't.

Spoiler alert: not there yet!

Make it compile → Test it → Ship it

# PyTorch on RISC-V

From Cross-Compilation to Native CI

PyTorch Conference EU 2026

# Hello!

## Ludovic Henry

- Software Engineer - [Meta](#)
- Co-Chair, Technical Steering Committee - [RISE](#)



# Enabling Open-Source Software on RISC-V

# RISE Members



MEDIATEK



北京开源芯片研究院  
BEIJING INSTITUTE OF OPEN SOURCE CHIP



Meta



# Where does PyTorch fit?



PyTorch is a flagship project for RISE

If it runs well on RISC-V, it proves the platform is ready for AI/ML workloads

# What is RISC-V?

## Open standard ISA



Royalty-free, open source

Anyone can implement a RISC-V processor

Community-led, billions of chips shipped



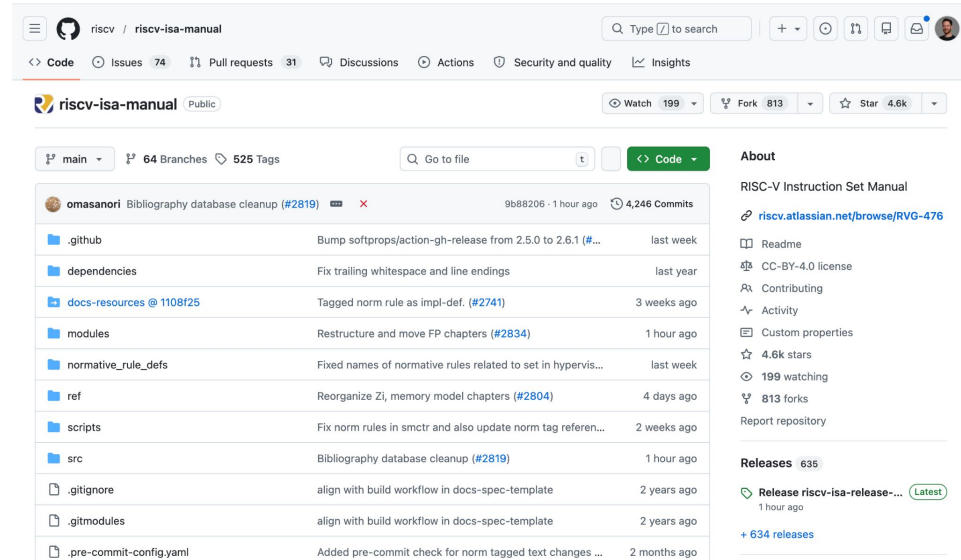
# From Embedded to Datacenter



# Why should you care?

## Open hardware & software

No single vendor controls the ISA  
Diversity of implementations and competition  
Natural fit for open-source AI/ML frameworks



riscv / riscv-isa-manual

Code Issues 74 Pull requests 31 Discussions Actions Security and quality Insights

riscv-isa-manual Public Watch 199 Fork 813 Star 4.6k

main 64 Branches 525 Tags Go to file Code

omasanori Bibliography database cleanup (#2819) 9b88206 · 1 hour ago 4,246 Commits

.github	Bump softprops/action-gh-release from 2.5.0 to 2.6.1 (#...	last week
dependencies	Fix trailing whitespace and line endings	last year
docs-resources @ 1108f25	Tagged norm rule as impl-def. (#2741)	3 weeks ago
modules	Restructure and move FP chapters (#2834)	1 hour ago
normative_rule_defs	Fixed names of normative rules related to set in hypervis...	last week
ref	Reorganize ZI, memory model chapters (#2804)	4 days ago
scripts	Fix norm rules in smctr and also update norm tag referen...	2 weeks ago
src	Bibliography database cleanup (#2819)	1 hour ago
.gitignore	align with build workflow in docs-spec-template	2 years ago
.gitmodules	align with build workflow in docs-spec-template	2 years ago
.pre-commit-config.yaml	Added pre-commit check for norm tagged text changes ...	2 months ago

**About**  
RISC-V Instruction Set Manual  
[riscv.atlassian.net/browse/RVG-476](https://riscv.atlassian.net/browse/RVG-476)  
Readme  
CC-BY-4.0 license  
Contributing  
Activity  
Custom properties  
4.6k stars  
199 watching  
813 forks  
Report repository

**Releases 635**  
Release riscv-isa-release-... (Latest)  
1 hour ago  
+ 634 releases

# Hardware Today: Scaleway EM-RV1



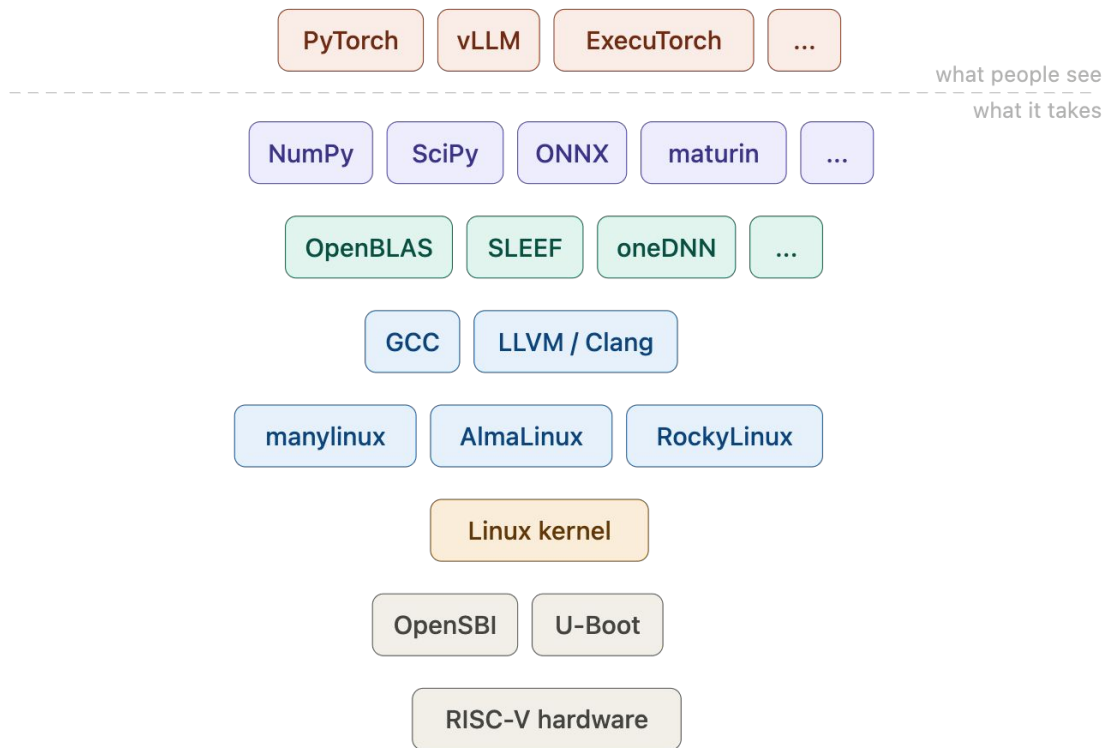
Real RISC-V bare-metal servers  
[Available today](#) in Scaleway's cloud



Step 1

Make it compile

# The Dependency Iceberg



# The Foundation: OS & Toolchain

AlmaLinux, RockyLinux, manylinux  
GCC, LLVM/Clang

RISE invested in getting these working on riscv64  
Without this layer, nothing above can even start

AlmaLinux OS Kitten 10 Begins Supporting RISC-V

Written by Michael Larabel in Operating Systems on 17 March 2026 at 01:45 PM EDT. 9 Comments

MAY 21, 2025

RockyLinux Official Support for RISC-V in RL10!

Alexia Stein, Community Lead

feat: manylinux\_2\_39\_aarch64, manylinux\_2\_39\_riscv64 &  
musllinux\_1\_2\_riscv64 #1743

Merged mayeut merged 1 commit into pypa:main from mayeut:musl-riscv64 on Jul 20, 2025

RISCV64 support #1425

Closed

# Math Libraries

## OpenBLAS

RVV support contributed by RISE 

## SLEEF

Vectorized math functions

## oneDNN

Major piece of work, in the RISE pipeline

Python

```
# your PyTorch code
C = torch.matmul(A, B) # A: [1024, 4096], B: [4096, 1024]
```

ATen dispatch

```
// aten/src/ATen/native/LinearAlgebra.cpp
Tensor mm(const Tensor& self, const Tensor& mat2) {
    auto result = at::empty({self.size(0), mat2.size(1)}, self.options());
    return at::mm_out(result, self, mat2);
}
// dispatches to backend-specific kernel via DispatchKey
// CPU → at::native::cpublas::gemm()
```

OpenBLAS  
kernel select

```
/* kernel/riscv64/KERNEL.RVV */
#define DGEMMKERNEL    dgemm_kernel_RVV
#define DGEMMINCOPY   dgemm_incopy_RVV
#define DGEMMITCOPY   dgemm_itcopy_RVV
#define DGEMMONCOPY   dgemm_ncpy_RVV
#define DGEMMOTCOPY   dgemm_otcopy_RVV
/* micro-kernel tile: M=VLEN/64 x N=4 */
/* packing: column-major → panel layout for L1 cache */
```

RVV micro-  
kernel (asm)

```
// dgemm_kernel_RVV.S - inner loop, 4 accumulators
vsetvli t0, a4, e64, m4, ta, mu // set vlen for f64
vle64.v v0, (a0) // load A column tile
fld ft0, 0(a1) // load B[k,0]
fld ft1, 8(a1) // load B[k,1]
fld ft2, 16(a1) // load B[k,2]
fld ft3, 24(a1) // load B[k,3]
vfmacc.vf v8, ft0, v0 // C[:,0] += A[:,0] * B[k,0]
vfmacc.vf v12, ft1, v0 // C[:,1] += A[:,1] * B[k,1]
vfmacc.vf v16, ft2, v0 // C[:,2] += A[:,2] * B[k,2]
vfmacc.vf v20, ft3, v0 // C[:,3] += A[:,3] * B[k,3]
addi a0, a0, 32 // advance A ptr
addi a1, a1, 32 // advance B ptr
bne a5, a0, .Lk_loop // loop over K dim
vse64.v v8, (a2) // store C[:,0]
vse64.v v12, (a3) // store C[:,1]
// ... 200+ lines of packing, tiling, edge cases
```

# The Python Last Mile




NumPy, SciPy, maturin, ONNX, conda

Every wheel in the dependency tree needs to be available on riscv64

2 years of RISE enablement work

```
luhenry@Ludovics-MacBook-Pro ~ % docker run --rm -it --platform linux/riscv64 ubuntu:24.04
root@9c7bd2903bee:/# uname -m
riscv64
root@9c7bd2903bee:/# python3 -m venv venv && source venv/bin/activate && pip3 install --upgrade pip
(venv) root@9c7bd2903bee:/# pip3 install maturin
Collecting maturin
  Downloading maturin-1.12.6-py3-none-manylinux_2_31_riscv64.musllinux_1_1_riscv64.whl.metadata (16 kB)
  Downloading maturin-1.12.6-py3-none-manylinux_2_31_riscv64.musllinux_1_1_riscv64.whl (10.0 MB)
----- 10.0/10.0 MB 25.0 MB/s 0:00:00
Installing collected packages: maturin
Successfully installed maturin-1.12.6
```

# Many Years of Ecosystem Work












# Status Check: Where Are We?



## Artifacts

Produced during runtime

Name	Size	Digest	
 linux-noble-riscv64-py3.12	258 MB	sha256:a6f7787b9d690da93854a13c598e0ce6893c8818e020a814e64...	  

-  linux-noble-riscv64-py3.12 
-  test (default, 1, 3,  
linux.riscv64.2xlarge.ephemeral  
)
-  test (default, 2, 3,  
linux.riscv64.2xlarge.ephemeral)
-  test (default, 3, 3,  
linux.riscv64.2xlarge.ephemeral)

```
>>> a = torch.randn(64, 128)
>>> b = torch.randn(128, 32)
>>> c = a @ b
>>> print(f"result shape: {c.shape}")
result shape: torch.Size([64, 32])
>>> print(f"sample value: {c[0, 0]:.4f}")
sample value: 9.0276
>>> print("ok")
ok
>>> print(c)
tensor([[ 9.0276, -2.8943,  1.2211, ...,  4.1037, -3.0044,  9.0749],
        [-2.1744, -3.6399, -6.0479, ..., 10.2100, -1.4276, -8.3336],
        [-16.5510, -4.3232, 13.3887, ..., -17.9506,  0.4866, -6.3470],
        ...,
        [-5.3377, -13.8820,  0.9861, ..., -12.7317, -12.7856, 16.3534],
        [ 9.7428,  0.6340, -11.6616, ..., -3.3935,  2.4640, -6.7308],
        [ 1.3742,  1.8036,  3.7884, ...,  8.5102, 13.0203, 10.2730]])
>>>
```

Honest snapshot

Step 2

# Test it

No CI = no official support

# RISE RISC-V Runners

Free, managed GitHub Actions runners on bare-metal [Scaleway EM-RV1](#)

Ephemeral, secure, just like x86/arm64

Install: [github.com/apps/rise-risc-v-runners](https://github.com/apps/rise-risc-v-runners)

Add `runs-on: ubuntu-24.04-riscv`

ggml-org / llama.cpp

Code Issues 554 Pull requests 856 Discussions Actions Projects Wiki Security and quality 13 Insights

CI (riscv)

ggml-webgpu: Add the support of MUL\_MAT\_ID (#21147) #244

Summary

All jobs

Triggered via push 12 hours ago

Status Success

Total duration 1h 22m 41s

Artifacts -

build-riscv.yml

on: push

Matrix: ubuntu-riscv64-native...

- ubuntu-riscv64-native... 50m 32s
- ubuntu-riscv64-native... 1h 2m
- ubuntu-riscv64-native... 1h 22m

Run details

- Usage
- Workflow file

# Runners: What's Next

## Next-gen RVA23 hardware

2-3x performance boost + VM-based runners

## Matching the GitHub runners experience

setup-python, setup-go, ccache, base image  
Everything should **just work** out of the box

## White-glove integration for flagship projects

PyTorch/Executorch, Kubernetes, Llama.cpp, and more



# Greening the Test Suite

## Contributors + contractors

Systematically working through test failures  
Fixing bugs, filing issues, upstreaming patches

The goal: green CI on RISC-V

- ✘ linux-noble-riscv64-py3.12 ^
- ✘ test (default, 1, 3,  
linux.riscv64.2xlarge.ephemeral  
)
- ✘ test (default, 2, 3,  
linux.riscv64.2xlarge.ephemeral)
- ✘ test (default, 3, 3,  
linux.riscv64.2xlarge.ephemeral)

# Better Hardware is Coming



# 2-3x

expected performance boost

Fully RVA23 compliant  
RVV (Vector) extensions enabled

Makes builds more tractable  
Enables testing with vector optimizations

Step 3

Ship it

# Enabling the Build



## Fork or upstream?

Working with the LF / PyTorch team to understand how we can collaborate

Goal: upstream, not a permanent fork

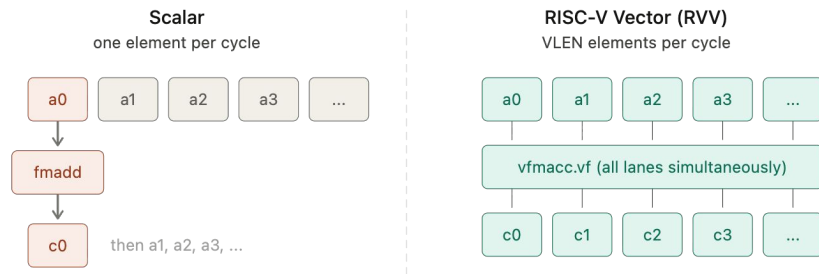
# Hardware Acceleration

## OpenBLAS + RVV

RISE-contributed RVV support, close to optimal for generic RVV  
*No uarch tuning yet*

## oneDNN

Big piece of work, in the RISE pipeline  
Critical for PyTorch performance on RISC-V



### Where RVV acceleration matters in PyTorch

- OpenBLAS + RVV** ✓ GEMM, GEMV, dot products  
Backbone of linear layers
- oneDNN + RVV** Ⓛ Convolutions, pooling, batch norm  
JIT-compiled kernels tuned per shape  
fp16, bf16, fp8 support

# ExecuTorch: The Embedded Angle

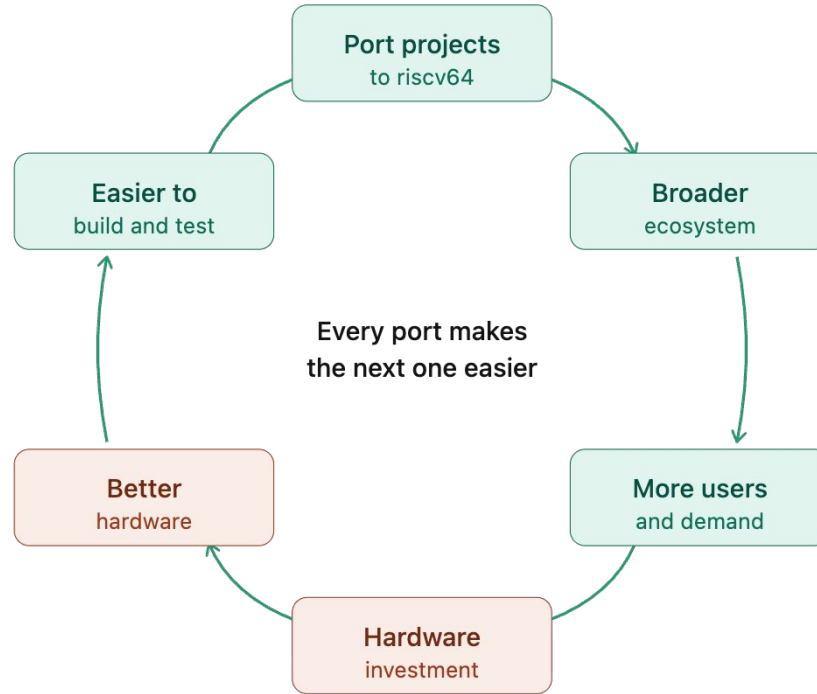
PyTorch for edge and embedded devices

RISC-V is already strong in embedded  
Natural convergence point

An important frontier for the ecosystem



# The Virtuous Cycle



# How You Can Help

Port your libraries to riscv64

Even basic builds + CI help the ecosystem

Use the RISE RISC-V Runners

Free CI for open source projects

File issues at [github.com/riseproject-dev/riscv-runner](https://github.com/riseproject-dev/riscv-runner)

Contribute upstream

Patches, bug reports, testing

Not if, but when.

# Q&A

# Get in Touch

Ludovic Henry  
[mail@ludovic.dev](mailto:mail@ludovic.dev)

<https://github.com/luhenry>  
[https://mastodon.social/@ludovic\\_dev](https://mastodon.social/@ludovic_dev)

RISE  
<https://riseproject.dev>

RISE RISC-V Runners  
<https://github.com/apps/rise-risc-v-runners>

