SOPHGO RISC-V Roadmap & RXU

SG2042 Empowering RISC-V in High-Performance Computing

RISC-V®

Our Vision:
Become the World's Leading Provider of General Computing Power

High-Performance RISC-V CPU RoadMap

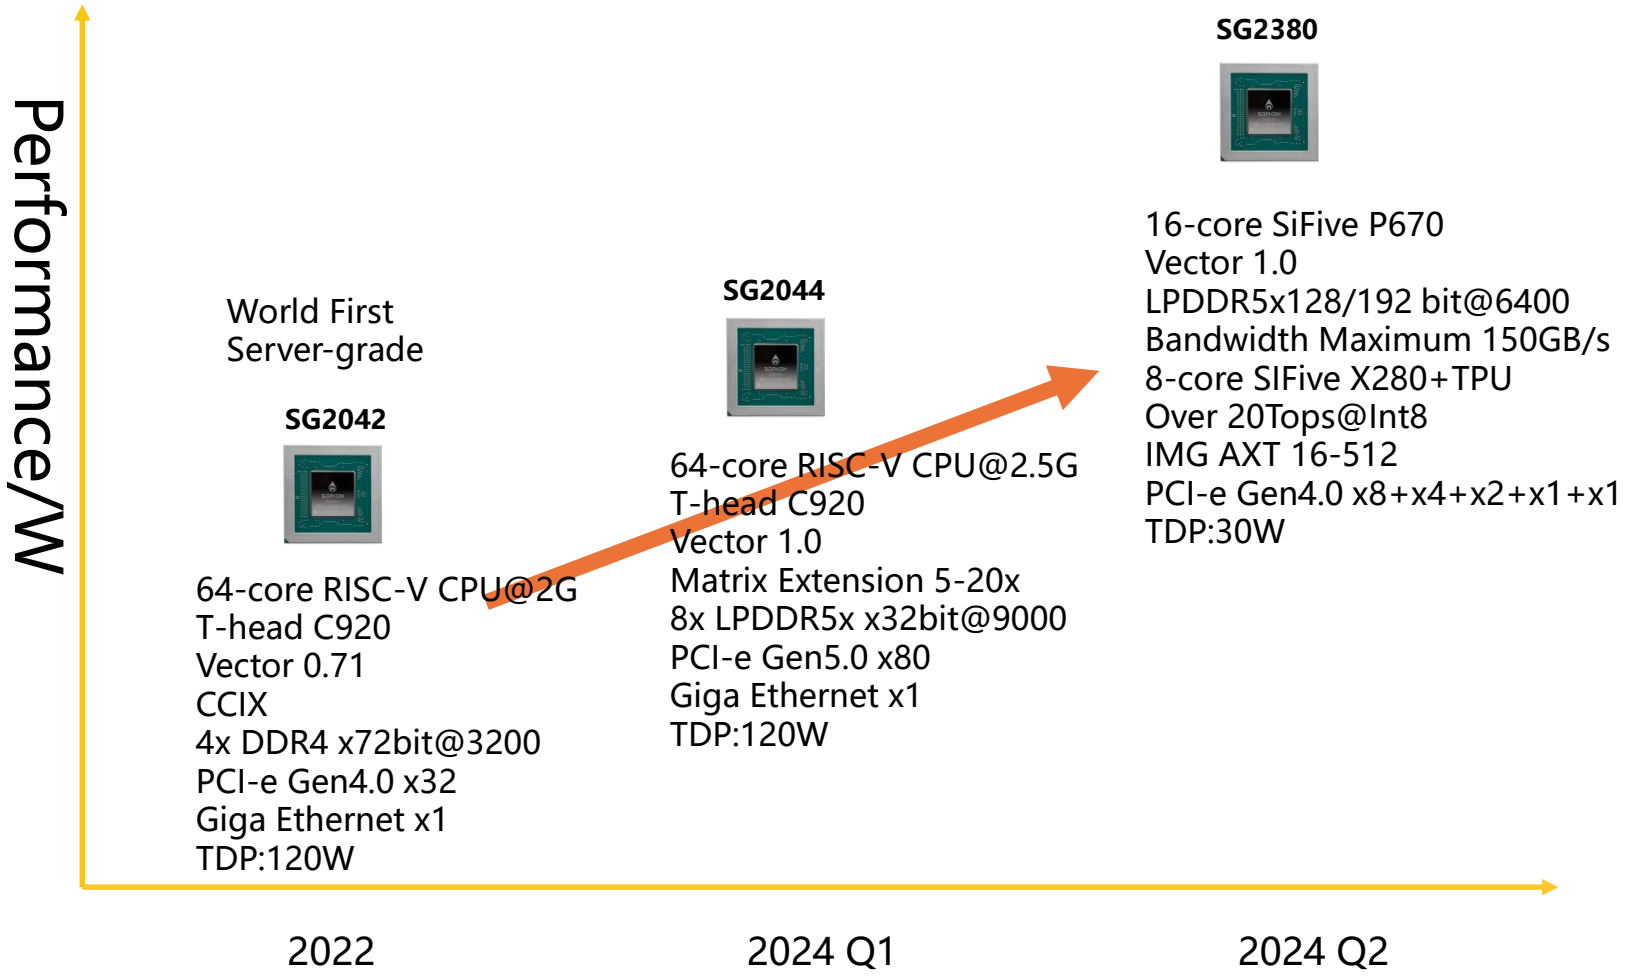

>1 TFLOPS(FP64)
Double floating-point computing power

64 Cores
Processor

2.0 GHz
CPU Frequency

120 W
TDP

1
1Gbps Ethernet RGMII

3200 MHz
Max DIMM Frequency

Released in Dec 2022
SG2042

0.7
RISC-V Vector

1
LPC

up to
256 GB
Max Memory Capacity

FCBGA
57mmx57mm
Package Size

2
SPI Flash Interface

4 MB
L1 I-Cache/L1 D-Cache

16 MB
Unified L2 Cache

64 MB
System Level L3 Cache

up to
16 GT/s
2 PCIe Gen 4 x16 Slots

2
General SPI Controller

support
2 eMMC5.1/SDIO3.0
4 Bit Data Width

4
PWM Generator for Fan control

support
4 I2C
00K/400/1M Clock Frequenc

4
Fan Speed control

4
UART

32
GPIO Pins

SG2042 Soc

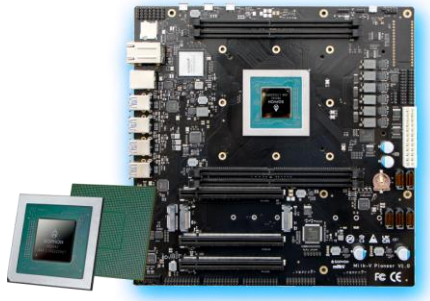

Milk-V Pioneer EVB

Milk-V Pioneer Box

SG2042
RISC-V 2U Server

SG2042
RISC-V 16U Cluster

RISC-V

About RISC-V | Membership | RISC-V Exchange | Technical | News & Events | Community

SOPHGO Donates 50 RISC-V Motherboards – Learn More About the Pioneer Box

By RISC-V Community News | June 5, 2023 | No Comments

New RISC-V International member SOPHGO is committed to the development and promotion of AI RISC-V CPU and other computing products. RISC-V member Milk-V delivers high-quality RISC-V products to developers, enterprises, and consumers, to promote the development of the RISC-V hardware and software ecosystem. Together, these two members have created a must-have RISC-V motherboard.

SOPHGO's 64-core RISC-V CPU codenamed SG2042 has been implemented into a motherboard by Milk-V and assembled as a RISC-V workstation named Pioneer Box.

The Pioneer Box is a standard mATX form factor. The PC-like interface, coupled with PC industrial compatibility, delivers native RISC-V development environments and one of the industry's most premier RISC-V desktop experiences. As RISC-V developers and hardware enthusiasts look for the latest motherboard, the Pioneer Box is a go-to option.

At this week's RISC-V Summit Europe, SOPHGO is donating 50 Pioneer Boxes to be shared with the RISC-V ecosystem. The company is committed to furthering the advancement of the open standard ISA and its supporting software ecosystem.

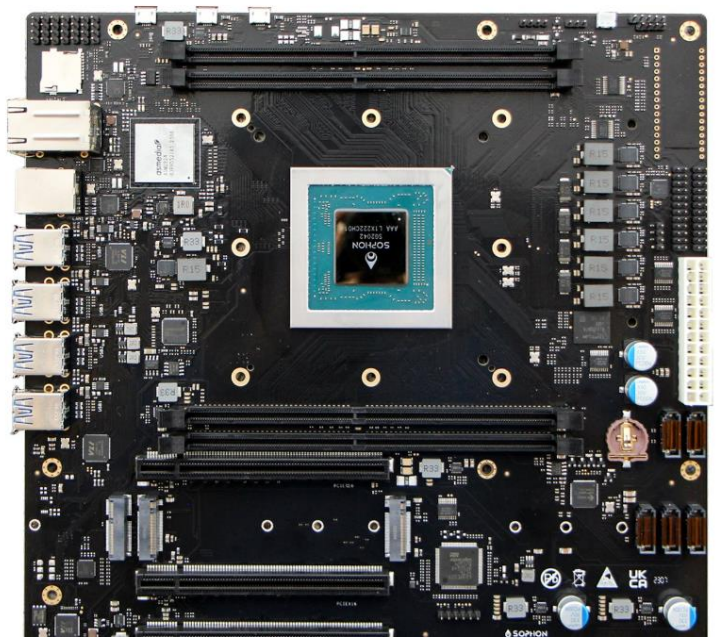

In the Press

Hackster News

"With promises to make schematics and design files available soon, the 64-core Milk-V Pioneer is an absolute monster of a board."

Tom's Hardware

"Milk-V...[is] building a new RISC-V powered line of computers... The Milk-V Duo looks like a Raspberry Pi Pico. There's also a Raspberry Pi shaped quad-core board and the current cream of the crop: a micro ATX system called the Milk-V Pioneer."

Tech Powerup

"Chinese RISC-V developers Milk-V Technology and SOPHGO recently announced their collaborative open source Milk-V Pioneer developer motherboard and workstation based on the SOPHON SG2042 RISC-V server SoC."

Sophgo is donating 50 Pioneer boxes to the **RISC-V foundation**.

SG2042 Newsletter

Upstream:

Linux 6.7

- Added support CV1800B chip found in the milkv-duo board
- Added SG2042 chip found in the milkv-pioneer 64-core developer workstation. (Currently only support booting into console with only UART)

Tutorials:

RISC-V Public Test Platform Released · Exploring Docker on SG2042
RISC-V public beta platform released · FFTW porting and performance comparison
RISC-V public beta platform released · CoreMark test report
RISC-V public beta platform release · SG2042 OpenMPI Test

Join Us!


IN 2024
SG2044 & SG2380
RISC-V®

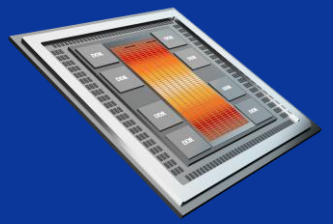

SG2044

Q1, 2024

Vector 1.0
RISC-V 64 core@2.5Ghz
T-head Martix Extension, LLaMA-7b 40 tok/s

Larger Bandwidth up to 300GB/s
8x DDR Control
PCI-e Gen5.0 x 80

Large Language Model

An open-source enthusiast：He Wang

added support for the RISC-V platform to the high-performance large model inference framework 'inferllm', enabling smooth operation of models like ChatGLM2, LLaMA2, Alpaca, Mistral and Baichuan-13B on the SG2042

Pure SG2042 CPU can run a llm model like LLaMA-7B at a speed of 5 tokens per second @INT8 while SG2044 at least 20 tokens per second

ChatGLM2

FALCON

LLAMA

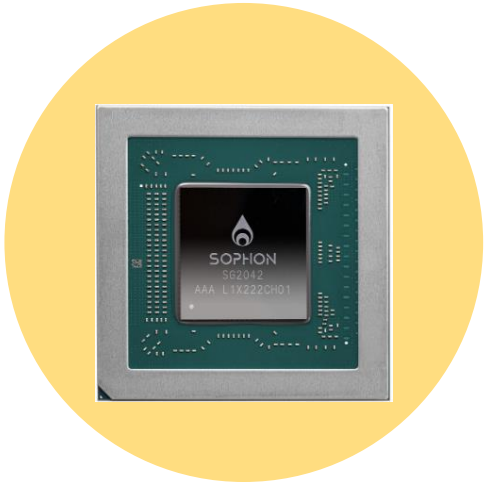

CodeGeeX2
A Multilingual Code Generation Model

Stable Diffusion XL

Whisper

RISC-V

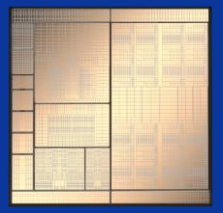

SG2380

Q4, 2024

Vector 1.0
SiFive 16-core P670, 12P, 4E
8-Core SiFive X280 and SOPHGO TPU
20 Tops@INT8
IMG AXT-16-512
Larger Bandwidth up to 150GB/s
PCI-e Gen4.0 x 16, Flexibility
TDP 30W, IDLE 15W

CPU: 12 P Cores@2.5Ghz & 4 E Cores@1.6Ghz
GPU: IMG AXT-16-512
TPU: 8-Core SiFive X280+SOPHGO TPU
Vector 1.0
SiFive Matrix Mul Extension
Max: 192/128 bit LPDDR5@6400
Min: 128 bit LPDDR4x@4266
TDP: Max to 30W
Temperature: -40℃ ~ +125℃

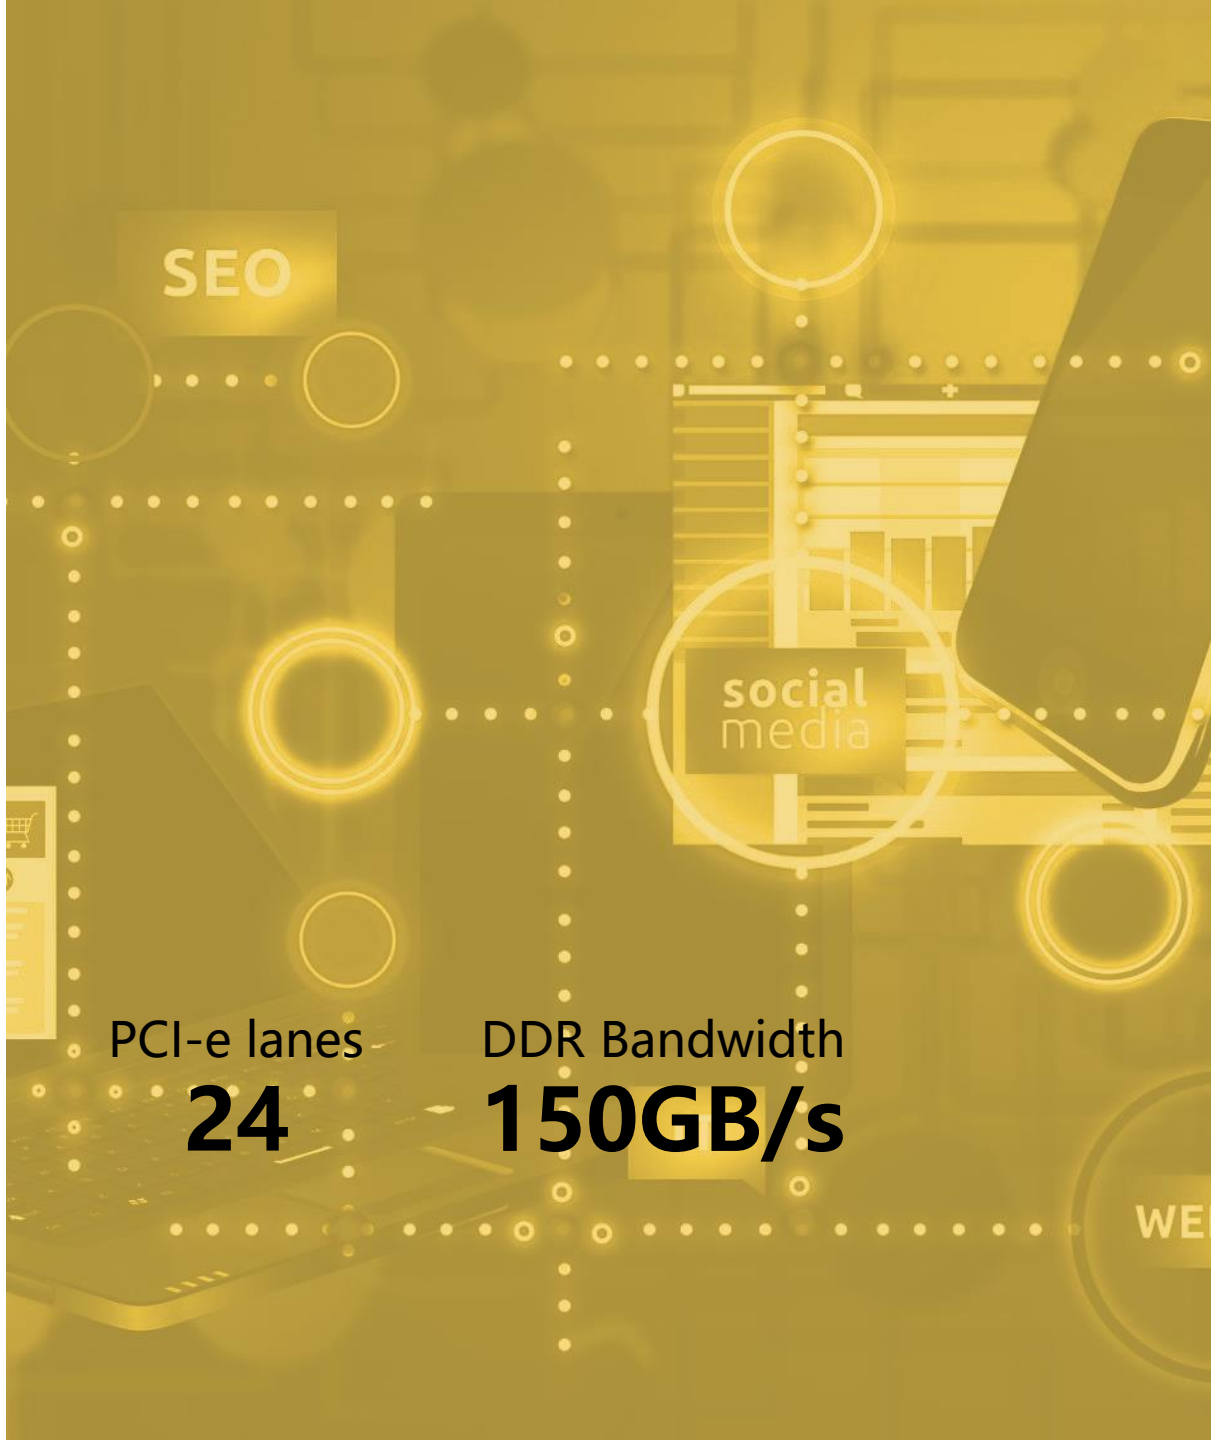

RISC-V Core	Vector	PCI-e lanes	DDR Bandwidth
16(12+4)	**1.0**	**24**	**150GB/s**

RISC-V®

PCIe 4.0 x 16

1x16 or 1x8 + 2x4
or 4x2 + 2x4 lanes PCIe Gen4 root complex, or 4 x2 + 1x4 PCIe Gen4 root complex + 1x4 PCIe Gen4 end point

PCIe 3.0 x 4

4x1 lanes high speed interface, configured as 4x1 lane Gen3 root complex, or ,or 2x 1/25/5/10/25Gbps Ethemet and 2 x 1/2.5/5 Ethernet

- 1/2.5/5/10Gbps Ethernet support IEEE1588 / AVB TSN
- Ethernet support 1000BASE-KX, 2.5GBASE-KX, 5GBASE-KR, 10GBASE-KR, 10GBASE-KX4, XAU, 25GAU C2C, 25GAUI C2C, 25GAUI C2M interface

PCIe 3.0 x 4

4 x1 lanes high speed interface, can be configured as 4 x1 lane PCIe Gen3 root complex, or 4 x SATA 3.0

- SATA 3.0 support AHC

User-friendly Toolchain

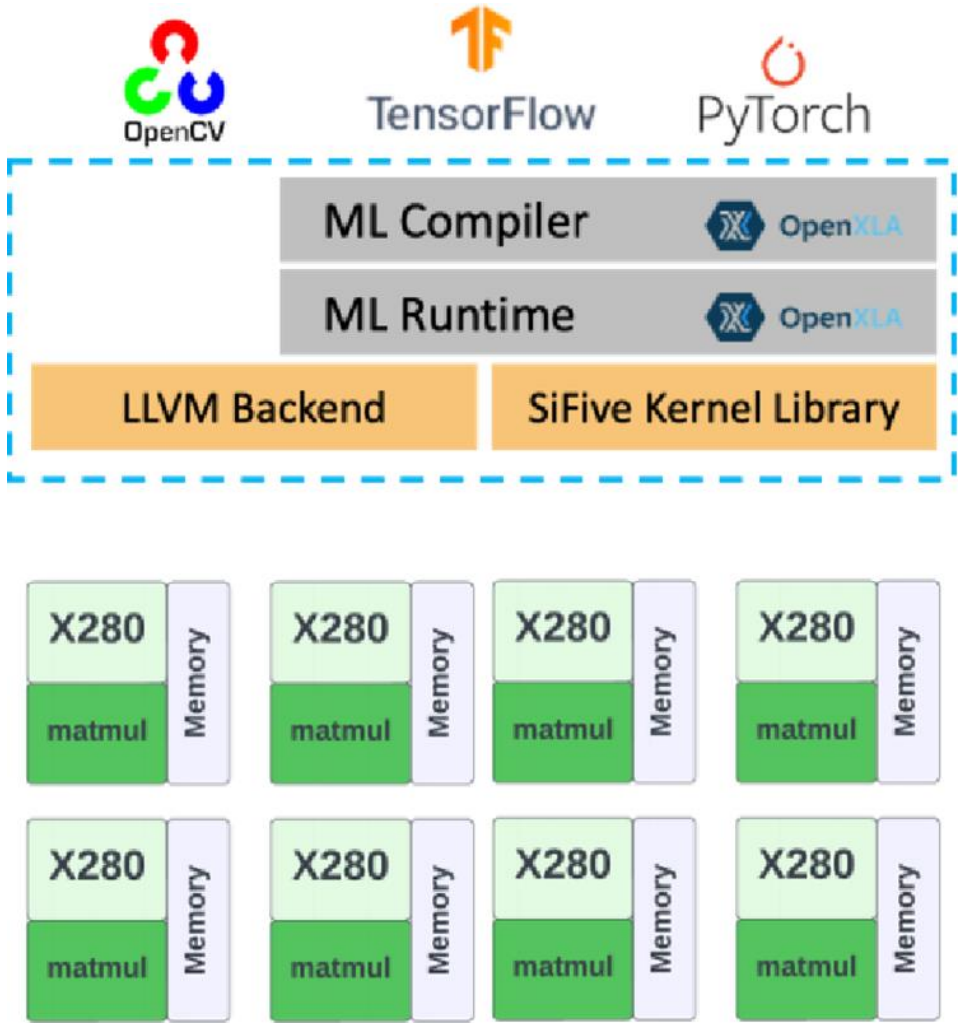

Deep Learning Compiler

OpenXLA MLIR Compiler/TVM

SiFive LLVM Backend

RISC-V®

Data Center

Cloud, HPC

New features specific to HPC such as Vector and larger Virtual address space，Hypervisor Support(AIA, H Ext and IOMMU)

Automotive

Customer algorithms (driver monitoring, driver assistance, etc.)
Multi-screen different display, CAN FD, support for 4G/5G/WiFi/BT
Wide temperature range (industrial grade), Android Support

AI/ML

Artificial intelligence is incorporated across 280s and TPU, over 20Tops@INT8
Especially For LLM(Large BW)

Telecom & Communications

4G/5G/WiFi 6/WIFI 7/BT and More
PCIE Gen 4.0 with 16 lanes
2 x Gbit Ethernet MAC, RGMII interface, support 10/100/1000 Mbps, support IEEE1588 / AVB / TSN

Consumer and IoT devices

Low Power Consumption, Power Switch, IDLE Power >10W and TDP 30W
Android Support
Fexible PCIE Configuration

Edge Computing

CPU、GPU and TPU
over 20Tops@INT8 and 8Tflops@FP16
LLM and CV Models
Useful and Easy Toolchain(LLVM backend)

2380 adoption spans industries

RISC-V®

RXU High Performance RISC-V Core From 2024.6.30

The first high-performance RISC-V core shared with members

Free access to SRC clusters and **MPW** services

Shared resources for chip design and verification platforms

Support for hardware-software co-design services, assisting university projects to fruition

Provision of backend layout technology services

RISC-V®

RXU Membership

RXU membership plan:
Membership fee set at 2.5 million US Dollars for 3 years

Release: June 30, 2024
(RXU2.6)
Benchmarked against single-core A72
Specifications: 8-Way Out-of-Order
RV64IMDFZcsr, supports **MMU**, **interrupts**,
Timer, G instruction set, capable of running a
full Linux OS

Release: June 30, 2025
(RXU3.5)
Benchmarked against X1
Specifications: 8-Way Out-of-Order,
RV64GCV**H**, SPECINT2K6 >**20**/GHz,
Vector 1.0, supports multicore,
IOMMU, AIA, RVA23, **Hypervisor**

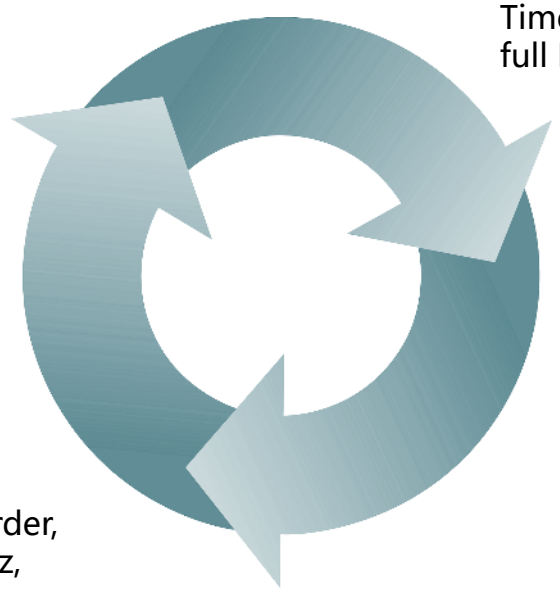

Release: December 30, 2024
(RXU3.0)
Benchmarked against A77
Specifications: 8-Way Out-of-Order,
RV64GCV, SPECINT2K6 >15/Ghz,
Vector 1.0, AIA, supports multicore

Release: May 30, 2025
(RXU3.2)
Benchmarked against A78
Specifications: 8-Way Out-of-Order,
RV64GCV, SPECINT2K6 >15/GHz,
Vector 1.0, supports multicore,
IOMMU, AIA, RVA23

RISC-V®

Graded Licensing

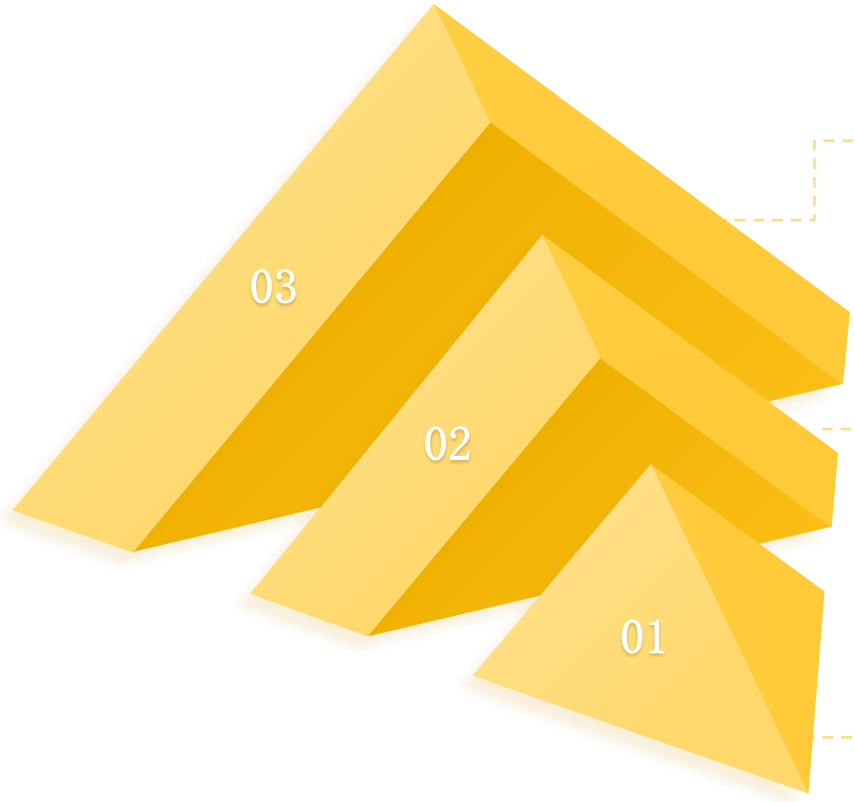

Others

L3 member pay licensing fees, support fees, and royalty fees to the L2 units based on the proportion of intellectual property ownership.

Universities

L2 can develop derivative IPs and industry-specific chips based on L1 technology, and they can agree on the terms of licensing.

SOPHGO

L1 Level Member, Provides IP, source code, and other materials along with free MPW support.

The world's first high-performance RISC-V core
shared with members
offering free MPW and SRC services.
Assisting universities in exploring the possibilities in directions such as
2.5D/3D Chiplet and in-memory computing.

Thank you

Info: zihan.wang@sophgo.com
qihang.fan@sophgo.com