Building an Open HPC Ecosystem

Special Interest Group for High Performance Computing

John D. Davis, Chair, SIG-HPC

Overview

- The HPC Vision
- SIG HPC Goals
- SIG HPC Activities
- RISC-V HPC Research in Europe
- Concluding thoughts

Open Ecosystem HW/SW Co-Design

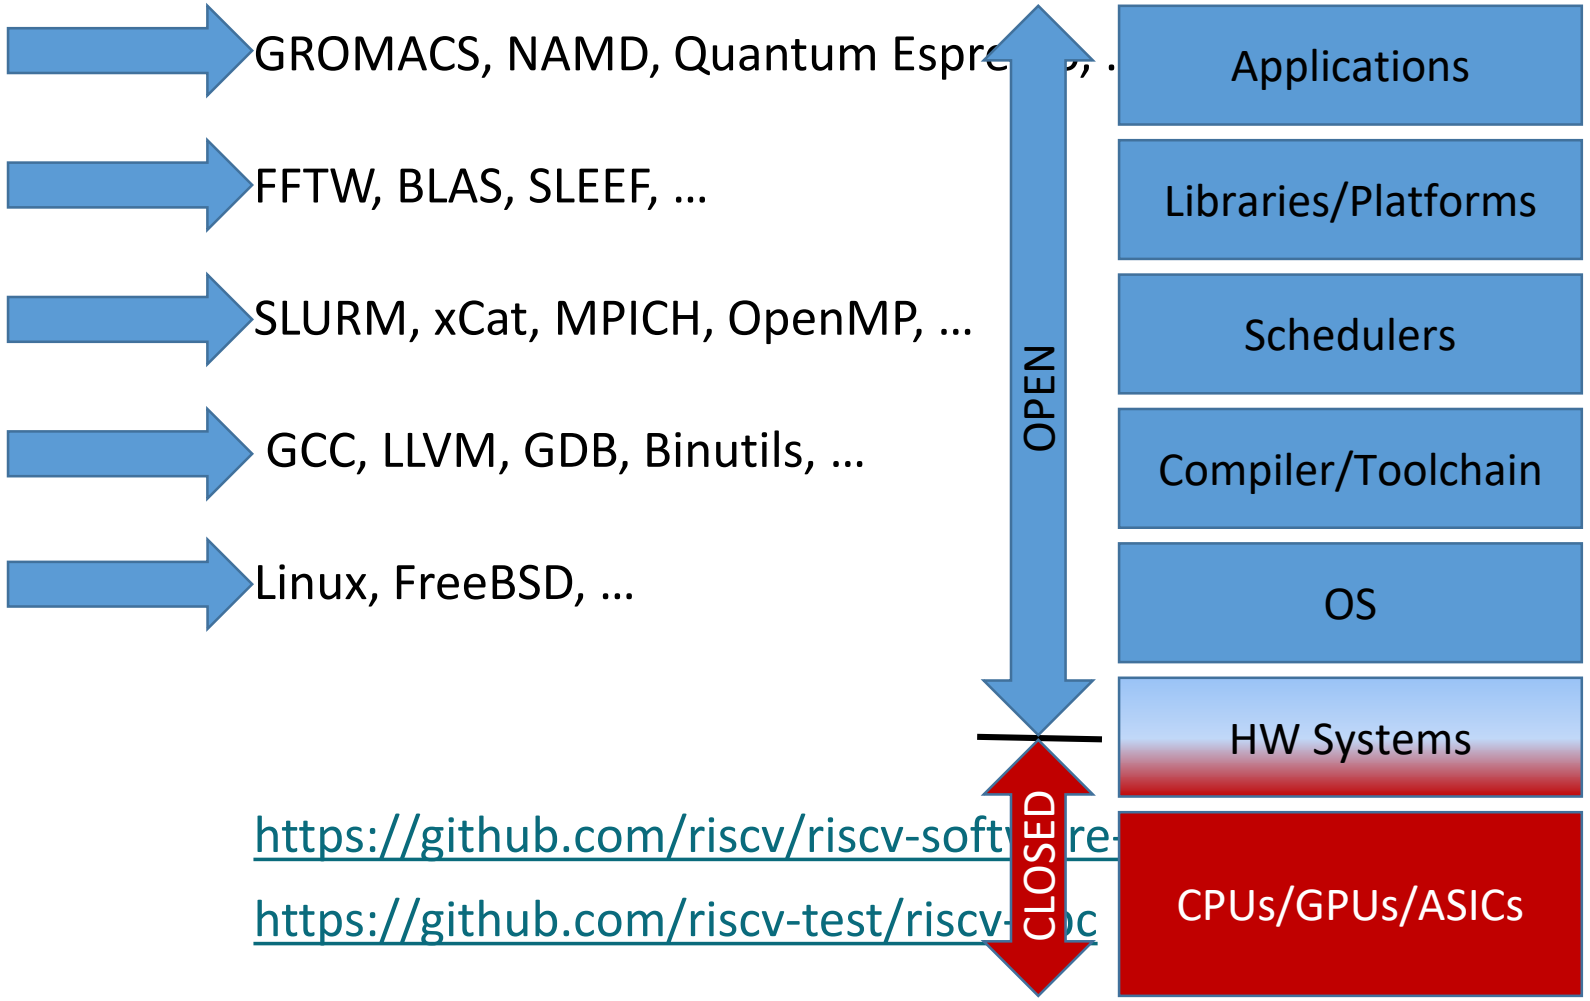

HPC Tomorrow

- Europe can lead the way to a completely **open SW/HW stack for the world**
- RISC-V provides the open source hardware alternative to dominating proprietary non-EU solutions
- Europe can achieve complete technology independence with these foundational building blocks
- Currently at the same early stage in HW as we were with SW when Linux was adopted many years ago
- **RISC-V can unify, focus, and build a new microelectronics industry in Europe.**

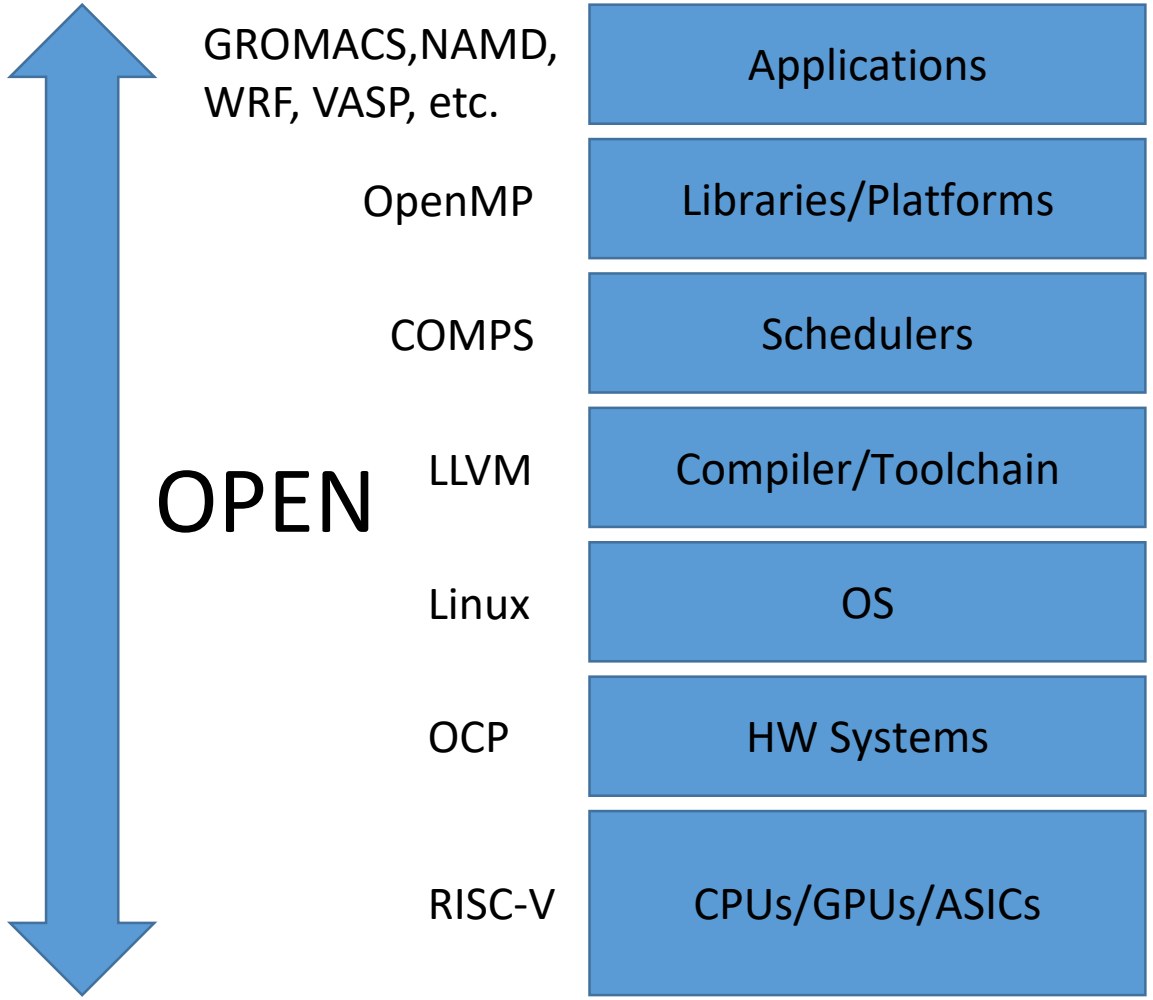

RISC-V®

HPC [Open] Source Software for Open Source Hardware

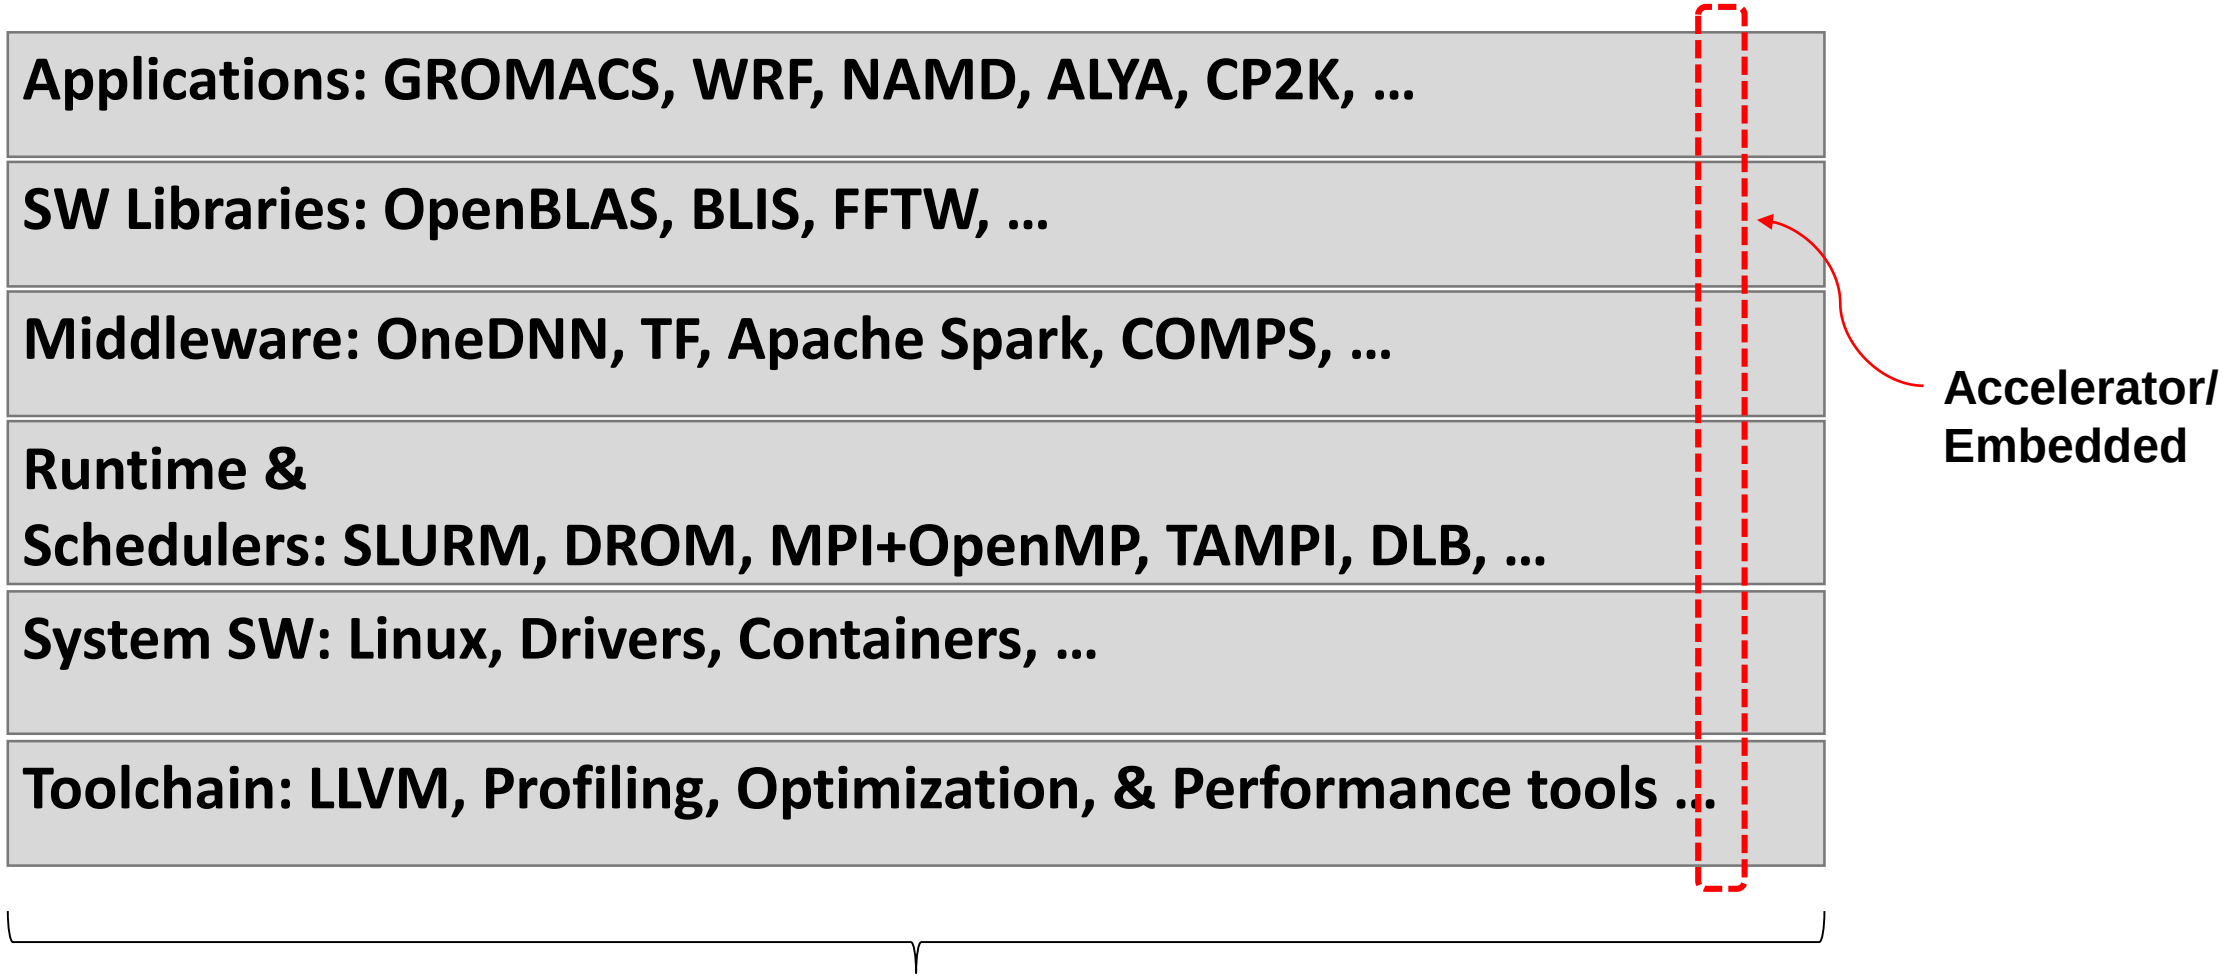

RISC-V®

SIG-HPC Vision & Mission: RISC-V: IoT to *HPC*

Vision:

The technical and strategic imperatives that guide the RISC-V ecosystem development to enable an Open HPC Ecosystem…

Mission:

…enable RISC-V in a broader set of new software and hardware opportunities in the High Performance Computing space, from the edge to supercomputers, and the software ecosystem required to run legacy and emerging (AI/ML/DL) HPC workloads.

RISC-V®

SIG-HPC: An Open era of HPC!

- CPUs, Accelerators, other hardware units, and coprocessors
- Verification and compliance infrastructure and methodologies specific to HPC
- Alignment and engagement and IP enablement.
- RISC-V software ecosystem alignment
- Engage and represent RISC-V in compute intensive industry and academic events
- Identify key industrial and academic partners.
- Support global technology independence with a RISC-V ecosystem roadmap and partners

SIG-HPC Initiatives

- Guide and enable the community
  - Virtual Memory
    - SV57, SV57K, SV64, SV128
  - HPC SW & HW ecosystem & roadmap
  - Accelerators
  - ISA Extensions
  - **HPC Software Stack**
    - Starting with HPC Libraries
  - **New Announcement!**

Tactical Computing Labs

HPC Software Testbed

John Leidel, Ph.D.

Tactical Computing Labs, RISC-V Technology HC Chair & SIG-HPC Co-Chair

Categories of Software

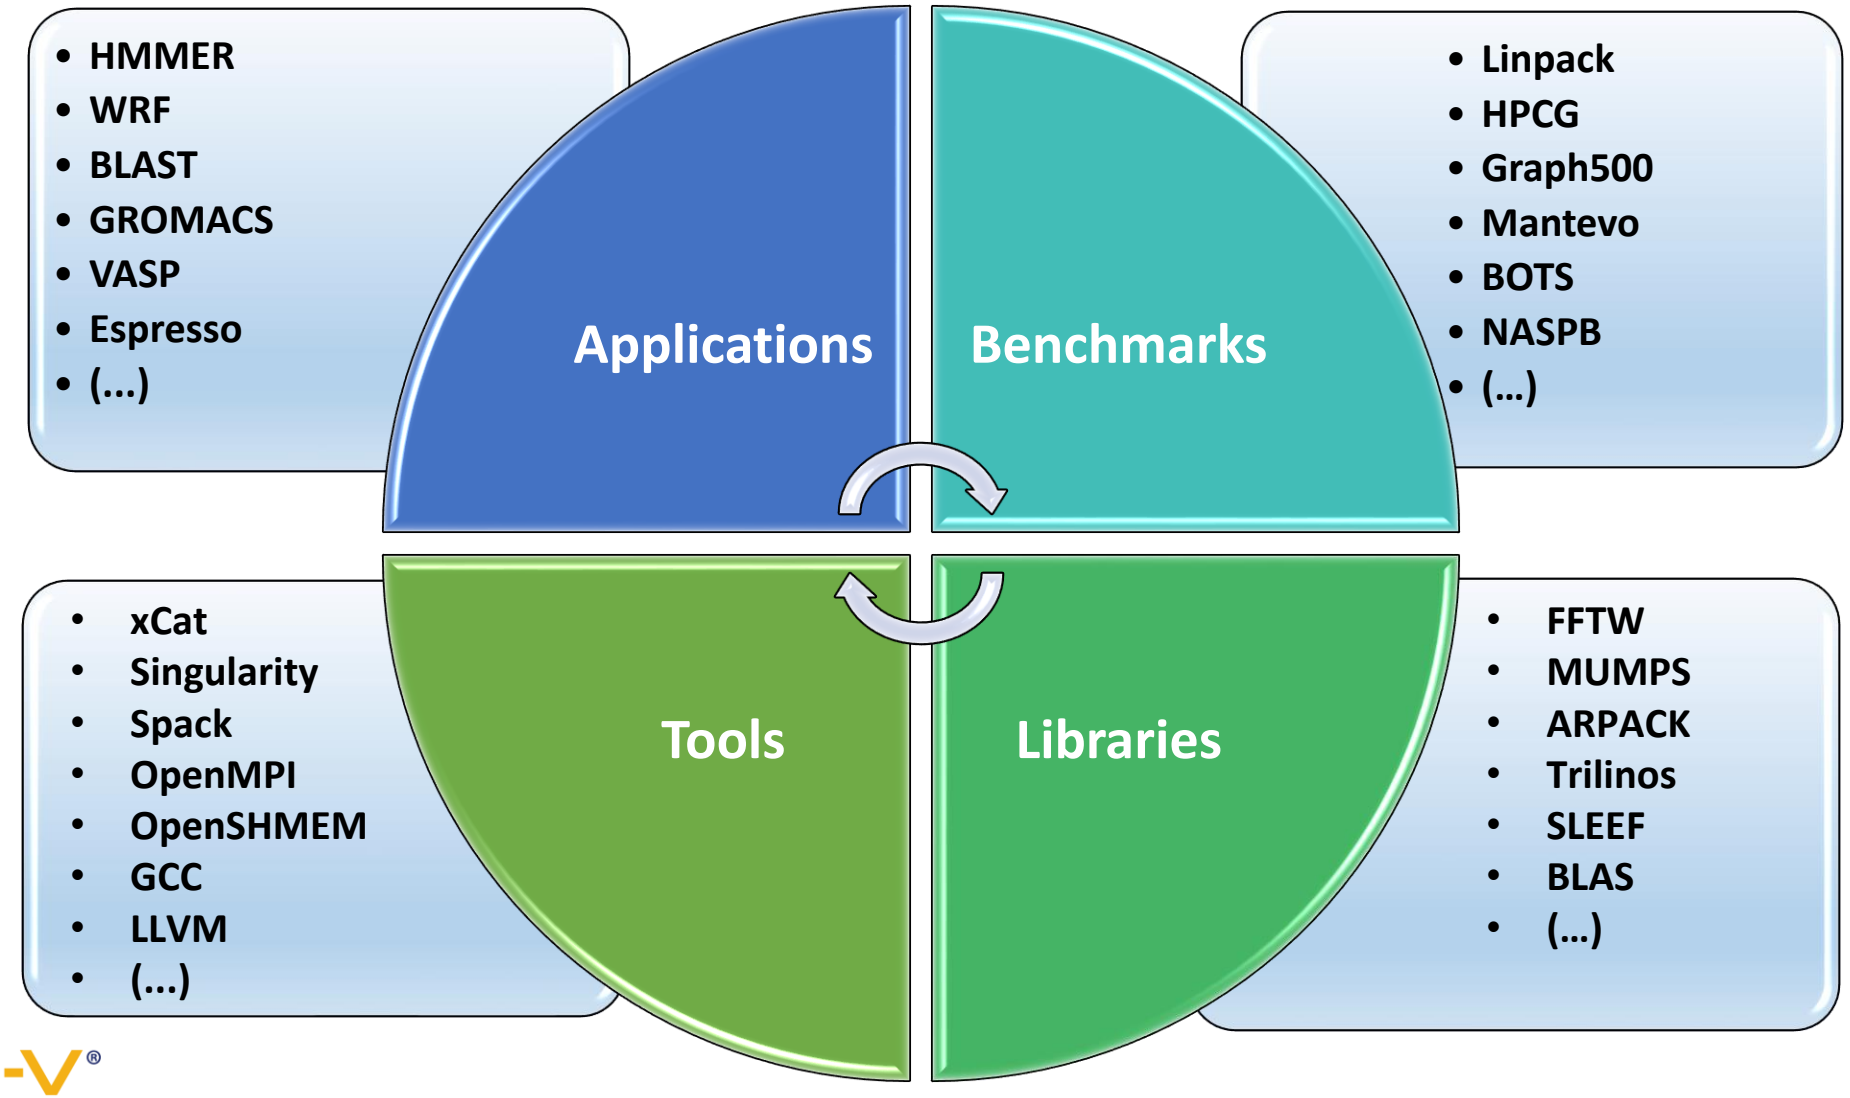

RISC-V®

HPC Software Testbed

- We are working with the RISC-V International group to drive requirements for adjacent working groups
- RISC-V HPC Tests
  - HPC-centric software test suite
  - Multi-version compiler centric: GCC, LLVM
  - Using each compiler, we cross compile each target library, benchmark and application suite for RISC-V compatibility

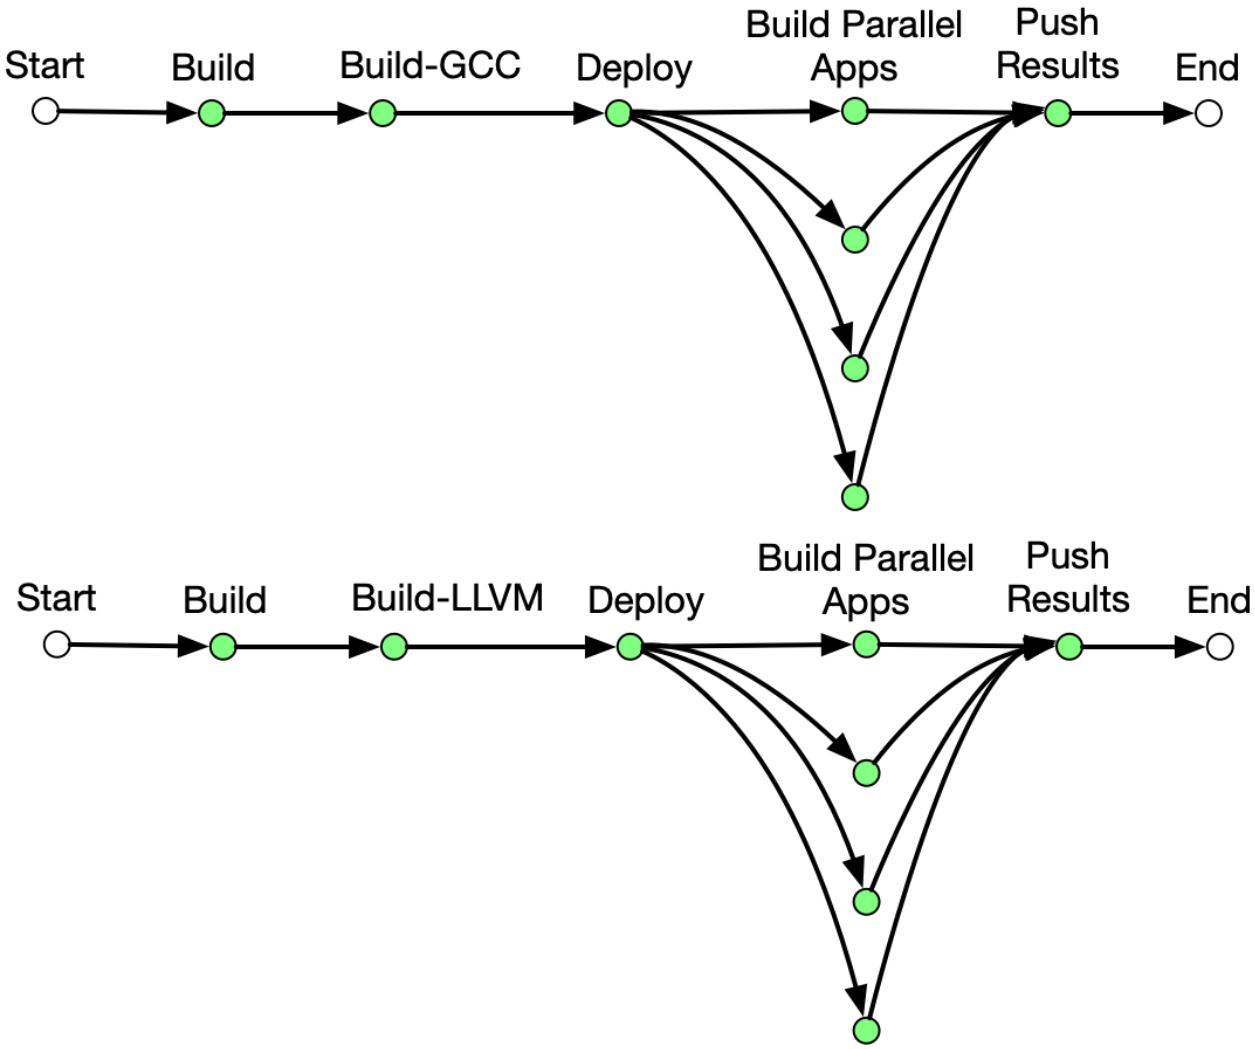

HPC Software Testbed: riscv-test.org

RISC-V Test

All

S	W	Name ↓	Last Success	Last Failure	Last Duration
✓		llvm-project-11.0.0	6 days 12 hr - #38	7 mo 9 days - #1	37 min
✓		llvm-project-11.0.1	6 days 12 hr - #33	6 mo 15 days - #1	3 hr 51 min
✓		llvm-project-12.0.0	6 days 13 hr - #6	1 mo 0 days - #1	4 hr 18 min
✓		llvm-project-12.0.1	6 days 12 hr - #5	N/A	4 hr 15 min
✓		llvm-project-master	6 days 12 hr - #62	1 mo 4 days - #58	7 hr 20 min
✓		riscv-gnu-toolchain-master	6 days 12 hr - #60	5 mo 21 days - #36	2 hr 47 min

Icon: S M L

Legend Atom feed for all Atom feed for failures Atom feed for just latest builds

RISC-V®

HPC Software Testbed: Public Results

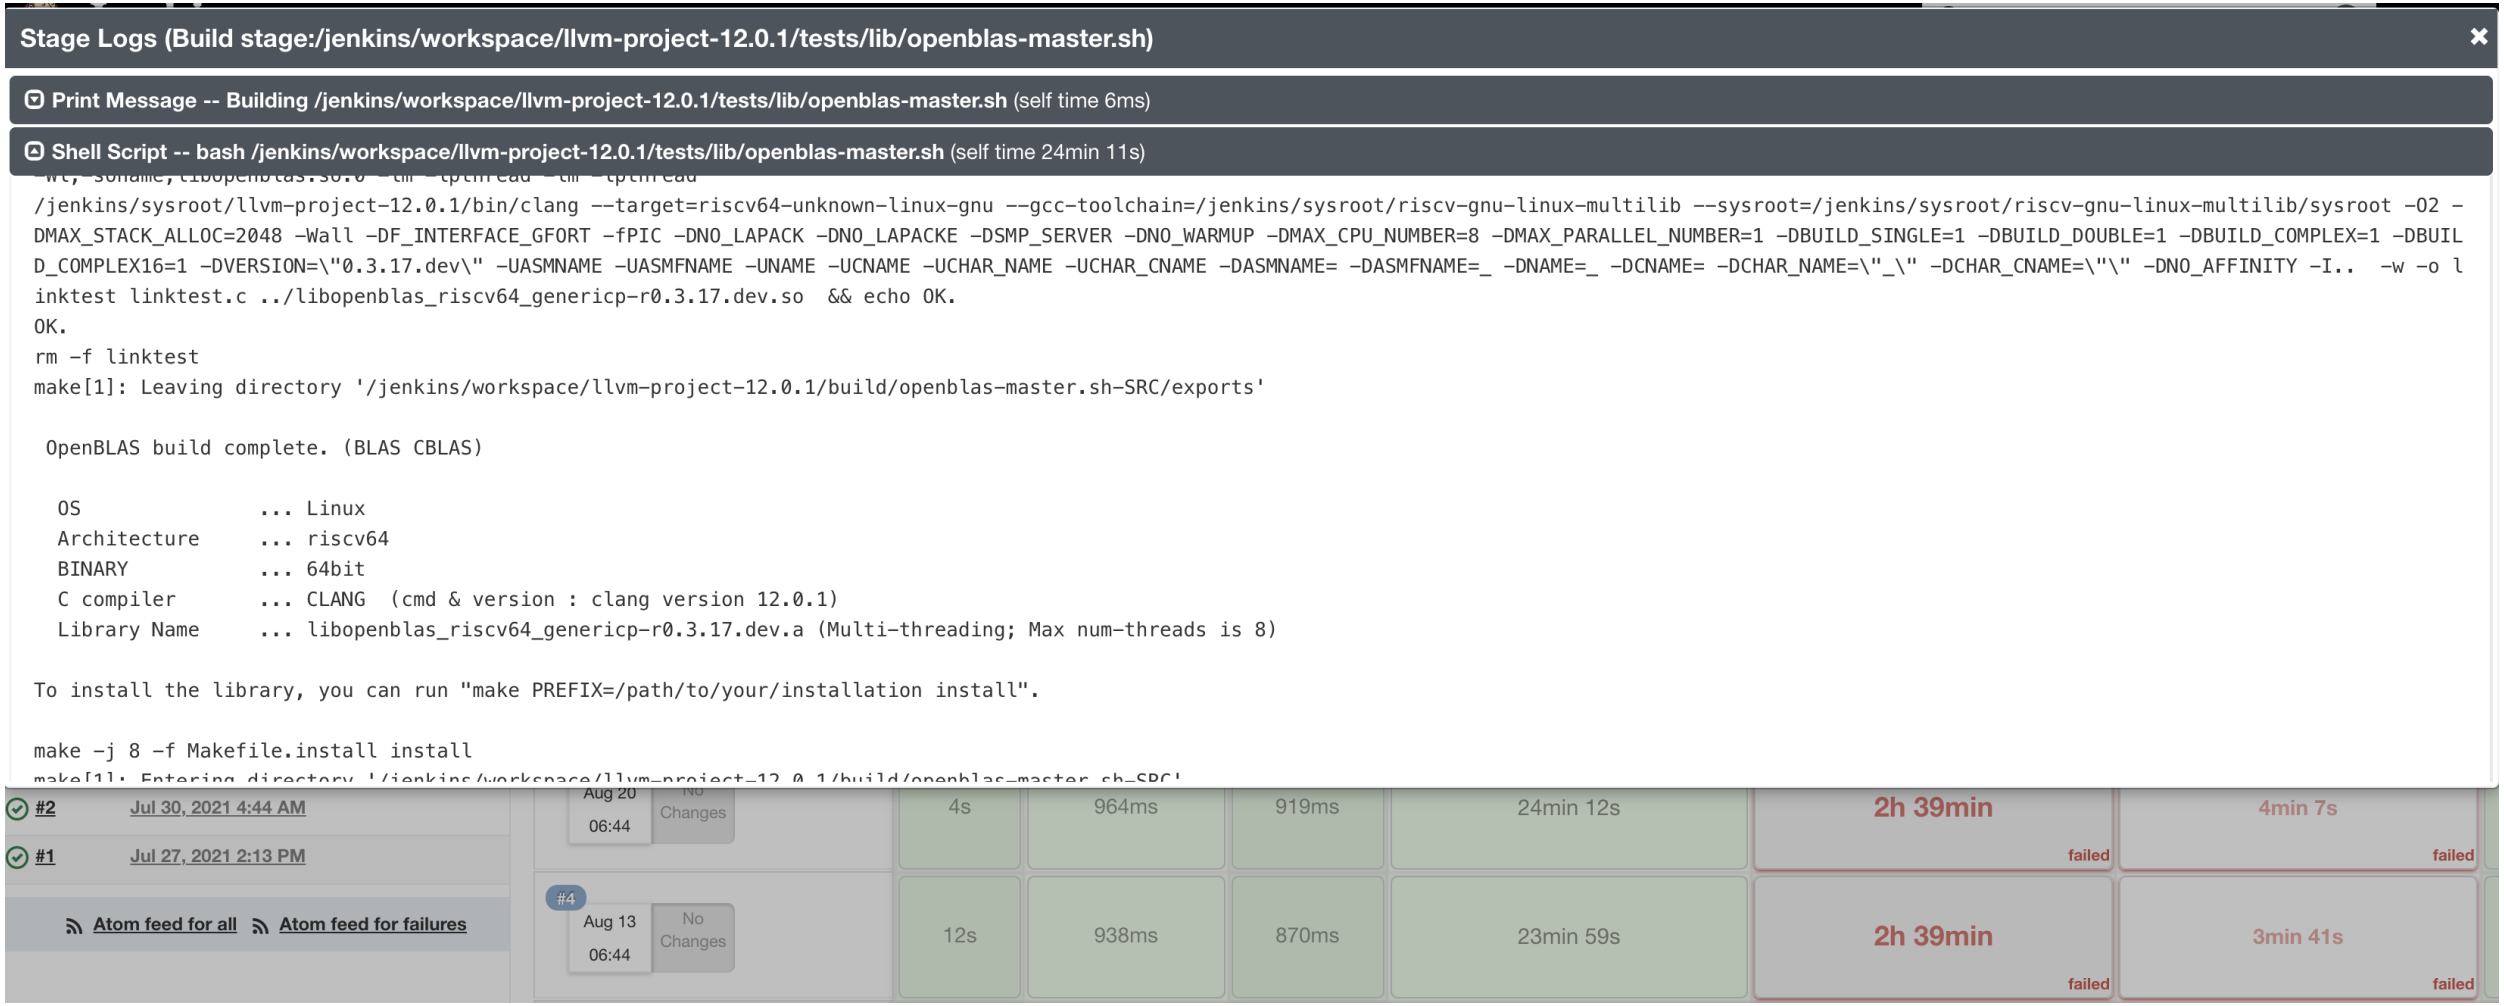

Public Test Harness

- The results are public!
  - https://riscv-test.org/
- The test harness has been created in a public Github repo
- https://github.com/riscv-test/riscv-hpc
- The harness includes a set of template scripts for each test
  - Each test pulls it own source code (archive or Git repo)
  - Unpacks the source, configures it and builds it
  - Adding tests is as simple as adding scripts to the Github repository
- The top-level harness provides common access to:
  - Directory structures
  - Compiler paths
  - Compiler flags
- **Pull requests are encouraged!**

riscv-hpc

license BSD

This repository serves as the basis for the official RISC-V HPC test suite that is executed via the Jenkins CI host at https://riscv-test.org.

License

The RISC-V HPC test suite is licensed under a BSD-style license - see the LICENSE file for details.

Architecture Overview

The RISCV-HPC test architecture is setup specifically to support testing various compilers, libraries benchmarks and applications for high performance computing against the RISC-V software ecosystem.
The test architecture is crafted specifically to support inclusion of a variety of different software tools that can be built using different compilers and/or compiler versions. Note that these software entities are NOT executed. We do not execute and/or track benchmark performance of any software package, benchmark or application. Rather, this infrastructure is designed to find the gaps in the RISC-V software ecosystem.

As we see in the figure below, the test infrastructure is crafted using a series of build scripts integrated into a Jenkins pipeline. Each pipeline instance is initiated at a specific cadence (documented in the Jenkins environment). The first stage of the pipeline initializes the test environment. The second stage of the pipeline initializes a set of global variables that are utilized by individual test scripts. These global variables initialize values such as the absolute installation path of the target compiler, the required compiler flags and other various paths. The third stage of the pipeline builds the target compiler. For compilers that are designated as "release" builds, we utilize a previously built compiler (as release builds rarely change).
For compiler builds that target the top of tree source code, we build and install the entire compiler from scratch. Once the compiler build has been deemed stable, we execute three sets of nested pipeline stages. The first stage builds and installs candidate libraries. This can be from release archives or from top-of-tree git repositories. We

InfiniBand on RISC-V

- Infiniband working on RISC-V Unmatched Board
  - https://network.nvidia.com/pdf/prod_adapter_cards/PB_ConnectX3_VPI_Card.pdf
  - System boots and the card shows up in the necessary configuration points
    - standard kernel drivers without manual intervention
  - Running Ubuntu 21.04 and 21.10 with factory kernels
    - 21.04+ on RISC-V even offers all the IB tools via the standard apt repositories
    - Subnet manager has the normal issues

`ibstat` shows this (for one of the two ports):

```
CA 'mlx4_0'
CA type: MT4099
Number of ports: 2
Firmware version: 2.40.7000
Hardware version: 1
Node GUID: 0x7cfe9003009ece10
System image GUID: 0x7cfe9003009ece10
Port 1:
State: Active
Physical state: LinkUp
Rate: 40
Base lid: 0
LMC: 0
SM lid: 0
Capability mask: 0x00010000
Port GUID: 0x7efe90fffe9ece10
Link layer: Ethernet
```

ConnectX3 Infiniband not out of the Box

- The latest Linux kernels are shipped with sufficient drivers to run most of the IB stack. This includes the basic drivers and IP over IB
- The latest RISC-V kernels lack RDMA driver support for Mellanox devices
- You can download the source for the RDMA drivers directly from Mellanox, but due to some recent kernel changes, they do not build on RISC-V Linux
- This may be solved by moving to more recent IB devices (ConnectX-4+), but there is also a planned fix coming in forthcoming kernel releases.

RISC-V BoF @ ISC ’22

- **Creating an Open HPC Ecosystem with RISC-V**
- **Monday, May 30, 2022 2:30 PM to 3:30 PM**
  - **David Donofrio, Tactical Computing Labs**
  - **Doug Norton, Inspire Semiconductor**
  - **Michael Wong, Codeplay Software**
    - **Also leads the RISC-V Datacenter Group**
- Discuss the SW/HW RISC-V Ecosystem
  - 20 minutes of presentation/40 minutes of discussion
- Come join us for the conversation

RISC-V®

RISC-V @ SC'22 in Dallas ???
November 13-18

Coming Soon: SUPERcomputing Risc-V LAB **SUPER-V** @ BSC

- Enabling the development of the HPC ecosystem for RISC-V based systems
- Variety of systems
  - RISC-V clusters running HPC software stack (i.e., Unmatched cluster)
  - RISC-V Experimental/research platforms for vector architectures
    - FPGA-based system
    - Software emulators
    - Hybrid software/hardware emulators
- HPC Software ecosystem development
- And more…
- Access information coming soon…
  - Easybuild and Gentoo first success story

RISC-V®

SUPER-V@ BSC

- Clusters:

Board	OS	Details
PolarFire	Fedora	4 cores w/ 2 GB
BeagleV	Fedora	2 cores w/ 8 GB
Unmatched	Fedora/Ubuntu	4 cores w/ 16 GB
Allwinner D1 (Vector extension)	Fedora	1 core w/ 2 GB

Emulators:

- A RISC-V soft vector core running in an FPGA.
- The Vehave RISC-V emulator on top of QEMU
- The Vehave RISC-V emulator on top of a native RISC-V core

- RISC-V Software Stack:
  - Linux, SLURM
  - Compilers:
    - go/1.17
    - openmpi/fedora/4.1.1_gcc10.3.1
    - llvm/EPI-0.7-development
    - openmpi/ubuntu/4.1.1_gcc10.3.0
    - llvm/EPI-development
    - python/fedora/2.7.16
  - Tools
    - extrae/3.8.3
    - papi/6.0.0
    - perf/5.11.10
    - singularity/3.8.2
  - Libraries
    - boost/1.77.0
    - glibc/fedora/2.33
    - openBLAS/0.3.15
    - fftw/3.3.9_gcc10.3.1_ompi4.1.1
    - libunwind/git
    - openBLAS/0.3.17

RISC-V®

Building Open European HPC CPUs & Accelerators

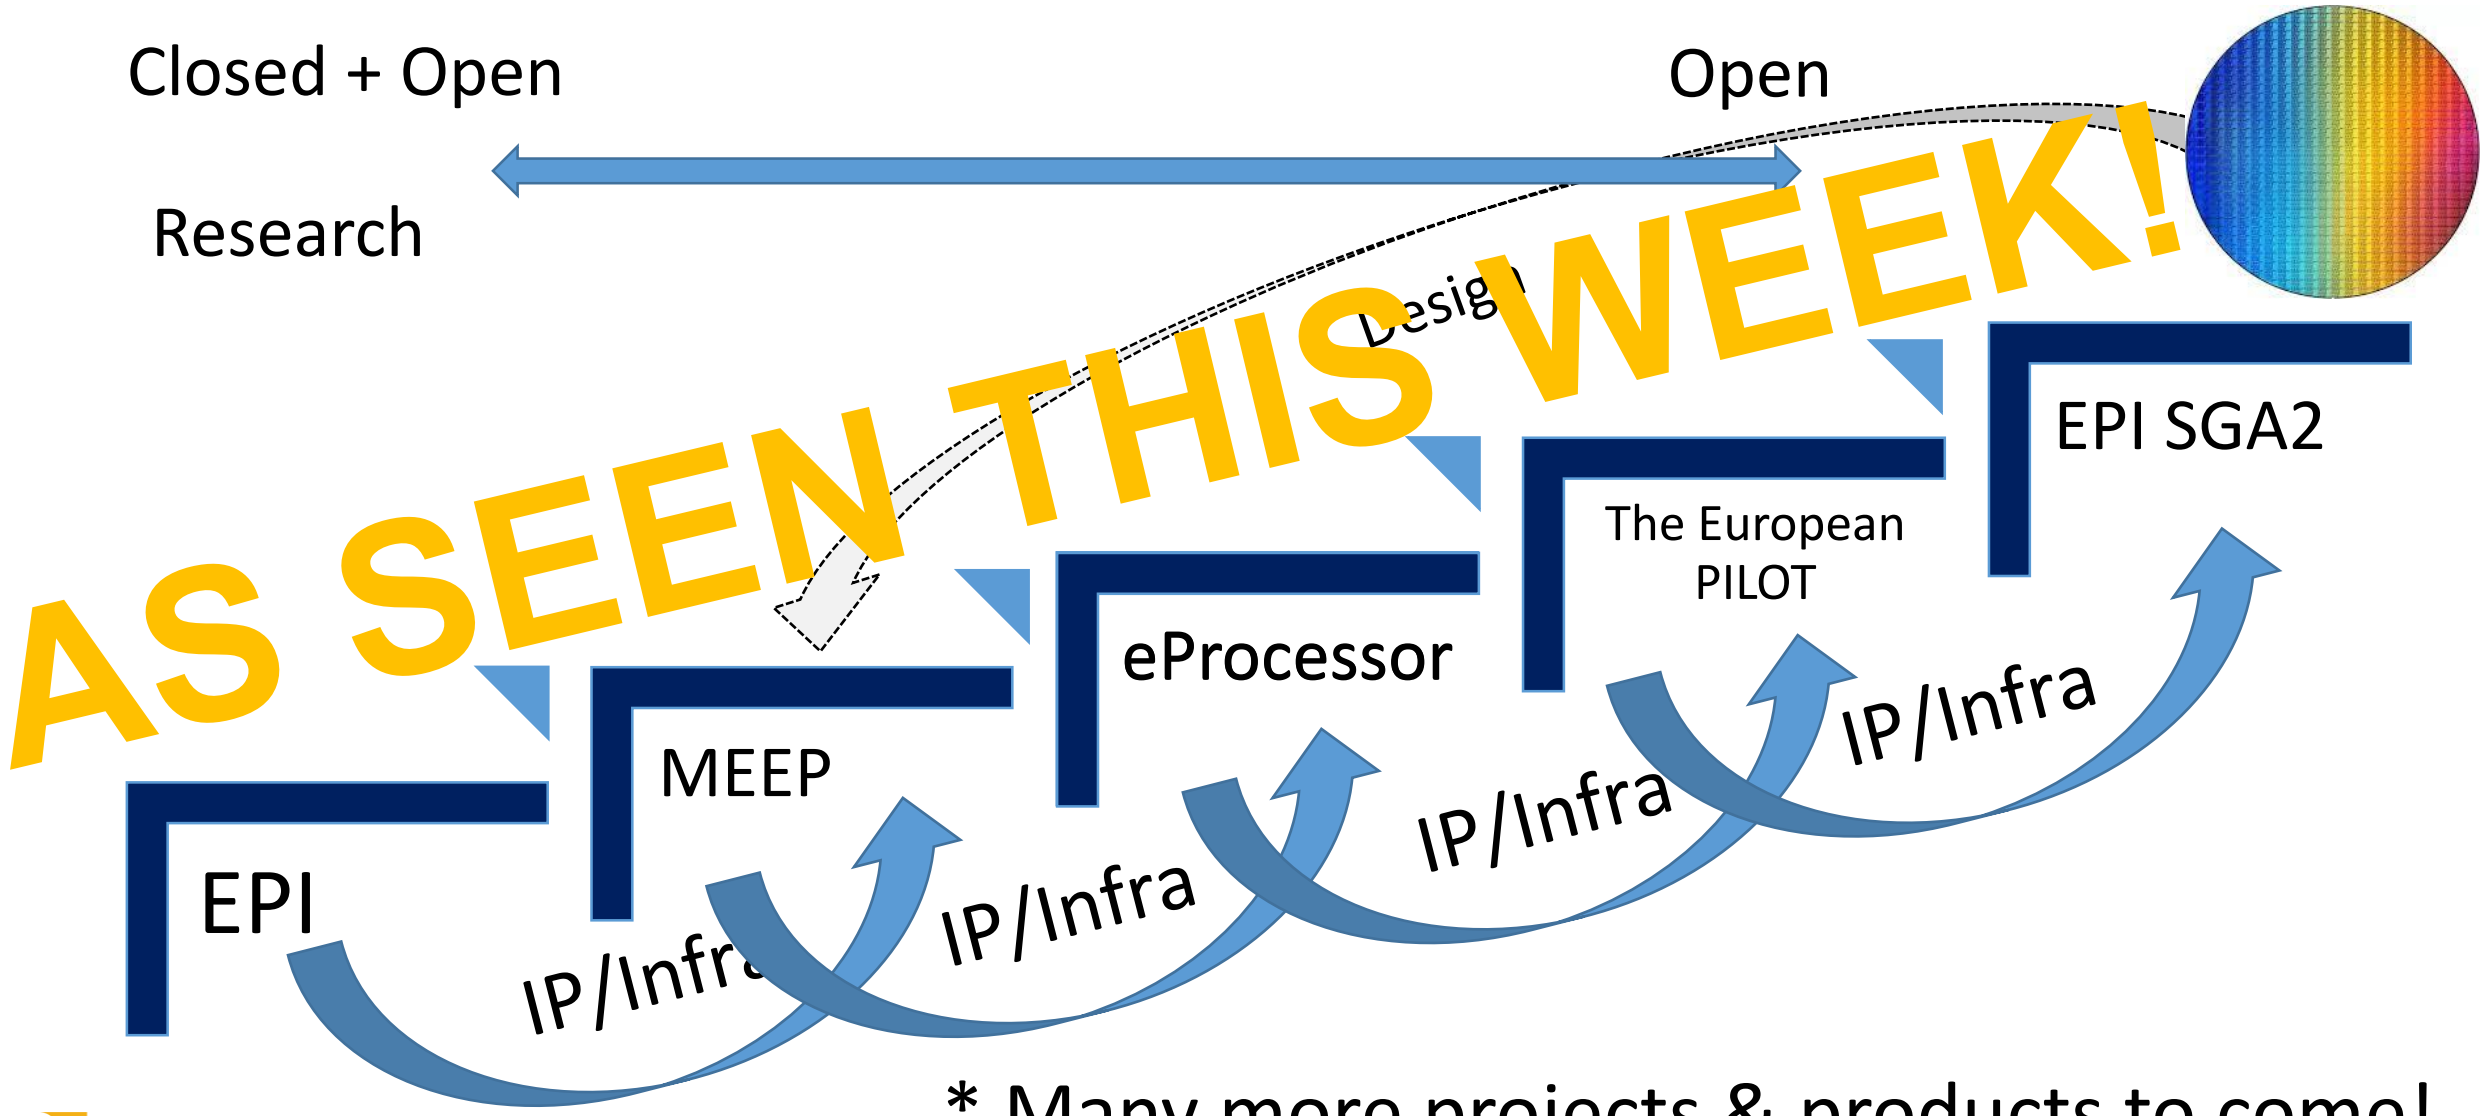

* Many more projects & products to come!

RISC-V®

Recap: RISC-V in EU RESEARCH

- 2019: What is Open Source Hardware??
- 2020: Open Source Hardware
- 2021: RISC-V? Roadmap?
  - November roadmap report
  - Horizon Europe Work Programme 2021-22:
    - Open Source Hardware (OSH) appears 6 times
    - CSA Roadmap
- 2022: Build RISC-V!!!
  - KDT JU Work Programme 2021 v13:
    - RISC-V appears 25 times / OSH appears 2 times
  - EU Chip Act (collection of documents)
    - RISC-V and OSH appears 5 each
  - **New KDT Call: Call 2022-1 Topic 3: Focus topic on Design of Customisable and Domain Specific Open-source RISC-V Processors (IA)**
- **More to come! (EuroHPC, KDT/Chip JU, etc.)**

RISC-V®

Building an Open HPC Ecosystem with RISC-V

- HPC requires customized hardware solutions for many HPC domains
- HPC is tackling grand challenges that benefit from the Global Technology ecosystem
- HPC is on the technology leading edge

- An Open ISA complements Open Source Software and combines to create an open Ecosystem

- RISC-V accelerates innovation both in Research and Industry
- The RISC-V community and ecosystem are rapidly growing

RISC-V®

You Can Help

- Get involved in SIG HPC
- https://lists.riscv.org/g/sig-hpc
- Subscribe:
  - Send email to: sig-hpc+subscribe@lists.riscv.org
- Monthly meetings
  - 3rd Thursday of the month @ 16:00 CET
  - Next meeting is May 19, 2022

Thank you

john.davis@bsc.es