Ushering in an Open Era of HPC

John D. Davis, Ph.D. @ BSC & John Leidel, Ph.D. @ TCL

An Open HPC Ecosystem?

- HPC
  - Past/Present/Future
- SIG-HPC
  - Intro
  - SW Ecosystem
- Building An Open HPC Ecosystem

Software History

Closed Priority SW/HW Systems: VMS, Lotus, AIX, etc.

Horizontal Platforms: Windows, Solaris, etc.

Open Source Software: Arduino, Android, Linux

Expensive, Rigid, Lock-in, Custom APIs, Additional services

Commodity, Standard APIs, Lock-in, More Apps → more users

Free-mium, Standard APIs, User Customized, Ubiquitous

Linux History

- 1991: Started development, release
- 1992: X Windows released
- 1998: Adopted by many major companies
- 2004: BSC Supercomputer OS
- 2009: Basis for many new business systems and the cloud
- 2013: Android in 75% of the world smart phones
- 2015: De facto OS for IoT, mobile, cloud, and supercomputers

Mont-Blanc HPC Stack for ARM

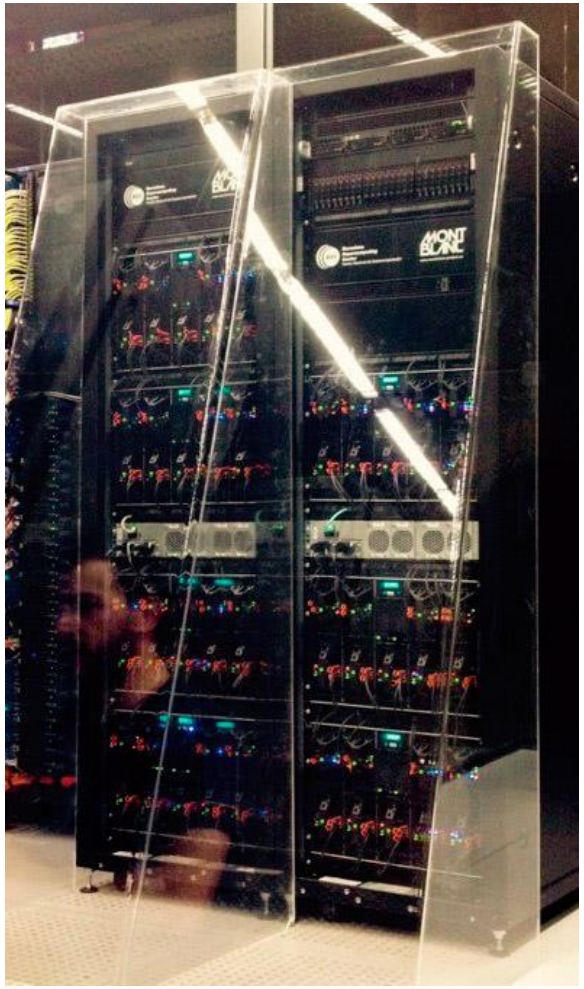

Industrial applications

Pharmacelera, MUREX, DASSAULT AVIATION, AVL, Termo Fluids, Cenaero, ROLLS ROYCE

Applications

GENCI, CINECA, UNI GRAZ, HLRS, University of BRISTOL

System software

ARM, ETH zürich, Inria, Barcelona Supercomputing Center, JÜLICH FORSCHUNGSZENTRUM, Bull atos technologies

Hardware

ARM, Barcelona Supercomputing Center, Bull atos technologies, cea, UC UNIVERSIDAD DE CANTABRIA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

The Exascale Race – The Japanese example

Co-design from Apps to Architecture

R-CCS

- **Architectural Parameters to be determined**
  - #SIMD, SIMD length, #core, #NUMA node, O3 resources, specialized hardware
  - cache (size and bandwidth), memory technologies
  - Chip die-size, power consumption
  - Interconnect
- **We have selected a set of target applications**
- **Performance estimation tool**
  - Performance projection using Fujitsu FX100 execution profile to a set of arch. parameters.
- **Co-design Methodology (at early design phase)**
  1. Setting set of system parameters
  2. Tuning target applications under the system parameters
  3. Evaluating execution time using prediction tools
  4. Identifying hardware bottlenecks and changing the set of system parameters

Target applications representatives of almost all our applications in terms of computational methods and communication patterns in order to design architectural features.

Target Application	
Program	**Brief description**
① GENESIS	MD for proteins
② Genomon	Genome processing (Genome alignment)
③ GAMERA	Earthquake simulator (FEM in unstructured & structured grid)
④ NICAM+LETK	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
⑤ NTChem	molecular electronic (structure calculation)
⑥ FFB	Large Eddy Simulation (unstructured grid)
⑦ RSDFT	an ab-initio program (density functional theory)
⑧ Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
⑨ CCS-QCD	Lattice QCD simulation (structured grid Monte Carlo)

9

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

RIKEN

LOCA

The Exascale Race – The Japanese example

"Post-K" Arm64fx Processor is…

- an Many-Core ARM CPU…
  - 48 compute cores + 2 or 4 assistant (OS) cores
  - Brand new core design by Fujitsu
  - Near Xeon-Class Integer performance core
  - ARM V8.2 --- 64bit ARM ecosyste
- …but also a GPU-like processo
  - SVE 512 bit vector extensions (ARM & Fujitsu)
    - Integer (1, 2, 4, 8 byte) + Float (16, 32, 64 bytes)
  - Ca + access localization (sector cache) – similar to scratchpad
  - HBM2 OPM – Massive Mem BW (1TByte/s, Bytes/DPF ~0.4 same as K)
  - Strea ing memory access, strided access, scatter/gather etc.
  - Intra-chip barrier synch. and other memory enhancing features
  - 40GByte/s Tofu-.D interconnect + PCIe 3
- GPU-like High performance in HPC, AI/Big Data, Auto Driving…

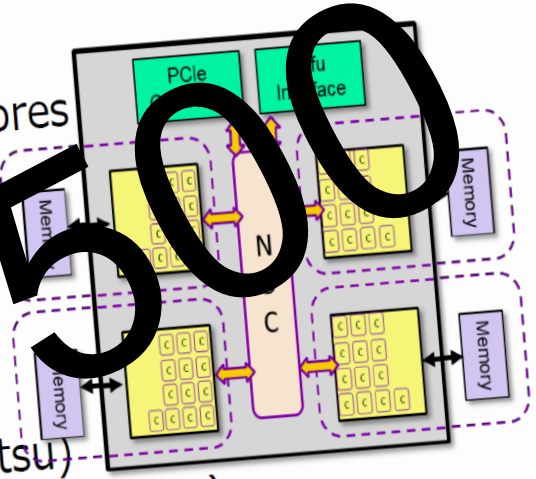

11

#1 TOP500

HPC Today

- **Europe has led the way in defining a common open HPC software ecosystem**
- **Linux** is the de facto standard OS despite proprietary alternatives
- Software landscape from Cloud to IoT already enjoys the benefit of open source
- Open source provides:
  - A common platform, specification and interface
  - Accelerates building new functionality by leveraging existing components
  - Lowers the entry barrier for others to contribute new components
  - Crowd-sources solutions for small and larger problems
- **What about Hardware and in particular, the CPU?**
- Inhibits opportunities in holistic co-design
  - Facing barrier to innovation
  - Being able to have a conversation or not

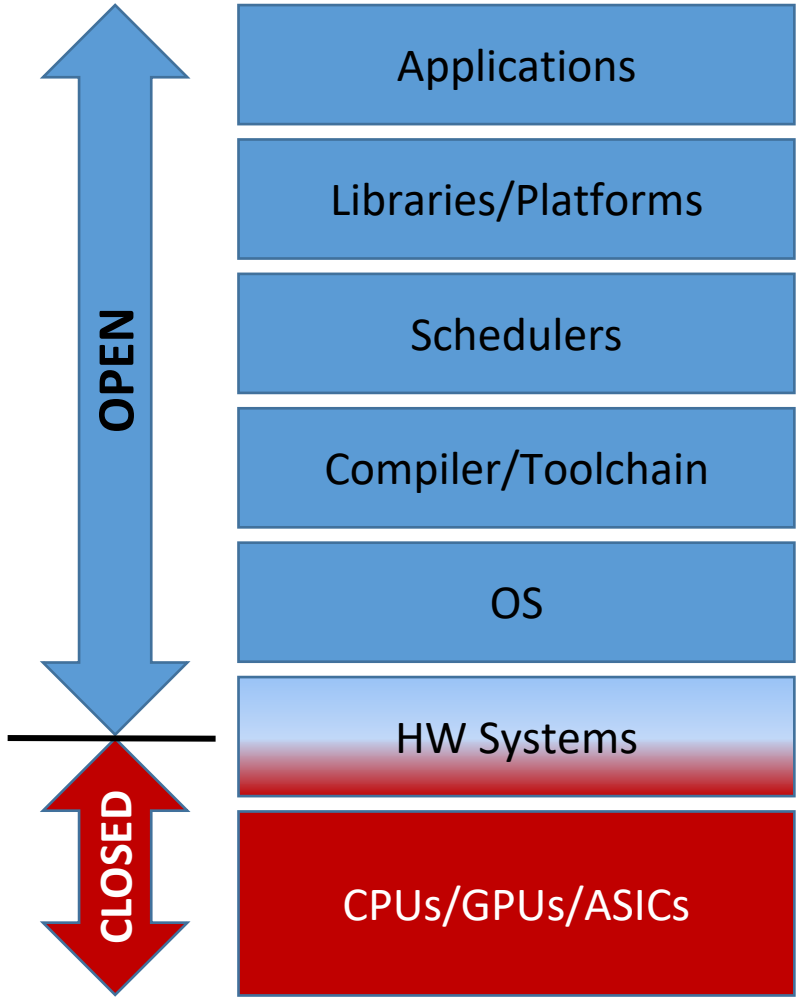

RISC-V®

Exascale: An Energy Efficiency Challenge

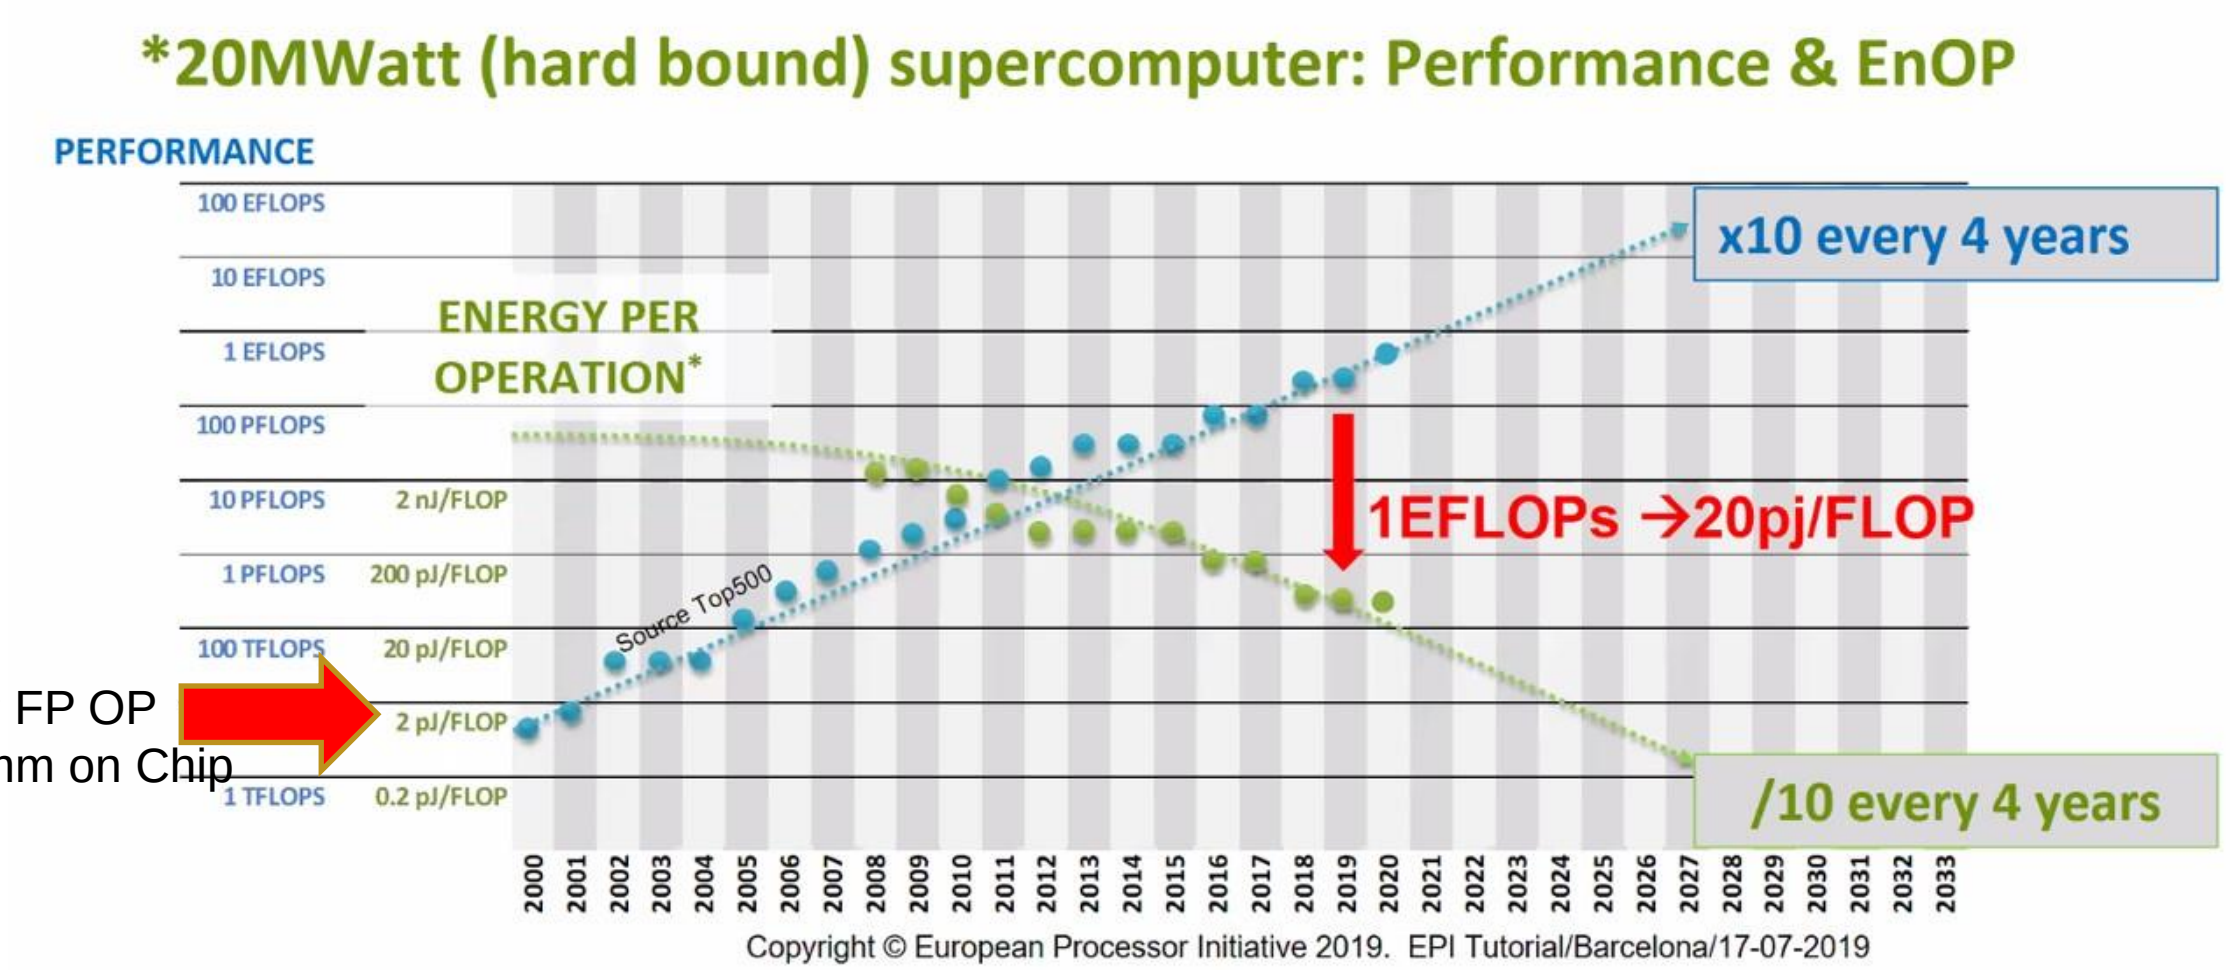

~DP FP OP
~1 mm on Chip

HPC is power-bound! 10X energy efficiency improvements every 4 years!

RISC-V®

Today´s technology trends

Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)

Moore´s Law + Power = **Specialization**

- More cost effective
- More performant
- Less Power

SOFTWARE/
HARDWARE
CO-DESIGN

RISC-V®

Today´s technology trends

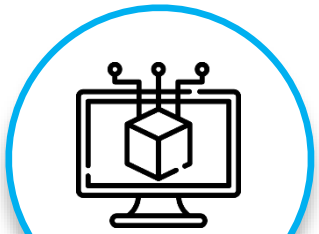

Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)

Moore´s Law + Power = **Specialization**

- More cost effective
- More performant
- Less Power

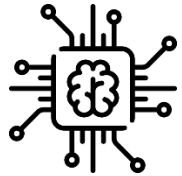

New Open Source Hardware Momentum from IoT and the Edge to HPC

- RISC-V
- OpenPOWER

RISC-V®

Open Ecosystem HW/SW Co-Design

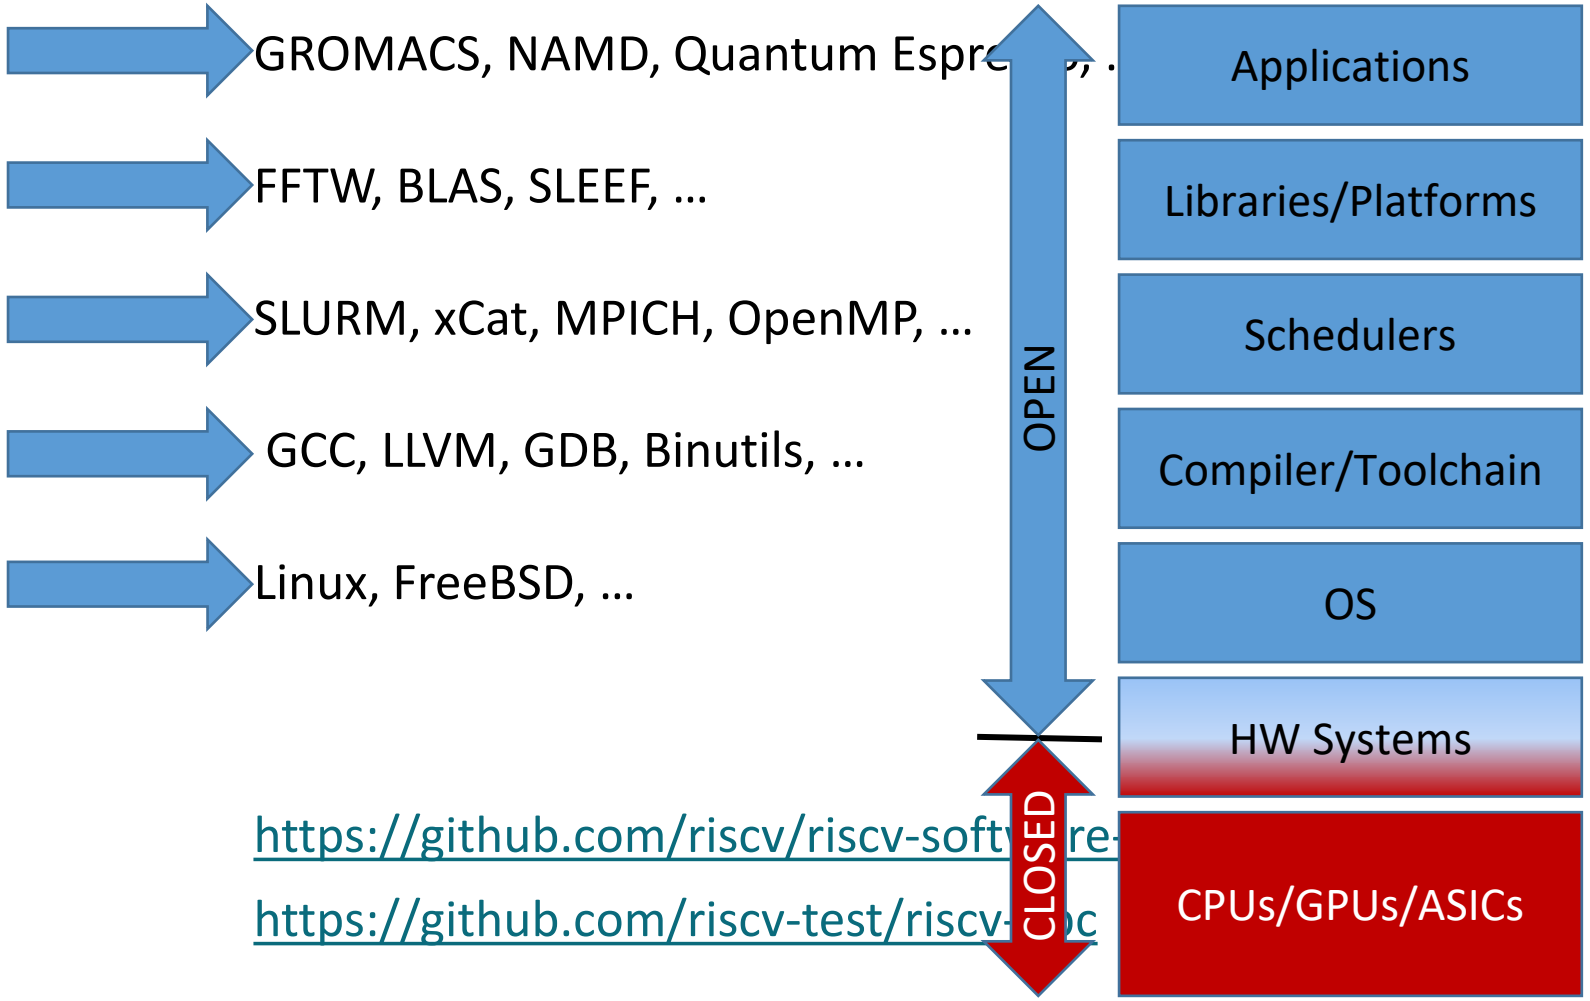

RISC-V®

Why Open Source Hardware?

Software: Leverage a large ecosystem compatible across implementations
Security: A fully auditable collection of IPs: processors, accelerators, etc.
Safety: No black-boxes
SWaP & Customization: SW/HW co-design for exact feature match
Performance: State-of-the-art implementations
No vendor lock-in: Ecosystem to enable custom develop from SME to large enterprise
Sovereignty: Freedom of access and implementation from design to production
Open Collaboration: Faster time to market, community, leverage existing open source

RISC-V Ecosystem

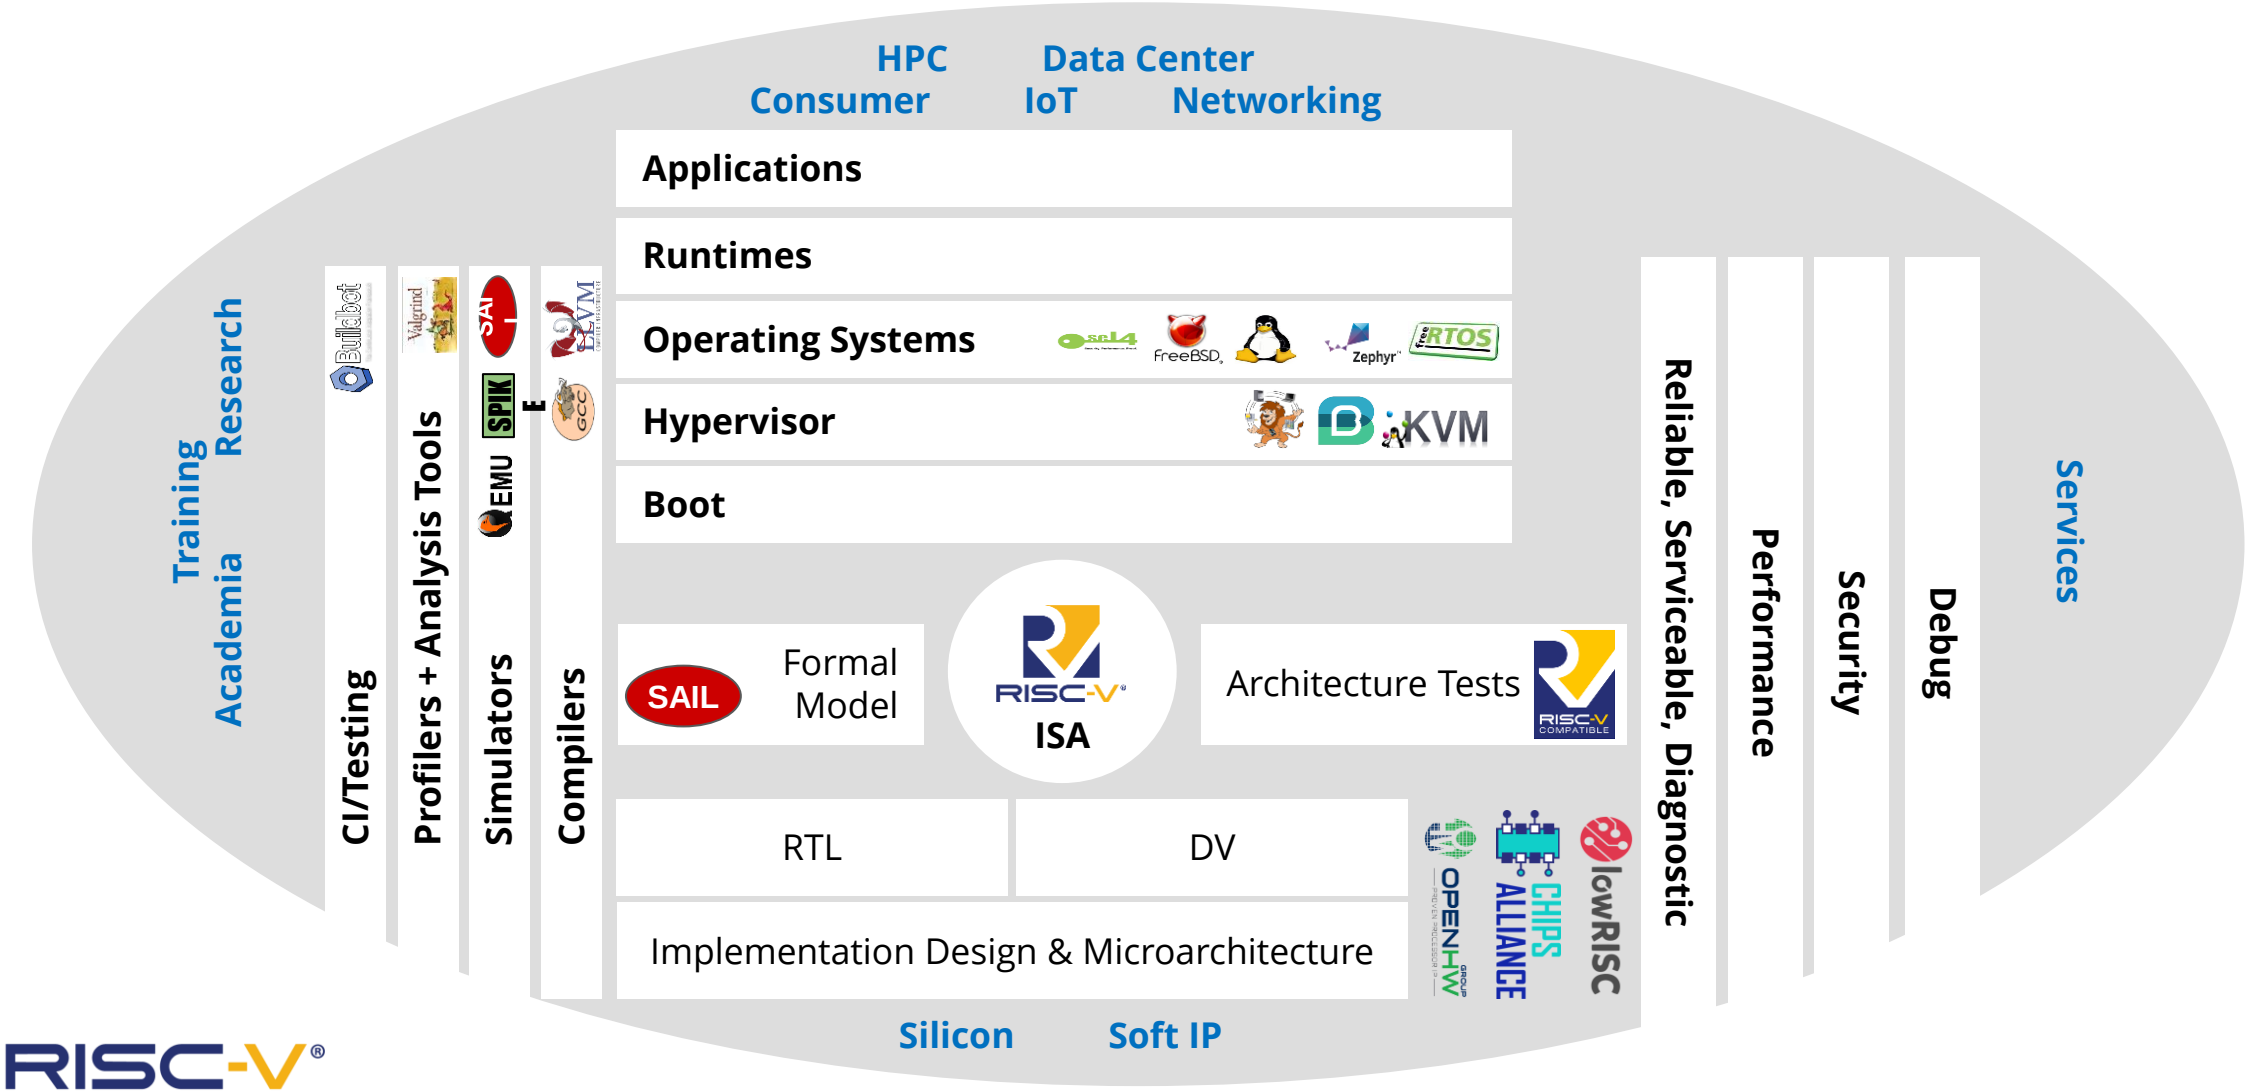

RISC-V®

More than 2,200 RISC-V Members

across 70 Countries

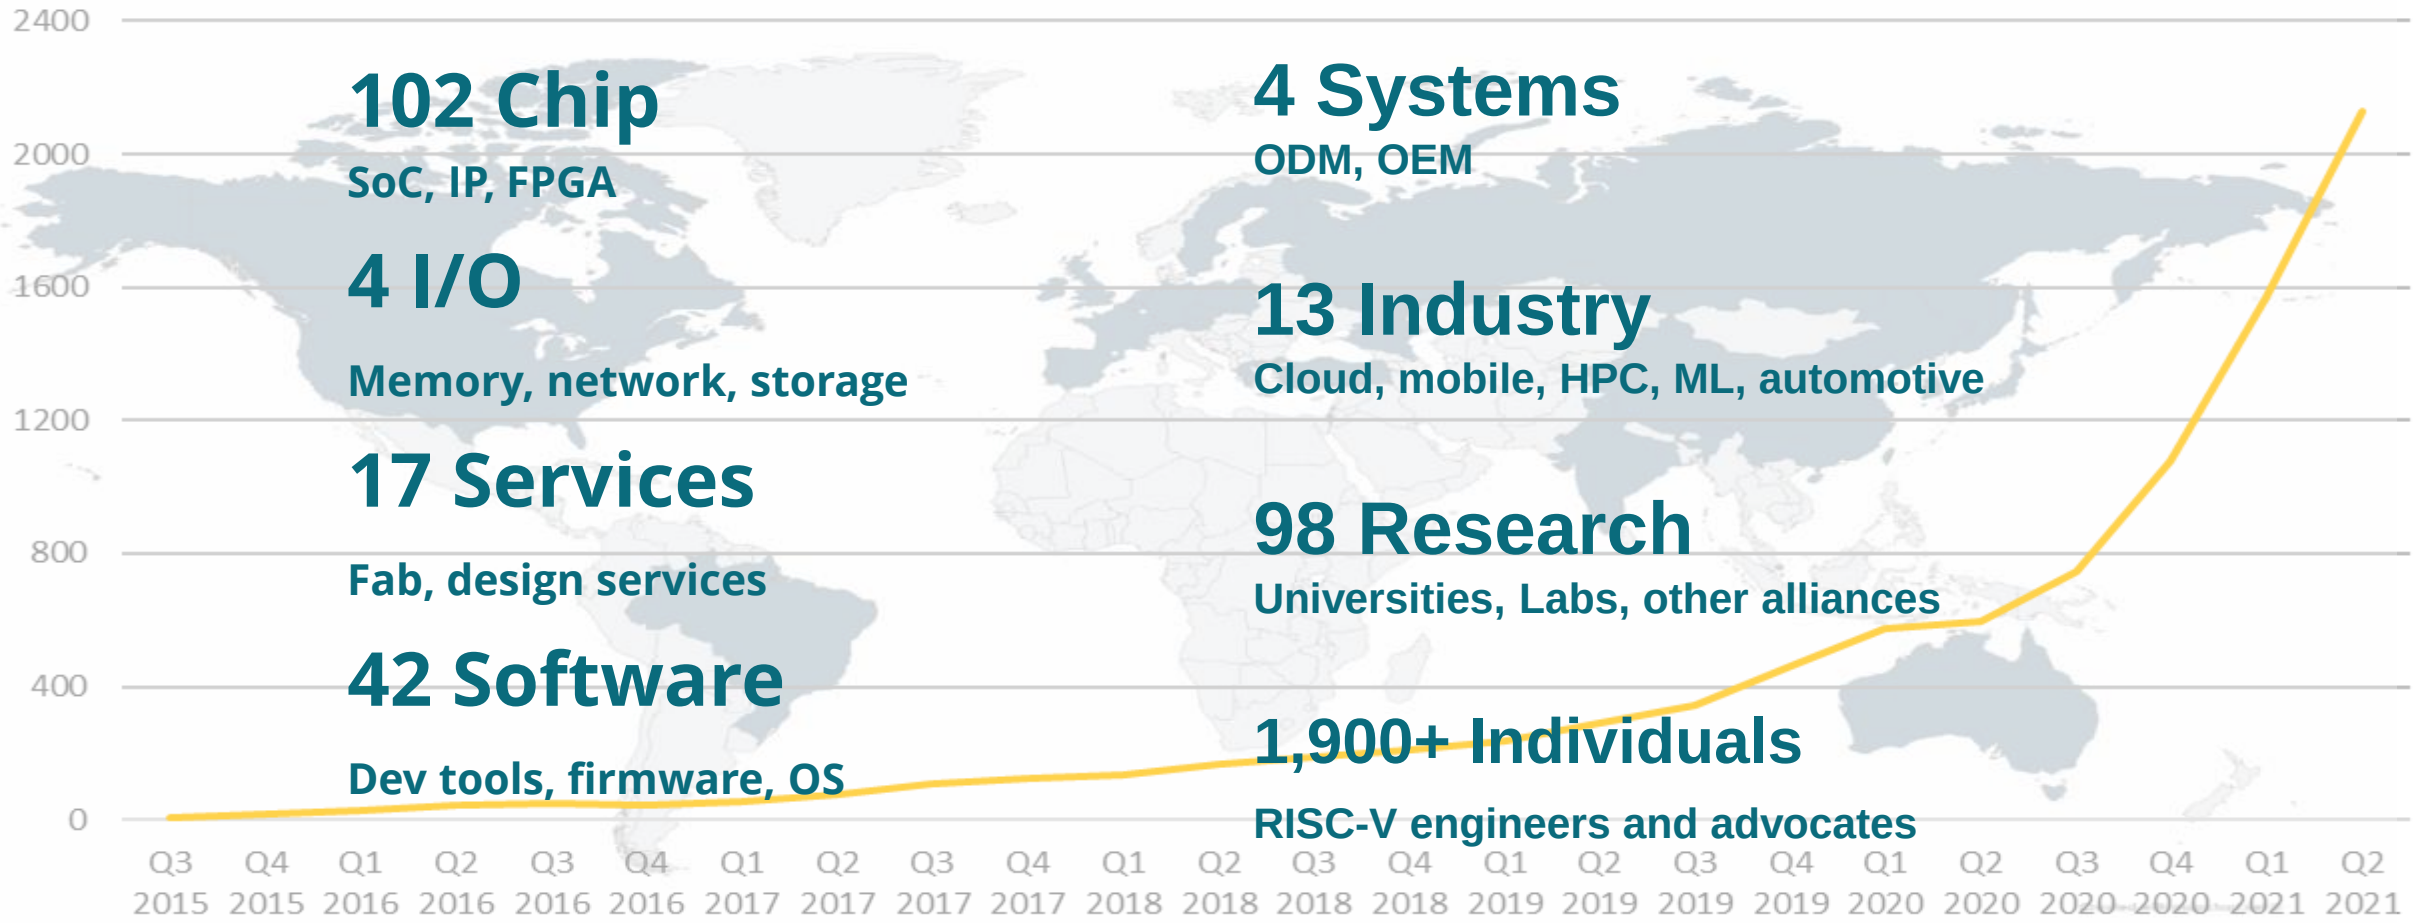

102 Chip
SoC, IP, FPGA

4 I/O
Memory, network, storage

17 Services
Fab, design services

42 Software
Dev tools, firmware, OS

4 Systems
ODM, OEM

13 Industry
Cloud, mobile, HPC, ML, automotive

98 Research
Universities, Labs, other alliances

1,900+ Individuals
RISC-V engineers and advocates

RISC-V®

RISC-V membership grew 133% in 2020.
In 2021, RISC-V membership has already doubled.

HPC Tomorrow

- A completely **open SW/HW stack for the world**
- RISC-V provides the open source hardware alternative to dominating proprietary solutions
- Enables complete technology independence using 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 Globally.**

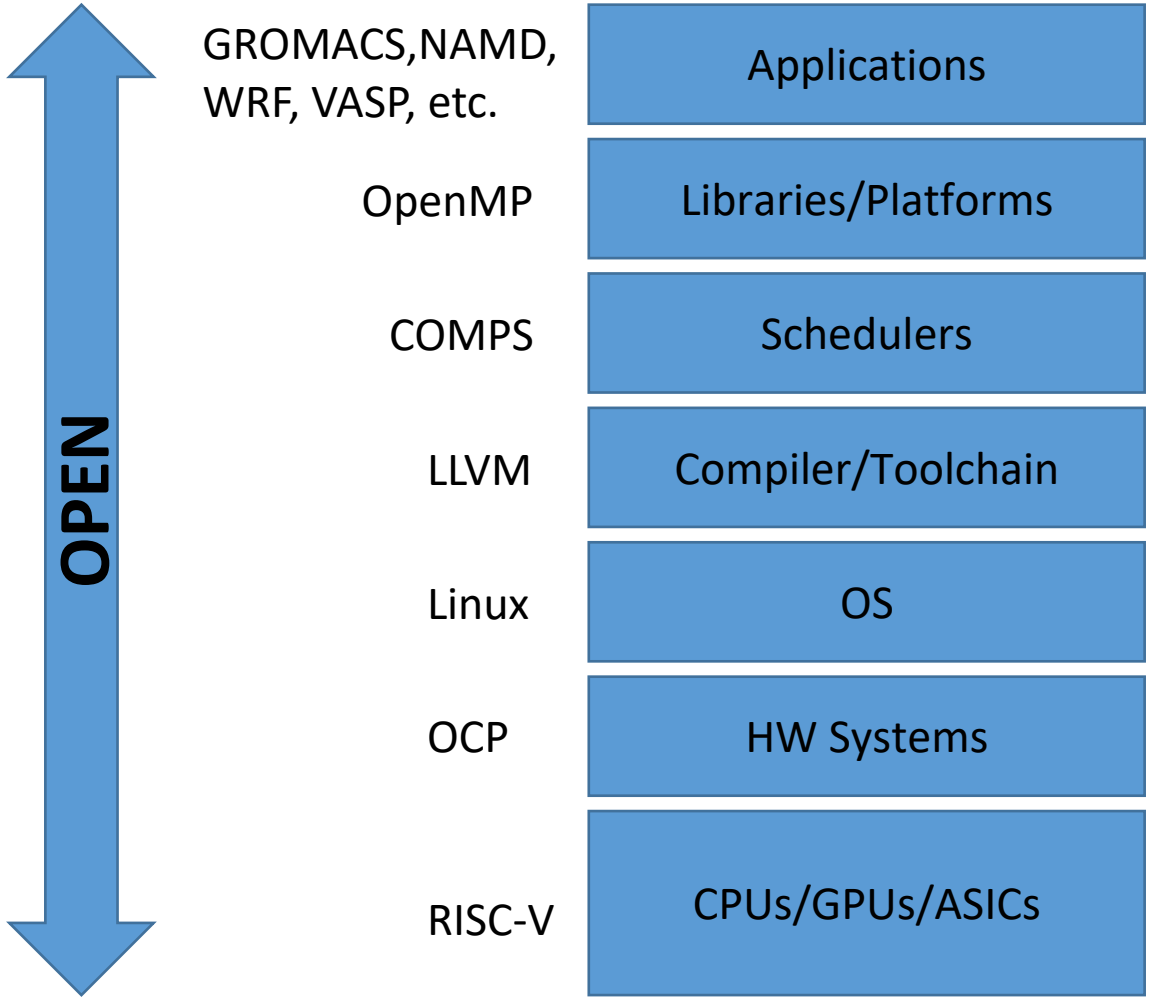

RISC-V®

Special Interest Group – High Performance Computing

SIG-HPC

https://lists.riscv.org/g/sig-hpc

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

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

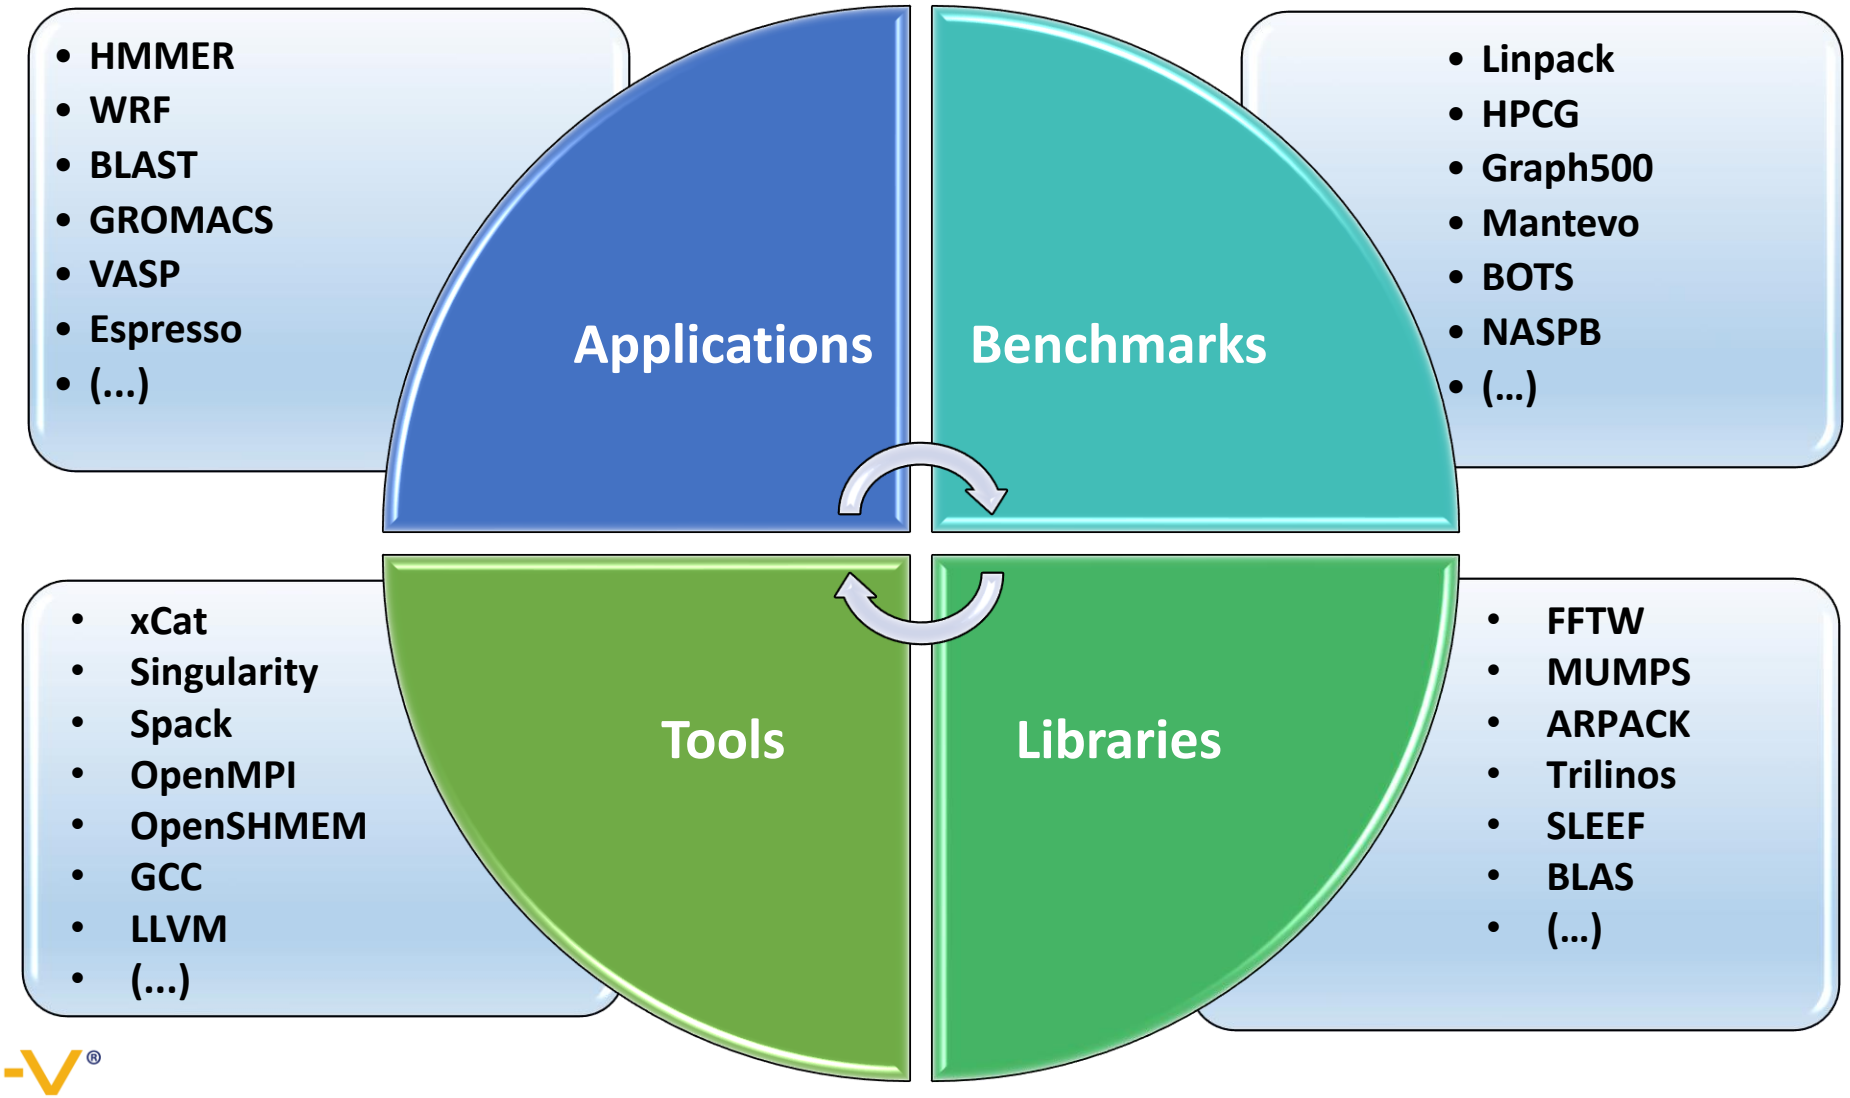

RISC-V®

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

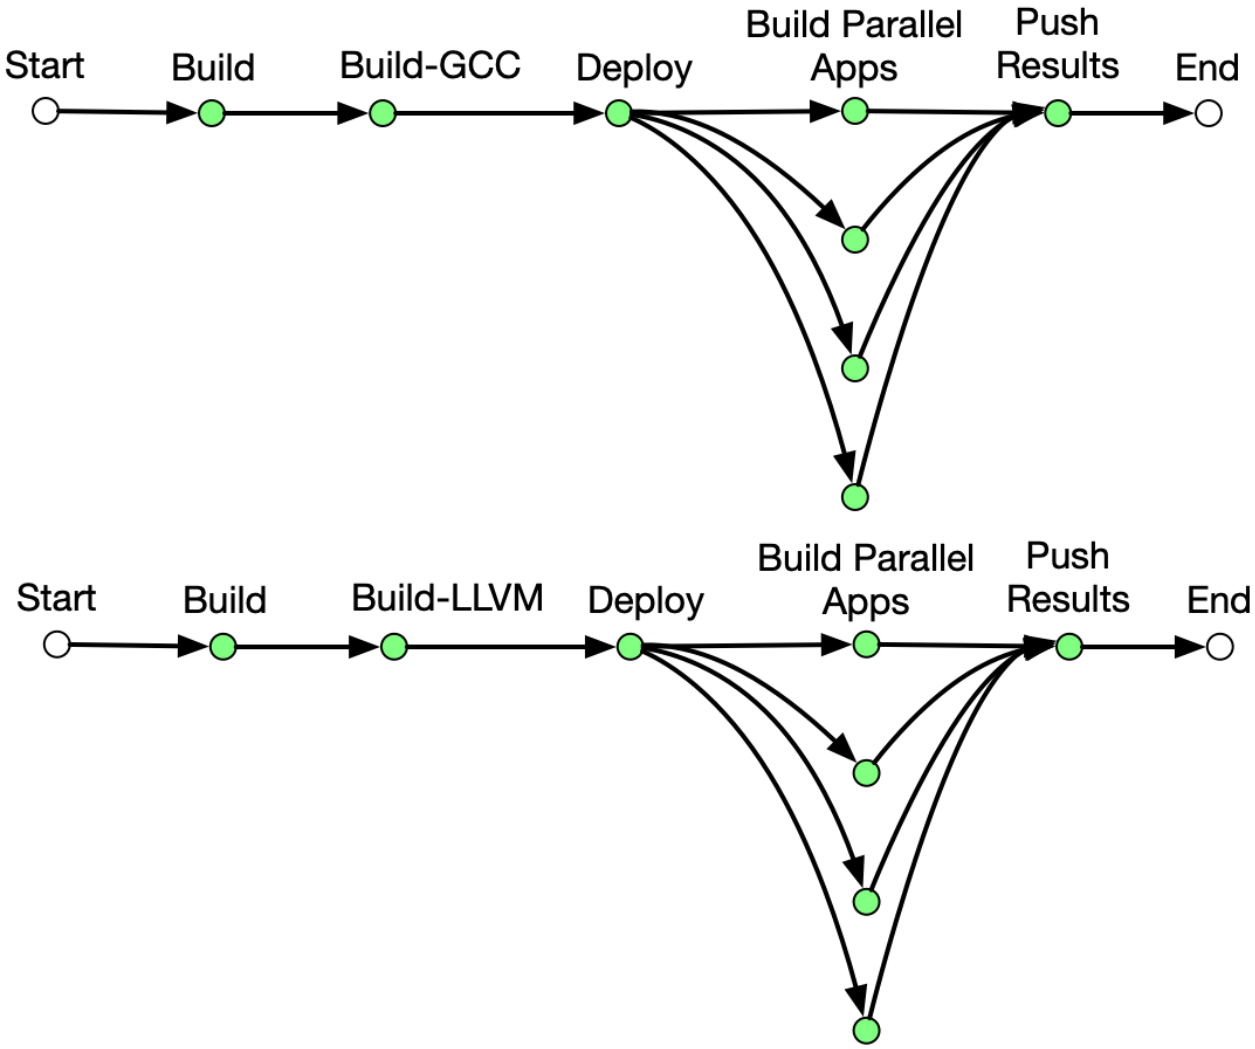

HPC Software Testbed: riscv-test.org

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

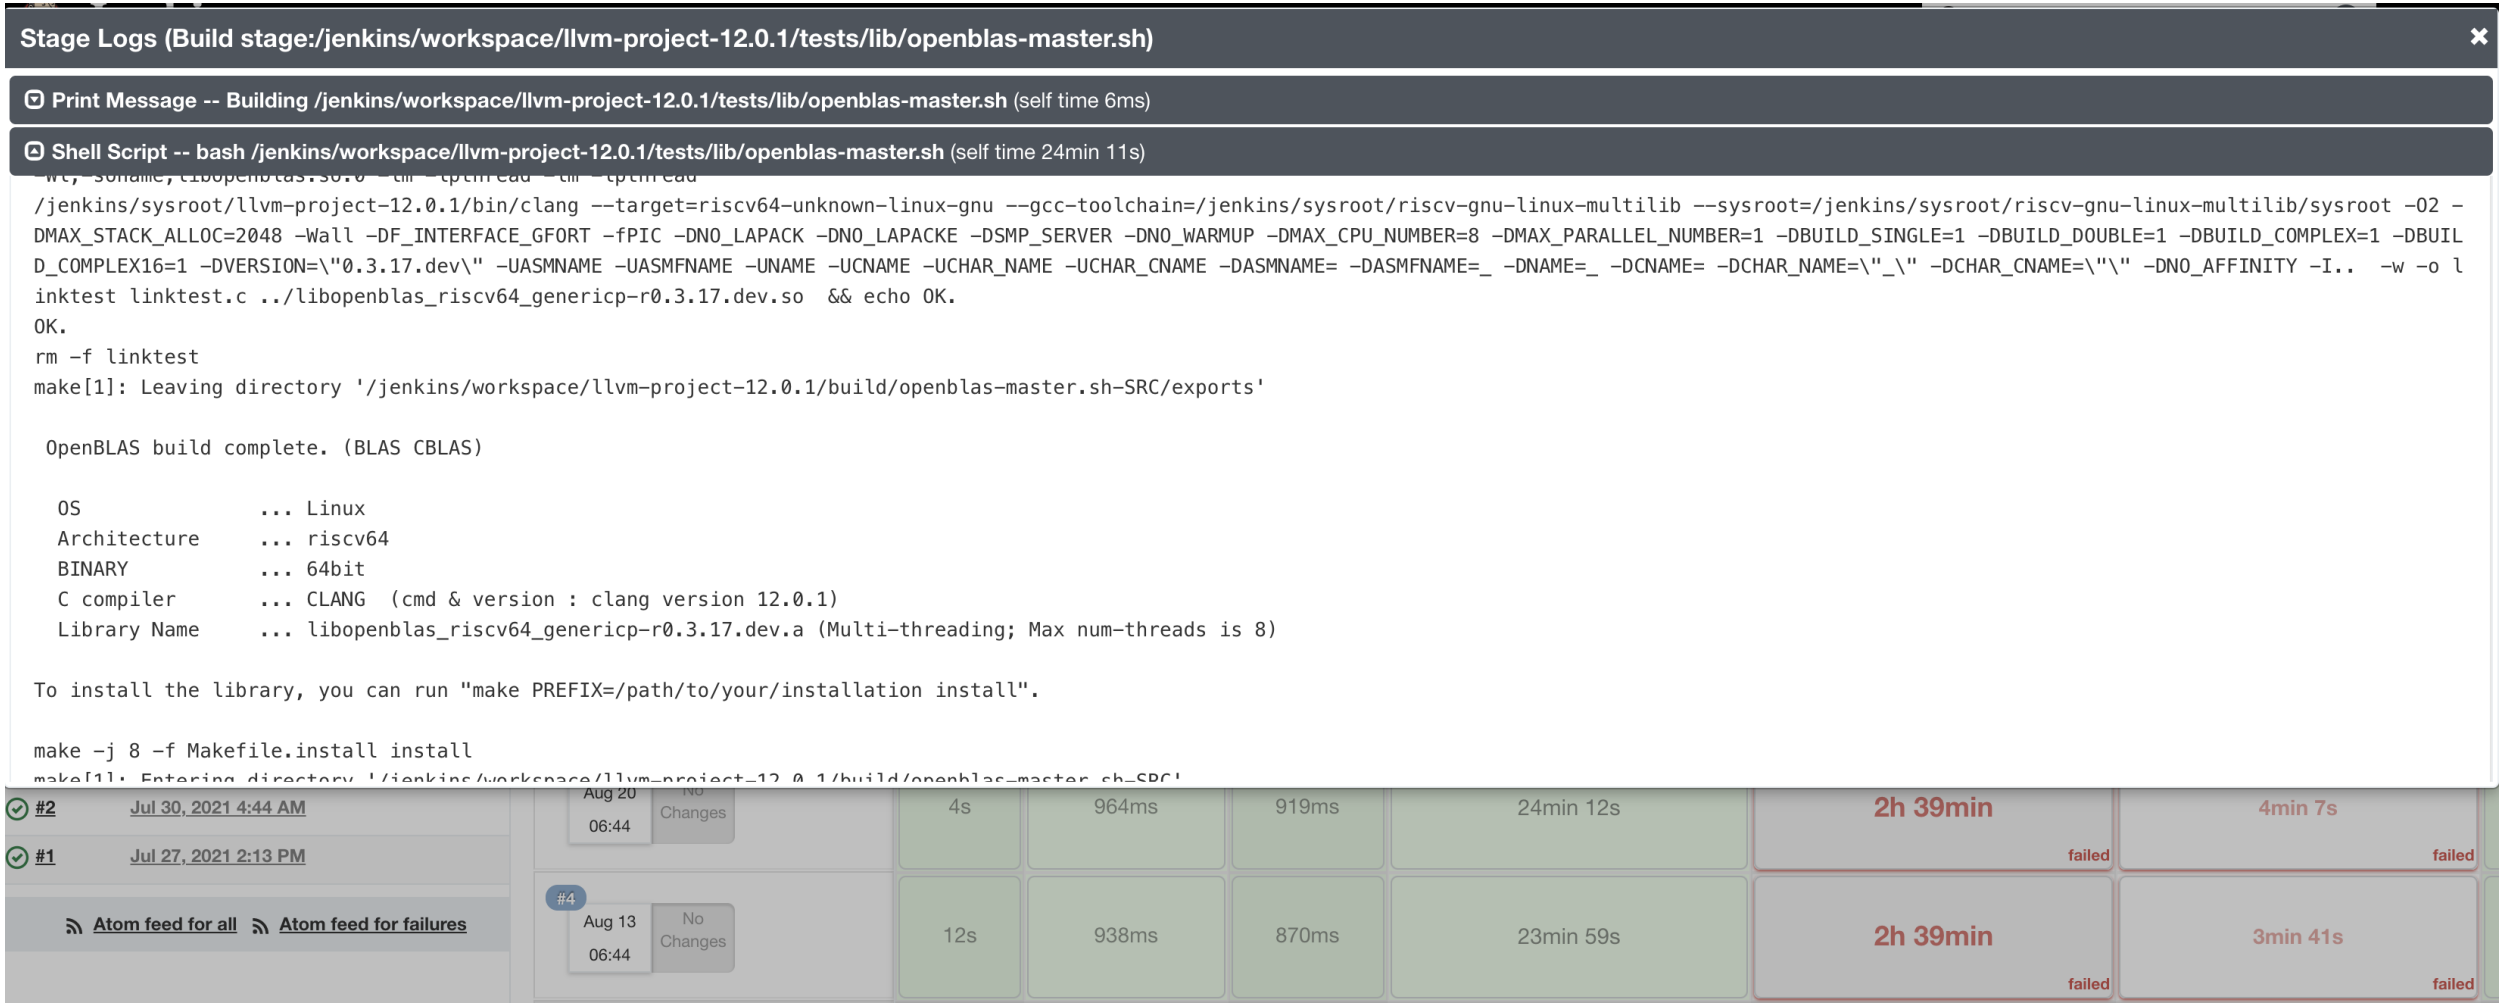

Public Test Harness

- 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

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®


Thank You!
Help define the future
Come Join Us
Subscribe: sig-hpc+subscribe@lists.riscv.org
https://lists.riscv.org/g/sig-hpc
RISC-V®