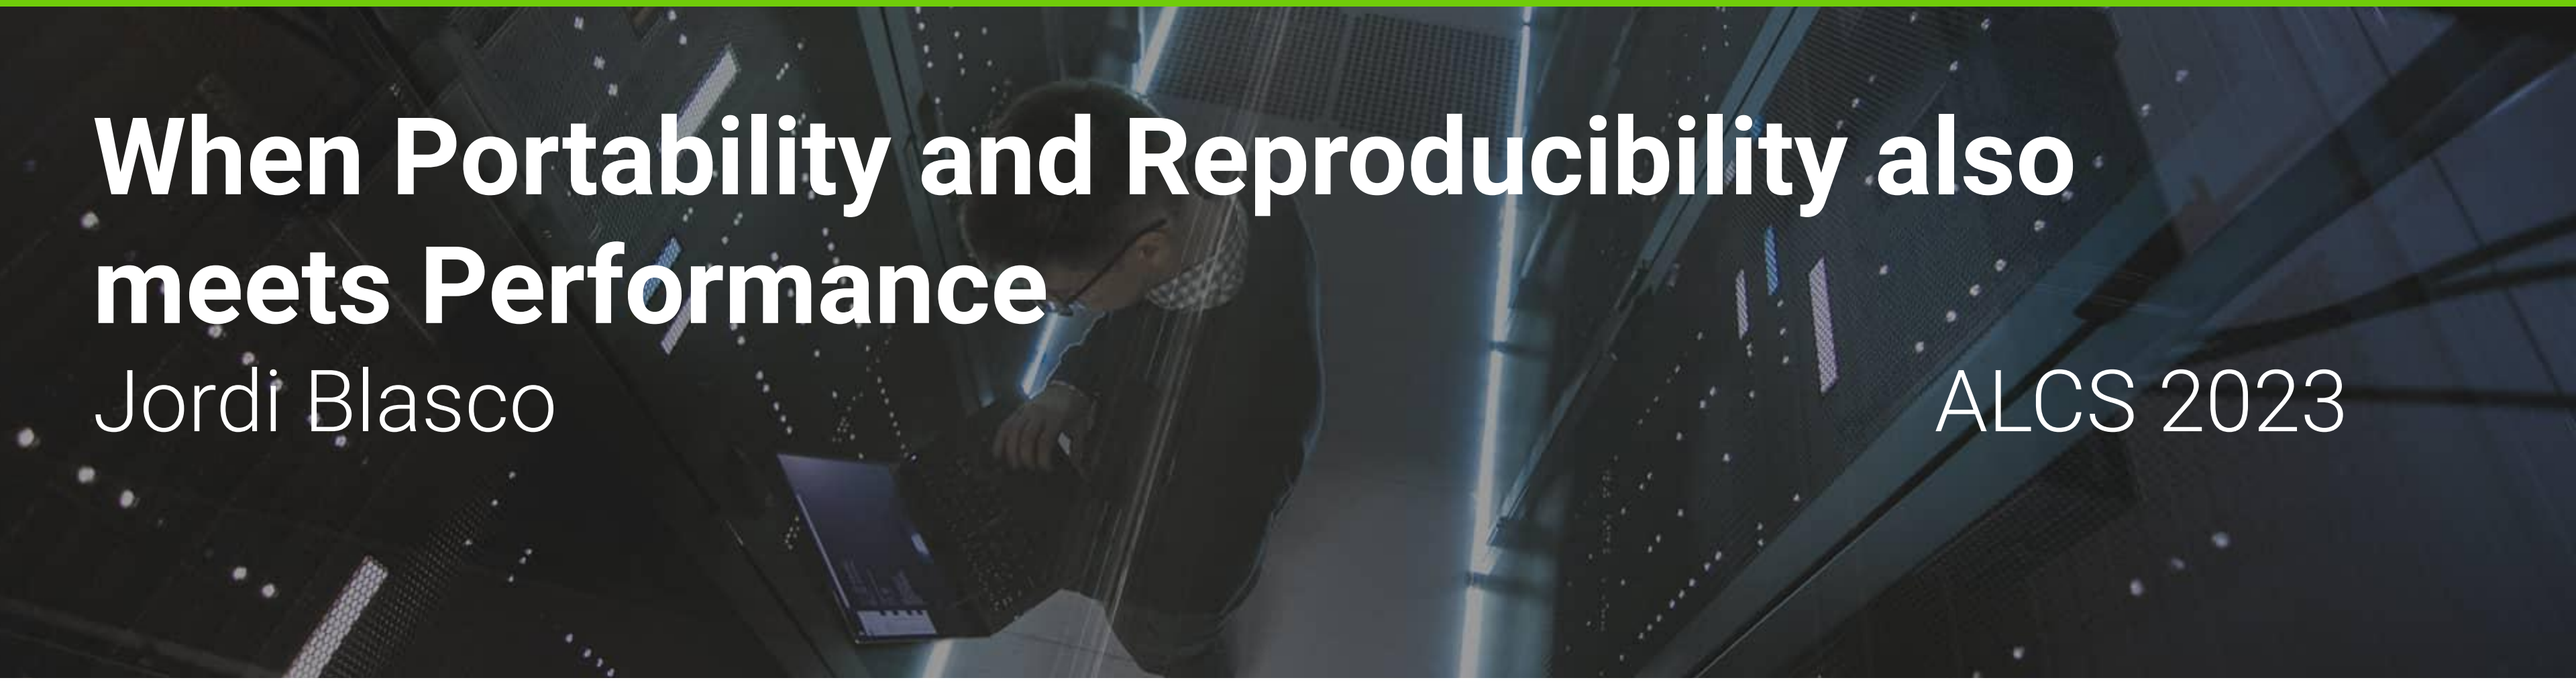

When Portability and Reproducibility also meets Performance

Jordi Blasco

ALCS 2023

15/06/2023

www.hpcnow.com


HPC Now!
Advanced supercomputing services for science and engineering
We plan, install & support
You run
www.hpcnow.com

Services and turnkey solutions adapted to customers needs

Planning

Consulting

Solution design

Installation

Infrastructure

Software

Training

Maintenance

Support

Managed services

We are passionate about new challenges and HPC technologies enthusiasts. Our goal is to take full advantage of supercomputing to provide solutions to our customers' scientific or engineering dares.

User-oriented company

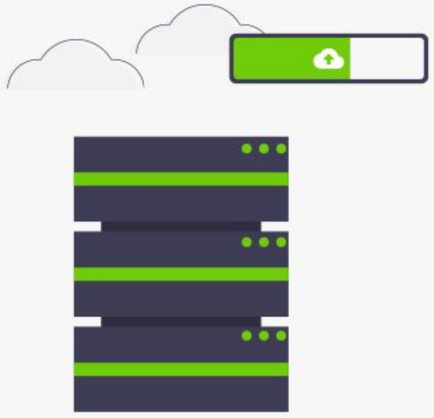

SW & HW & Cloud agnostics

IT + scientific background

HPC services and solutions

- **Customers**















and many more...

- **Contribution to HPC community**

openHPC

easybuild

hpckp'19 Barcelona

ETP 4 HPC EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

edci European Extreme Data & Computing Initiative

HPC ADVISORY COUNCIL NETWORK OF EXPERTISE

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING

BioNT

Multiscale

EESSI EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

HEROES

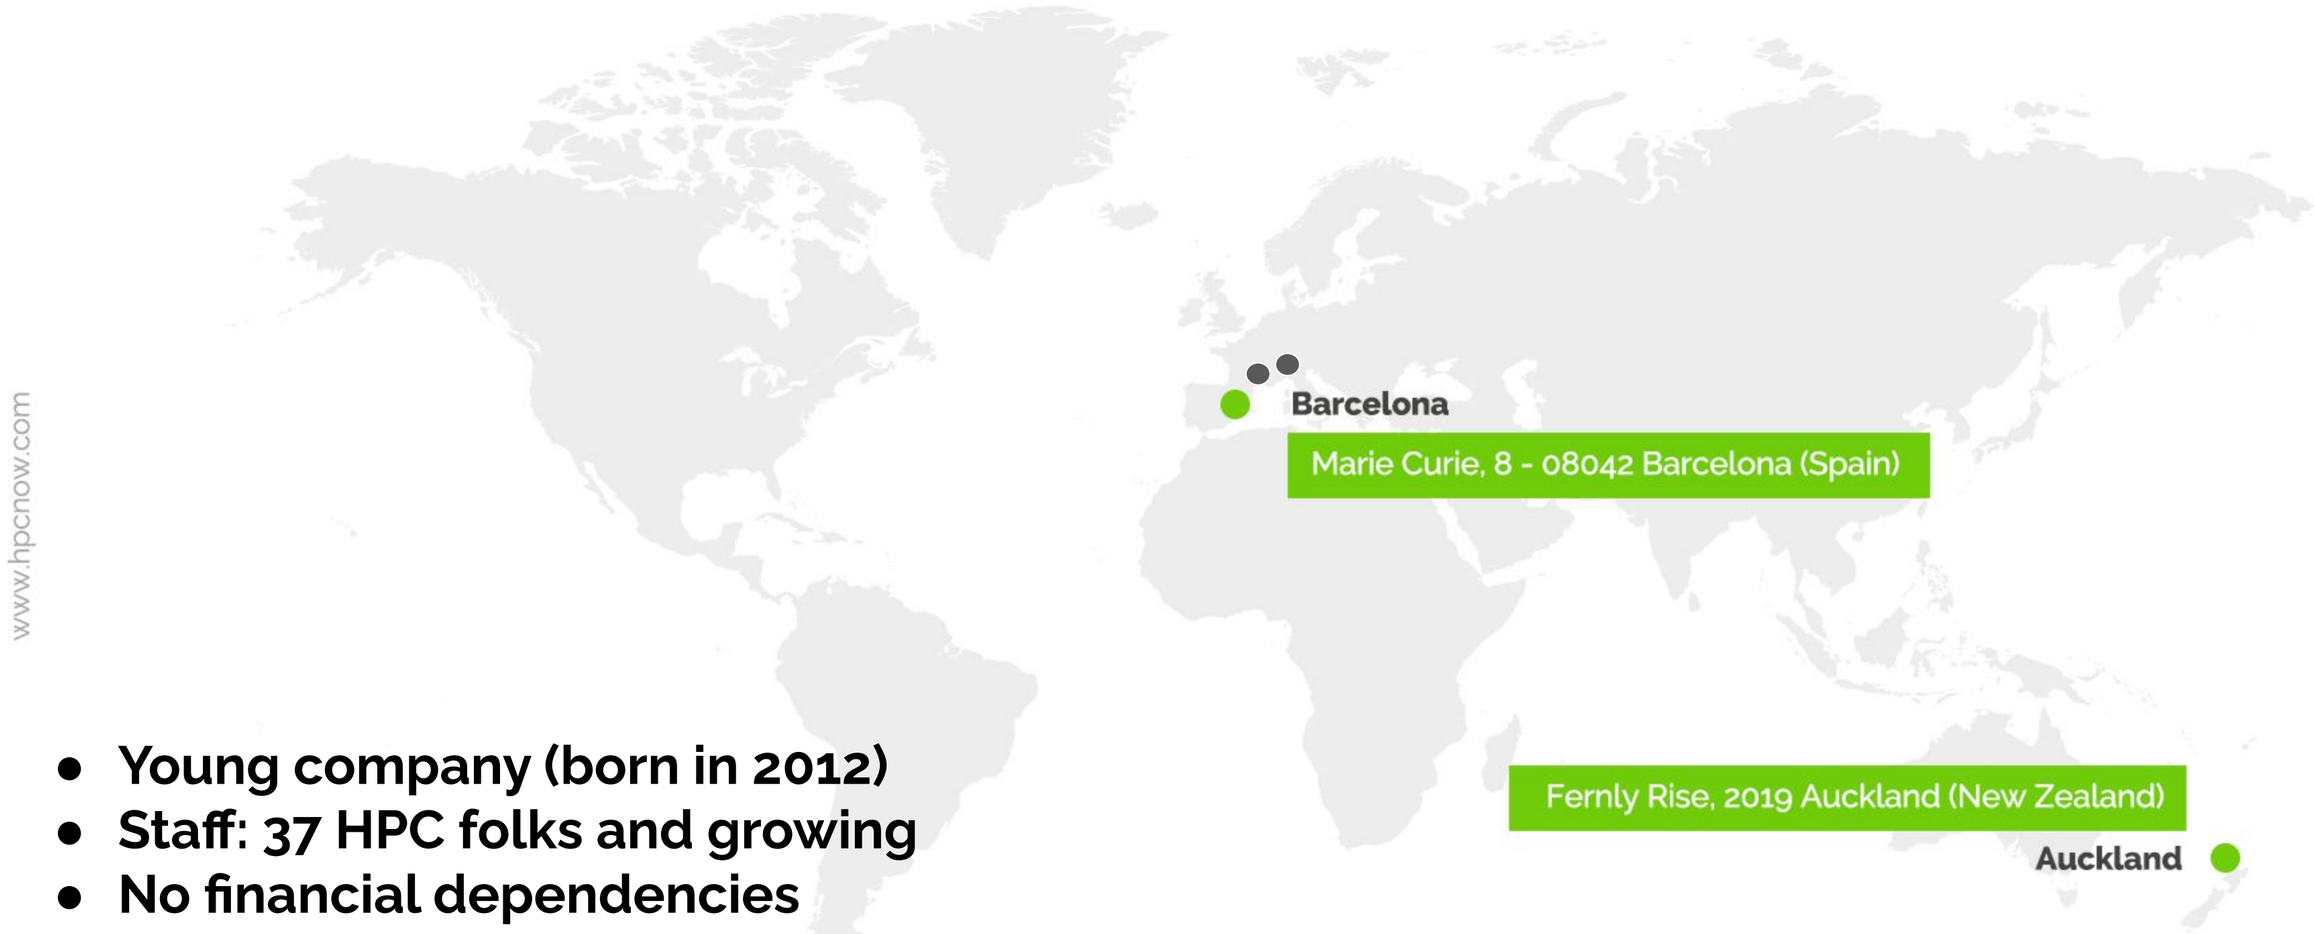

- Young company (born in 2012)
- Staff: 37 HPC folks and growing
- No financial dependencies
- Strong growth
- EU joint venture

doitnow
HPC Services

Needs and Challenges

HPC Now!

Micro Architecture Awareness

Keep HPC in capital letters

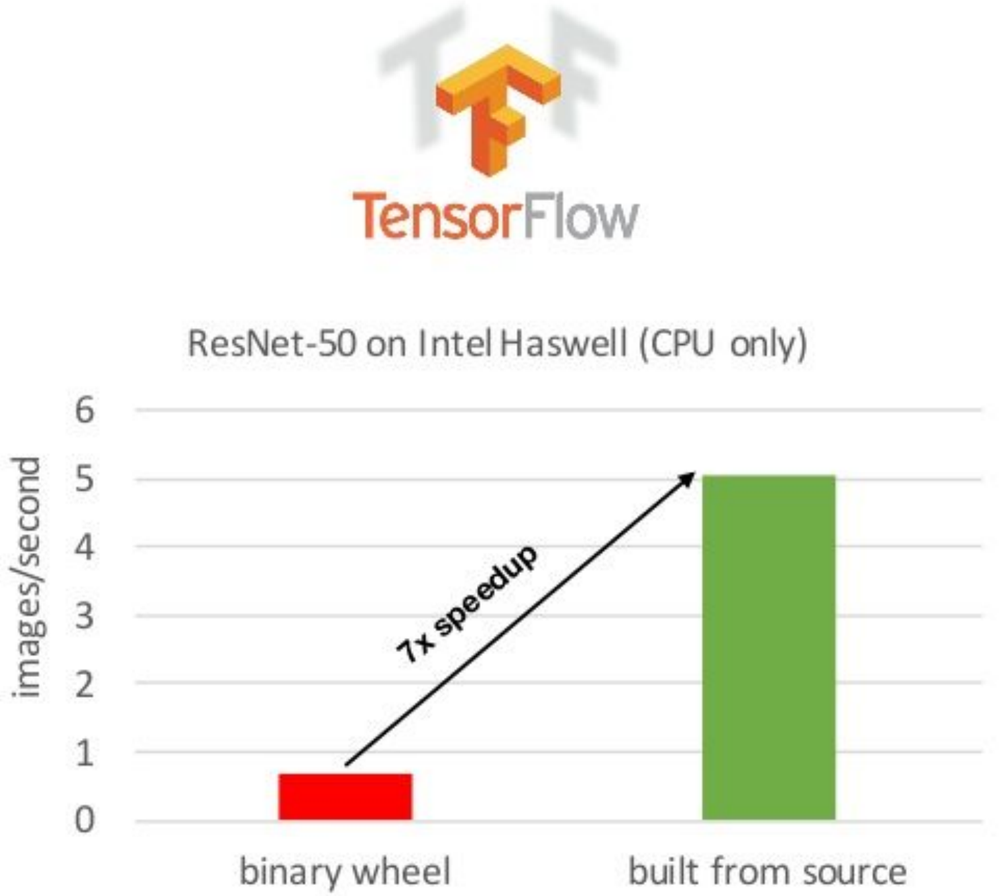

Micro Architecture Awareness

- Binaries from repositories have no microarchitecture optimisation for portability.
- Binaries compiled in old architectures will work in new ones but with a performance penalty.
- Binaries compiled in new architectures will not work in old ones due to lack of hardware instructions.

HPC Now!

Micro Architecture Awareness

Vectorial Instructions

New instruction sets allow achieving more operations per cycle. The following table provides an overview of the performance increase in terms of DP FLOPs in the course of x86 SIMD evolution.

Instruction Set Details

Microarchitecture	Year	Instruction Set	DP FLOPs/Cycle
Diamond Rapids	2025?	AVX-1024	64
Skylake	2015	AVX-512	32
Haswell	2013	AVX2	16
Sandybridge	2011	AVX	8
Nehalem	2008	SSE	4

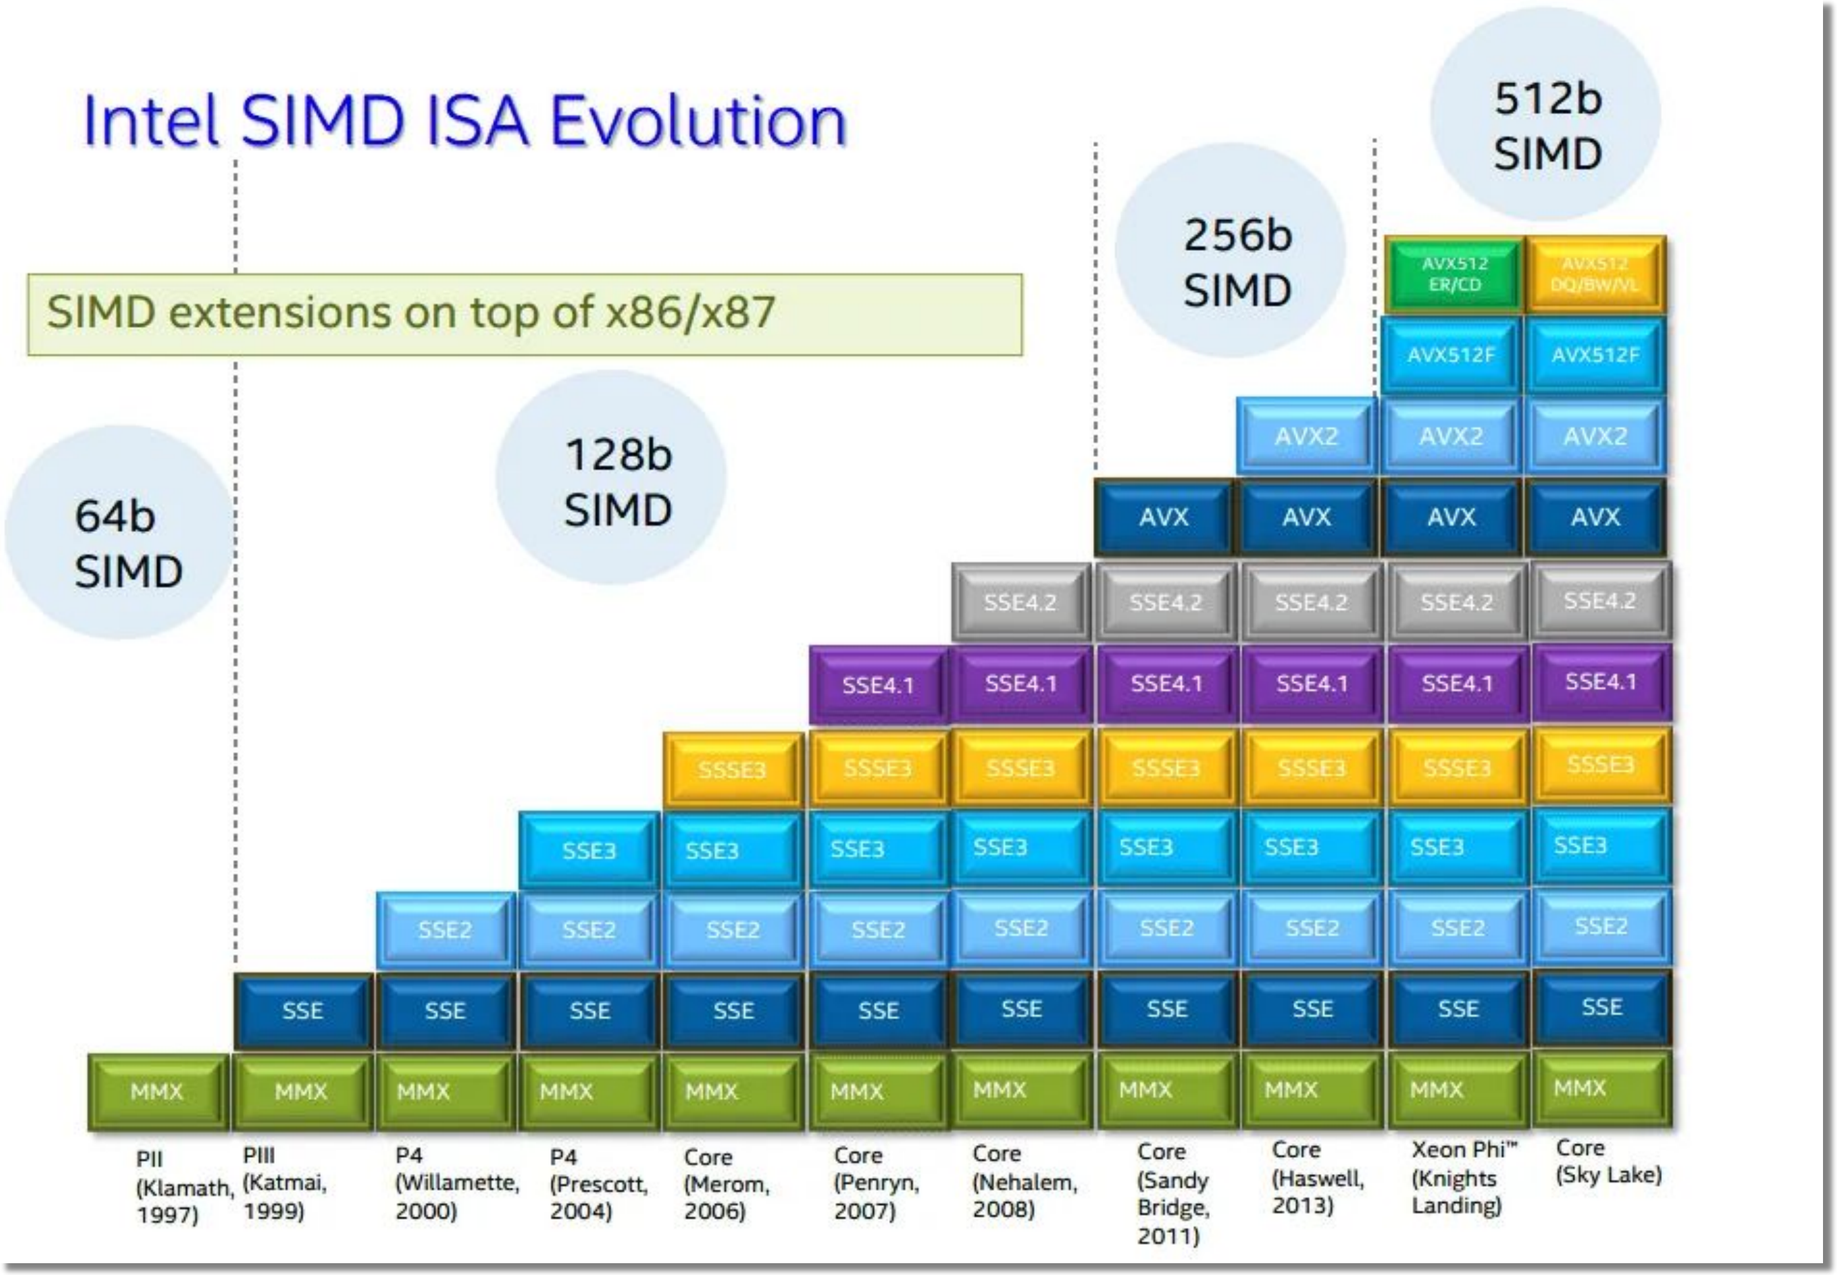
Intel SIMD ISA Evolution
SIMD extensions on top of x86/x87
64b SIMD
128b SIMD
256b SIMD
512b SIMD
MMX
SSE
SSE2
SSE3
SSSE3
SSE4.1
SSE4.2
AVX
AVX2
AVX512F
AVX512 ER/CD
AVX512 DQ/BW/VL
PII (Klamath, 1997)
PIII (Katmai, 1999)
P4 (Willamette, 2000)
P4 (Prescott, 2004)
Core (Merom, 2006)
Core (Penryn, 2007)
Core (Nehalem, 2008)
Core (Sandy Bridge, 2011)
Core (Haswell, 2013)
Xeon Phi™ (Knights Landing)
Core (Sky Lake)

It’s time to go back to the future!

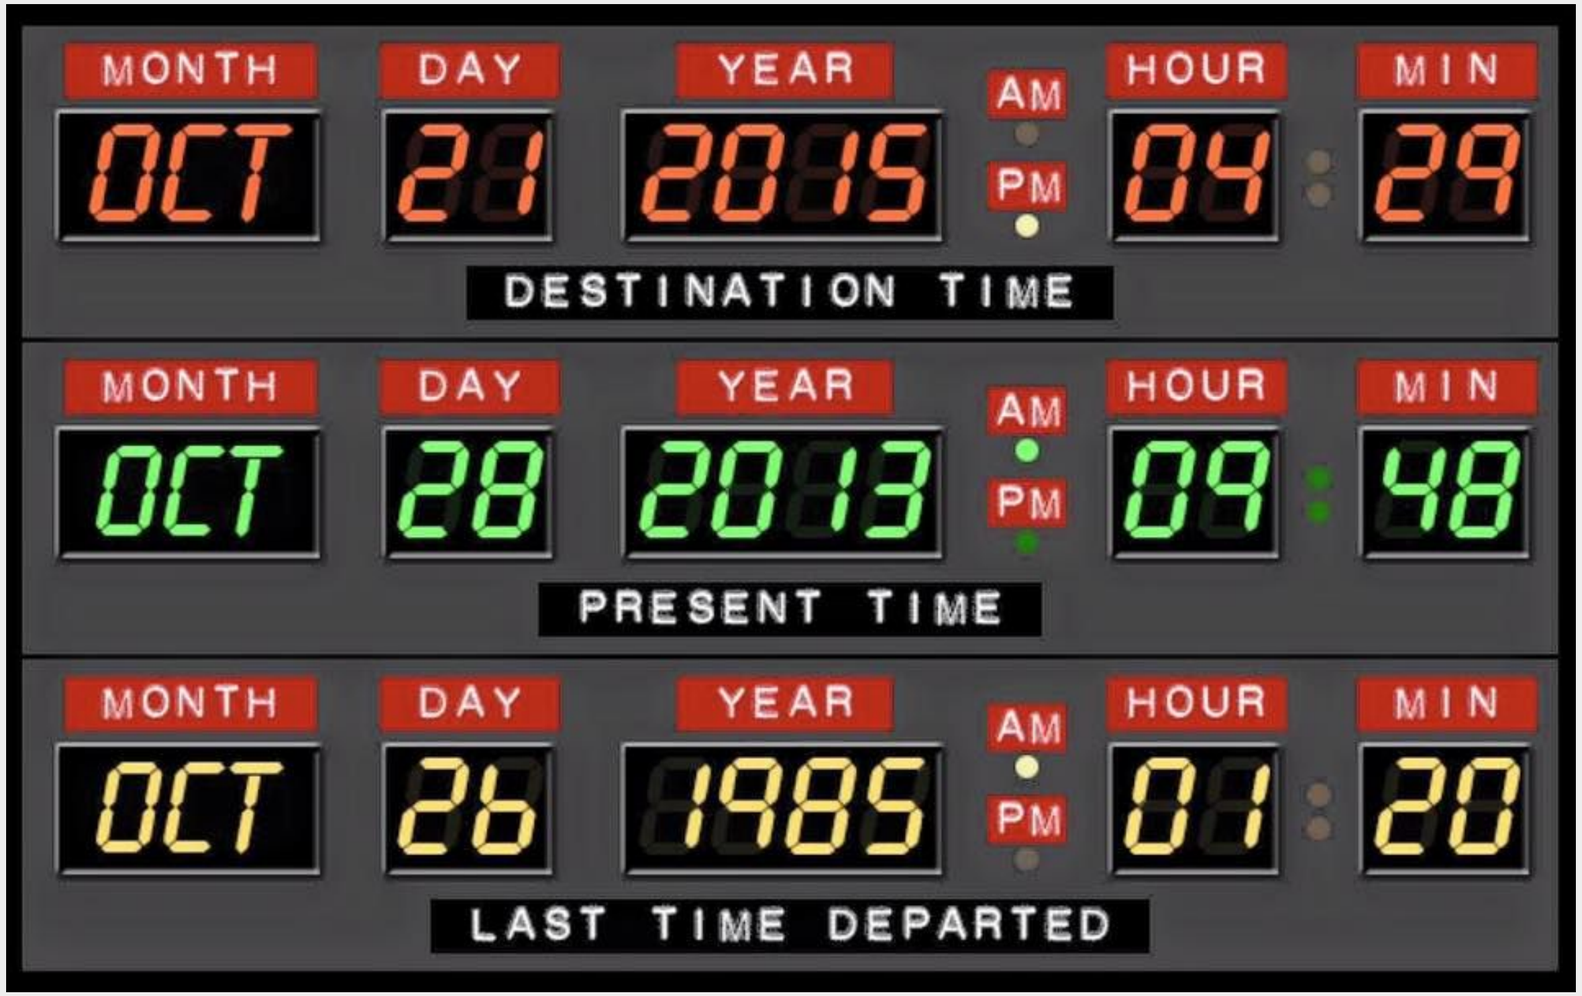

EESSI overview

HPC Now!

Overview

1. Collaboration between different European partners in HPC community.
2. Goal: building a scientific software stack for HPC systems & beyond.
3. Requirements: reproducibility, portability and performance.
4. Inspired by Compute Canada software stack.

EESSI

EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

HPC Now!

EESSI Partners

EESSI Architecture

- The bottom layer is the filesystem layer, which is responsible for distributing the software stack across clients.
- The middle layer is a compatibility layer, which ensures that the software stack is compatible with multiple different client operating systems.
- The top layer is the software layer, which contains the actual scientific software applications and their dependencies.

EESSI Architecture

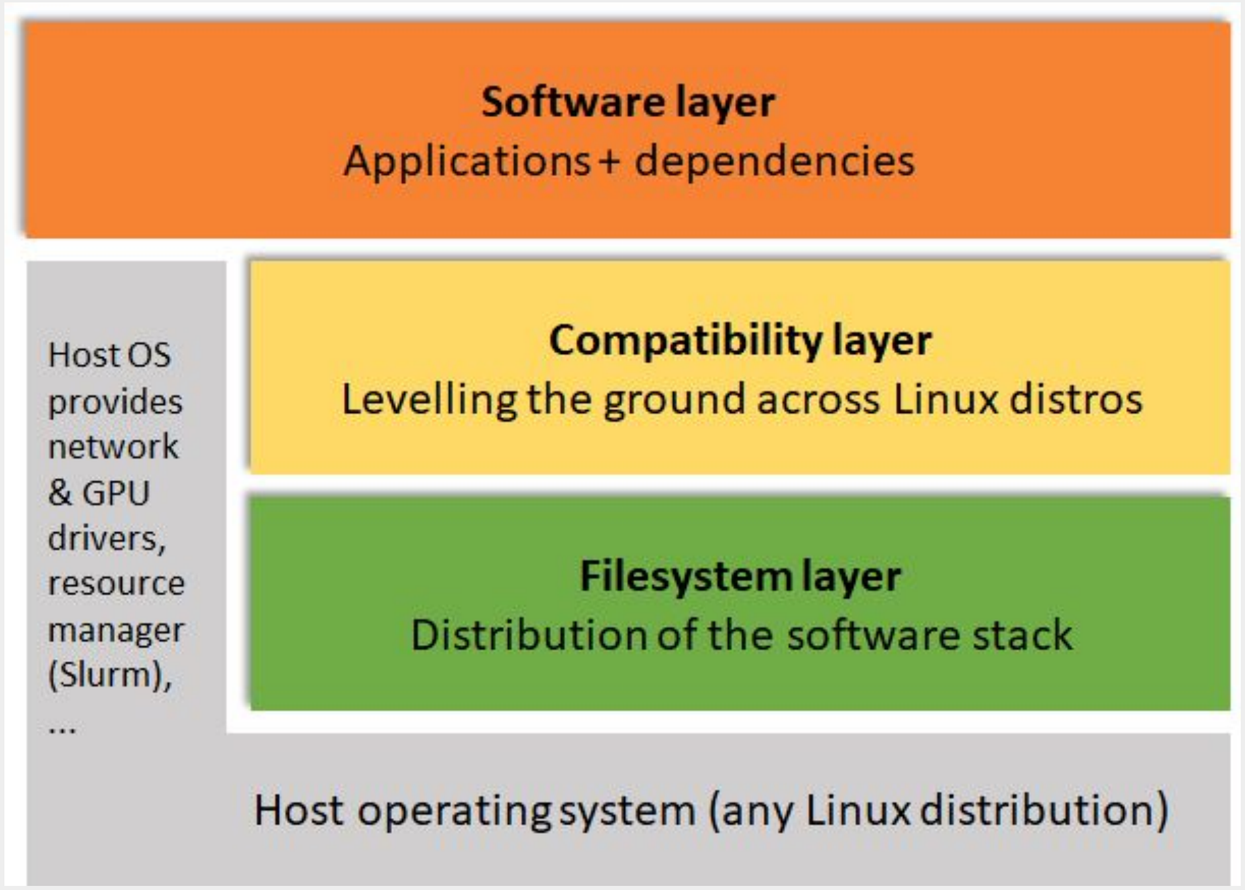

Lmod

arch spec

easybuild

gentoo

EESSI Architecture

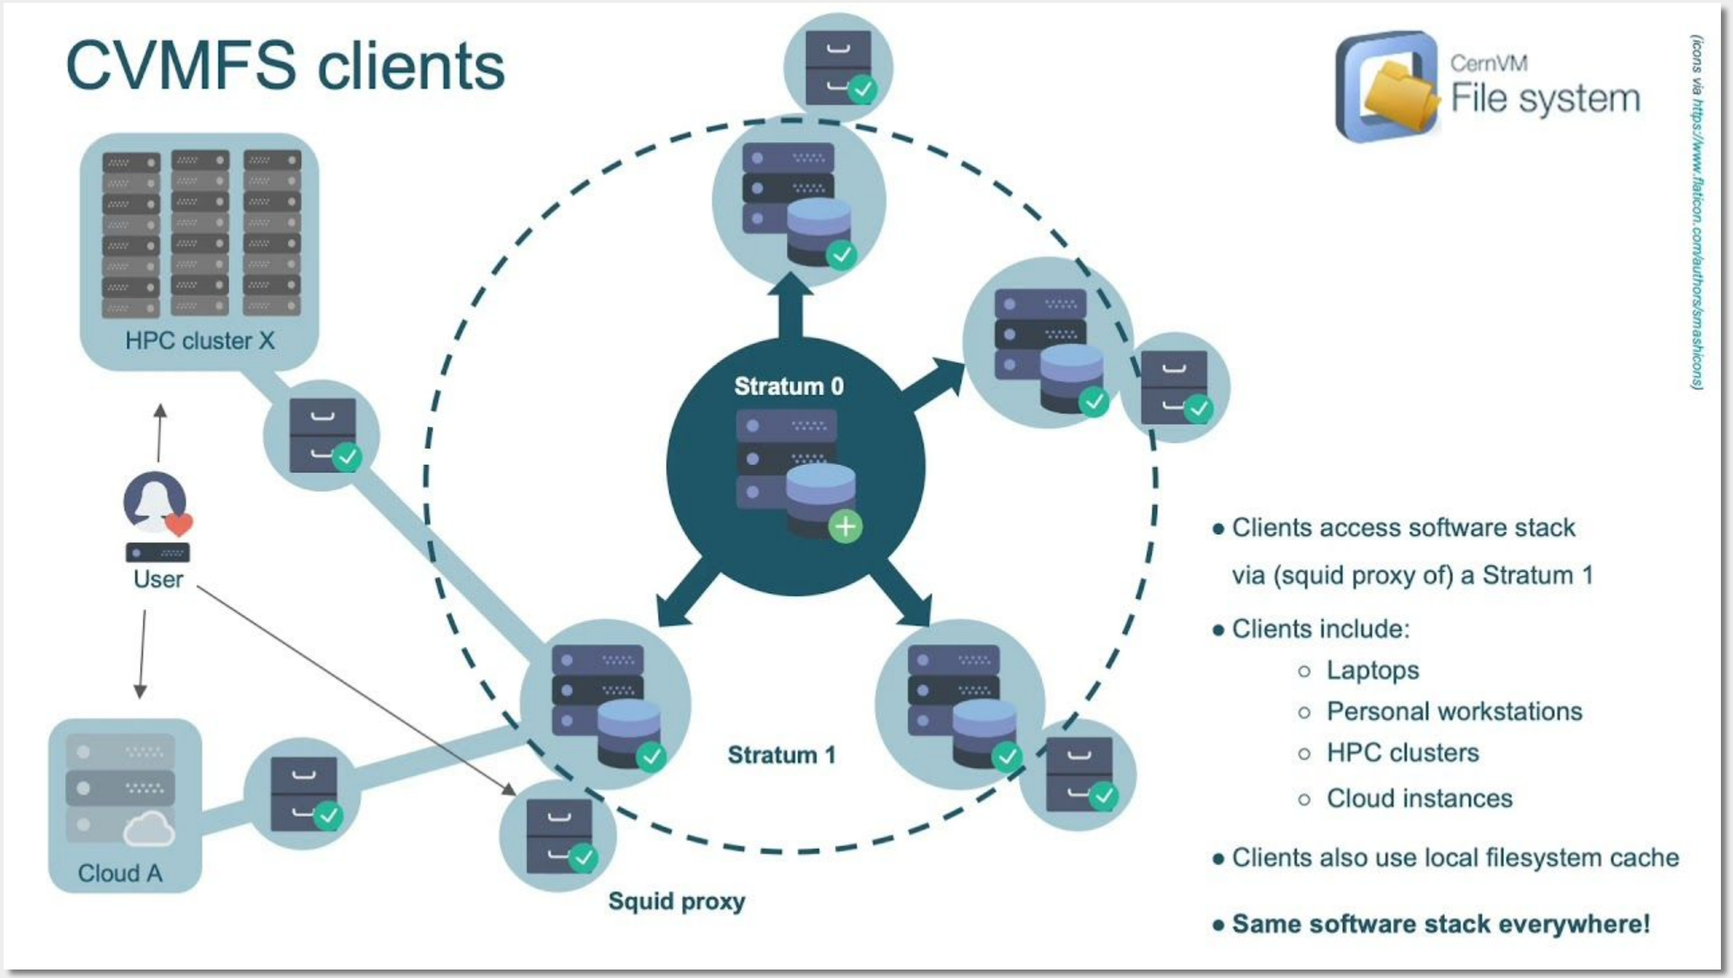

HPC Now!

More information

- EESSI website: https://www.eessi-hpc.org
- Slack channel: https://www.eessi-hpc.org/slack-channel/
- EasyBuild + EESSI workshop: https://hpckp.org/talks/ (available soon)

HPC Now!

Case Studies

Case Studies

Long-term Reproducibility

Pharmaceutical industry has a need for at least 10 years **reproducibility**. HPCNow! provides ongoing support to adopt EESSI architecture in order to cover this need, and also achieve **portability**, **performance**, and CI/CD for scientific software.

Use case: Pharmaceutical Company

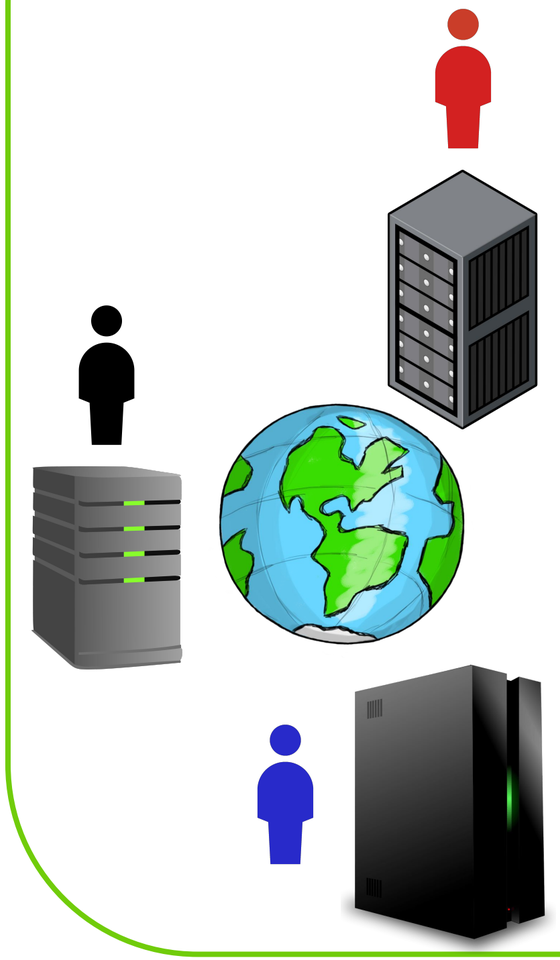

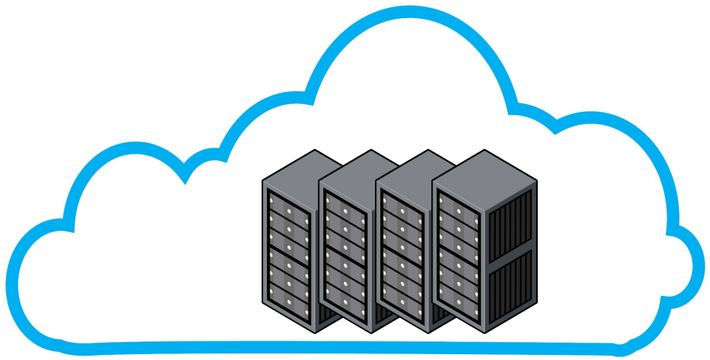

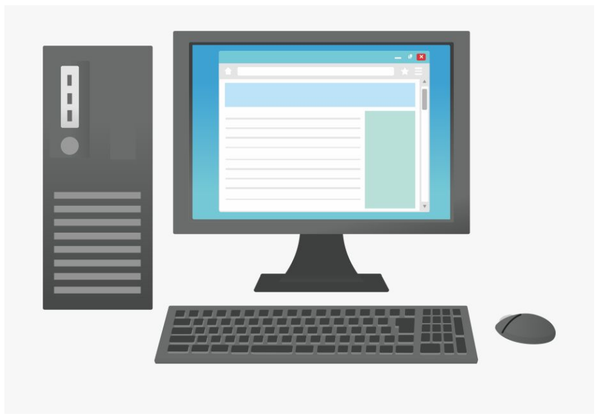

Public distributed appstacks will not match

Self developed databases

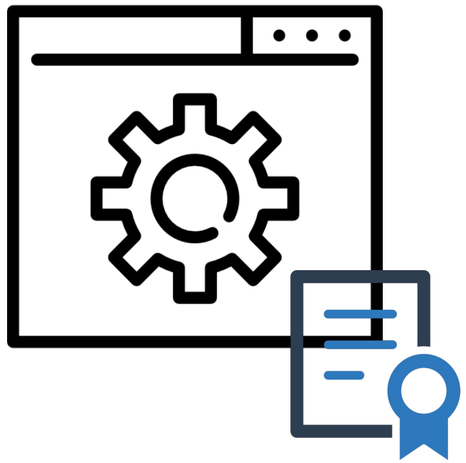

“Restricted”, custom and internal-developed software

Security concerns

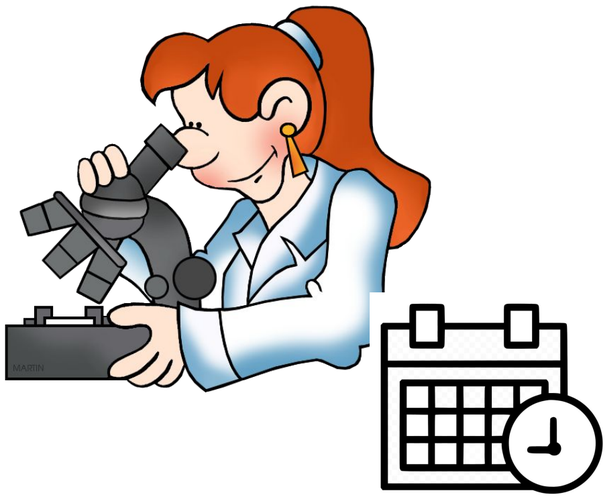

Reproducibility

Applications are built locally, and distributed using automatized workflows

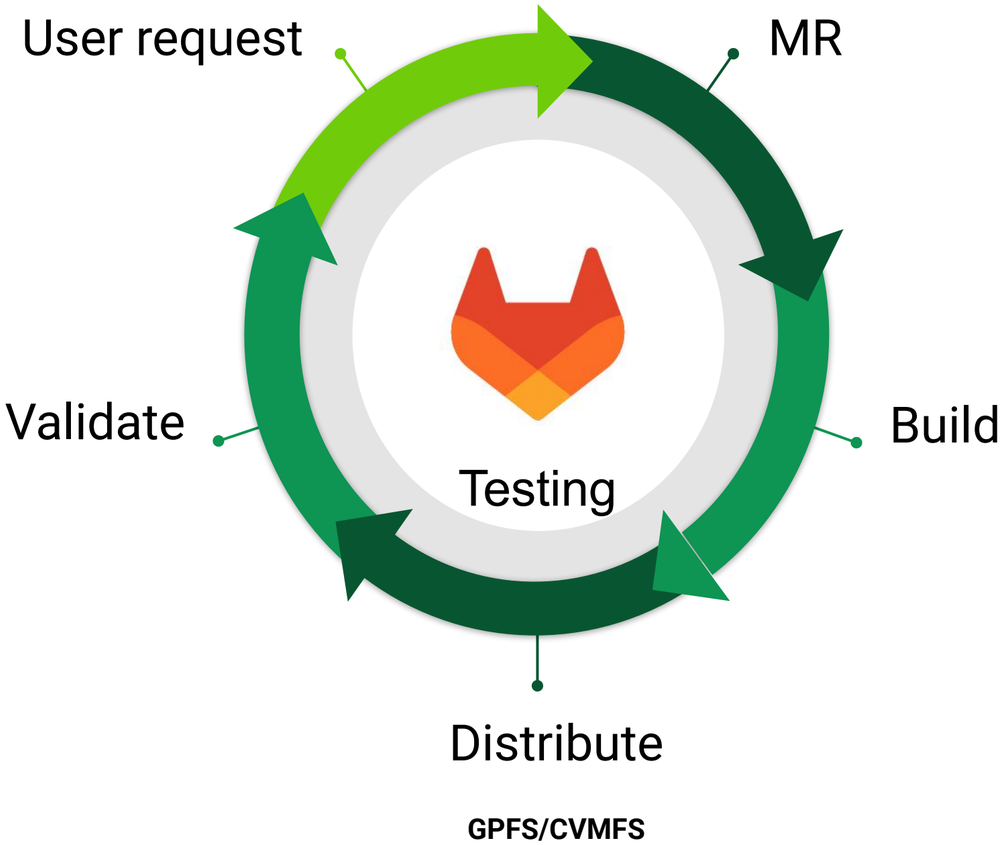

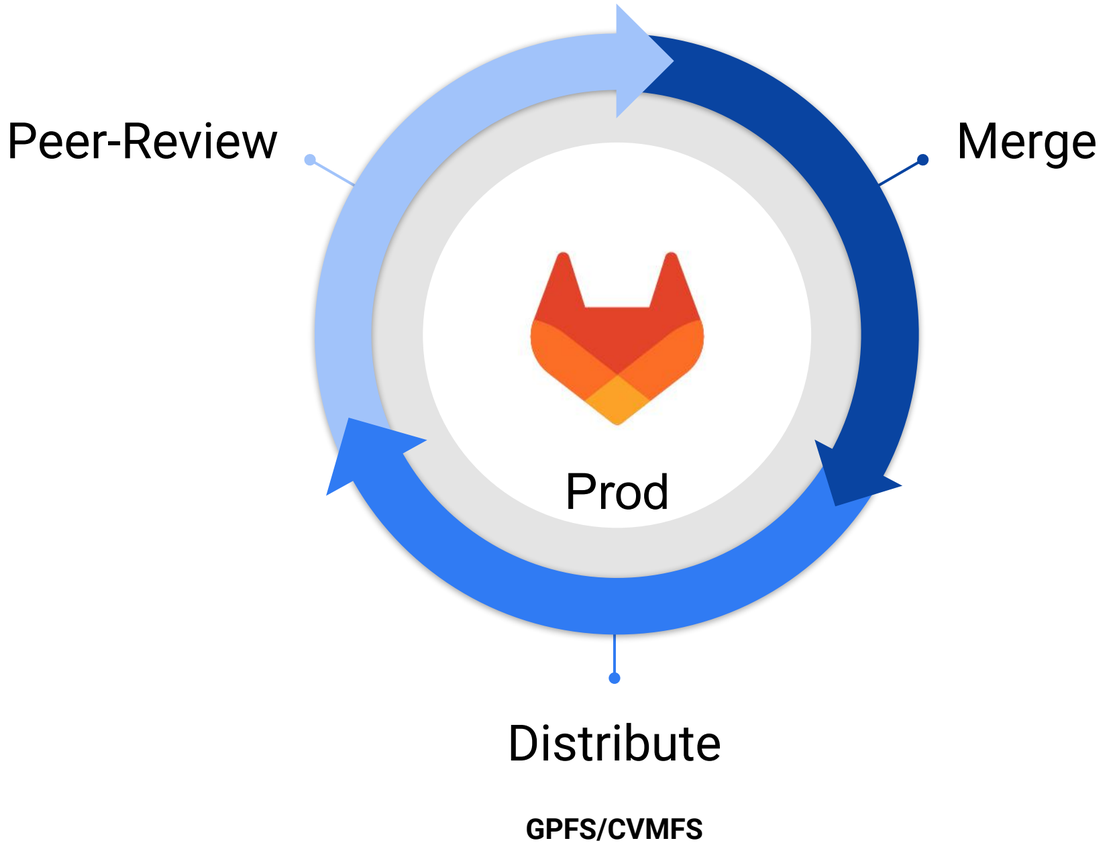

~200 apps + dependencies!

What about “restricted” software?

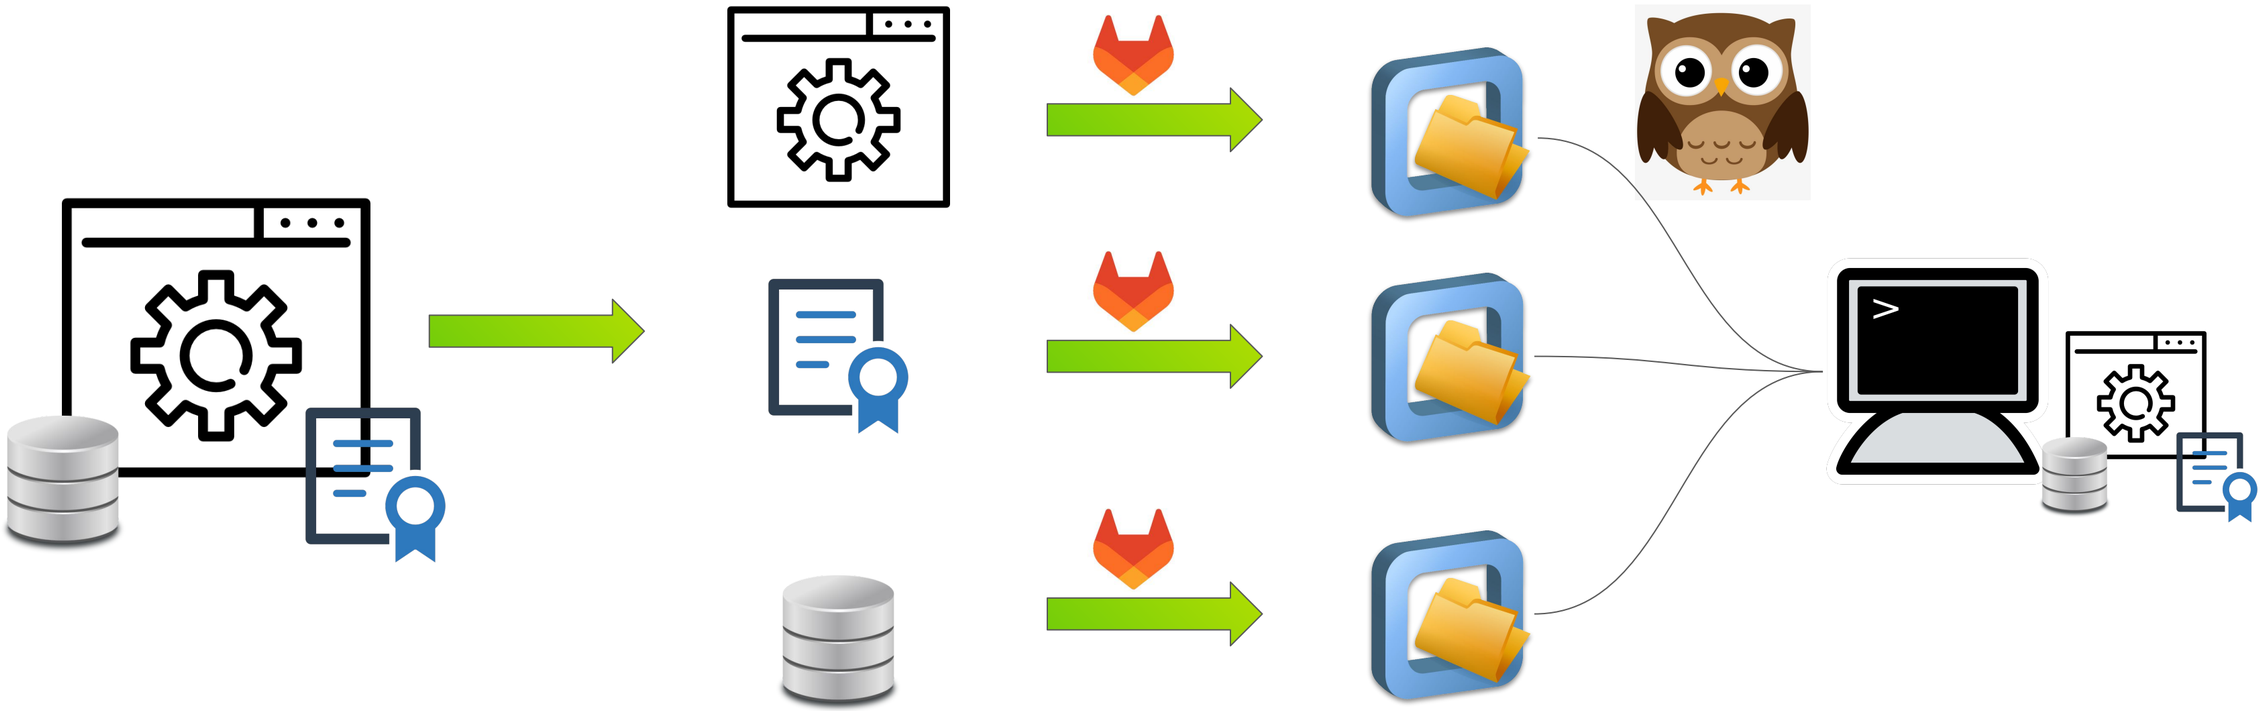

Access control

Layered approach provides a user suited appstack

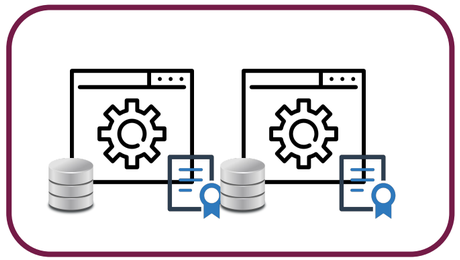

Org. Restricted

multiple repositories
apps.org1…
restricted.org1…
apps.org2…
restricted.org2…

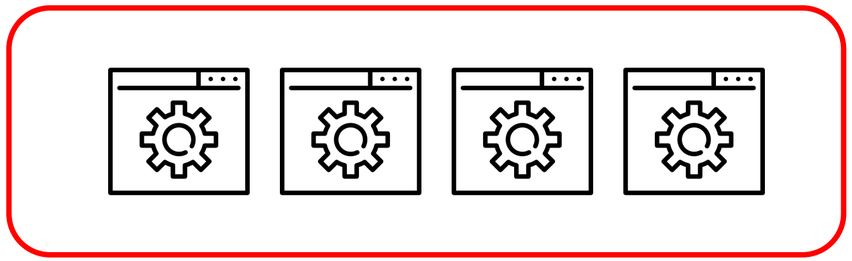

Org. Specific

Common

In production over a year ago!

We are part of Multiscale project!

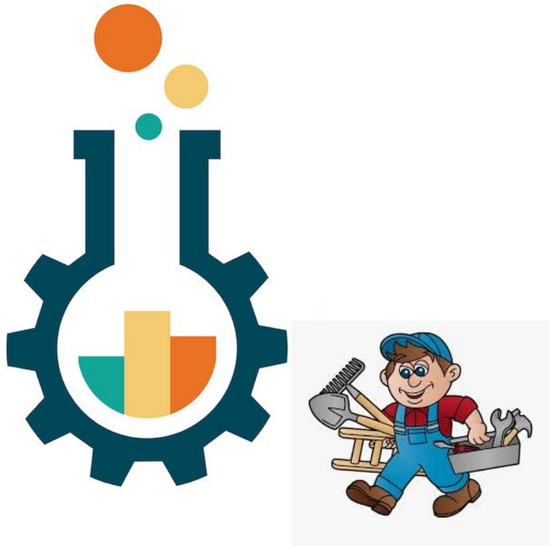

Appstack
Maintenance

Training

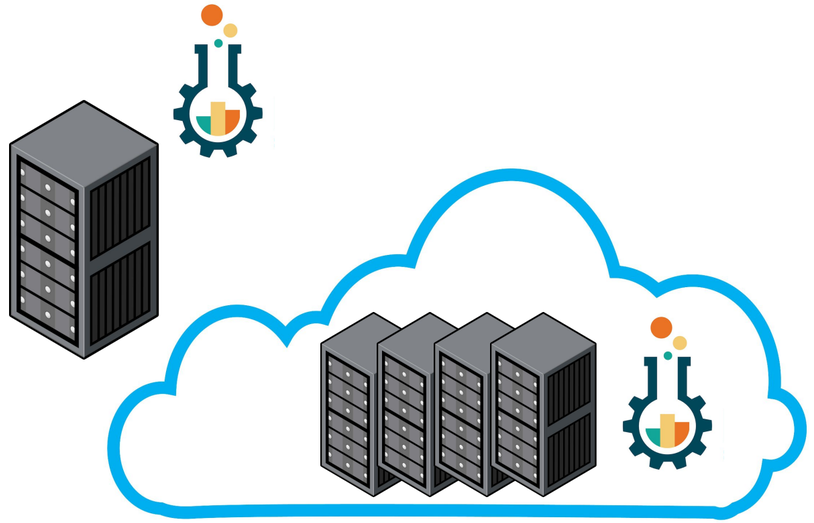

Customer support

and more!

We are hiring now!

careers@hpcnow.com

Thank you

Contact us:
info@hpcnow.com