



HPC&P

# HPC ANNUAL MEETING '23

Book of abstracts



# Contents.

Contents.	2
HPC Annual Meeting '23.	3
Schedule.	8
17 May 2023 - Wednesday	8
Schedule.	9
18 May 2023 - Thursday	9
Abstracts.	10
New Era for HPCKP	10
IBM Quantum roadmap and next developments	11
Simplification of quantum circuits with toolchains from the Eviden Quantum Learning Machine (QLM) Quantum Simulator	12
HPC in quantum era	13
Qibo, a full-stack quantum computing framework	14
Managing Complex Workflows Across Cloud Platforms with Covalent	15
Cloud and HPC Convergence: Flux for Job Management on Kubernetes	16
Cloud-Native Supercomputing: Exploring New Technologies for High-Performance Computing	17
Running Slurm cluster within Kubernetes environment	18
Container Technologies for HPC	19
Optimize the use of your IB NDR network	20
HPC Infrastructure for Energy: Past, Present, and Beyond	21
HPC for a weather information company	22
Cornelis Omni-Path Express: The open-source software approach to High Performance Fabrics	23
Bridging the Gap: Introducing Fusion - A POSIX Distributed File System for Cloud-Native and On-Premises Storage	24
Reducing the HPC Skills Gap with OpenFlightHPC	25
DPU Offloading with OpenMP Programming Model	26
Streaming scientific software: an introduction to EESSI	27

# HPC Annual Meeting '23.

HPC Knowledge Portal annual meeting is a key global event for High-Performance Computing able to attract relevant professionals and main developers of leading software projects widely used by the HPC community.

## Location

The meeting will be hosted by  
Barcelona Advanced Industry Park,  
Marie Curie, s/n,  
08042 - Barcelona (Spain).

## Social Media

**Twitter:**  
<https://twitter.com/HPCKP>  
**LinkedIn:**  
<https://www.linkedin.com/company/hpckp/>  
**Youtube:**  
<https://www.youtube.com/HPCKP>

## Website

An up-to-date agenda, slides and  
training material are available in  
HPCKP  
<https://hpckp.org/annual-meeting/>

## Adamantium Sponsors








# EVIDEN

Expanding the  
possibilities of  
data and  
technology

An Atos business

## LENOVO - High Performance Computing Solutions

In today's competitive environments, accelerated research is the key to success. But what's the best way to manage the huge amounts of data and heavy compute workloads—on a tight budget? Lenovo high-performance computing (HPC) solution can make a huge difference. Powered by the Intel®



Xeon® Platinum processor, Lenovo HPC solutions can boost performance, accelerate innovation, and drive breakthroughs, all while minimizing costs and risks to address business and technical challenges across many industries.

Lenovo is helping engineering firms accelerate simulation and analytics, life sciences research and development organizations speed drug discovery and genomics sequencing insights, and financial services firms boost high-frequency trading and risk analytics applications. Also, it has completed the delivery and implementation of the world's largest, next-generation Intel-based Supercomputer at the Barcelona Supercomputing Center (BSC) in 2017. The fast delivery, installation and optimization of the MareNostrum 4 system at BSC, showcases Lenovo's end-to-end,

high-performance computing strength. It is the third leading-edge HPC system that Lenovo has installed at the Partnership for Advanced Computing in Europe (PRACE), making Lenovo the largest provider of leading-edge HPC systems to this high impact scientific community, and solidified Lenovo's position as the fastest growing supercomputer company in Europe.

Lenovo is a US\$43 billion global Fortune 500 company and a leader in providing innovative consumer, commercial, and data center technology. Our portfolio of high-quality, secure products and services covers PCs (including the legendary Think and multimode Yoga brands), workstations, servers, storage, networking, software (including ThinkSystem and ThinkAgile solutions), smart TVs and a family of mobile products like smartphones (including the Motorola brand), tablets and apps. Join us on LinkedIn, follow us on Facebook or Twitter (@Lenovo) or visit us at [www.lenovo.com](http://www.lenovo.com)







## Interconnecting a Data-driven World

Omni-Path delivers unprecedented ease of use while our community-first approach ensures the interoperability you need.

### Cornelis Omni-Path Fabric Advantage

World-class End-to-End High-Performance Fabric Solution for HPC, AI & HPDA

**Performance** Lowest scalable latency, highest message rate, high data throughput and bandwidth

**Power** Leading energy efficiency end-to-end fabric at all scale

**Price** Cost effective cluster solutions for better budget use

**Product Innovation** Innovate high-performance link-layer using FLITS: 100G/400G/800G

**Partner Friendly** Eco-system friendly, no bundling, no bid conditions no exclusivity requirements

**Product Availability** Short OPA100 product leadtime

[cornelisnetworks.com](https://cornelisnetworks.com)

 @qilimanjaro

 qilimanjaro@qilimanjaro.tech



**Qilimanjaro** makes you quantum ready with our quantum algorithmic consultancy service and our analog quantum computer with high-quality flux-qubit chips, to be released soon on the cloud! Stay tuned!

**We're interested  
in your hard compute challenges!**

[WWW.QILIMANJARO.TECH](http://WWW.QILIMANJARO.TECH)

# Schedule.

## 17 May 2023 - Wednesday

08:15 - 09:00 Registration

09:00 - 09:30 Welcome  
**Jordi Blasco** (CTO @HPCNow!)  
New Era for HPCKP

### Quantum

09:30 - 10:00 Talks  
**Albert García** (Quantum Community Engineer Lead @ IBM)  
IBM Quantum roadmap and next developments

10:00 - 10:30 Coffee Break

10:30 **Andrés Bravo** (QLM Presales Quantum Engineer @ Eviden)  
Simplification of quantum circuits with toolchains from the Eviden Quantum Learning Machine (QLM) Quantum Simulator

11:00 **Elisabeth Ortega** (R&D&I Manager @ HPCNow!)  
HPC in quantum era

11:30 **Albert Solana** (Chief Business Officer @ Qilimanjaro)  
Qibo, a full-stack quantum computing framework

12:00 **William J. Cunningham** (Head of High Performance Computing @ Agnostiq)  
Managing Complex Workflows Across Cloud Platforms with Covalent

12:30 Panel discussion

13:00 - 14:00 Lunch Break

### Convergence of HPC/AI and Cloud

14:00 **Claudia Misale** (Researcher @ IBM) + **Vanessa Sochat**  
(Computer Scientist @ LLNL)  
Cloud and HPC Convergence: Flux for Job Management on Kubernetes

14:30 **Carlos Arango** (Senior Systems software engineer @ NVIDIA)  
Cloud-Native Supercomputing: Exploring New Technologies for High-Performance Computing

15:00 **Hussein Harake** (HPC System Manager @ CSCS)  
Running Slurm cluster within Kubernetes environment

15:30 **Alberto García** (HPC Architect @ Do IT Now)  
Container Technologies for HPC

16:00 - 17:00 Trip by bus to BSC

17:00 - 18:30 Visit to MareNostrum (2 groups)

18:30 - 21:00 Beers & Pizzas at **AltaglioBcn**



---

# Schedule.

---

18 May 2023 - Thursday

---

## HPC Innovation

---

- 09:00 Miguel Terol** (Senior HPC Architect @ Lenovo)  
Optimize the use of your IB NDR network
- 09:30 Felipe Portella** (IT Advisor @ Petrobras)  
HPC Infrastructure for Energy: Past, Present, and Beyond
- 10:00 Kim Serradell** (Senior Site Reliability Engineer @ Meteomatics AG)  
HPC for a weather information company
- 10:30 - 11:00** Coffee Break
- 11:30 Robert Bollig** (HPC Solution Architect @ Cornelis Networks)  
Cornelis Omni-Path Express: The open-source software approach to High Performance Fabrics
- 12:00 Jordi Deu-Pons** (Senior Software/Data Engineer @ Seqera Labs)  
Bridging the Gap: Introducing Fusion - A POSIX Distributed File System for Cloud-Native and On-Premises Storage
- 12:30 Cristin Merritt** (CMO @ Alces Flight Limited)  
Reducing the HPC Skills Gap with OpenFlightHPC
- 13:00 - 14:00** Lunch Break
- 14:00 Muhammad Usman** (Research Engineer @ BSC)  
DPU Offloading with OpenMP Programming Model
- 14:30** Panel discussion
- 15:00 - 15:10** Coffee Break

---

## EESSI workshop

---

- 15:10 - 17:00 Alan O'Cais** (Research Fellow, University of Barcelona; Research software engineer, CECAM)  
Streaming scientific software: an introduction to EESSI
- 17:00 - 17:30** Wrap up

# Abstracts.

---

## New Era for HPCKP

### Abstract

I am thrilled to announce the inauguration of our new online forum and articles section on our HPC knowledge website. Our team has been working diligently to create a platform where experts in High-Performance Computing can come together and share their knowledge and ideas with the broader community.

The online forum will allow users to engage in discussions on a variety of topics related to HPC, including hardware and software optimization, parallel programming, and performance analysis. We hope this will facilitate the exchange of ideas and foster collaboration within the community.

Additionally, our new articles section will feature in-depth articles and tutorials on various HPC topics, written by experts in the field. These articles will provide valuable insights in areas like performance, efficiency, technology adoption, reliability, and end-user education.

We are excited to offer these new resources to the HPC community and look forward to seeing the exciting discussions and collaborations that will arise from them.



**Author:** Jordi Blasco  
**CTO @ HPCNow!**  
**Date:** 17 May 2023

# IBM Quantum roadmap and next developments

## Abstract

Quantum computing is creating a world of possibilities for new technologies. To better understand them, we will provide an in-depth analysis of the current state of IBM quantum, with a particular focus on its potential impact on several industries, including healthcare, finance, and transportation. It is important to understand also the future trends and how we will arrive there through an overview of the future of IBM Quantum technology and the HW and SW roadmap that will guide its development.



**Author:** Albert García  
**Quantum Community Engineer Lead @ IBM**  
**Date:** 17 May 2023

---

# Simplification of quantum circuits with toolchains from the Eviden Quantum Learning Machine (QLM) Quantum Simulator

## Abstract

Through the tools provided by the Eviden QLM quantum simulator, it is possible to rewrite circuits. This technique consists of automatically compressing sequences of consecutive one-qubit gates into a universal pattern. In this intervention, it will be detailed how this technique has been applied in a quantum circuit and how it has been possible to reduce the number of quantum gates of a qubit. In addition, the results obtained after the execution in different IBM backends will be shown and how the reliability of the circuit is improved. Everything from a single environment, that of QLM.



**Author:** Andrés Bravo Montes  
**QLM Presales Quantum Engineer @ Eviden**  
**Date:** 17 May 2023

---

# HPC in quantum era

## Abstract

This 2023, quantum computing is evolving faster than ever. Last year, Europe selected six sites to host the first European Quantum Computers, which will be added to other projects like the future quantum computers in CESGA, BSC and Ikerbasque. However, the quantum community is not waiting to have real quantum hardware before investigating new algorithms and solutions.

Quantum emulators can be used to start experimenting with quantum software. On a small scale, algorithms can be executed on a laptop, but complicated solutions will need the use of CUDA or MPI acceleration in an HPC environment. In this talk, we will focus on this last case.



**Author:** Elisabeth Ortega  
**R&D&I Manager @ HPCNow!**  
**Date:** 17 May 2023

---



# Qibo, a full-stack quantum computing framework

## Abstract

In this talk, Qilimanjaro will introduce its full-stack quantum computing framework, Qibo [1], putting a special focus on its novel approach to maximize the use of local computing resources (both quantum and classical) when executing quantum algorithms.

Qibo is an open-source full stack API for quantum simulation and quantum hardware control with several contributors, including Qilimanjaro, who integrates it to its Quantum as a Service solution, including its own low-level control, named qililab.

We will present the newest design of qililab solution to run quantum algorithms for both gate-based and annealing quantum computers taking into consideration the interaction with classical hardware (multi-CPU, multi-GPU) and QPUs to ease the integration of our quantum solutions in HPC centers.

[1] <https://qibo.readthedocs.io/en/stable/>



**Author:** Albert Solana  
**Chief Business Officer @ Qilimanjaro**  
**Date:** 17 May 2023

# Managing Complex Workflows Across Cloud Platforms with Covalent

## Abstract

Covalent is an open-source Pythonic tool designed to orchestrate heterogeneous HPC and quantum workflows, enabling computational research to be reproducible and manageable. By attaching metadata to workflows as they execute, Covalent provides users with a way to inspect everything from source code to data dependencies to software and environments in real time. With a browser-based user interface, users can easily visualize large and complex workflows distributed across multiple clouds, including AWS, Azure, and GCP Batch services.

In this talk, we will explore how Covalent can be used in multi-cloud and hybrid-cloud environments to enable seamless workflow management across various platforms. Specifically, we will discuss the benefits of using Covalent with AWS compute resources, as well as some of the challenges involved in joining cloud and on-prem HPC resources. I will also highlight the upcoming release of Covalent Cloud, a SaaS offering of Covalent which expands the tool's capabilities and makes it accessible to a broader range of users.



**Author:** William J. Cunningham  
**Head of High Performance Computing @ Agnostiq**  
**Date:** 17 May 2023

# Cloud and HPC Convergence: Flux for Job Management on Kubernetes


## Abstract

Adoption of cloud technologies by high performance computing (HPC) is accelerating, and HPC users want their applications to perform well everywhere. While container orchestration frameworks provide advantages like resiliency, elasticity, and declarative management, they are not designed to natively enable application performance to the same degree as HPC workload managers and schedulers. In this talk, we present our effort in promoting Cloud and HPC convergence, an effort that spans computing platforms and institutions. Specifically, we target Kubernetes platforms with the introduction of the Flux Operator and the Fluence plugin scheduler, both based on the Flux open-source HPC resource manager and job scheduler. We also describe the challenges we faced when shifting from traditional HPC to Kubernetes, and what it takes to run workflow natively on such platforms.



**Author:** Claudia Misale (**Researcher @ IBM**) & Vanessa Sochat (**Computer Scientist @ LLNL**)

**Date:** 17 May 2023



# Cloud-Native Supercomputing: Exploring New Technologies for High-Performance Computing

## Abstract

As cloud-native technologies continue to advance, researchers and engineers in high-performance computing (HPC) are beginning to explore how these technologies can be used to build more scalable, reliable, and efficient supercomputers. In this talk, we will explore some of the latest technologies being used in cloud-native supercomputing, and discuss how they are being integrated into Kubernetes, the de facto standard for container orchestration.

One of the most exciting new technologies for cloud-native supercomputing is dynamic resource allocation (DRA), which has recently been merged into Kubernetes upstream. DRA allows users to dynamically allocate resources based on their application's requirements, making it possible to optimize resource usage and reduce costs.

Another important technology is the container device interface (CDI), which provides a standardized way for containers to access devices like GPUs and FPGAs. CDI is critical for running high-performance computing workloads on cloud-native infrastructure, as it allows applications to take advantage of specialized hardware.

We will also discuss Kueue, a Kubernetes-native queuing controller that enables users to easily manage complex workloads with dependencies and priorities. Finally, we will provide updates on the MPI operator, which automates the deployment and scaling of MPI applications in Kubernetes clusters.

By leveraging these new technologies, cloud-native supercomputing promises to unlock new levels of scalability and efficiency for HPC workloads, making it possible to tackle even the most demanding scientific and engineering challenges.



**Author:** Carlos Arango  
**Senior Systems software engineer @ NVIDIA**  
**Date:** 17 May 2023

# Running Slurm cluster within Kubernetes environment

## Abstract

This talk will discuss the process of containerising and running a Slurm cluster within a Kubernetes environment. The discussion will primarily focus on the challenges encountered while building the container images and how to orchestrate them effectively. During the talk, we will also demonstrate how pods can be run as a complete operating system, including all necessary services. Additionally, we will might be able to benchmark results of running MPI (Message Passing Interface) jobs within this containerised environment.



**Author:** Hussein Harake  
**HPC System Manager @ CSCS**  
**Date:** 17 May 2023

---



# Container Technologies for HPC

## Abstract

Containers have been growing in popularity all across the information technologies landscape, including specialized infrastructures, like HPC, but the constraints of architecture, performance, and security make its adoption difficult. This presentation provides an overview of the current technologies in competition on HPC, from a technical as well as non-technical perspective, highlighting their differences and providing an assessment according to experiences extracted from real use cases.



**Author:** Alberto García  
**HPC Architect @ Do IT Now**  
**Date:** 17 May 2023

---

# Optimize the use of your IB NDR network

## Abstract

The new Infiniband NDR technology allow for bandwidths of up to 400 Gbps per host port. Although some applications use as much bandwidth as available, for many other these numbers are overkilling. We will present how to optimize your IB topology, leveraging Lenovo SharedIO concept for dense systems to allow up to 256 compute nodes in a single 1U IB switch.



**Author:** Miguel Terol  
**Senior HPC Architect @ Lenovo**  
**Date:** 18 May 2023

---

# HPC Infrastructure for Energy: Past, Present, and Beyond

## Abstract

Petrobras is one of the world's leading energy companies. As such, it relies heavily on high-performance computing (HPC) infrastructure for various workloads, ranging from seismic processing to reservoir simulation and, more recently, artificial intelligence. This presentation will provide an overview of Petrobras' HPC infrastructure history, from its early beginnings in 1969 to the present day. During its early years, Petrobras employed a mix of mainframe computers, RISC machines, and Beowulf clusters to perform its computing tasks. Throughout the years, the company has upgraded its infrastructure, and today it operates several TOP500-ranked supercomputers, including the Pegasus and Dragão systems, which are currently ranked as the 33rd and 65th most powerful supercomputers in the world, respectively.

We are already looking toward the future at Petrobras, and this presentation will provide insights into cutting-edge HPC architectures, such as cloud and quantum computing. The presentation will also explore how serverless computing infrastructures could push HPC boundaries for geoscience applications.



**Author:** Felipe Portella  
**IT Advisor @ Petrobras**  
**Date:** 18 May 2023

# HPC for a weather information company

## Abstract

Meteomatics is a Swiss company, global leader in weather intelligence. The main mission is to provide the most accurate weather data for any location at any time to improve our customers' business.

Meteomatics data is delivered using the powerful weather API, combining massive amounts of data from external sources and also from the EURO1k, an in-house high-resolution European weather model for precise forecasts. In this talk, we will present the work done to deploy this model in our on-premise high-performance cluster, the technical challenges experienced to efficiently manage an intense I/O workflow, future directions and we will compare the previous experiences of the company using a private cloud provider.



**Author:** Kim Serradell  
**Senior Site Reliability Engineer @ Meteomatics AG**  
**Date:** 18 May 2023

---

# Cornelis Omni-Path Express: The open-source software approach to High Performance Fabrics

## Abstract

As HPC and AI cluster architecture becomes increasingly heterogenous, Cornelis Networks have adopted the Open Fabrics Alliance (OFA) model for host software development to drive enhanced application performance and scalability.

This, coupled with a fully open-source software model, fully leverages development work done in the community.



**Author:** Dr. Robert Bollig  
**HPC Solution Architect @ Cornelis Networks**  
**Date:** 18 May 2023

---



# Bridging the Gap: Introducing Fusion - A POSIX Distributed File System for Cloud-Native and On-Premises Storage

## Abstract

Fusion is a virtual, lightweight, distributed file system that bridges the gap between cloud-native storage and pipelines. Is used by Nextflow pipeline manager, a popular tool for developing and executing data-intensive pipelines.

It enables seamless filesystem I/O to object stores via a standard POSIX interface, resulting in simpler pipeline logic and faster, more efficient pipeline execution. But that's not all. Fusion can also be used on-premises and coexist with traditional HPC parallel file systems, allowing for a seamless data management experience across hybrid and multi-cloud environments.

This talk proposal introduces Fusion as a versatile solution for managing data in hybrid and multi-cloud environments. It will delve into the technical aspects of Fusion, how it integrates with object stores, and its benefits in terms of performance and efficiency.

<https://seqera.io/>



**Author:** Jordi Deu-Pons - Software Engineer at Seqera Labs  
**Senior Software/Data Engineer @ Seqera Labs**  
**Date:** 18 May 2023

# Reducing the HPC Skills Gap with OpenFlightHPC

## Abstract

OpenFlightHPC is an open-source project focused on platform agnostic HPC cluster deployment and use. The project's aim is to reduce the skills gap in building and using HPC systems for a range of different use cases. Initially developed to support construction of supercomputers on public-cloud platforms, OpenFlightHPC has since evolved into covering bare-metal hardware and private cloud deployments as well.

With the ability to handle jobs requiring low-latency fabrics, parallel filesystems, big-memory systems, and modern container-based AI and ML workloads, there is plenty of scope for building both agile, single-use and more complex feature-full environments. In this talk we will focus on our newest features designed for novice HPC user enablement and how, from template use-cases, they are able to step gently into the world of supercomputing, avoiding a steep learning curve.

Utilising a small demo cluster we will demonstrate how an individual can develop their skills, starting with the guidelines you can put into place, up to the user being able to confidently amend or reconfigure their environment. We will also discuss how to use OpenFlightHPC as a form of user empowerment, and how this project can aid in promoting self-improvement and help foster greater inclusivity in our industry.

OpenFlightHPC site: <https://www.openflighthpc.org/>



**Author:** Cristin Merritt  
**CMO @ Alces Flight Limited**  
**Date:** 18 May 2023

# DPU Offloading with OpenMP Programming Model

## Abstract

Data Processing Units (DPUs) are taking high-speed SmartNIC development further by integrating performant processing cores it. We are enabling BlueField DPU capabilities in the OpenMP programming model. The enablement of OpenMP Target Offload features for BlueField DPUs will contribute to accessibility to a wider range of users. It will be an opportunity for the HPC community to leverage DPU features for a wider range of existing and emerging applications.



**Author:** Muhammad Usman  
**Research Engineer @ BSC**  
**Date:** 18 May 2023


---

# Streaming scientific software: an introduction to EESSI


## Abstract

Have you ever wished that all the scientific software you use was available on all the resources you had access to without having to go through the pain of getting them installed the way you want/need? The European Environment for Scientific Software Installations (EESSI - pronounced "easy") is a common stack of scientific software installations for HPC systems and beyond, including laptops, personal workstations and cloud infrastructure. In many ways it works like a streaming service for scientific software, instantly giving you the software you need, when you need it, and compiled to work efficiently for the architecture you have access to.

In this tutorial, we'll explain what EESSI is, how it is being designed, how to get access to it, and how to use it. We'll give a number of demonstrations and also give access to a resource where you can try EESSI out for yourself.



**Author:** Alan O'Cais  
**Research Fellow, University of Barcelona; Research software engineer, CECAM**  
**Date:** 18 May 2023



+34 931640488

[info@hpcnow.com](mailto:info@hpcnow.com)