



CEPP Training
Accelerate workload,
give value to simulation!

The CEPP (Centre for Excellence in Performance Programming) supports organisations worldwide in tackling industrial and scientific challenges by harnessing the value from vast amounts of data. By combining deep HPC proficiency, data sciences, and extensive industry knowledge, CEPP accelerates HPC simulations, optimises cluster performance and operations, and reduces the cost-to-innovation.

With a modular and flexible service offering, CEPP can customise its services to adapt to your goals, offering training, workshops, webinars, or dedicated resources for your specific project.



CEPP welcome pack

This Welcome Pack is designed for new HPC, AI, & QC customers and is offered free of charge. It serves as an introductory session to introduce customers to Bull projects and research activities, different user groups, technology roadmaps, technology preferences, specific needs, and evaluate interests. Additionally, it introduces the Bull User group for eXtreme computing (BUX). Following this, we discuss potential joint collaborations and, if suitable, define a plan of activities. The session also delves into methodological details.



Why is this important?

Regardless of your activity or the production use of your supercomputer, fostering a community is crucial. This community can offer a deeper technological understanding of the supercomputer, enrich knowledge of current trends, and provide guidance on best practices. It also serves as a platform to submit useful project proposals.

The 'Welcome Pack' aims to understand customer expectations and present the ecosystem surrounding your operations. This approach allows us to initiate appropriate discussions, take relevant actions, and establish fruitful collaborations throughout the lifespan of your supercomputer.

How does it work?

The initial step of this Welcome Pack involves a discussion between the Bull collaboration coordinator and key stakeholders on your end: users, scientists, managers, and project leaders. This initial discussion typically occurs shortly after the machine's acceptance. If collaboration is part of the offer, it is then discussed and refined.

A possible agenda might include:

- Presentation of your activities, ecosystem, and projects.
- Presentation of Bull activities, ecosystems, and projects.
- Introduction to user groups (HPC, AI, Quantum, etc.).
- Discussions regarding collaboration, training, and service offerings.
- Open discussion to identify next steps.

Offerings (non-exhaustive)

- Technology surveillance, training, vision, and roadmaps.
- User groups like BUX (Bull User group for eXtreme computing).
- Tailored services that flexibly align with customer needs, with detailed taxonomy and methodology presented during the initial meeting.
- Proof of Concept (POC), code optimisation, and training sessions.
- Access to our project office: Bull actively participates in numerous funded projects, both national and European.
- Co-design opportunities: Collaborations with our R&D teams often result in co-designed solutions. An excellent illustration is the HDEEM project conducted with TU-Dresden, focusing on Energy Efficiency.
- Collaborative engagements: Use-case interactions frequently lead to structured collaborations. The ECMWF collaboration stands as a testament to this form of engagement.



CEPP trainings

Exploitation & innovation

Objective	This level aims to furnish participants with adequate information necessary for compiling and submitting code on a new cluster. The primary objective is to equip participants with essential knowledge, providing them with the necessary insights to comprehend HPC processes, from the initial compilation stages to the optimisation phase
Dedicated to	Beginners
Prerequisite knowledge	Linux, Compilation workflow

Webinar: What is HPC (1h)?

This webinar is offered free of charge.

During this session, we delve into HPC, commencing with an exploration of fundamental laws such as Moore's Law and Amdahl's Law. Our aim is to draw insights from the past to illuminate the path towards the future.

This webinar offers a detailed explanation of cluster construction from a hardware perspective. We systematically cover the hardware elements, starting from CPUs and progressing through networks and file systems, providing a comprehensive understanding of cluster building from its foundational components.

The subsequent segment focuses on software aspects tailored for end-users. Covering everything from code compilation to optimisation, this session elucidates crucial steps that developers and end-users need to grasp in order to effectively utilise a modern HPC cluster.

Training: How to start using your cluster?

This comprehensive training course caters to end users over a span of 3 days, accommodating up to 12 participants. The training targets two distinct categories of end users: beginners and individuals needing to adapt previous workflows to fit the new architecture. Our structured approach revolves around an end user's typical workflow, covering compilation, execution, debugging, and optimisation.

By the end of these three days, participants will gain proficiency in code compilation, scripting for job submissions, utilisation of debugging tools, and application optimisation techniques.

The initial phase of the training initiates with an overview of HPC, focusing on the cluster's architecture from CPU functionalities to network utilisation. This hardware-centric segment lays the groundwork for comprehending the subsequent software architecture. The first software section introduces tools for generating binaries, libraries (such as mathematical and communication libraries), and explores compiler usage with its performance implications.

In the subsequent software segment, we delve into the resource manager and batch scheduler. In alignment with system administrators, we elucidate cluster usage regulations encompassing queues, priorities, and related guidelines.

The final software segment centers on analysis. Participants will explore debugging applications and tools to enhance performance. This includes an introduction to profiling, covering basic bottleneck identification and progressively delving into complex concepts involving micro-architecture tuning.

Webinar pack

With this series of 6 webinars, each spanning 1 to 2 hours, we offer a comprehensive overview of your HPC environment:

Step 1	Cluster Architecture: In this webinar, we provide end users with an understanding of the architecture, from processors to the entire cluster infrastructure, encompassing I/O and communication systems.
Step 2	Compilation Workflow: This webinar exposes the intricacies of compiling parallel applications (including OpenMP, OpenACC, CUDA, MPI) tailored for your cluster. We recommend a set of libraries compatible with your hardware.
Step 3	Application Execution: Users will learn to utilise a resource and batch scheduler to efficiently run applications on the cluster. This session offers insights into submitting and optimising applications for optimal performance without necessitating code modification.
Step 4	Input/Output Optimisation: Focusing on I/O as a specific topic, this webinar explores understanding I/O bottlenecks and best practices.
Step 5	Debugging Tools: Given the increasing complexity of new architectures, this session introduces tools designed to aid end users in identifying bugs within applications featuring multi-level parallelism.
Step 6	Profiling Techniques - Understanding and identifying bottlenecks in applications is vital for performance enhancement. This webinar delves into profiling codes, presenting tools to characterise hot spots in both code and HPC architectures. Various tools for analysing I/O, computation, and communications will be covered, providing a comprehensive view of profiling tools.



One day training

Programming language

When new hardware emerges, it often introduces new programming languages, novel access methods for the hardware, and fresh approaches to compiling, debugging, and optimisation.

This training program is tailored for users who are newcomers to the new hardware (12 users per session) and will cover the foundational aspects of programming languages. It encompasses classical languages aligned with the hardware specifications (like CUDA for NVIDIA® hardware) as well as the utilisation of standards-compatible languages such as OpenMP5, OpenACC, and SYCL. The training includes practical labs to reinforce the concepts covered during the course.

It's important to note that this training initiates at a second-level understanding, assuming users possess some foundational knowledge of programming.



Usage of debugger

As HPC grows in complexity, scaling applications to match the cluster's size becomes essential. Debugging codes using traditional methods has become increasingly challenging, necessitating the use of specialised tools to assist users.

Our training session provides an introduction to methods aimed at alleviating the complexities associated with bug identification, outlining various types of bugs. We introduce parallel debugging techniques and hardware debugging specifically tailored for specialised hardware. Through hands-on labs, participants gain practical insight into utilising these methods effectively on the cluster.

We provide support for the Linearo DDT debugger throughout the training sessions to aid users in mastering efficient debugging practices.

Usage of Profiler

There are three primary facets when it comes to profiling HPC applications:

- Node-level profiling: This encompasses classical profiling techniques within the node.
- Task communication profiling: A more recent profiling method focused on analysing communication between tasks.
- Input/Output profiling: The last aspect involves profiling input/output operations.

Advance Training Sessions:

For new programming paradigms, we provide advanced training courses catering to end users seeking deeper insights into programming models. For instance, our advanced training

We offer specialised training for each of these elements, leveraging state-of-the-art tools from various partners, such as Intel Advisor or Linearo MAP. Our approach aligns with our debugging methodology, starting with an introduction to profiling and bottleneck identification, followed by an in-depth exploration of the tools. These sessions are complemented by hands-on labs conducted on your cluster.

on CUDA programming delves extensively into low-level configurations and the advanced utilisation of CUDA.

“ Our CEPP Training programs, tailored from welcome packs, webinars, to advanced boot camps, cater to diverse user needs, ensuring optimal alignment with project objectives. By closely collaborating with stakeholders and partners, we craft customised solutions adapted to specific customer requirement that accelerate HPC utilisation and drive transformative results. ”



Christophe Berthelot
CEPP Manager, Bull
christophe.berthelot@bull.com

Connect with us



bull.com

Bull is a registered trademark © Copyright 2026, Bull SAS - All rights reserved.