

EVIDEN

Eviden's Center for Excellence in Performance Programming (CEPP) – Accelerate Workload, Add Value to Simulation!

Leading the Way in HPC, AI and Quantum Efficiency:
CEPP's Solutions for Performance & Application Optimization

Authors

Bruno Lecoïnte

Group VP Business Support HPC, AI & Quantum, Eviden

Christophe Berthelot

CEPP Manager, Eviden

This white paper spotlights Eviden's Center for Excellence in Performance Programming (CEPP), which empowers organizations to take on industrial and scientific challenges by unlocking the potential of vast datasets.

The Market


Since the advent of digital simulation after World War II, significant advancements have occurred in the HPC arena, spanning both technological developments and application expansions. While quantum computing continues to be explored for its stability and industrial applications, Exascale computing has already become a reality worldwide. The emergence and widespread adoption of LLMs (Large Language Models) and machine learning have unlocked new avenues and possibilities for classical HPC. HPC applications are entering an era filled with both excitement and challenges.

The Challenges



Data Deluge: Meeting the Demands of Complex HPC Workloads

In the current industry landscape, data sets entering HPC systems are reaching unprecedented enormity and complexity. This surge in data volume and diversity poses significant challenges, leading to increasingly intricate configurations of HPC systems. Additionally, the arrival of AI (Artificial Intelligence) and generative AI technologies further complicates this landscape. Consequently, optimizing code and applications to harness these vast datasets effectively has become paramount. As organizations navigate this data deluge, they are compelled to adopt innovative strategies to ensure the efficiency and efficacy of their HPC infrastructure and applications.



Addressing System and Applications Efficiency Challenges Amidst Rapid Technology Evolution

The rapid evolution of hardware technologies, including multi-core CPUs, many-core processors, GPUs, and FPGAs, often outpaces the ability to optimize code accordingly. The heterogeneous systems built with these technologies present challenges in system management and application efficiency.

Furthermore, energy consumption has emerged as a critical concern, especially considering the evolving geopolitical landscape and the growing corporate focus on sustainability. Achieving significant reductions in energy consumption remains elusive without expertise in managing complex systems and optimizing code and algorithms.



Challenging Talent Pool with Skilled Data Scientists and Related Cost

With the rapid evolution of technology and applications, particularly the surge in AI adoption, specialized AI data scientists have become indispensable. This necessitates a refreshment of organizations' talent pools to keep pace with the complexity of HPC systems. Only by pooling together these specialized experts and leveraging best practices, can organizations reap the benefits of enhanced speed and cost-effectiveness for HPC and AI applications.

CEPP value proposition - Accelerate workload, give value to simulation!

The CEPP enables organizations globally to address industrial and scientific challenges by unlocking the potential of vast datasets. By integrating extensive expertise in HPC, data sciences, and industry insights, the CEPP team enhances simulations, optimizes cluster performance and operation, and minimizes time-to-result and cost-to-innovation.

Our adaptable service offerings cater to organizations' unique goals, providing customized solutions through training, workshops, webinars, or dedicated resources tailored to specific projects.

CEPP one+: A Joint Effort with Intel to Facilitate Seamless Code Evolution




We are living in an exciting era today. Ensuring the versatility of codes across multiple architectures, including emerging ones, is critical for organizations, making code migration essential.


As we move forward from legacy Intel® compilers to newer tools within the oneAPI framework, CEPP one+ stands ready to facilitate a smooth transition for our valued customers. Our dedicated team works closely with Intel, leveraging pre-release versions of Intel oneAPI to ensure that crucial fixes are implemented effectively. With a comprehensive range of customizable services, including training, workshops, webinars, and dedicated resources, CEPP tailors its support to meet the unique needs of each project. Using a task force approach, we assist clients with code migration and modernization, optimizing efficiency while minimizing costs and time investment.

CEPP Acceleration: Accelerate Time-to-result and Reduce Cost-to-innovation

Our approach focuses on assisting users in efficiently adopting applications in new hardware and software environments, such as GPU deployment, during the early lifecycle of the cluster. This involves facilitating a swift transition to the new cluster, which is crucial for an early commencement of production on the fresh system. Moreover, it often yields significant energy-efficiency improvements, hence, substantial Total Cost of Ownership (TCO) advantages, by optimizing the utilization of a significant corporate investment. The CEPP team includes specialists dedicated to ensuring a smooth and successful application transition.



CEPP Training: Foster Community, Enhance Knowledge, and Best Practice



Regardless of the type of activity or the production use of the supercomputers, fostering a community is crucial. This community can offer a deeper technological understanding of the supercomputer, enrich knowledge of current trends, and guide best practices. It also serves as a platform to submit useful project proposals. The CEPP training programs, tailored from welcome packs, and 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 customized solutions adapted to specific customer requirements that accelerate HPC utilization and drive transformative results.

CEPP AI4SIM: Accelerate and Improve Numerical Solvers with AI

AI is now an integral part of HPC and our lives. We are exploring with our clients in the simulation world whether the introduction of AI can help them save time and energy compared to conventional methods. Client conversations enable us to characterize their problem and we implement the most suitable solution.



Customer Cases – CEPP’s Impact on Real-world Applications

We have established the CEPP one+ offering in collaboration with a leading meteorological simulation provider in Europe. As part of this initiative, we migrated the production code from an older version of Parallel Studio to the latest version of oneAPI. This effort was conducted in collaboration with Intel compiler developers. Additionally, a migration of compute kernels initially implemented in OpenACC to OpenMP5 was also executed.

As part of MésoNET, we have implemented the CEPP Acceleration and Training offering at CALMIP. We collaborated with CALMIP to facilitate the onboarding process for the new Turpan machine by providing education sessions through the Training offering. To ensure optimal utilization, the CEPP Acceleration offering was selected, thereby enabling support for optimal machine usage through specific projects.

In a tight collaboration between ECMWF and CEPP, the Center of Excellence in Weather & Climate Modelling was created in 2020 to support ECMWF's roadmap with HPC, AI, and quantum capabilities. As part of this Center of Excellence, one of the projects is to develop Machine learning solutions for applications across the numerical weather prediction workflow that are customized

towards the needs of Earth system modeling. A second project is to help develop a CPU-GPU-based version of ECMWF's Integrated Forecasting System (IFS) and to prepare the ECMWF product-generation pipeline and data-centric workflows for new technologies. The Center is also supported by expert staff from AMD, NVIDIA, NVIDIA Mellanox, and DDN.

CEPP experts are actively involved in numerous European research projects. This enables them to remain at the forefront of technological developments and even to influence these technologies through co-design (EPI and EUPEX projects). Through these participations, CEPP collaborates with the best teams from academia and industry, in a wide range of application areas, such as earth system modeling, weather, and climate prediction (ESIWACE3 CoE - Center of Excellence), astrophysics (SPACE CoE), solid earth (seismology, volcanology, geodynamics in ChEESE-2P CoE), or materials simulations (MAX3 CoE). As a result, CEPP's experts can provide our customers with highly specialized assistance in optimizing and improving the efficiency of their applications on many different system architectures.

Conclusion

In the dynamic realm of high-performance computing (HPC) and new workload simulations, we witness an era of unprecedented growth, complexity, and opportunities. From the dawn of digital simulation to the present reality of Exascale computing and AI, the journey has been marked by remarkable technological strides and expanding horizons for HPC applications. In this ever-changing landscape, the CEPP emerges as a robust resource and ally in tackling complexity. Through strategic collaborations and tailored solutions, CEPP empowers organizations to navigate the complexities of HPC with confidence, accelerating time-to-insight and driving transformative outcomes.

Contact us



Christophe Berthelot
christophe.berthelot@eviden.com

Connect with us

in /in/eviden

✂ @EvidenLive

@evidenlive

▶ /EvidenLive

eviden.com

Eviden is a registered trademark © Copyright 2024, Eviden SAS – All rights reserved

ECT-240619-SJ-WP-Eviden's Center for Excellence in Performance Programming (CEPP) –Accelerate Workload, Add Value to Simulation!

