



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Computing in Science and Engineering Converged Computing Special Issue Introduction

V. Sochat, D. Milroy, C. Misale, J. Luetzgau, E.
Bollig, W. Magro

November 4, 2024

Institute of Electrical and Electronics Engineers Computing in
Science and Engineering

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

CiSE Converged Computing Introduction

Authors:

- Vanessa Sochat*, Lawrence Livermore National Laboratory
- Daniel Milroy, Lawrence Livermore National Laboratory
- Claudia Misale, IBM T.J. Watson Research Center
- Jakob Luettgau, Inria
- Evan F. Bollig, Amazon Web Services
- William Magro, Google

Final word count: 1,886

This special issue had an open call for proposals (CFP).

* Corresponding author: sochat1@llnl.gov

Abstract

Collaboration between Cloud and High Performance Computing (HPC) communities has accelerated in the last half decade. A common goal to run batch workloads combined with a desire for reproducibility, automation, and optimization has led to successful projects that range from container technologies to workload management and security. This span of current and future work defines a novel "Converged Computing" paradigm that aims to combine the best of both worlds, both from a technological and cultural standpoint. In this Special Issue, we review common themes in the space, showcasing current work and encouraging a continued effort toward innovative ideas that will enable the next generations of scientific discovery.

Introduction

Cloud and High Performance computing (HPC) trace their origins to different problems. HPC was created to solve the largest computational problems in science, engineering, and mathematics. The HPC community was built by domain scientists and engineers whose problems needed endlessly greater computational efficiency and scale. Cloud computing's beginning was slightly more practical and business-minded, having a purpose to deliver global access to infrastructure hosting under a sustainable economic model. Cloud emphasizes reliability, security, and breadth of services to meet infrastructure needs for any use-case. The two areas have been tremendously successful at advancing science and providing valuable services to billions of people. The latter has a more immediate and vastly larger economic impact, with revenues projected to exceed \$1.1T USD by 2027 and annual growth estimated at 20% [1]. Such

monetary cloud draws diverse talent and translates to rapid innovations, and when coupled with the demise and waning of tried-and-true scaling laws (Dennard and Moore, respectively), Cloud will soon be the leading sector of computing by revenue [2]. In the early 1990s, the scientific community understood the benefits of geographically distributed computing, which were realized in Grid Computing. These general concepts have been taken and augmented into what is now Cloud, backed by a business model to ensure its survival. Within the last decade, community demand pushed Cloud to innovate HPC capabilities layered atop service APIs, and backed by consistent (converged) hardware and software stacks found in HPC. As the HPC community increasingly moves to create more complex, heterogeneous workflows that require an ability to scale, there is a new requirement for orchestration and dynamism that traditional HPC does not well support. This heterogeneity results from more interdisciplinary interplay between workflows, where multiple scales are orchestrated to model more holistic behavior of a system [3]. An environment is needed that combines the performance and efficiency of HPC with the resiliency, portability, dynamism, and automation of the Cloud. While Cloud has been able to provide and innovate on many of these features, the environment is not without gaps. The HPC community can supplement the technological landscape with further expertise in scheduling, storage and filesystems, and communication protocols. The innovations of Cloud can also further trickle back into HPC setups. This technological environment, together with the representative communities, is called Converged Computing.

Converged Computing acknowledges that HPC and Cloud are neither disjoint nor orthogonal in mission. Cloud's shift to support HPC embeds it as part of the HPC community – i.e., HPC can take on features of Cloud and vice versa. Both serve as tools to progress science, engineering, and many other industries. With a common goal to efficiently run complex, batch workloads, much of the expertise of the HPC community continues to be relevant, creating opportunities for collaboration and knowledge sharing. The HPC community brings decades of experience optimizing application performance and increasing resource efficiency, which complement Cloud's strength in automation, modularity, portability, and pace of innovation. While the HPC and larger scientific communities may struggle to match Cloud's pace, the pillars of Converged Computing, namely integration, participation, and shared curiosity, offer a path forward to translate ideas between the spaces. The transformation of not just technology, but also policy to support the change, represent more nuanced, cultural underpinnings of the Converged Computing movement.

A part of this cultural change is related to incentives and behavior. For example, the service-oriented nature of Cloud can influence users to better consider resource utilization. The focus of HPC is the performance of the application and the hardware, while Cloud often emphasizes the performance to price ratio, which suggests achieving the best performance given a reasonable cost [4]. Having direct accountability for money spent compels the customer to think more critically about their Cloud hygiene, which can range from efficient application design to weighing the tradeoff between cost and time to results. The change in thinking would have a positive impact on HPC, where there has traditionally been little incentive to consider efficiency due to lack of transparency about costs.

Communities have also emerged that enable easy collaboration between practitioners of Cloud and HPC. The Cloud Native Computing Foundation (CNCF), a part of the Linux Foundation, is one such example [5]. The CNCF provides regular conferences and a streamlined process for collaborative work on projects innovating in the cloud-native space that also capture the interest of the HPC community. The HPC community similarly recognized the benefits to having such a foundation, and followed suit by establishing the High Performance Software Foundation (HPSF) recently [6]. It will be maximally advantageous for both communities to make an effort to participate in or attend the community events [7,8] from each side of the landscape, as convergence and collaboration can lead to changes that benefit both spaces. A final example comes from the security space, where NIST Special Publication 800-223, authored collaboratively by both sides, provides guidance on standardizing architectural security in HPC [9]. These examples suggest that the catalyst for any kind of cultural or behavioral change starts on the level of individuals that are working in the space, and often starts on the level of the technologies themselves.

This Special Issue includes five examples of work that illustrate the convergence of different domains and technologies such as workload managers and compute provisioning, artificial intelligence and machine learning, workflows, and security.

Workload managers paired with clusters of compute and storage have led innovation in the space, primarily driven by the need to run batch workloads. In Cloud, Kubernetes [10], has become the de-facto Cloud container orchestration framework. The HPC community, having developed workload managers for decades, has recognized the opportunity for improving the batch submission process and access to resources. The **Jetstream2** project presents the idea of an on-premises research Cloud, considering the trade-offs and needed balance to afford performance, and consider resource availability and utilization. The model of dynamic, on-demand creation of user-scoped compute resources and software offers an alternative to permanent, slower moving multi-tenancy software stacks. Such a dynamic model is being tested with projects like the versatile software-defined cluster (**vCluster**), a framework that packages HPC resources in a more service-oriented architecture. These setups are supported by application programming interfaces (APIs) to both orchestrate infrastructure, workflows, and access to resources, whether that is within a single research Cloud, or across Clouds, as is the case with the **GA4GH Task Execution API**. As contention for highly demanded resources like GPU increases, approaches to effectively and economically use resources that are available will be paramount [11].

Convergence is important not just to embody features of the Cloud that are highly desired, but to maintain strong needs for scientific computing that are often lost in the public sector. Such an example is performance and consistency. With Cloud leading the innovation space to develop hardware oriented toward AI, associated tooling and even chip-level floating point precision is being targeted to best address the needs of model training, which often means more bang for the buck – achieving comparable training outcomes at the lowest cost. The drive toward AI/ML workloads also affords a design for compute to maximally pack processing power on fewer

nodes, favoring vertical scaling with fewer nodes offering larger and more powerful GPU. A push toward using surrogate models is also emerging as a complement to traditional approaches to reduce simulation time of multi-physics models. While these approaches are promising, it is unknown how long it will take to validate the stability and accuracy of these new models. It also remains problematic if the HPC community cannot afford the GPU resources [12]. Converged work in the space of AI/ML offers early ideas here, such as **XaaS**, a shared execution platform that offers computing resources and accelerators as a service, tackling the resource availability problem by focusing on maximizing utilization of resources available, and also providing a focus on communication paradigms that are required for HPC workloads today. Of a similar thread, the Advanced Privacy-Preserving Federated Learning (**APPFL**) framework has a focus on security, offering streamlined federated learning for the HPC community.

The work here reflects only a subset of a larger expanse of exciting and innovative ideas in Converged Computing that extend to workflows, workload managers and schedulers, container technologies, kernel monitoring, operating systems, and package managers. For example, supporting elasticity, portability, automation, and efficient resource usage requires convergence of resource management and scheduling. HPC has focused development on providing powerful mapping and binding capabilities, allowing applications to execute on exactly the hardware desired under the job topology required for best performance. The fine-grained mapping and binding imposes rigidity on the execution environment that extends to the resource specifications themselves. HPC uses “imperative” resource specification that provides performance through specificity but can burden users with defining complex job shapes and fixed requirements. Shifting some of the burden of task definition from the user to the system and enabling resource mapping flexibility are capabilities offered by convergence with Cloud. Cloud technologies feature “declarative” resource management and specification, which reduces mapping complexity, automates management, and improves portability at the expense of outright performance. Eliminating the trade-off between performance and dynamism and portability by supporting both imperative and declarative paradigms is crucial to support scientific workflows.

In the domain of workflow tools, the Kubeflow (<https://kubeflow.org>) community is taking proactive action to unify interfaces for machine learning training, using common abstractions and accounting for the needs of the HPC community. Such an effort will provide an automated, declarative means to run scaled AI/ML in Cloud, and potentially HPC applications given user-space Kubernetes environments. The unit of work for those workflows, the container technology, has also taken interesting directions to address needs from both communities – user space or rootless approaches for HPC, and better strategies for caching, registry pulling, and distribution for Cloud. Alongside these workflows, monitoring and a general understanding of what is happening on the level of the system is increasingly important. The eXtended Berkeley Packet Filter (eBPF) [13], an interface that allows for running sandboxed programs in the kernel, provides this level of visibility. eBPF has become a leader in Cloud for a wide range of tools that range from networking, load-balancing, security, performance analysis, and observability use cases, and the HPC community has yet to fully realize its benefits.

There has never been a more exciting time to be a developer in this nascent space, and there has never been a better time to take calculated risks, think creatively, and imagine a future for High Performance Computing that looks very different from what exists today. Exploring this space and working to integrate new and transformative, but status-challenging ideas and techniques is both the responsibility and reward of those inspired to contribute to Converged Computing.

References

- [1] "Worldwide HPC in the Cloud Forecast, 2022-2027." 2023. Hyperion Research. Accessed May 29, 2024. <https://hyperionresearch.com/product/worldwide-hpc-in-the-cloud-forecast-2022-2027/>.
- [2] Reed, Daniel, Dennis Gannon, Jack Dongarra, R. Colin Johnson, Marc Rotenberg, and Robin K. Hill. 2023. "HPC Forecast – Communications of the ACM." Accessed October 21, 2024. <https://cacm.acm.org/research/hpc-forecast/>.
- [3] Vetter, Jeffrey S., et al. *Extreme heterogeneity 2018-productive computational science in the era of extreme heterogeneity: Report for DOE ASCR workshop on extreme heterogeneity*. USDOE Office of Science (SC), Washington, DC (United States), 2018.
- [4] Kutzner, Carsten, et al. *GROMACS in the Cloud: A Global Supercomputer to Speed Up Alchemical Drug Design*. Journal of Chemical Information and Modeling. 2022. 62 (7), 1691-1711. <https://pubs.acs.org/doi/10.1021/acs.jcim.2c00044>
- [5] "Cloud Native Computing Foundation." 2015. CNCF. Accessed October 23, 2024. <https://www.cncf.io/>.
- [6] "HPSF." 2024. High Performance Software Foundation. Accessed October 23, 2024. <https://hpsf.io/>.
- [7] "SC Conference Series." 2020. Supercomputing. SC Conference Series. August 13, 2020. <https://supercomputing.org/>.
- [8] "KubeCon + CloudNativeCon North America." 2024. LF Events. Accessed October 23, 2024. <https://events.linuxfoundation.org/kubcon-cloudnativecon-north-america/>.
- [9] Guo, Yang. 2024. "High-Performance Computing Security: Architecture, Threat Analysis, and Security Posture." NIST Special Publication (SP) 800-223. Gaithersburg, MD: National Institute of Standards and Technology. <https://doi.org/10.6028/nist.sp.800-223>.
- [10] Honeypot. 2022. "Kubernetes: The Documentary [PART 1]." Youtube. January 21, 2022. <https://www.youtube.com/watch?v=BE77h7dmoQU>.
- [11] Strati, Foteini, Paul Elvinger, Tolga Kerimoglu, and Ana Klimovic. 2024. "ML Training with Cloud GPU Shortages: Is Cross-Region the Answer?" In *Proceedings of the 4th Workshop on Machine Learning and Systems*. New York, NY, USA: ACM. <https://doi.org/10.1145/3642970.3655843>.
- [12] Lange, Jack, Thomas Papatheodore, Thomas Papatheodore, Todd Thomas, Chad Effler, Aaron Haun, Carlos Cunningham, et al. 2023. "Evaluating the Cloud for Capability Class Leadership Workloads." ORNL/TM-2023/3083. Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States). Oak Ridge Leadership Computing Facility (OLCF). <https://doi.org/10.2172/2000306>.

[13] Soldani, David, Petrit Nahi, Hami Bour, Saber Jafarizadeh, Mohammed F. Soliman, Leonardo Di Giovanna, Francesco Monaco, Giuseppe Ognibene, and Fulvio Riso. 2023. "EBPF: A New Approach to Cloud-Native Observability, Networking and Security for Current (5G) and Future Mobile Networks (6G and Beyond)." *IEEE Access: Practical Innovations, Open Solutions* 11: 57174–202. <https://doi.org/10.1109/access.2023.3281480>.

Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-JRNL-871107

Biographies

Vanessa Sochat is a Computer Scientist in Livermore Computing. Her work focuses on container technologies and High Performance Computing (HPC) and Cloud convergence. She maintains the Flux Operator and Fluence custom scheduler as part of the Flux Framework, and collaborates closely with industry and several open source communities, leading work and initiative to innovate in the space. She received her Ph.D and M.S. in Biomedical Informatics from Stanford University, and her B.A. from Duke University.

Daniel Milroy is a Computer Scientist at the Center for Applied Scientific Computing. His research focuses on graph-based scheduling and dynamic resource management for High Performance Computing (HPC) and Cloud converged environments, and he collaborates closely with industry and academia. Dr. Milroy is a developer of the Fluxion directed graph-based scheduler in the Flux Framework and leads research on flexible resource representation and scheduling. He holds a Ph.D. and M.S. in Computer Science from the University of Colorado Boulder, and an A.B. in Physics from the University of Chicago.

Claudia Misale is a Staff Research Scientist in the Hybrid Cloud Infrastructure Software group at IBM T.J. Watson Research Center (NY). Her research is focused on Kubernetes for IBM Public Cloud and bare metal, and targets monitoring, observability and scheduling for HPC and AI workloads. She is mainly interested in cloud computing and container technologies, and her background is on high-level parallel programming models and patterns, and big data analytics on HPC platforms.

Jakob Luetzgau is a researcher at the French Institute for Research in Computer Science and Automation (Inria). His research interests include HPC/Cloud convergence, high-performance storage, holistic performance engineering for scientific workflows, machine learning applications in instrumentation, and optimization of distributed systems. He earned his PhD in Computer Science in 2021 from the University of Hamburg (Germany) after graduating with his BSc and MSc from the same University in 2014 and 2016.

Evan F. Bollig is a principal specialist solutions architect for HPC with Amazon Web Services (AWS) and a senior member of the Institute of Electrical and Electronics Engineers (IEEE). His experience in HPC spans all roles and all workloads, from performance tuning multi-GPU physics codes for large-scale systems, to developing secure cloud-native infrastructures for production clinical genomics pipelines. Evan champions for greater good in the AWS community through open science and open source (GitHub: bollig, ORCID: 0000-0001-9254-4598).

William Magro is General Manager of High Performance Computing at Google, where he drives HPC strategy and customer success for Google Cloud. He joined Google in 2020, after 20 years at Intel, where he was Intel Fellow and Chief Technologist for HPC. A recognized leader in the InfiniBand industry, Magro helped found the OpenFabrics Alliance and served as InfiniBand Trade Association Technical Working Group co-chair from 2007-2020. He has been a prominent voice in the HPC community for over two decades and regularly participates and presents in HPC conferences, advisory boards, and panels. He holds a bachelor's degree in applied and engineering physics from Cornell University and a master's degree and Ph.D. in physics from the University of Illinois at Urbana-Champaign.