

Quick Take

Open Source and Standards Thriving within the HPC Community

Mark Nossokoff and Bob Sorensen
January 2022

HYPERION RESEARCH OPINION

Open source and standards have long been a hallmark of the HPC community. Each encourages broad industry participation, wide-ranging collaboration, and a thriving ecosystem with a common goal of robust, interoperable solutions. Key to the success of standards and open source is a stable, neutral governance and support structure.

Two new unrelated industry efforts (Apptainer evolving under the auspices of the Open Linux Foundation and DMTCP establishing a public/private open-source partnership) have been embraced by the HPC community to strengthen the standards and open source activities of the HPC ecosystem.

CURRENT SITUATION

Openness and collaboration are seen by many as key elements for the long-term success for any project that spans a diverse and dispersed population of users and suppliers. Users want to know that the product they're using is stable, reliable, has a long life with continued investment in new features and support, and is available from more than one source. Suppliers want to provide their products to the widest audience as possible with minimal variations and customizations to optimize their investments. Open source software and standardization of interfaces support and promote both the product requirement needs of users and the business interests of the vendors.

Independently and unrelated, two new open source and standardization efforts have recently emerged. One (Apptainer) addresses the growth in adoption of a new technology (containers) and the other (DMTCP) extends the implementation of a stalwart HPC method for resilient and fault-tolerant system design (checkpointing).

Apptainer

Containers are a ready-to-run software package bundled with everything needed to run an application quickly and reliably from one computing environment to another. In this case, containers may include code, runtime, system libraries, default settings, and other dependencies. Containers aim to make running HPC workloads easier to run with consistent and reliable performance across diverse, heterogeneous operating infrastructures, such as between on-premises, hybrid-cloud, and multi-cloud environments.

Formerly known as Singularity, Apptainer is a widely used container system for HPC workloads with a focus on bare-metal performance while being secure, portable, and reproducible. Being under the governance of the Open Linux Foundation (also home to the Cloud Native Computing Foundation,

Open Container Initiative, and OpenHPC) promotes the long-term health of the project and a sustainable roadmap of new, robust features. Supporters of the project and new governance structure includes both government organizations like Berkeley Lab, Fermilab, and Sandia National Lab, as well as corporate entities including AMD and Intel.

DMTCP

Checkpointing is a commonly used method with HPC systems that allows the recovery from system failures during complex simulation runs with long execution times to occur seamlessly and orderly. It is a well-understood technique for saving the states of independent node memory during a failure mode and restoring that state when the machine is back up and running.

As HPC system architectures evolve to support data-intensive computing and complex systems with pooled, distributed memory across multiple nodes operating on extremely large datasets, traditional checkpointing implementations that were designed to support more monolithic systems are not up to the task. Distributed MultiThreaded Checkpointing (DMTCP) technology is an emerging standard to implement the functionality of traditional checkpointing within the new constructs of data-intensive, big memory architectures.

The DMTCP project currently consists of Northeastern University, MemVerge, and NERSC, representing the academic, industry, and government sectors, respectively. Northeastern University has been working on the underlying technology since 2005. NERSC has been contributing testing efforts to optimize MPI-Agnostic Network-Agnostic (MANA), a plugin on top of DMTCP that has been used for transparent checkpointing of MPI on the Cori and Perlmutter supercomputers. MemVerge is providing the commercial stewardship and support of the project, as well as directly investing in and providing resources to support the effort.

FUTURE OUTLOOK

While some organizations take the approach that a closed, vertically-integrated supply chain of solution elements provides a more reliable and innovative long-term environment for robust products and roadmaps, they are in the minority as few have the resources necessary to sustain the required investments. Open source and standardization should continue to be key tools in promoting a healthy architecture and solutions environment for the HPC community. Supporting a "rising tide lifts all boats" approach to the market, they allow users to trust that solutions will be reliable, have a long life, and are available from multiple sources. Vendors also benefit from being able to support a broader range of customers while optimizing investments in as few product variations as possible.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue

St. Paul, MN 55102

USA

612.812.5798

www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2022 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.