

ISC 2020 SuperCompCloud Workshop Proposal

Title

SuperCompCloud: Workshop on Interoperability of Supercomputing and Cloud Technologies

Abstract

Exascale computing initiatives are expected to enable breakthroughs for multiple scientific disciplines. Increasingly these systems may utilize cloud technologies, enabling complex and distributed workflows that can improve not only scientific productivity, but accessibility of resources to a wide range of communities. Such an integrated and seamlessly orchestrated system for supercomputing and cloud technologies is indispensable for experimental facilities that have been experiencing unprecedented data growth rates. While a subset of high performance computing (HPC) services have been available within a public cloud environments, petascale and beyond data and computing capabilities are largely provisioned within HPC data centres using traditional, bare-metal provisioning services to ensure performance, scaling and cost efficiencies. At the same time, on-demand and interactive provisioning of services that are commonplace in cloud environments, remain elusive for leading supercomputing ecosystems. This workshop aims at bringing together a group of experts and practitioners from academia, national laboratories, and industry to discuss technologies, use cases and best practices in order to set a vision and direction for leveraging high performance, extreme-scale computing and on-demand cloud ecosystems. Topics of interest include tools and technologies enabling scientists for adopting scientific applications to cloud interfaces, interoperability of HPC and cloud resource management and scheduling systems, cloud and HPC storage convergence to allow a high degree of flexibility for users and community platform developers, continuous integration/deployment approaches, reproducibility of scientific workflows in distributed environment, and best practices for enabling X-as-a-Service model at scale while maintaining a range of security constraints.

Keywords

Supercomputing; Cloud computing; Virtualisation; Extreme Data workflows; Data analytics; Parallel and distributed storage systems; Software defined infrastructure; Exascale Computing Project¹; Exascale Centers of Excellence²

Workshop Scopes

This workshop will cover topics related to interoperability of supercomputing and cloud computing, networking and storage technologies that are being leveraged by use cases and research infrastructure providers with a goal to improve productivity and reproducibility of extreme-scale scientific workflows.

Topics included within the scope:

¹ <https://exascaleproject.org/>

² <https://ec.europa.eu/digital-single-market/en/news/ten-new-centres-excellence-hpc-applications>

- Virtualization for HPC e.g. virtual machines, containers, etc.
- Storage systems for HPC and cloud technologies
- On-demand and interactivity with performance, scaling and cost efficiencies
- Resource management and scheduling systems for HPC and cloud technologies
- Software defined infrastructure for high-end computing, storage and networking
- Application environment, integration and deployment technologies
- Secure, high-speed networking for integrated HPC and cloud ecosystems
- Extreme data and compute workflows and use cases
- Research infrastructure deployment use cases
- Resiliency and reproducibility of complex and distributed workflows
- Isolation and security within shared HPC environments
- X-as-a-Service technologies with performance and scalability
- Workflow orchestration using public cloud and HPC data centre resources
- Authentication, authorization and accounting interoperability for HPC and cloud ecosystems
- Workforce development for integrated HPC and cloud environments

Papers Submission Guidelines:

Workshops will accept unpublished material on topics within the scope. A call for papers is planned to be released immediately after receiving the workshop notification. All submitted manuscripts will be peer reviewed by the technical program committee for relevance, technical contributions, originality, impact and correctness. The review process will not be double-blind. We will use ISC workshop paper publishing service that is offered by ISC 2020 workshop, Springer's Lecture Notes in Computer Science (LNCS) series. The call for papers will be publicized through different outreach channels and mailing lists. It will highlight that at least one author of an accepted paper is expected to register and attend the ISC 2020 workshop program.

Relevance and impact of the workshop for ISC

This workshop builds on a successful SC19 workshop³ "SuperCompCloud: Workshop on Interoperability of Supercomputing and Cloud Technologies"⁴ and SC18 panel⁵ titled "HPC in Cloud or Cloud in HPC: Myths, Misconceptions and Misinformation". Both events had over 200 participants. At the same time, it is noted that the SC and ISC exhibit programs have experienced a steady growth in participation of exhibitors that identify themselves with keyword "cloud computing" over the last 3 years. In fact, "supercomputing" and "cloud computing" are among the most widely used keywords that exhibitors have listed during SC19, SC18 and SC17 conferences. Hence, there is great deal of interest among the Supercomputing technical program such as ISC as well as exhibitors in topics that are covered by this workshop. For the proposed workshop, we would like to include a diverse group of speakers have been involved in a number of large multidisciplinary, international research

³ https://sc19.supercomputing.org/proceedings/workshops/workshop_pages/wksp115.html

⁴ <https://sites.google.com/view/supercompcloud>

⁵ <https://sc18.supercomputing.org/panel-focus-hpc-in-cloud-or-cloud-in-hpc-myths-misconceptions-and-misinformation/>

and development projects as well as consortiums and user groups related to operational aspects of national and institutional data centers as well as public cloud providers.

Format for the workshop

In order to ensure a narrative around the European efforts related to Open Science and Exascale programs, the workshop will be in part mini-symposium i.e. invited contributions and in part peer-reviewed technical program with papers. Optionally, there could be a panel with speakers. Short lightening talks may be included to improve engagement of the participants.

Tentative program for a full-day workshop with sessions and speakers:

Morning session:

- Keynote (45 minute): EuroHPC⁶ program
- Invited Talk (25 min): Exascale Centre of Excellence
- Submitted Paper presentation (25 min): peer reviewed material
- Submitted Paper presentation (25 min): peer reviewed material
- Panel with speakers (optional)

Afternoon session:

- Keynote (45 minute): European Open Science Cloud (EOSC)⁷ initiative
- Submitted Paper presentation (25 min): peer reviewed material
- Submitted Paper presentation (25 min): peer reviewed material
- Submitted Paper presentation (25 min): peer reviewed material
- Lightening talks (optional-based on extended abstracts)

Expected outcome from the workshop

This full day workshop is expected to develop a community forum and initiative for bringing together researchers, practitioners and service providers to share experiences and expertise. We hope to continue in form of a full day workshop with an open call for contributions for subsequent SCxy and ISCxy conferences. Another option would be to reach out to ongoing initiatives and standards committees such as OpenHPC and OpenStack scientific working group. Presentation material will be available for SC archive.

How the organizers plan to advertise and attract attendees

In addition to ISC and HPC communication channels, the workshop information will be distributed through online, email and social media channels of different projects. Workshop organizers have participated in previous ISCxy and SCxy at several occasions. User groups and activities include Cray and IBM user groups, national infrastructure activities such as XSEDE and PRACE and various academic and research organizations. Workshop announcement will be sent to the following mailing lists: hpc-announce, siam-sc, tcpc-announce, tcsc-announce,

⁶ <https://eurohpc-iu.europa.eu/>

⁷ <https://ec.europa.eu/research/openscience>

swiss academic channels, swiss hpc-ch, OpenStack scientific working group, etc. Workshop organizers and invited speakers, who often give talks on this topic to other events and funding agencies meetings, will communicate workshop announcements at these meetups.

Planned Schedule for Call for Papers and Review Process

- Call for papers published: December 20, 2019
- Submission deadline: March 15, 2020
- Notifications: April 2, 2020
- Program building: April 20, 2020

Deadlines can be adjusted to ensure ISC2020 timelines for the final list of speakers and early-bird registration deadline.

Expected Number of Papers

We expect around 10 papers because this is the first workshop with call for papers.

Expected Number of Participants

SC18 panel and SC19 workshop had 200+ participants (people were turned away due to lack of space). We anticipate 80-120 participants for the first ISC workshop.

Tentative Technical Program Committee

- Deborah Bard, NERSC, USA
- Eli Dart, Energy Science Network (ESnet), USA
- Giuseppe Fiameni, NVIDIA
- Maria Girone, CERN Open Lab, Switzerland
- Stathis Papaefstathiou, Cray-HPE, USA
- Tiago Quintino, ECMWF, UK
- Stig Telfer, StackHPC, UK
- Georgia Tourassi, ORNL, USA
- Allan Williams, NCI, Australia

SuperCompCloud Steering Committee

- Victoria Stodden, University of Illinois, USA
- Dirk Pleiter, JUELICH Supercomputing Centre , Germany
- Steve Poole, Los Alamos National Laboratory, USA
- Roy Campbell, DOD, USA
- Kevin Kissell, Google Cloud Platform

For questions, comments, or suggestions contact the workshop organizers: Sadaf Alam and David Hancock at supercompcloud@googlegroups.com