

HPCW 2020 - Proposal

Title: 6th Annual High Performance Containers Workshop
Short Title HPCW20
Workshop Website <http://qnib.org/isc/>

Abstract

Linux Containers have become an industry standard for sharing and distributing software. This workshop will create a space of interaction between field experts, end-users, and newcomers to discuss container technologies and their current and future challenges.

This workshop presents the current state of Linux Containers in HPC/AI, what challenges and research opportunities currently exist in the adoption in HPC/Data Analytics and how containers can foster improvements when applied to the field of HPC, Big Data and AI in the mid and long term.

The ecosystem is dissected into different layers (see below) within the workshop and presented as a curated set of lightning talks. This provides a holistic and state-of-the-art overview so that participants can make informed discussions on how to start, improve, or continue the adoption containers in HPC.

Topic Areas

Architectures & Networking

Programming Models & Systems Software

AI/ML

HPC Applications

Emerging Technologies

Keywords

Clouds and Distributed Computing, Containers in HPC, AI, ML, DL, MPI, Reproducibility, Resiliency

Target Audience

System Operators and Administrators, System Engineers/Architects, Software Developers, Data Engineers/Scientists

Keywords

HPC, cloud, Kubernetes, virtualization, container runtimes, software stack, MPI, RDMA, monitoring, orchestration, AI, big data, accelerators

Format

This workshop will follow the concepts of the last [5 successful workshops at ISC](#) as well as the [CANOPIE workshop](#) at SC19 to provide a holistic overview of the current state-of-the-art in High Performance Containers (HPCCon).

The workshop will be segmented into five parts, to provide guidance for the attendees and limit the scope of each presentation. Each segment will start with a short introduction to define the session's scope. Next, domain experts and power-users will share their experience in a condensed, crisp lightning talk. Following each set of talks, a panel will be formed to field questions from the audience and discuss what was presented. This unique session organization helps coalesce and organize key aspects of container technologies into a more consumable format, enabling the latest research topics and community trends to be discussed. The segmentation with defined scope will help newcomers to ramp up quickly and follow the workshop discussions.

The segments are:

- **RUNTIME**: Discussion about different runtimes within the container ecosystem and how they differ
- **BUILD & TEST**: How are container images structured, best-practices and how to test the artifact
- **DISTRIBUTION**: Once images are built, how are those images distributed efficiently
- **SECURITY**: Different aspects of security within the container lifecycle and at scale
- **ORCHESTRATION**: While running one container is easy, orchestrating a cluster of containers is complicated. This segment will dive into different complexity levels and implementations of orchestration.
- **HPC SPECIFICS**: What makes HPC workloads different and how might they converge with non-HPC workloads
- **USE-CASES**: This segment leads to an open discussion about what use-cases, experiences and limitations are.
- **FUTURE DIRECTIONS**: Looking out into the future what are the challenges ahead and what new technologies are emerging at the edge of the container landscape?

Agenda

Segment	Duration	Title
INTRO	15min	Workshop Introduction
		Speaker Introductions
RUNTIME	1h	Introduction and Scope
		Container Runtimes Overview
		The state of container runtimes
		Panel Discussion
BUILD & TEST	1h	Introduction and Scope
		Building Container Images (via runtime)
		Build Tools (SPACK/EasyBuild)
		Verify and Test Container Artifacts
		Panel Discussion
BREAK	11:00	Coffee Break
DISTRIBUTION	30min	Introduction and Scope
		Container Image Formats (OCI/SIF)
		Distribution mechanisms and tools
		Panel Discussion
SECURITY	30min	Introduction and Scope
		Security at Scale
		Reproducibility and signing
		Runtime Security
		Panel Discussion
ORCHESTRATION / SCHEDULING #1	1h	Introduction and Scope
		Orchestration Overview
		Workflow Orchestration
		Kubernetes Ecosystem
BREAK	13:00	Lunch Break
ORCHESTRATE /	30min	Other Schedulers and Intersections

SCHEDULE #2		Panel Discussion
HPC SPECIFIC	1h	Introduction and Scope
		Storage Solutions and Trends
		HPC Device Integration
		HPC-Containers at Scale
		MPI / PMI Update
		Panel Discussion
BREAK	16:00	Coffee Break
USE-CASES	45min	6y HPCW Containerization ReCap
		NERCS/CSCS Looking Back
		Mellanox/NVIDIA Journey
		Panel Discussion
FUTURE DIRECTIONS	45min	New compute models (Lambda, Inferencing)
		Does K8s take over?
		ECP Supercontainer Project
		Panel Discussion
END	18:00	Workshop Ends