

# Application

## **Title of the workshop**

15th Workshop on Virtualization in High-Performance Cloud Computing (VHPC '20)

## **Names, affiliations, and e-mail addresses of the workshop organizers**

(a typical workshop should count no more than three co-chairs affiliated with different organizations)

Michael Alexander (chair), University of Vienna  
<mfalexand@gmail.com>

Anastassios Nanos (co-chair), SunLight, UK  
<ananos@sunlight.io>

## **Abstract/Scope**

Containers and virtualization technologies constitute key enabling factors for flexible resource management in modern data centers, and particularly in cloud environments. Cloud providers need to manage complex infrastructures in a seamless fashion to support the highly dynamic and heterogeneous workloads and hosted applications customers deploy. Similarly, HPC environments have been increasingly adopting techniques that enable flexible management of vast computing and networking resources, close to marginal provisioning cost, which is unprecedented in the history of scientific and commercial computing. Most recently, Function as a Service (Faas) and Serverless computing, utilizing lightweight VMs-containers widens the spectrum of applications that can be deployed in a cloud environment, especially in an HPC context. Here, HPC-provided services can become accessible to distributed workloads outside of large cluster environments.

Various virtualization-containerization technologies contribute to the overall picture in different ways: machine virtualization, with its capability to enable consolidation of multiple underutilized servers with heterogeneous software and operating systems (OSes), and its capability to live-migrate a fully operating virtual machine (VM) with a very short downtime, enables novel and dynamic ways to manage physical servers; OS-level virtualization (i.e., containerization), with its capability to isolate multiple user-space environments and to allow for their coexistence within the same OS kernel, promises to provide many of the advantages of machine virtualization with high levels of responsiveness and performance;

lastly, unikernels provide for many virtualization benefits with a minimized OS/library surface. I/O Virtualization in turn allows physical network interfaces to take traffic from multiple VMs or containers; network virtualization, with its capability to create logical network overlays that are independent of the underlying physical topology is furthermore enabling virtualization of HPC infrastructures.

### **Tentative Program Committee**

Stergios Anastasiadis, University of Ioannina, Greece  
Jakob Blomer, CERN, Europe  
Ron Brightwell, Sandia National Laboratories, USA  
Eduardo César, Universidad Autónoma de Barcelona, Spain  
Julian Chesterfield, OnApp, UK  
Stephen Crago, USC ISI, USA  
Christoffer Dall, Columbia University, USA  
Patrick Dreher, MIT, USA  
Robert Futrick, Cycle Computing, USA  
Balazs Gerofi, RIKEN Advanced Institute for Computational Science, Japan  
Maria Girone, CERN, Europe  
Kyle Hale, Northwestern University, USA  
Brian Kocoloski, University of Pittsburgh, USA  
Nectarios Koziris, National Technical University of Athens, Greece  
John Lange, University of Pittsburgh, USA  
Che-Rung Lee, National Tsing Hua University, Taiwan  
Giuseppe Lettieri, University of Pisa, Italy  
Qing Liu, Oak Ridge National Laboratory, USA  
Nikos Parlavantzas, IRISA, France  
Kevin Pedretti, Sandia National Laboratories, USA  
Amer Qouneh, University of Florida, USA  
Carlos Reaño, Technical University of Valencia, Spain  
Thomas Ryd, CFEngine, Norway  
Na Zhang, VMWare, USA  
Borja Sotomayor, University of Chicago, USA  
Craig Stewart, Indiana University, USA  
Anata Tiwari, San Diego Supercomputer Center, USA  
Kurt Tutschku, Blekinge Institute of Technology, Sweden  
Yasuhiro Watashiba, Osaka University, Japan  
Nicholas Wright, Lawrence Berkeley National Laboratory, USA  
Chao-Tung Yang, Tunghai University, Taiwan

**Expected number of papers:** 30 submissions, 8 presented papers

**Proceedings:** Springer Workshop Volume

**Targeted audience**

The Workshop on Virtualization in High-Performance Cloud Computing (VHPC) aims to bring together researchers and industrial practitioners facing the challenges posed by virtualization, containerization, and orchestration techniques for deploying HPC and HPC-like applications in large data centers.

**Expected outcome from the workshop**

VHPC aims to foster discussion, collaboration, mutual exchange of knowledge and experience, enabling research to ultimately provide novel solutions for virtualized computing systems of tomorrow.

Diffusion of new cloud and HPC virtualization results to researchers and practitioners.

**Estimated attendance:** > 30

**One or more keywords**

container, virtualization, cloud, hpc, orchestration

**If the workshop has been held before, please indicate the dates, locations, and approximate number of attendees**

Prior History:

VHPC 7 had a 33% accept rate and drew 28 attendees peak in 2012 (EuroPar '12).

VHPC 8 had 45 attendees peak, min  $\geq 30$  (SC '13)

VHPC 9 had a 44% acceptance rate with 20 attendees peak (min  $\geq 18$ ). (EuroPar '14)

VHPC 10 had a 40% acceptance rate, 18 attendees peak. (EuroPar '15)

VHPC 11 had a 40% acceptance rate, 30 attendees peak. (ISC '16)

VHPC 12 had a 40% acceptance rate, 29 attendees peak. (ISC '17)

VHPC 13 had a 29% acceptance rate, 27 attendees peak. (ISC '18)

VHPC 14 had a 44% acceptance rate, 26 attendees peak. (ISC '19)

**Format of the workshop**

Duration: Workshop Duration is one day

Format: Talks, some demonstrations and lightning talks.

**Tentative list of talks/invited talk(s)/panel session(s)/keynote talk(s)/combinations of these**

Not set at this time; prior invited speakers included an:

NASA IT CTO, CERN IT group leader, ARM director

**Short CV for each workshop organizer** (1,000 characters maximum)

Michael Alexander holds degrees in electrical engineering (TGM), business administration (University of Southern California) and economics (University of Vienna). He is currently performing large data analytics for multiple domains and supporting architecting of clusters. His professional experience includes education and product management at IBM, and Alcatel. Prior he was a HPC Specialist at TU Wien and a Product Line Manager for Alcatel ADSL and Optical Access Networks. He is the author of a textbook on networks and network security published by Hüthig/Verlagsgruppe Süddeutsche, and editor of a special issue on mathematical methods in network management of the Wiley International Journal of Network Management. For the last fourteen years, he has served as the Program Committee Chair for VHPC, Workshop on Virtualization in High-Performance Cloud Computing. His current research interests include content centric networking, distributed databases, virtualization system-network management.

With over 12 years of experience in Virtualization technologies, Anastassios Nanos is currently working on the lower-level parts of the stack to attack issues related to performance, scalability, power-efficiency, and security in hypervisors. Previously, he was a post-doc at CSLab, NTUA, working on bridging the gap between common HPC practices and virtualization. His research interests include I/O Virtualization, systems software for high-performance I/O in virtualized environments, systems support for heterogeneous platforms, communication architectures for clusters, and scalable storage architectures based on clusters. He holds a Diploma in Engineering (2006) from ECE, NTUA and a PhD in Computer Engineering (2013) from NTUA. He has been involved in EU-funded projects, conducting research in emerging, power-efficient micro-server architectures on scalable network and storage I/O, and energy-driven resource management in cloud architectures.

**If known: list of all workshop speakers**, including first name, last name, e-mail address and short CV (1,000 characters maximum) for each speaker: expected in March

## Call for Papers

=====

CALL FOR PAPERS

### **145h Workshop on Virtualization in High-Performance Cloud Computing (VHPC '20)**

held in conjunction with the International Supercomputing Conference - High Performance, June 21-25, 2019, Frankfurt, Germany.

(Springer LNCS Proceedings)

=====

Date: June 25, 2019

Workshop URL: <http://vhpc.org>

Paper Submission Deadline: April 19, 2019

Springer LNCS, rolling abstract submission

Abstract/Paper Submission Link: <https://edas.info/newPaper.php?c=26973>

## **Call for Papers**

Containers and virtualization technologies constitute key enabling factors for flexible resource management in modern data centers, and particularly in cloud environments. Cloud providers need to manage complex infrastructures in a seamless fashion to support the highly dynamic and heterogeneous workloads and hosted applications customers deploy. Similarly, HPC environments have been increasingly adopting techniques that enable flexible management of vast computing and networking resources, close to marginal provisioning cost, which is unprecedented in the history of scientific and commercial computing.

Various virtualization-containerization technologies contribute to the overall picture in different ways: machine virtualization, with its capability to enable consolidation of multiple underutilized servers with heterogeneous software and operating systems (OSes), and its capability to live-migrate a fully operating virtual machine (VM) with a very short downtime, enables novel and dynamic ways to manage physical servers; OS-level virtualization (i.e., containerization), with its capability to isolate multiple user-space environments and to allow for their coexistence within the same OS kernel, promises to provide many of the advantages of machine virtualization with high levels of responsiveness and performance; lastly, unikernels provide for many virtualization benefits with a minimized OS/library surface. I/O Virtualization in turn allows physical network interfaces to take traffic from multiple VMs or containers; network virtualization, with its capability to create logical network overlays that are independent of the underlying physical topology is furthermore enabling virtualization of HPC infrastructures.

## **Publication**

Accepted papers will be published in a Springer LNCS proceedings volume.

## **Topics of Interest**

The VHPC program committee solicits original, high-quality submissions related to virtualization across the entire software stack with a special focus on the intersection of HPC, containers-virtualization and the cloud.

Major Topics:

- HPC on Containers and VMs
- Containerized applications with OS-level virtualization
- Lightweight applications with Unikernels
- HP-as-a-Service

each major topic encompassing design/architecture, management, performance management, modeling and configuration/tooling:

Design / Architecture:

- Containers and OS-level virtualization (LXC, Docker, rkt, Singularity, Shifter, i.a.)
- Hypervisor support for heterogeneous resources (GPUs, co-processors, FPGAs, etc.)
- Hypervisor extensions to mitigate side-channel attacks  
([micro-]architectural timing attacks, privilege escalation)
- VM & Container trust and security models
- Multi-environment coupling, system software supporting in-situ analysis with HPC simulation
- Cloud reliability, fault-tolerance and high-availability
- Energy-efficient and power-aware virtualization
- Containers inside VMs with hypervisor isolation
- Virtualization support for emerging memory technologies
- Lightweight/specialized operating systems in conjunction with virtual machines
- Hypervisor support for heterogeneous resources (GPUs, co-processors, FPGAs, etc.)
- Novel unikernels and use cases for virtualized HPC environments
- ARM-based hypervisors, ARM virtualization extensions

Management:

- Container and VM management for HPC and cloud environments
- HPC services integration, services to support HPC
- Service and on-demand scheduling & resource management
- Dedicated workload management with VMs or containers
- Workflow coupling with VMs and containers
- Unikernel, lightweight VM application management
- Environments and tools for operating containerized environments (batch, orchestration)

- Novel models for non-HPC workload provisioning on HPC resources

#### Performance Measurements and Modeling:

- Performance improvements for or driven by unikernels
- Optimizations of virtual machine monitor platforms and hypervisors
- Scalability analysis of VMs and/or containers at large scale
- Performance measurement, modeling and monitoring of virtualized/cloud workloads
- Virtualization in supercomputing environments, HPC clusters, HPC in the cloud

#### Configuration / Tooling:

- Tool support for unikernels: configuration/build environments, debuggers, profilers
- Job scheduling/control/policy and container placement in virtualized environments
- Operating MPI in containers/VMs and Unikernels
- Software defined networks and network virtualization
- GPU virtualization operationalization

The Workshop on Virtualization in High-Performance Cloud Computing (VHPC) aims to bring together researchers and industrial practitioners facing the challenges posed by virtualization in order to foster discussion, collaboration, mutual exchange of knowledge and experience, enabling research to ultimately provide novel solutions for virtualized computing systems of tomorrow.

The workshop will be one day in length, composed of 20 min paper presentations, each followed by 10 min discussion sections, plus lightning talks that are limited to 5 minutes. Presentations may be accompanied by interactive demonstrations.

#### **Important Dates**

Apr 19th, 2019 - Abstract (extended), Paper submission deadline (Springer LNCS)

May 3, 2019 - Acceptance notification

June 20th, 2019 - Workshop Day

July 10th, 2019 - Camera-ready version due

#### **Chair**

Michael Alexander (chair), University of Vienna, Austria

Anastassios Nanos (co-chair), SunLight.io, UK

## **Program committee**

Stergios Anastasiadis, University of Ioannina, Greece  
Jakob Blomer, CERN, Europe  
Eduardo César, Universidad Autonoma de Barcelona, Spain  
Taylor Childers, Argonne National Laboratory, USA  
Stephen Crago, USC ISI, USA  
Tommaso Cucinotta, St. Anna School of Advanced Studies, Italy  
Christoffer Dall, Columbia University, USA  
Patrick Dreher, MIT, USA  
Kyle Hale, Northwestern University, USA  
Bob Killen, University of Michigan, USA  
Brian Kocoloski, Washington University, USA  
John Lange, University of Pittsburgh, USA  
Giuseppe Lettieri, University of Pisa, Italy  
Qing Liu, Oak Ridge National Laboratory, USA  
Nikos Parlavantzas, IRISA, France  
Kevin Pedretti, Sandia National Laboratories, USA  
Amer Qouneh, Western New England University, USA  
Carlos Reaño, Queen's University Belfast, UK  
Borja Sotomayor, University of Chicago, USA  
Jonathan Sparks, Cray, USA  
Joe Stubbs, Texas Advanced Computing Center, USA  
Anata Tiwari, San Diego Supercomputer Center, USA  
Kurt Tutschku, Blekinge Institute of Technology, Sweden  
John Walters, USC ISI, USA  
Yasuhiro Watashiba, Osaka University, Japan  
Chao-Tung Yang, Tunghai University, Taiwan  
Na Zhang, VMware, USA

## **Paper Submission-Publication**

Papers submitted to the workshop will be reviewed by at least two members of the program committee and external reviewers. Submissions should include abstract, keywords, the e-mail address of the corresponding author, and must not exceed 10 pages, including tables and figures at a main font size no smaller than 11 point. Submission



of a paper should be regarded as a commitment that, should the paper be accepted, at least one of the authors will register and attend the conference to present the work. Accepted papers will be published in a Springer LNCS volume.

The format must be according to the Springer LNCS Style. Initial submissions are in PDF; authors of accepted papers will be requested to provide source files.

Format Guidelines:

<ftp://ftp.springernature.com/cs-proceeding/lncs/lncs2e.zip>

Abstract, Paper Submission Link:

<https://edas.info/newPaper.php?c=26973>

### **Lightning Talks**

Lightning Talks are non-paper track, synoptical in nature and are strictly limited to 5 minutes. They can be used to gain early feedback on ongoing research, for demonstrations, to present research results, early research ideas, perspectives and positions of interest to the community. Submit abstract via the main submission link.

### **General Information**

The workshop is one day in length and will be held in conjunction with the International Supercomputing Conference - High Performance (ISC) 2019, June 21-25, Frankfurt, Germany.