EuroHPC Containers Forum

Members

Site name	Representative(s)	Emails
UiO/Sigma2	Abdulrahman Azab (coordinator)	azab@uio.no
IJS/Sling	Barbara Krasovec	barbara.krasovec@ijs.si
IJS/Sling	Dejan Lesjak	dejan.lesjak@ijs.si
IZUM/VEGA	Teo Prica	teo.prica@izum.si
UL	Leon Kos	leon.kos@lecad.fs.uni-lj.si
UL	Matjaz Pancur	matjaz.pancur@fri.uni-lj.si
QNIB Solutions	Christian Kniep	christian@qnib.org
CINECA/Leonardo	Francesco Cola	f.cola@cineca.it
HPE	Alfio Lazzaro	alfio.lazzaro@hpe.com
HPE	Jonathan “Bill” Sparks	jonathan.sparks@hpe.com
CSC	Henrik Nortamo	henrik.nortamo@csc.fi
HPE	David Brayford	dbrayford@gmail.com
Karolina	Lukas Krupcik	lukas.krupcik@vsb.cz
Karolina	Jakub Kropacek	jakub.kropacek@vsb.cz
Karolina	Radovan Pasek	radovan.pasek@vsb.cz

Background

Google Drive:
https://drive.google.com/drive/folders/1Mj4lv-aa5EcQ8h4BtMRpWTdAANEnMmni?usp=sharing
WS site: https://events.prace-ri.eu/event/1410/

EuroHPC Supercomputers

https://eurohpc-ju.europa.eu/about/our-supercomputers_en

Presented topics

- Checkpoint/restart (Adrian reber) - TBC October 23
- HPC Container conformance paper (Christian/Bill) - November
- Draft for the HPCW (Christian) - Candidate October 23
- Recap of SC23 (Violeta Sikaleska, Tykierko Mateusz, Nina Mujkanovic) - December 23
- Charlicloud (David) - January 24
- SME users experience (Barbara Krasovec, Tykierko Mateusz) - February 24
    - User namespaces and security (Barbara)
    - Agenda for the HPC containers WS
- Software containers in HPC (History of containers on the slovenian clusters) - Dejan (May)
- “ORAS (artefacts in OCI registries) - May 2025

Topics Blacklog

- Meta-hub for AI Factories - Christian - May June 2025
- LUMI Containers workflow - May/June
- HPC Containers in the LEM and BEM cluster - Hubert Mazur - August 2024
- Container registry for AI Factories - Christian - Jan 2025

Meetings

2025-05-15

Presence:

- Abdulrahman Azab
- Lukas Hejtmanek
- David Brayford
- Teo Prica
- Viktoria
- Christian Kniep
- Dejan Lesjak

- Alja Prah
- Barbara Krasovec

Topics:

Jonathan Sparks presents “ORAS (artefacts in OCI registries)

https://github.com/langflow-ai/langflow
I saw ragflow for building rag chats
https://ragflow.io/docs/dev/

2025-03-13

Presence:

- Alfio Lazzaro
- Abdulrahman Azab
- Christian Kniep
- Dejan Lesjak
- Alja Prah
- Barbara Krasovec
- Mateusz Tykierko

Topics:

Who is attending EuroHPC summit:
- Abdulrahman
- Mateusz Tykierko
- Leon Kos
- Dejan Lesjak
- Jan Jona Javoršek
-

Proposed Agenda:
- Containers for AI Factories: Software, training, runtimes
- Repositories/registries for AI containers and AI models
- Other …

ISC workshop:

- Apply for a talk: https://docs.google.com/forms/d/1-Be1uypwDeLkA9rQPFukPINZ686mgxruGtYIpfiGcrc/viewform?edit_requested=true#responses

2025-02-20

Presence:

- Alfio Lazzaro
- Abdulrahman Azab
- Christian Kniep
- Dejan Lesjak
- Alja Prah
- Barbara Krasovec
- Mateusz Tykierko

Topics:

- ISC Container workshop: the Google form is out to submit a talk, in March we will meet to choose the talks and establish the agenda
- 2 BoFs at ISC, will be known shortly if they are accepted
  ISC2025-HPCW_w/o_proceedings
  - Software distribution for AI Factories
  - Bioinformatics workflows on the HPC
- Software distribution registry - we still need a common image registry, maybe also in the context of Epicure or later for AI factories
- AI factory: moving toward hybrid setups, a combination of traditional HPC, private cloud and containerized environment (Kubernetes), containerized environments will be predominant
- Epicure collaboration, it was in the proposal that they would collaborate with EuroHPC Container Forum to provide also container images for EuroHPC users. Currently, some containers were done, but for a specific case. Maybe a joint meeting on the 13th of March with Epicure to discuss this?
- F2F meeting at EuroHPC Summit, Mateusz will try to get a room with remote participation

2025-01-16

Presence

- Abdulrahman Azab
- Alja Prah
- David Brayford
- Barbara Krasovec
- Dejan Lesjak
- Christian Kniep
- Anton Tustanowski
- Kos, Leon

Software Distribution for AI Factories: Bridging AI and HPC Workflows at Scale

As AI Factories emerge as critical hubs for advancing AI and HPC convergence, efficient software distribution becomes pivotal to scaling operations, enhancing collaboration, and delivering reproducible results. This Birds of a Feather (BoF) session will bring together practitioners, researchers, and industry stakeholders to explore challenges and solutions for software distribution in AI Factories. The discussion will focus on containerization, cross-platform compatibility, best practices in software management, and fostering interoperability between national and international AI Factory infrastructures, such as those aligned with the LUMI AI Factory and EuroHPC.

Session Objectives

- Discuss Challenges: Explore barriers in software distribution, including multi-platform compatibility, performance optimization, and user accessibility for AI and HPC workflows.
- Share Best Practices: Highlight success stories and methodologies for effective software distribution using tools like containerized environments (e.g., Singularity, Docker), package managers, and version control systems.
- Foster Collaboration: Connect attendees from diverse AI Factory ecosystems to identify synergies and foster cross-border collaborations.
- **Propose Standards:** Evaluate opportunities to establish standardized practices for software distribution across AI Factories.

Proposed set of invitees:
CSC – IT Center for Science (Finland)
IT4I
DeiC
BSC
CINECA
SURF

ETH Zurich (Switzerland)
Other ….

2024-12-19

Presence

- Abdulrahman Azab
- Alja Prah
- David Brayford
- Alfio Lazzaro (HPE)
- Barbara Krasovec
- Dejan Lesjak
- Christian Kniep
- Anton Tustanowski
- Kos, Leon

Topics

ISC workshop

AI Factories

Introduction to AI Factories in EuroHPC

The European High-Performance Computing Joint Undertaking (EuroHPC JU) has set ambitious goals to establish Europe as a leader in artificial intelligence (AI) by integrating advanced AI capabilities with its high-performance computing (HPC) infrastructure. AI Factories are central to this vision, acting as hubs for cutting-edge AI research, development, and application. These factories aim to create a cohesive ecosystem where HPC resources, high-value datasets, and innovative AI tools converge to empower researchers, industries, and public institutions. By fostering collaboration across borders, EuroHPC's AI Factories enhance

Europe's competitiveness in the global AI landscape while addressing societal challenges through trustworthy and impactful AI solutions.

AI Factories are designed to address key challenges in AI development, such as access to computing resources, secure data handling, and the need for interdisciplinary collaboration. They enable the development of general-purpose AI models, facilitate the adoption of AI in industry and academia, and promote the growth of AI talent across Europe. The factories also align with EU initiatives on digital transformation and sustainability, ensuring that AI development adheres to ethical and environmental standards.

The LUMI AI Factory Proposal

The LUMI AI Factory is a cornerstone initiative under the EuroHPC umbrella, led by a Finland-based consortium and involving partners from six European countries. Building upon the success of the LUMI supercomputer, the proposal introduces an advanced AI ecosystem that integrates world-class computing power, high-value data, and top-tier AI talent. The LUMI AI Factory's vision is to create an open, accessible platform that empowers startups, SMEs, researchers, and public institutions to develop innovative AI models and applications.

This proposal aims to address critical challenges in AI development, such as data reachability, talent shortages, and access barriers. By leveraging the existing LUMI ecosystem and expanding it with new AI-focused features, the LUMI AI Factory seeks to foster a thriving cross-border AI community. Key objectives include advancing generative AI, supporting industry-specific AI solutions, and promoting the development of trustworthy AI aligned with European values.

The consortium behind the LUMI AI Factory includes prominent organizations such as CSC – IT Center for Science (coordinator), Aalto University, the University of Helsinki, and international partners like Sigma2 (Norway), IT4Innovations (Czech Republic), DeiC/DTU (Denmark), and the University of Tartu (Estonia). Together, these partners contribute expertise in HPC, AI research, industrial collaboration, and training, ensuring the LUMI AI Factory's success as a pan-European initiative.

Slovenia will apply to the February call

Ideas to link the forum to AI Factories

- EuroHPC Container forum can contribute to a common way of submitting workloads, e.g. using SLURM API
- Containerising services

Next meeting: Jan 16th 2025

2024-04-11

Presence

- Abdulrahman Azab
- Alja Prah
- David Brayford
- Teo Prica
- Kos, Leon
- Tykierko Mateusz
- Dejan Lesjak
- Christian Kniep

Topics

Presentation: Scientific software and containers in Slovenian HPC

Slides:
https://www-f9.ijs.si/~lesi/slides/Scientific_software_and_containers_in_Slovenian_HPC.pdf

Other links:

- Easy-Config repo for LUMI: https://github.com/Lumi-supercomputer/LUMI-EasyBuild-contrib
- LUMI software stack: https://github.com/Lumi-supercomputer/LUMI-SoftwareStack
- Containers: https://github.com/Lumi-supercomputer/LUMI-EasyBuild-containers/tree/main/easybuild/easyconfigs

2024-04-11

Presence

- Abdulrahman Azab
- Christian Kniep
- Alja Prah
- Alfio Lazzaro (HPE)
- Barbara Krasovec
- Tiziano Müller (HPE)
- Dejan Lesjak
- Teo Prica
- Violeta Šikaleska
- Radostin Stoyanov
- Matjaz Pancur
- Lebar Bajec, Iztok

Topics

HPCW 2024

HPCW24 Planning KickOff

NeIC Conference CFP

https://equinocs.springernature.com/service/NeIC2024

Container Security

Comments:

- The following is another recently discovered CVE for runc: https://github.com/advisories/GHSA-xr7r-f8xq-vfvv
- The recently discovered xz backdoor is another interesting example of supply chain attack: https://en.wikipedia.org/wiki/XZ_Utils_backdoor

2024-02-15

Presence

- Abdulrahman Azab
- Christian Kniep
- Alfio Lazzaro (HPE)
- Barbara Krasovec
- Tiziano Müller (HPE)
- Jani Pogačar

Topics

HPCW 2024

HPCW24 Planning KickOff

EuroHPC Summit week 2024

Attendees:

- Abdulrahman
- Barbara
- VEGA team
- Jani Pogačar (zoom)
- Christian (zoom)(TBC)

https://www.eurohpcsummit.eu/programme

Thursday 21.03 12:30 - 14:00

Last year: https://indico.neic.no/event/248/

Topics:

- Software containers in HPC (History of containers on the slovenian clusters) - Dejan Lesjak
- Container related EuroHPC JU - Barbara Krasovec

    - AI Factory as the 7th pillar
    - 20% of the computers will be used for AI especially for SMEs
    - Containers are the option to provide reproducibility (same environment)
- Interoperability of Cloud and HPC - ??
- HPC Containers and Security - Barbara

NeIC 2024

https://indico.neic.no/event/259/

AP Abdulrahman: Ask the Springer Editors about the possibility of having papers/abstracts from the HPCW sessions

2024-01-25

Presence

- Abdulrahman Azab
- Christian Kniep
- Alfio Lazzaro (HPE)
- Nina Mujkanovic (HPE)
- David Brayford
- Alja Prah
- Barbara Krasovec
- Dejan Lesjak
- Tiziano Müller (HPE)
- Wygladacz Konrad
- Teo Prica
- Radostin Stoyanov
- Lukierska Urszula
- Mateja
- Jani Pogačar

Topics

Charlicloud

Presentation recording: CharlieCloud.mp4
Unprivileged mode: https://dl.acm.org/doi/10.1145/3458817.3476187

How to use IB libraries: Bind-mount from the host or install them inside the container

Most of the exploits in the last two years where related to enabling both network and user namespaces

EuroHPC Summit

https://www.eurohpcsummit.eu/

Alja Prah
Barbara Krasovec
Abdulrahman Azab
Leon Kos

AP Barbara: Ask for a possible slot for the EuroHPC containers forum

EuroHPC machines will be used for AI workloads for the industry

- Press release: https://ec.europa.eu/commission/presscorner/detail/en/IP_23_5739
- https://ec.europa.eu/commission/presscorner/detail/en/ip_24_383

E.g. LUMI will need to offer 20% for AI workloads

There will be an amendment to the regulations for the EuroHPC to add AI pillar (getting more computing power to AI). They couldn't invest in data access.

Amendment:
https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-amending-regulation-eu-2021/1173-regards-eurohpc-initiative-start-ups-boost

- AI users are not fond of modules but more containers
- The investment in the next two years will be in data and AI

HPCW@ISC

Container boat-trip
Full-day workshop
Hybrid setup
Setup:

- Presentations
- Lightning talks

Planning to go to ISC:
- David
- Christian
- Abdulrahman

AP All: Contact Christian if you have something interesting to present

FOSDEM

HPC DevRoom: https://fosdem.org/2024/schedule/track/hpc-big-data-data-science/
If you are around, there is a dinner on Saturday. Ping me or Kenneth Hoste

NeIC 2024 conference

https://indico.neic.no/event/259/

HPC Containers workshop agenda and dates

Dates:
- **AP Teo:** Provide preferences on dates by next week

Next meeting: **Feb 15th**

2023-12-07

Presence

- Christian Kniep
- Abdulrahman Azab
- Alfio Lazzaro (HPE)
- Nina Mujkanovic (HPE)
- David Brayford
- Alja Prah
- Barbara Krasovec
    - Dejan Lesjak
    - Tykierko Mateusz
    - Violeta Šikaleska

- Mazur Hubert
- Teo Prica
- Tiziano Müller (HPE)
- Adrian Reber
- Matjaz Pancur
    - Mateusz Tykierko

Topics

SC23 recap

Canopie WS was full

- Issues with the labeling of containers
- Giving access to users in a system

A lot of talks about Frontier and other exascale-related topics

Panel sessions:

- Getting Python into HPC

Open HPC Booth:

- Questions about RedHat and how to get more involved with HPC users
- Building something and losing access to the code if you distribute it

HPC Software Foundation was announced

- https://hpsfoundation.github.io/#welcome
- High Performance Software Foundation Pitch Deck

ISC24 workshop submission

Submitted a WS without proceedings

OCI Image compatibility WG

Discussing the use cases, broader not just HPC specific

HPC Conformance paper

https://gitlab.com/qnib-metahub/hpc3-paper

Contributors:

- David
- Alfio
- Abdulrahman
- Tiziano Müller

Future conferences

FOSDEM 2024
https://discuss.linuxcontainers.org/t/fosdem-2024-containers-devroom-call-for-papers/18253
NeIC 2024 conference
https://indico.neic.no/event/259/
Possibility to have a HPC containers Satilite event

LUMI Container building

Next meeting

EuroHPC containers training

Best time: June next year
AP Teo: Ask about the possibility to move the course to January
Platform/Infrastructure:

- Confirm the workshop contributions and possibility for platform accounts on VEGA - Confirmed
- Admin WS: Container security Confirmed (if online, also possible to travel)
    - Namespaces/cgroups
- https://doc.sling.si/en/workshops/advanced-supercomputing/01-intro/01-course/
- https://indico.cern.ch/event/1297500/contributions/5454541/
- AP Matjas and Barbara: Check with the CoE workshop holders about the possible dates for a Spring WS

EuroHPC Summit week

Containers on HPC for industry: make it easy for industry to use HPC:

- Christian: knows someone from Volxvagen

AI and containers:

- We can prepare something about AI and Containers
- Mention it to CJ NVIDIA

2023-11-09

Presence

- Christian Kniep
- Abdulrahman Azab
- Nina Mujkanovic (HPE)
- David Brayford
- Barbara Krasovec
- Dejan Lesjak
- Tykierko Mateusz
- Violeta Šikaleska
- Mazur Hubert
- Teo Prica
- Tiziano Müller (HPE)
- Adrian Reber
- Matjaz Pancur
- Mateusz Tykierko

Topics

SC23 attendees

- Mateusz Tykierko
- Hubert Mazur
- David Brayford
- Violeta Sikaleska
- Nina Mujkanovic
- Tykierko Mateusz

OCI WG

Image compatibility group: Christian will be our ambassador

EuroHPC Summit week

Presentation about how to provide goood users experience using containers

EuroHPC containers training

AP Teo: Ask about the possibility to move the course to January
Platform/Infrastructure:

- Confirm the workshop contributions and possibility for platform accounts on VEGA - Confirmed
- Container security Confirmed (if online)

2023-10-12

Presence

- Christian Kniep
- Abdulrahman Azab
- David Brayford
- Barbara Krasovec
- Dejan Lesjak
- Adam Hough
- Radostin Stoyanov
- Teo Prica
- Tiziano Müller (HPE)
- Adrian Reber
- Matjaz Pancur
- Alfio Lazzaro

Topics

Development Call for EuroHPC was announced for developing applications for the end-users

https://eurohpc-ju.europa.eu/eurohpc-federation-platform_en

EuroHPC Summit week

https://www.hpcwire.com/off-the-wire/eurohpc-summit-2024-will-be-in-belgium-booth-presence-at-isc-sc/#:~:text=March%2031%2C%202023%20%E2%80%94%20The%20EuroHPC,of%20EuroHPC%20Summit%20in%202023.

Applications to fully exploit EuroHPC systems

AP Abdulrahman and Barbara: ask how to propose a topic about portable HPC containers

EuroHPC containers training

Platform/Infrastructure:

- Confirm the workshop contributions and possibility for platform accounts on VEGA - Confirmed
- Container security Confirmed (if online)
- Containers and K8s Abdulrahman to contact Mattias

Checkpoint/Restore in userspace CRIU

Presentation by Adrian Reber
Recording:
https://drive.google.com/file/d/1-a356pvDQ-cAy8rif_GIo9BdpgYf1egT/view?usp=drive_link

CRIU for research: https://criu.org/Articles

OpenVZ: community driven, moving a container from one machine to another
Integrated with Podman and Docker

Integration with Podman

- podman container checkpoint
- podman container restore

K8s based integration offers more advanced options

- Working on: kubectl checkpoint
- Pod Checkpoint/Restore: Ckpt/restore containers in a pod and probably moving them somewhere else: PoC is made

Migrating container from one system to another:

- Similar to VM migration
- The CPU cannot be old

Integration with Docker:

- No clear view about the status of it
- Advice: use Podman

HPCW at ISC24

Due date is tomorrow

Compatibility WG

Container Packaging paper (no progress)

Christian will share the template and initial project
https://gitlab.com/qnib-metahub/hpc3-paper

Next meeting November 2nd

2023-09-14

Presence

- Christian Kniep
- Abdulrahman Azab
- Bill Sparks
- Dejan Lesjak
- Tiziano Müller
- Barbara Krasovec
- Teo Prica
- Violeta Šikaleska
- Lebar Iztok
- Iulia Ibanescu

Topics

Container Packaging paper (no progress)

Christian will share the template and initial project
https://gitlab.com/qnib-metahub/hpc3-paper

HPCW at ISC24 (10th)

The submission is open for workshops with proceedings (deadline mid October - 10/13)

- We might submit for a workshop with proceedings?
- Bill: Anker the workshop around ‘hot topics’ within HPC where container would help (make it catchy and current).
- How can we stand out from the VHPC workshop? (one is what Bill said)
- Bill: Container runtimes is somewhat an obsolete topic, but the direction of environment integration e.g. Cloud/HPC convergence is more valid and alive topic.
- Focusing on container packaged applications/solutions, i.e. use cases for research/industry communities and how to create portable and easy deployable solutions
- Containers for ISV (Independent Software Vendor) applications
- https://github.com/opencontainers/tob/pull/128

christian.qnib@gmail.com Create a draft, fetch review from last years rejection
Abdul: We might aim for short-paper and extended abstracts.

Compatibility WG

- https://github.com/opencontainers/tob/pull/128
- Revisit container and host OS compatibility
- Bill: There are different directions in the implementation of container compatibility including container annotation
- Working subgroup from the forum:
    - Coordinator: Bill
    - Contributors: Christian, Abdulrahman, Alfio?, David?, Lebar?

NIST Proposal

https://www.federalregister.gov/documents/2023/08/24/2023-18263/request-for-information-regarding-file-specification-for-findable-accessible-interoperable-and

- This document has a comment period that ends in 84 days. (12/07/2023)

EuroHPC containers training

Platform/Infrastructure:

- Confirm the workshop contributions and possibility for platform accounts on VEGA - Confirmed
- Container security - AP Barbara
- Containers and K8s - TBC

Development Call for EuroHPC was announced for developing applications for the end-users

- All applications developed by the CoEs needs to be there
- Copy all applications and create a registry/repo for these applications
- EuroHPC summit will in March 2024
- Will be discussed in the next governance board meeting in October
- Will be out in early 2024

Next meeting October 12th

2023-08-03

Presence

- Abdulrahman Azab
- David Brayford
- Lebar Bajec, Iztok
- Christian Kniep
- Jan Wender
- Dejan Lesjak
- Teo Prica
- Alfio
- Alja Prah

Topics

Container Packaging paper

Christian will share the template and initial project
https://gitlab.com/qnib-metahub/hpc3-paper

Interesting contributors:
- Abdulrahman
- David
- Alfio
- Lebar

HPE project: Working on Podman and SLUM integration with K8s

CANOPIE workshop

Christian is submitting a lightning talk
Abdulrahman may submit a paper about containerised HPC

EuroHPC containers training

Platform/Infrastructure:

- Vega: it is possible to create temp user accounts which will be available throughout and 1 week after the training, then deleted (this should be ok)
    - Aim for max 50 participants
- Singularity runtime

APs

- Confirm the workshop contributions and possibility for platform accounts on VEGA
- Check if possible to book an auditorium @ UL Or Maribor

Next meeting

- Container checkpoint/restart
- HPE work on K8s/Slurm?
-

2023-07-06

Presence

- Abdulrahman Azab
- David Brayford
- Lebar Bajec, Iztok
- Christian Kniep
- Jan Wender
- Barbara Krasovec
- Dejan Lesjak
- Teo Prica

- Alfio

Topics

EuroHPC containers training

Previous training: https://events.prace-ri.eu/event/1229/

Targeted Date: October/November/December
Setup: Online/Hybrid (Oslo/UiO is possible)

User and Provider sessions:
- Users: 2.5 Days
- Providers: 2.5 Days

Platform/Infrastructure:
- Vega: it is possible to create temp user accounts which will be available throughout and 1 week after the training, then deleted (Ask Teo)
    - Singularity Pro
    - Podman?
    - enroot?
- https://sling.si/event/seminar-virtualizacija-in-kontejnerji-pod-pokrovom/
- K8s
- Container runtimes

User training:
Target: How to use a container runtime for HPC workloads
Container platforms
- Docker (Abdulrahman)
    - Container concept
    - Container building
    - Container building - advanced, e.g. multistage build
    - Overview of the container runtimes
- Container packaging and distribution
    -
- Apptainer (https://ciq.com/) (Abdulrahman, VEGA folks (Matjaž)?)
    - Introduction (including the main pros, cons, and comparison) and basic commands
    - Container building, rootless building
    - MPI
    - GPU
- Podman (NERSC/HPE? Alfio asks)

        - Introduction (including the main pros, cons, and comparison) and basic commands
        - Container building, rootless building
        - MPI
        - GPU
        - Checkpoint-restart (David asks)
            - https://www.youtube.com/watch?v=tEyH2ogFKeg
- Sarus
    - Alberto?
- K8s
    - Ask Matjaž (Teo will ask)
    - https://indico.ijs.si/event/1400/

Provider training:
Target: How to install and maintain a container runtime, and manage user privileges, for HPC workloads

- Container security (2.5 hrs)
    - https://indico.cern.ch/event/1106023/contributions/4915297/ (hands-on)
    - https://arxiv.org/abs/2104.07508 (webinar)
- Container packaging and distribution
    -
-
- K8s?
- Apptainer: best practices, MPI, and GPU support
- Podman: best practices, MPI, and GPU support
- Sarus: best practices, MPI, and GPU support
-

Duration: 1 week?
Date

2023-05-11

Presence

- Abdulrahman Azab
- David Brayford
- Lebar Bajec, Iztok
- Christian Kniep
- Jan Wender
- Barbara Krasovec
- Dejan Lesjak

- Teo Prica
- Alfio

Topics

HPCW workshop

Who is attending:
- Christian
- Abdulrahman

Online attendance:
- Barbara
- Alfio
- Teo Prica

Container Boat Trip:

For folks interested in container, please join (and invite more folks):
Info: https://container-in-hpc.org/isc/2023/friends-of-container-trip/index.html
SignUp: https://forms.gle/mjx9ymEw6hTfQUmJ9

EuroHPC containers session

3 slides for Vega
3 slides for LUMI
- Support for Open MPI and PMIx on a system that has MPICH
- Container wrapper

Alberto's presentation:
- https://www.hpcadvisorycouncil.com/events/2023/swiss-conference/pdf/day3/CSCS_Madonna.pdf
- https://youtu.be/-EECnPA5oCc
- Full event: https://www.hpcadvisorycouncil.com/events/2023/swiss-conference/agenda.php

CSC applied for the EuroHPC call "EuroHPC containers forum" is mentioned in multiple places

Barbara will provide link to the submission

Next meeting topics

K8S in Slurm from LUMI/VEGA (if the solution is ready to present)
PMIx and Open MPI support on Cray (MPICH) systems

2023-04-13

Presence

- Abdulrahman Azab
- David Brayford
- Leon Kos
- Lebar Bajec, Iztok (virtual)
- Christian Kniep
- Jan Wender

Topics

HPCW workshop

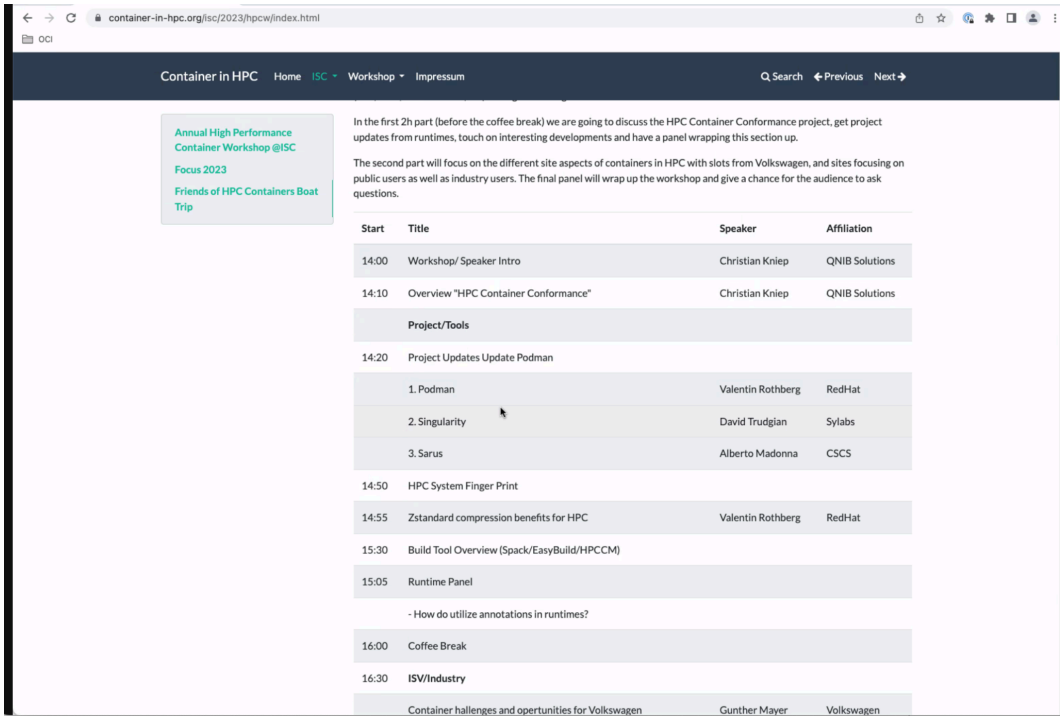

Contributions from the forum:

- 7 mins presentation for EuroHPC systems
- Propose use cases by the industrial users of EuroHPC systems

Boat Trip

Please sign up for the boat trip this ISC:
https://container-in-hpc.org/isc/2023/friends-of-conthttps://container-in-hpc.org/isc/2023/friends-of-container-trip/index.htmlainer-trip/index.html

2023-03-21

Presence

- Abdulrahman Azab
- David Brayford (virtual)
- Lebar Bajec, Iztok (virtual)
- Christian Kniep

- Jan Wender

Topics

Intro

Slides

Projects of interest

- Efficient container runtimes: Apptainer/Singularity, Sarus, Podman, Charliecloud, …
- Portable HPC containers: Software stack that can be used efficiently on multiple EuroHPC systems for HPC applications running on CPUs and GPUs
    - Share the binaries or recipes
    - Handling the heterogeneity in GPUs, MPI, etc.
- Efficient sharing platform
    - Which sharing platform/registry to use
- Provide a portable Container cloud solution (K8S etc).
    - How to control the cloud application: K8s in control or Slurm in control
    - VMs? How to preserve the performance: Virtualization to be Lightweight

Building containers outside the HPC system might be ok but in many cases, build on premises is needed

HPC Container Conformance

Slides

Activities

- Monthly community meetings: first thursday of every month
- Training activity: One workshop/school per year
- Hackathons: Two per year
- Projects

2023-02-02

Presence

- Christian Kniep
- Dejan Lesjak
- Alfio Lazzaro
- David Brayford
- Abdulrahman Azab
- Pancur Matjaz
- Christian Kniep
- Lebar Bajec, Iztok
-

Upcoming EuroHPC call

Application Support in High Performance Computing
https://eurohpc-ju.europa.eu/application-support-high-performance-computing_en

AP All: Define interests by next week

Fakeroot presentation

K8s on Vega

Tests are ongoing

Container behaviour

GROMAC container: https://github.com/amd/InfinityHub-CI/tree/main/gromacs-docker
PyTorch: https://pytorch.org/tutorials/intermediate/dist_tuto.html

For the next meeting

K8s on Vega (Pancur Matjaz)

2023-01-19

Presence

- Christian Kniep
- Bill Sparks
- Dejan Lesjak
- Alfio Lazzaro
- David Brayford
- Abdulrahman Azab
- Barbara Krasovec

Possible outlets for a workshop

- euroHPC Workshop 20 March -> 23 March
  -> Program Website
    - Planned attendees:
      - Abdulrahman Azab (TBC)
      - Christian Kniep (if we put something together)
      - Bill Sparks (maybe - depending on budget/agenda)

- CUG23 (DL: Jan 23rd): May 7-13
- ISC23: May 21-25
    - Maybe a workshop the weekend before? 20 and/or 21 May
- Previous workshop: https://events.prace-ri.eu/event/1229/

HPC Container Conformance

Christian is presenting the first iteration on the project 'HPC Container Conformance', which aims to provide guidance on how a container should behave and what annotations we might converge towards.

- Slides (2023-01-19, v9)
- Usage of container annotations
- Good troubleshooting
- Configuere hooks and runtimes to tweek containers correctly
- Can recommend a suitable image if you don't have one
- Hardware annotations:

Hardware Annotations

CPU/GPU/...

Information about what the application in the containers user-land is compiled for.

- Will the application segfault due to architecture mismatch (beyond the platform specification ARM/x86)?
- What CUDA version and GPU architecture is the application build against?

org.supercontainers.hardware.cpu.optimized.mode	architecture, genericMicro, microarchitecture
org.supercontainers.hardware.cpu.optimized.version	x86_64 / x86_64_v4 / skylake / skylake_avx512
org.supercontainers.hardware.gpu.nvidia.driver.version	346.34
org.supercontainers.hardware.gpu.nvidia.cuda.version	12.1
org.supercontainers.hardware.gpu.nvidia.architecture	sm_35 (kepler), sm_86 (ampere)

- MPI annotations

MPI/Interconnect Annotations

Information about what the user-land is compiled for and what methods to tweak the container is the container designed for?

org.supercontainers.mpi.implementation	(openmpi,mpich,threadmpi)
org.supercontainers.communication.framework	(ucx, libfabrics)
org.supercontainers.openmpi.version	1.16.1
org.supercontainers.libfabric.abi.version	1.6
org.supercontainers.mpi.portability.optimization	stock, cray-xc-cnl10
org.supercontainers.mpi.portability.mode	mpi_replace, libfabric_inject, ucx_replace

- System annotations
- Documentation annotations
- How to annotate:

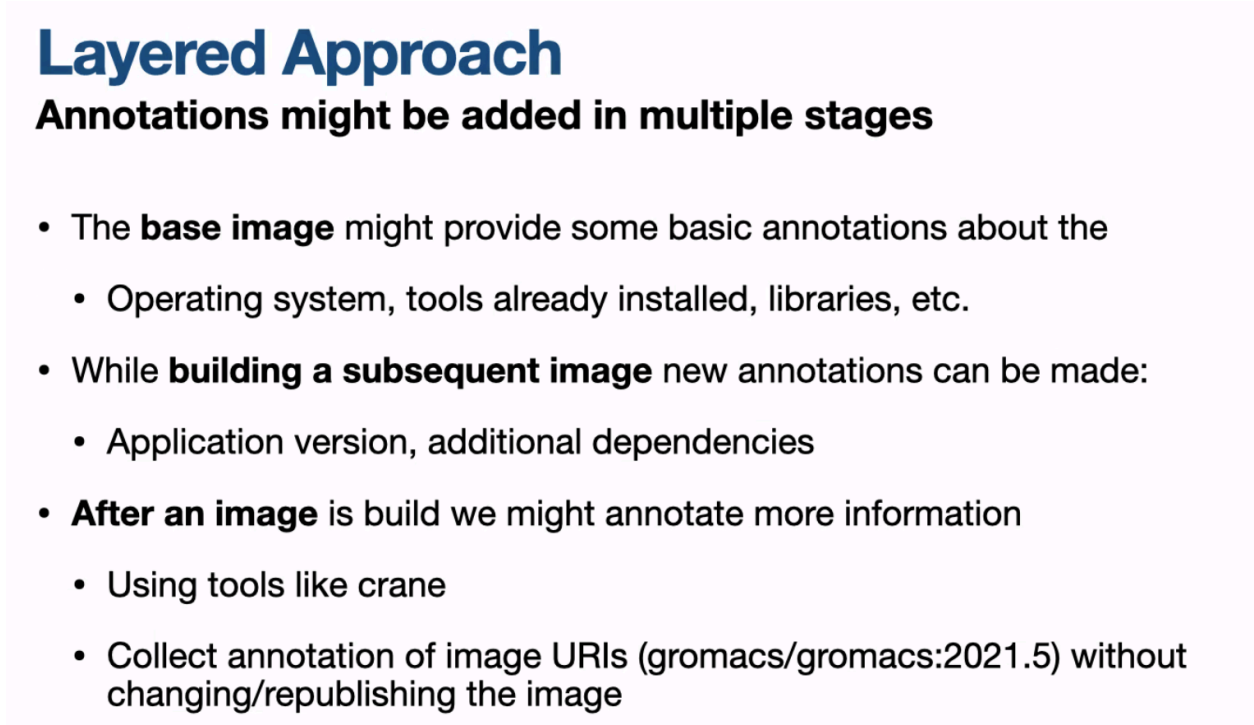

Relevant work by HPE:
https://www.sodalite.eu/sites/sodalite/files/public/content-files/articles/D3.4_Full%20application%20and%20infrastructure%20performance%20models.pdf

Next meeting topic:

- Presentation of rootless container by Alfio

2022-12-01

Presence

- Abdulrahman Azab
- Teo Prica
- Barbara Krasovec
- Alfio Lazzaro
- Christian Kniep
- Jakub Kropacek
- David Brayford
- Radovan Pasek
- Pancur Matjaz

Site Updates

Vendor Communication

2022-12-01

Presence

Present

- Abdulrahman Azab
- Teo Prica
- Barbara Krasovec
- Alfio Lazzaro
- Christian Kniep
- Dejan Lesjak
- Radovan Pasek (IT4I CZ)
- Jakub Kropáček (IT4I CZ)

AP Abdulrahman: ask about the possibility to have a HPC container workshop

AP Christian: discuss with Shane and CJ and initiate the connection regarding the NVIDIA containers

- Pytorch and GROMACS

AP Barbara: ask the Karolina guys to present next time
AP Christian: Update on Meta-hub (https://app.metahub-registry.org/)
AP Alfio: Present/demo fakeroot for singularity without user namespaces (https://github.com/apptainer/singularity/issues/5941)
Item: Prepare a workshop proposal in Gothenburg 2023: https://www.etp4hpc.eu/events/eurohpc-summit-week023-sweden_624.html

Karolina presentation

-

- For apptainer: the registries need to be manually configured

Metahub update

OCI Compience

Action points

AP Radovan: Prepare a brief presentation about Puhuri
AP Abdulrahman and Alfio: Present/demo fakeroot for singularity without user namespaces
AP Christian: Prepare a webinar with Meta-hub and the new functionality
AP Abdulrahman and Christian: Recap from the last HPC advisory council meeting

2022-11-11

Presence

Present

- Abdulrahman Azab
- Teo Prica
- Barbara Krasovec
- Alfio Lazzaro
- Christian Kniep
- Dejan Lesjak

Site Updates

Vega

Teo: cloud-hypervisor, instances have access to inter-node interconnect (SR-IOV VF on IB Mellanox adapter), Lustre was tested, air-gapped installation of Kubernetes (use snap/microk8s), 2 nodes are reserved for testing. Security concerns still remain.

Who has access send PubKey to: support@sling.si

The resource allocation options:

- Separate partition

- Slurm allocation/reservation for each K8s cluster
- Internal meeting at UL will take place to evaluate different options

High Level support Teams for EuroHPC systems (EuroHPC) call

Another call to:

- Make sure that all applications that come from the EuroCC centers can run on all EuroHPC systems
- Several projects for supporting several software packages (the setting to be decided and)

Karolina joined the forum

Vendor Communication

Base-Image initiative, action items from CJ for EuroHPC (Christian took the item)

- nominate an app each from HPC and DL that are on NGC now
    - Consider a NCCL app from DL and an MPI/NCCL app from HPC (maybe VASP, which benefits from NCCL)
    - HPC apps @ NGC are built with HPCCM. We can work with Scott McMillan to enable HPCCM for Slingshot, at least wrt answering questions and helping cleaning up a PR against the HPCCM code base
- ID at least a couple of sites (CSCS, LANL, NERSC, …) at which we'd try to test these
- We'll try to collaborate on build recipes that will ease the pain of making those work at various sites. That could include creating a base image layer that is known to work at multiple sites using Slingshot, that those apps can they be re-built on top of.

ToDo:

- 11/10 Christian send an email to Shane (NERSC), Jordan (LANL), Abdul (UiO), Alberto (CSCS) to check in at SC22 to see what we can agree on.

Candidates: PyTorch and GROMACS

For the next meeting

AP Christian: discuss with Shane and CJ and initiate the connection regarding the NVIDIA containers
AP Barbara: ask the Karolina guys to present next time
AP Christian: Update on Meta-hub (https://app.metahub-registry.org/)
AP Alfio: Present/demo fakeroot for singularity without user namespaces (https://github.com/apptainer/singularity/issues/5941)

Item: Prepare a workshop proposal in Gothenburg 2023:
https://www.etp4hpc.eu/events/eurohpc-summit-week023-sweden_624.html
AP Abdulrahman: ask about the possibility to have a HPC container workshop

2022-10-06

Presence

Present

- Abdulrahman Azab, UiO/Sigma2 (organizer)
- Dejan Lesjak
- David Brayford, LRZ
- Matjaz Pancur, UL FRI
- Christian Kniep, QNIB Solutions
- Barbara Krasovec, IJS/Sling
- Teo Prica, IZUM/VEGA
- Alfio Lazzaro, HPE
- Francesco Cola, CINECA

.

Excused

- .

Actual topics

Future meetings

Monthly basis: Every 1st Thursday of the month: 14:00 - 15:30 CET

Site updates

- VEGA:
    - Tested Cloud hypervisor
        - Disk access may need to be tested
        - Network is ok with IB SR-IOV VFs, tested on a similar server but not yet on VEGA compute nodes

- LUMI:
    - Two directions for K8s
        - LUMI-K: external K8s cluster where services run and submit jobs to the slurm cluster
        - On-demand K8s: on the compute nodes
        -
- **AP Matjaz, Teo:** Prepare a demo and an installation recipe

Communication with the vendors

LUMI had a productive discussion with AMD on the development of infinity-hub, and also for the development of containers which are not on infinity-hub

For NVIDIA NGC, we are not sure whether all images are maintained by NVIDIA
What can the vendors help in:

- Get the right drivers
- Updates on the drivers and Compute environments etc
- Optimize the performance
- TODO: Invite CJ and the rest to attend next month meeting

EuroHPC JU

Not clear how SW support on HPC systems:

The advisory RIAG (Research and Innovation Advisory Group) groups:

- They would like to support existing initiatives, e.g. EuroHPC containers forum

Governance board:

- The container forum was proposed and it is found to be a good idea
- NB! All EuroHPC centers of excellence providing software, the software needs to be supported to run on ALL Hosting entities of EuroHPC not only the national one

Call for high level support teams (December 2022)

- Offer high level support on the EuroHPC level, not the hosting entity level

Call for federation (December 2022)

1 Federation of supercomputing and data resources (EuroHPC)

Deployment and operation of a platform for federating resources (including High Performance Computing, quantum computing and data resources) providing Union-wide, cloud-based secure services for a wide range of public and private users across Europe. A solution will be deployed on top of specific dedicated resources in order to create a federated EuroHPC infrastructure. Indicative EU budget for the topic is EUR 40 million with an **EU funding rate of up to 50%**. A Call may be agreed

by the Governing Board in 2022. Public procurement 1. Delivering the EOSC core infrastructure and services

2 International Activities (EuroHPC)

The EuroHPC JU Regulation gives a mandate to the EuroHPC JU to develop strategic research and innovation partnerships in HPC with third countries like Japan, Brazil, USA or India that enables advancing the work on HPC applications in domains of common interest, including facilitating access for researchers to EuroHPC JU resources and co-development of HPC applications. The European HPC ecosystem will be further reinforced by enabling European stakeholders to develop novel algorithms, implement them in state-of-the-art codes and architectures, and test the applications and codes in academic and industrial cases to benefit both Europe and like-minded third countries. Indicative EU contribution for the topic is EUR 5 million and will be 100% EU funded. A Call for tender/proposal and other activities may be agreed by the Governing Board in 2022 and launched in 2022.

3 High Level support Teams for EuroHPC systems (EuroHPC)

The EuroHPC JU will provide grants to consortia who will establish High Level Strategic support Teams (HLSSTs) for EuroHPC systems in collaboration with EuroHPC JU Hosting Entities to provide specific services adapted to the needs of academia and industry users (i.e. benchmarking, code optimisation and scaling-out of applications). Indicative EU contribution for the topic is EUR 5 million and will be 100% EU funded. A Call may be agreed by the Governing Board in 2022. **(expected December 2022)**

State of MetaHub

Others

1. CK: CANOPIE-HPC at SC22 is going to have nice papers: libfabric-injection, podman@scale

2022-09-01

Presence

Present

- Abdulrahman Azab, UiO/Sigma2 (organizer)
- Matjaz Pancur, UL FRI

- Christian Kniep, QNIB Solutions
- Barbara Krasovec, IJS/Sling
- Teo Prica, IZUM/VEGA
- Alfio Lazzaro, HPE
- Francesco Cola, CINECA

.

Excused

- .

Actual topics

Future meetings

Monthly basis: Every 1st Thursday of the month: 14:00 - 15:30 CET

Wrap up from the EuroHPC containers workshop

https://docs.google.com/document/d/1krvaejetoZvfFhnwH3cHEYDGYdpjbS9CVizzq3PRJPk/edit#

Shared activities follow up:

- **K8s on HPC nodes (LUMI, VEGA, maybe LEONARDO)**
- Flexible/secure deployment of K8s nodes: Cloud hypervisor to be tested as the VM monitor:
    - Why not rootless K8s: not reliable, instead regular deployment will be used in a VM
    - Lightweight
    - Statically compiled binary
    - **AP: VEGA team (Next meeting):**
        - Reservation on ~ 2 nodes to perform the testing
        - Start the activity, sharing the source code on Github/Gitlab
- **How do we share containers among EuroHPC systems: Meta-hub registry**
    - Harbour is well known in the Grid community
    - **AP Christian:** organize a webinar next meeting
    - Hackaton can be organized in the month after
- **BK: Report from RIAC(**Research and Innovation advisory group**):**
    - BSC is pushing for E4S: https://e4s-project.github.io/
        - LUMI, VEGA, KAROLINA: use easybuild
        - LEONADO: Spack
    - Two aspects:
        - Attract more users from the industry (industrial users are used to the cloud-like resources) → disaster recovery, etc

            - Not very satisfied with the level of support from the HPC centers, e.g. they don't want to build their own containers
            - Barbara Suggested idea: provide a template, and mentioned about the EuroHPC containers forum
        - Formal RIAC meeting Next week (with all the working groups).
- **CK: Build:**
    - Spack/EasyBuild to build a single container
        - Dockerfile as the standard for the build script
    - E4S/EESSI to define the software stack (which is build using to feed Spack/EB)
    - OCI Registry to store the images
        - OCI image is the reproducible artifact for containers
    - CVMFS to store prefetch**:** If we use CVMFS for container snapshot distribution: as a foundation for sharing
- **Technical-Vendor forum for official images**
    - **AP Abdulrahman:** Bring it up for discussion in the HPC Containers advisory council meeting today
- **Gitlab for building containers on some compute nodes with the same HW as the compute nodes**
    - You may have the possibility of sharing builds and recipes
    - If there are security issues, use a VM
    - Follow up:
        - Workshop/hackathon for container building via gitlab

EuroHPC Containers Workshop UL June 2022

Minutes - Day 1 (June 30th 09:00 - 15:00)

WS site: https://events.prace-ri.eu/event/1410/

EuroHPC SupercomputersDas

https://eurohpc-ju.europa.eu/about/our-supercomputers_en

LUMI
LUMI is a pre-exascale EuroHPC supercomputer located in Kajaani, Finland. It is a Cray EX supercomputer supplied by Hewlett Packard Enterprise (HPE) and hosted by CSC – IT Center for Science.

LEONARDO

Leonardo is a pre-exascale EuroHPC supercomputer currently built in the Bologna Technopole, Italy. It is supplied by ATOS, based on a BullSequana XH2000 supercomputer and hosted by CINECA
Search for available translations of the preceding link
EN●

VEGA
IZUM
Vega

is a petascale EuroHPC supercomputer located in Maribor, Slovenia. It is supplied by Atos, based on the BullSequana XH2000 supercomputer and hosted by IZUM

- Team of 5 and have root access + IJS + Atos support
- Slovenian supercomputing network: https://doc.sling.si/en/

K8S

- Slurm Reservation on some of the nodes
- Each node has KVM/qemu//firecracker(cloud hypervisor) and VMs
    - firecracker has an OCI-compliant runtime shim
    - take into account any performance degradation: some hypervisors have very little overhead (2%)
        - firecracker/cloud-hypervisor has a very little performance degradation
    - Flexibility vs performance
    - If you care much about performance, down prioritize such tasks and put them in a queue
- Each user gets their own K8s controller (maybe per group of users)
- Use the K8s control plane to launch several k8s clusters
    - Main admin K8s cluster to host the control planes
- Need to integrate the network platform with VMs
- Users from the industry are used to cloud and K8s APIs: this use case is very interesting
- Other experience:
    - K8s for courses
    - issue: clean pods up regularly and not them be there after done using them

Runtime plans

- Plans on VEGA to install enroot and SARUS (Singularity by default, Podman for internal services)
- LUMI: Only singularity for now
- LEONARDO: Only singularity for now
    - enroot has been tested and can be enabled in the future
- Important to provide good documentation

Documentation

LUMI: https://docs.lumi-supercomputer.eu/software/
CINECA: https://wiki.u-gov.it/confluence/display/SCAIUS/Container
Vega: https://doc.vega.izum.si
https://docs.archer2.ac.uk/

Portability vs performance

- VEGA: AMD CPUs, A100 GPUs
- Leonardo: Intel CPUs, A100 GPUs
- Options:
    - **One Image w/ single fat binaries:** Multiple instruction sets in a single binary
        - PRO: submit one container and you are good

        - CONS:
            - need to build a new image when a platform is added
            - Reproducibility is reduced, b/c of runtime decisions made
    - **One Image w/ multiple targeted binaries**: Detect the CPU Arch at runtime and use the associated instruction set
        - PRO: submit one container and you are good
        - CONS:
            - need to build a new image when a platform is added
            - Reproducibility is reduced, b/c of runtime decisions made
    - **Multiple images, one image for each target**: One tag for each CPU Arch
        - PRO:
            - Image for each target system
            - less complex container
            - reproducibility
        - CONS:
            - So many image tags: maintaining a lot of images
        - Example:
            - gromacs:leonardo, gromacs:vega

Container Distribution

Once an image for each target is build, the image needs to be made available for different clusters. We discussed two ideas:

- **Use CVMFS to distribute (singularity) images via POSIX mount:**
    - Share a CVMFS: Directory structure for version control
        - Sysadmins define which directories can be seen on your cluster
        - Currently on VEGA: only easybuild modules are shared on CVMFS
- **Distribute Images w/ dynamic Registry**: docs.metahub-registry.org/ (WiP from Christian)
        - PRO
            - OCI compatible
            - Image for each target
            - decision which exact image is used is made centrally
            - Reproducibility
        -

Registry:

- Harbour, CVMFS etc.

Singularity:

- Security issues with the kernel:
    - CVEs were related to network namespaces not user namespaces (eg. CVE-2021-3715)
    - One CVE for user namespaces, but to be able to use the vulnerability you needed network namespaces enabled (cve-2022-25636)
- For Open Science Grid (USA) and EGI, the recommended settings for Singularity are written here: https://osg-htc.org/docs/worker-node/install-singularity/index.html
- fake-root building:
    - VEGA enables it on the login nodes (upon user request at support@sling.si)
    - MARCONI100 (A100, PowerPC) has it on a compute node (outside the slurm cluster, but it mounts the filesystem so that users has access to their home directory)
    - In both cases: users are added manually upon request
    - Julish has some way building images on behalf of the users: https://apps.fz-juelich.de/jsc/hps/juwels/container-runtime.html#container-build-system
        - The building of images with JSC's Build System takes place on a dedicated system that is external to the clusters. The dedicated system has different characteristics compared to the HPC machines (different CPU type, no GPUs); created images might not be optimized to the fullest extent of the targeted system.
            - The only advantage, in this case, is to build the container close to the cluster not having to upload it from laptop
    - One idea:
        - Enable building via GitLab and set down some compute nodes where users submit their dockerfiles via GitLab PRs and they are built on the nodes
        - You gain the possibility to do a lot of checks
        - https://github.com/NVIDIA/pyxis (plugin for slurm srun)

Shared activities candidates:

- **K8s on HPC nodes (LUMI, VEGA, maybe LEONARDO)**
- Follow up:
    - K8s on HPC workshop
    - Testing on VEGA
    - Workshop:
        - Slurm/K8s/firecracker
    - **AP: VEGA Admins:** to investigate how to get LUMI and LEONARDO staff access and whether some training can be made for other admins on the usage of K8s and firecracker on Slurm.
        - What VEGA gets: more hands-on for testing and piloting the solution

- **Gitlab for building containers on some compute nodes with the same HW as the compute nodes**
    - You may have the possibility of sharing builds and recipes
    - If there are security issues, use a VM
    - Follow up:
        - Workshop/hackathon for container building via gitlab
- **How do we share containers among EuroHPC systems: Meta-hub registry**
    - Test the second version of meta-hub
    - Test the integration with CVMFS, storing images as singularity files
    - **AP Christian:** organize a webinar/hackathon after summer

Ideas and action points

- Shared training/hackathon events
- Slack/rocketchat channel: https://hpc-containers.slack.com/archives/C03MR904FU5
-

Minutes - Day 2 (July 1st 09:00 - 12:00)

Container building

Dockerfile template to build user containers

Native use-case is a user who wants to run stuff from his home-directory. In a container world, the user picks a system specific Dockerfile template (like the python-one for Vega/LUMI) and adds her/his application to it.

Support containers that are widely used (support team maintains)

Example: Relion
NVIDIA image: https://catalog.ngc.nvidia.com/orgs/hpc/containers/relion

A new version is requested by the users:

A new version is needed:

- Get the CUDA base image
- Learn how to install

Official containers (by the vendors)

NVIDIA:
NVIDIA cloud: https://catalog.ngc.nvidia.com
These containers have been tested on VEGA: LAMMPS, Quantum Espresso, pyTorch, Tensorflow, NAMD, CUDA

- Added some python packages to pyTorch and Tensorflow
- Use the CUDA base image for compiling own containers

NVIDIA base images (with dockerfile): https://gitlab.com/nvidia/container-images

NVIDIA guide: https://docs.nvidia.com/ngc/pdf/NGC-User-Guide.pdf
page 27
As mentioned earlier in this document, the Docker containers on nvcr.io also provide a sample Dockerfile that explains how to patch a framework and rebuild the Docker image.

Comment: a nice thing of ngc containers (cuda, pytorch, tensorflow) is that typically the containers are optimized for nvidia HW (e.g. they include nvccl, melanox drivers, etc.), so if something like that is achievable from AMD that would of course be a huge plus; as a user one starts on a official base image and installs only the required toolsets for working on their problem

One example for CP2K: there can be a bug/issue in the official images. So it is still needed to have the original Dockerfile

Why do we need the Dockerfile/build-script for the official application images:

- In case of bugs in the official application images
    - There are cases where there are bugs
- To build the software in the way to get the best performance (arguments for compiling/building)
    - There are cases where there have been performance differences between official containers and user built containers, due to e.g. compilation/build options
- To get the best size optimized images
    - Many official images are size optimized which is an additional advantage
- To get the best integration with MPI

The wishlist for the hardware vendor (ordered by priority - lower is better)

1. Official application images (for the user requested software packages) that work out of the box on the targeted HPC systems - for applications requested by large number of users
2. Official application images (for all already published base images on the vendor's repo) that work out of the box on the targeted HPC systems
3. Dockerfiles for building the official application images

    - Example: when the official application image for the latest release is not out yet
4. Official base images and dockerfiles together with some examples to build application images on the top of those, e.g CUDA images
    - Example: Dockerfiles with the GPU/Network drivers and MPI base libraries

Comments:

- Licencing is an issue that needs to be considered
- Registries with proxies should be investigated
- CJ: Some official containers are tested on specific platforms but not on other platforms, e.g. Cray
    - Community to share the build instructions/recipes
- Bill: HPE is interested in the collaboration.
- CJ: Having resources available for different parties to do testing and validation

Action

Build a EuroHPC vendor forum

What do we have in common:

- Several EuroHPC systems have AMD CPUs
- All EuroHPC systems supporting GPUs have NVIDIA GPUs (except LUMI)
- Several EuroHPC systems are HPE (Cray and others)

How to approach:

- All user support teams have connections with the vendors as part of the service contracts
- Construct a “technical” EuroHPC collaboration forum which includes representatives of the user support teams and the vendors

Benefits:

- User support teams get better support for application containers
- Vendors get their application containers well tested, more hands-on for better performance. The may add more application containers to their catalog/repositories

Kick off:

- Shoot with the idea in the HPC containers advisory council meeting

MPI Containers

Fully containerised, Optimised MPI installation

Replicating the network stack:

- Check what kind of operating system the target is running
- Check what kind of network stack the target system has
  - Type of interconnect, software distribution type and version
- Check for any additional kernel modules enabled for shared memory transport (e.g CMA,knem,xpmem)

During container construction

- Install all the network stack components
- Install any other required libraries
- Build the MPI library against all the installed components
  - It might be required to first build some lower level communication libraries
  - Or there might be a pre-built package available from the same source as the network components.
- Compile program with installed MPI wrappers

Hybrid MPI installation

- Install MPI inside the container
  - It is safe to install MPICH for ABI compatibility (experience from NERSC)
- Build the program using the installed MPI
- Bind mount the host MPI (and all needed dynamic dependencies)
  - Mounting libfabric from the host gives a better better performance
- Start the program
  - using either `srun` or `mpirun` outside the container depending on site recommendation.
    Using this approach we will also be compatible with the startup environment (assuming that we mounted any sockets if needed.)
  - OFED container from Mellanox: https://github.com/Mellanox/ofed-docker

srun with enroot+pyxis:
https://archive.fosdem.org/2020/schedule/event/containers_hpc_unprivileged/attachments/slides/3711/export/events/attachments/containers_hpc_unprivileged/slides/3711/containers_hpc_unprivileged.pdf

Challenges

- **Containerised:** Maintenance effort
- **Hybrid:** Need to define all necessary mount points (and probably update them some times)

LUMI

- Both contained and hybrid
- Performance difference for osu (lower BW for the contained)
- Described in the documentation

VEGA

- Both contained and hybrid
- Performance tests PMI2 vs PMIx (better performance for PMIx)
- Not described in the documentation

LEONARDO

- Plan to support both
- Prefer the hybrid model

Conclusion

All systems will support both ways with different priorities

Action

Hybrid MPI:

- Develop a best practices guide for writing the bind-mount script/module
  - Better to use metadata annotation: what to annotate, e.g. MPI version
- Provide the necessary tests
- **Collaboration between LUMI, VEGA, and LEONARDO:**
  - **Define the metadata annotation**
  - **Making the best practices guide and develop the modules for binding**

Future follow up activities

EuroHPC JU plan:
https://drive.google.com/file/d/1ASx66Bzj52zxHd9kw2rVYCA902Fijt09/view?usp=sharing

- PRACE sites interests:
  https://docs.google.com/presentation/d/1iW5v6rP48AKloHoCVQpA-8cv-lV9Vlrg/edit?usp=sharing&ouid=117095049849694613185&rtpof=true&sd=true

2.1 High Level support Teams for EuroHPC systems

In the RIAG (Research and Innovation advisory group) in EuroHPC defined that this activity should be inter-site instead of having separate teams on each EuroHPC system

Next governing board meeting is in the beginning of September

Barbara is a member in RIAG representing Slovenia

Represented countries (50/50 Academia/industry) (voting members in bold): **Slovenia**, Italy, **Norway**, **Poland**, **Spain**, **Germany**, Sweden, Netherlands, Turkey, France + 2 members from the **EC**

All members:
https://eurohpc-ju.europa.eu/system/files/2022-03/EuroHPC%20JU%20Decision%204.2022.pdf

Who is interested to join a consortium:

- Slovenia: UL, VEGA
- Norway: Sigma2/UiO
- Italy: TBD (AP Francesco to ask)

Other future calls:

HPC Open software stack technologies
The aim is to develop a source of innovative HPC open software stack technology and support an
ambitious research and innovation agenda for developing a competitive and innovative supercomputing ecosystem addressing software technologies, including their integration into computing systems. A key aspect is to ensure that all the IP necessary to produce the solutions remains in the EU, effectively creating an independent European source of critical technology. This action should be linked to the action of the EuroHPC JU to develop the HPC Open Hardware
Technologies.
Indicative EU contribution for the topic is EUR 45 million.
A call may be agreed by the Governing Board in 2022 and launched in 2022.

Communication:

- Slack
- Trello board for follow up actions

Follow up after summer:

- Meeting after mid August
    - **25/08 14:00 CEST**