

Pawsey Webinar Series: *Containers in HPC* (with OpenFOAM workshops)

The use of containers has become an attractive framework for several areas of research supported by Pawsey (including bioinformatics and machine learning, among others).

Now, Pawsey supports the usage of OpenFOAM containers. You can read more in the [Pawsey documentation](#), re: **The most recent supported versions of OpenFOAM**.

Invitation for Online Training:

If you are interested in the general applicability of containers as an emerging technology, or in the usage of OpenFOAM containers at Pawsey, **we invite you to join Pawsey staff in a series of online events, to skill up with containers!**

The training is designed to be modular and fit readily into your schedule, offering short daily contacts over the course of a week. It starts with 3 online webinars of essential content about containers in HPC and finishes with 2 hands-on online workshops focused on OpenFOAM usage at Pawsey:

Webinars	1. Getting Started with Containers (Core)	Mon 25 10am	90 min
	2. Customising your Containers Experience	Tue 26 10am	60 min
	3. Scaling Research using Containers	Wed 27 10am	60 min
Workshops	1. OpenFOAM. Common production workflows	Thu 28 10am	2.5 hr
	2. OpenFOAM. Special cases and focus on your needs	Fri 29 10 am	2.5 hr

Registration:

Please register for the training at:

<https://pawsey.org.au/event/containers-openfoam/>

Contacts:

ann.backhaus@csiro.au

marco.delapierre@pawsey.org.au

alexis.espinosa@pawsey.org.au

Containers in HPC (Webinar 1): Core: "Getting Started with Containers"

Containers = software portability, data reproducibility, and improved collaboration.

Sounds great, right? But many researchers find themselves asking: "Are containers a right fit for the work I do?"

The "Getting Started with Containers" Webinar answers key questions about containers. What are they? Who uses them? When, and why?

This 90-minute webinar presents core information about containers, talks about their use on Supercomputers and the Cloud, and presents domain-agnostic examples of usage. This webinar is applicable for attendees across research areas.

More specifically, when you leave this webinar, you will be able to:

- Explain what containers are and the benefits to using them in HPC
- Download and run container images
- Transfer data between a host computer and a container
- Describe the high-level process of building a container image

The webinar is a mix of presentation, discussion, demonstrations and (optional) hands-on exercises.

This is a core webinar in a multi-part webinar series on *Containers in HPC*.

Session Prerequisites:

Participants require no experience with containers.

Basic familiarity with UNIX shell concepts is required (e.g., familiarity with command line interactions).

Containers in HPC (Webinar 2): Intermediate: "Customising your Containers Experience"

After the first workshop, you know how to run basic workflows with containers. What comes next?

This 60-minute webinar equips you with the skills needed to set up custom application scenarios using containers. The session's topics are particularly relevant to Pawsey users and include advanced container building and running interactive graphical applications. Examples are domain agnostic, and make use of popular software platforms, such as RStudio, Python/Jupyter.

More specifically, when you leave this webinar, you will be able to:

- Build customised container images

- Understand the pros and cons of building containers using different container technologies
- Run graphical applications using containers

This is an intermediate-level webinar in a multi-part webinar series on *Containers in HPC*.

Session Prerequisites:

Participants need to have achieved the learning outcomes from Webinar 1 of the series before joining this webinar.

[Containers in HPC \(Webinar 3\): Intermediate: "Scaling Research using Containers"](#)

Do you need to run applications that are characterised by distributed computing, GPU acceleration, or intensive Input/Output (I/O)?

This 60-minute webinar shows you how to set up containers to make the most out of applications in a high-performance computing (HPC) scenario. We discuss practical, real-world examples taken from a variety of science domains. No domain-specific knowledge is required.

More specifically, when you leave this webinar, you will be able to:

- Build and run MPI applications from a container
- Run GPU-enabled applications in containers
- Reduce filesystem overload by running I/O intensive applications (like OpenFOAM) using sandboxed filesystems mounted on a container

This is an Intermediate-level webinar in a multi-part webinar series on *Containers in HPC*.

Session Prerequisites:

Participants need to have achieved the learning outcomes from Webinar 1 of the series before joining this webinar.

[Containers in HPC \(Workshops x 2\): "Virtual Workshops for OpenFOAM containers"](#)

Having joined us for the 3 webinars of the *Containers in HPC* series, you are well positioned for hands-on practice with CFD research scenarios that can make use of OpenFOAM containers.

This workshop is run entirely online over 2 sessions of 2.5 hours each, providing a forum to discuss and practice the most common situations that you may face when applying OpenFOAM containers.

During these workshops, you will:

- Practise, using real use cases,

- Review important details about the use of sandboxed file systems for the reduction of I/O load for old OpenFOAM versions (avoiding the limitations of the 1M file quota)
- Build your own OpenFOAM solvers/tools and use them with an existing container
- Review important details in the definition files for building your OpenFOAM container
- Add community tools (or your own) inside a new customised container
- Participate in Q&A, including discussing tips and best practices, and
- Have opportunity to work on your own workflow – if desired – to get advice on how you can best containerise it.

Workshop Requirements:

- As these workshops apply what you learned in the previous webinars, you need to have watched Webinars 1, 2 and 3 to participate in this workshop.
- Make sure you have good internet connectivity so you can 1) visualise remote video, and 2) communicate two-way with audio,
- (Session 2 (Friday) only) If desired, bring along your own research scenario for discussion (optional)

Trialling novel formats of training delivery, this workshop will be run entirely online, with virtual connections nationwide. The lead trainer will be joined by multiple facilitators located around Australia. This means you will be supported by individuals well versed in using containers in the HPC environment, giving you the opportunity to draw from expertise both deep and broad.