

Pawsey Training Goes Remote: Experiences and Best Practices

Ann Backhaus¹

ann.backhaus@pawsey.org.au

Sarah Beecroft¹

sarah.beecroft@pawsey.org.au

Lachlan Campbell¹

lachlan.campbell@pawsey.org.au

Maciej Cytowski¹

maciej.cytowski@pawsey.org.au

Marco De La Pierre¹

marco.delapierre@pawsey.org.au

Luke Edwards¹

luke.edwards@pawsey.org.au

Pascal Elahi¹

pascal.elahi@pawsey.org.au

Alexis Espinosa Gayosso¹

alexis.espinosa@pawsey.org.au

Yathu Sivarajah¹

yathunanathan.sivarajah@pawsey.org.au

ABSTRACT

The Pawsey Supercomputing Centre training has evolved over the past decade, but never as rapidly as during the COVID-19 pandemic. The imperative to quickly move all training online — to reach learners facing travel restrictions and physical distancing requirements — has expedited our shift online. We had planned to increase our online offerings, but not at this pace or to this extent.

In this paper, we discuss the challenges we faced in making this transition, including how to creatively motivate and engage learners, build our virtual training delivery skills, and build communities across Australia. We share our experience in using different learning methods, tools, and techniques to address specific educational and training purposes. We share trials and successes we have had along the way.

Our guiding premise is that there is no universal learning solution. Instead, we purposefully select various solutions and platforms for different groups of learners.

Keywords

Online training, Virtual training, Remote training, HPC training, Engagement, Interactivity, Containers, Visualization, Australia

1. INTRODUCTION

Pawsey offers Australian researchers a diverse range of training. We provide basic computer science concepts, through to introductory and intermediate supercomputing, cloud and visualization, to parallel programming courses, GPU hackathons, and customized training for specific scientific domains and/or groups. Basic UNIX/Linux skills were also taught to prepare attendees for the hands-on training activities.

Pawsey has offered in-person training for over 12 years, at universities and research institutions across Australia. We have also offered virtual training ad-hoc as well as hybrid training (a

combination of face-to-face and virtual). While we have dabbled in different delivery modalities, our go-to approach has been in-person.

The advent of the COVID-19 pandemic dramatically changed Pawsey training. In January and February 2020, before COVID-19 restrictions in Australia, we had a fully (over)booked schedule of in-person user training, teacher professional development, student hands-on activities and on-site events, Internship Program events, and community outreach. However, in March and April 2020, physical distancing and Australian border restrictions and closures meant that all Pawsey in-person activities ceased. Pawsey staff were asked to work remotely.

Working off-site increased the complexity of creating new, online training, as we use a hands-on “whiteboarding” approach to training design and development.

Pawsey faced two main challenges when moving training online:

- Re-purposing training content. The existing two-day “roadshow,” which included hundreds of PowerPoint slides and dozens of Carpentry-esque episodes, required re-purposing.
- Re-focusing Pawsey’s trainers. In parallel with content re-creation, Pawsey trainers needed to develop or refine their virtual training skills, such as online engagement and community building. This was a non-trivial task, considering training is an add-on to the staff’s main role of working with researchers on code optimization and Pawsey resource uptake.

Pawsey set out on its training change journey.

2. RE-PURPOSING TRAINING

Pawsey staff saw the requirement to re-purpose training content into a virtual format as an opportunity to improve the content, ensuring alignment with learning objectives and learning outcomes, and incorporating best practices in (virtual) learner interaction and engagement. In this section, we briefly describe how this process was implemented for an example “core” course, *Introduction to Nimbus*, later renamed *Using the Nimbus Research Cloud*.

¹ Pawsey Supercomputing Centre, Kensington WA, Australia

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

2.1 Introduction to Nimbus: In-Person

To start, we listed the episodes of the existing, in-person *Introduction to Nimbus* training (Table 1) to assess the flow of content at a high level. We also noted additional supporting materials, such as online documentation.

Table 1. Introduction to Nimbus in-person training episodes.

<i>Introduction to Nimbus</i> training (3 hours)
Let’s talk about cloud computing
Nimbus: Cloud computing at Pawsey
How does Nimbus work for users?
First steps: making keypairs
Simplified security and networking
Launching an instance
Attaching volume storage
Maintaining your instance
Using snapshots to save time

Next, a new staff member attended the in-person training, reviewed the open code resources, and scanned the documentation to see if they could successfully launch an instance. They could not. They noted confusion and contradiction about the overall flow (what to do first, second, third), asked questions about security and jargon, etc. This feedback aligned with another new team member’s feedback, received informally, earlier.

2.2 Using the Nimbus Research Cloud: Virtual

The reviewer’s input on the existing, in-person content opened productive discussion and debate about what should and should not be in an introductory Nimbus Cloud training.

After several iterations, the team arrived at a newly designed series of high-level steps based on the flow of a new user’s tasks (Apply > Set Up > Use > Manage > Develop > Optimize > Retire) and accompanying, detailed flowcharts for each step. Figure 1 shows a working draft of the sample flowchart for the Set Up step.

Apply > **Set Up** > Use > Manage > Develop > Optimise > Retire




Figure 1. Early flowchart for the Set Up step for Pawsey’s *Using the Nimbus Research Cloud* training.

From the flowcharts, the team outlined those topics that were core versus those that were advanced or non-essential to getting an instance up and running. Only the essential topics would be included in the foundational training.

Table 2 shows the revised, online episodes for the new, two-part *Using the Nimbus Research Cloud* training as well as the modularized video recordings derived from the “live” virtual conducts.

Table 2. Nimbus virtual training & modularized recordings.

Virtual conduct (3 hours)	Short modularized recordings
Management	Introduction to Nimbus What is a Nimbus instance?
Authentication	Log into the Nimbus dashboard Review the Nimbus dashboard Create keypairs to access the Nimbus instance
Instance Creation	Create a Nimbus instance
Instance Access	Access your Nimbus instance
Storage	Set up your instance storage
Data	Transfer data to/from your instance
Software	Manage instance software

Short, modularized (edited) recordings are available for viewing after a virtual training conduct. The entirety of the workshop is usually not viewable, as the large file size and lengthy duration make them difficult to readily access and use by learners.

2.3 Design Considerations & Approach to Moving Online

During the reconstruction, consideration was given to feedback areas, including too few learner activities (beside hands-on coding), too quickly progressing from basic to intermediate concepts, and unexplained use of jargon. These latter comments are common manifestations of “expert blind spots” or “expert awareness gaps” [1]. Having forgotten what it is like to be a novice learner — unfamiliar with the language or concepts of the topic — experts may inadvertently overlook explanations and teach at what seems to be “breakneck speed” to a new learner.

A brief, but important, mention must be made here about Learning Outcomes (LOs) and backward design. LOs sit at the core of Pawsey training. Everything — training activities, discussions, coding, etc. — must contribute to and support the trainee in reaching LOs.

When designing a training course, we start with LOs, then work backwards to the activities, discussions, etc. In this way, we follow the backward design approach: LOs are formulated first and then course design follows, to determine how to assess (or validate / confirm learning) and how to teach (activities to use, etc.) [2].

During reconstruction, we rewrote LOs for each episode. Anything extraneous to those LOs were removed from content. New content was added as needed, with significant consideration and thought given to making the online content interactive and engaging [3], [4] and [5].

3. SUPPORTING SELF-GUIDED LEARNING

In July and August, Pawsey launched its suite of core online training to replace its national roadshow. For 90 minutes, each Monday morning, learners were offered free, significantly revised trainings on basic concepts in supercomputing, data, cloud, and visualization.

Pawsey offers these new virtual trainings in several ways: open enrolment, institutional requests, and domain requests. After a virtual conduct, we release videos of the events in short (5–15 minute), topical “chunks”. This modularization enables learners to readily find a specific topic or watch an entire series of recordings (See Table 2).

Learners can access all training content (PowerPoint slides, code samples (GitHub), video recordings, etc.) from the Pawsey Training Portal [6].

4. CREATING NEW INTERMEDIATE AND ADVANCED CONTENT: CONTAINERS

With an online repository of core training accessible 24/7, Pawsey staff can turn their attention to the creation of intermediate and advanced training. Topics include such areas as effective use of compute infrastructure (parallel and accelerated computing) architecture, reproducible science (containers), and advanced visualization. Focusing on intermediate and advanced training was very difficult pre-COVID, when much of the trainers’ time was spent traveling for in-person conducts of core trainings.

One advanced topic that Pawsey has focused on since moving entirely online is Containers. The *Using Containers in X* is a multi-day, webinar-workshop series focused on addressing specific needs for the targeted scientific domain.

This series has provided staff with rich opportunities to partner with domains, experiment with various online teaching techniques, and trial different peer-to-peer and community building approaches. At the time of writing, Pawsey has partnered with three domains to create tailored Containers training: computational fluid dynamics, bioinformatics, and radio astronomy. For each of these we offer core (“generic”) webinars followed by domain-specific (“bespoke”) workshops.

4.1 Creating a Baseline of Knowledge: Generic Container Webinars

To ensure that all learners have the requisite baseline to participate meaningfully in the hands-on workshops, we offer a three-part webinar series. In these 90–120-minute sessions, Pawsey trainers introduce concepts using illustration, discussion, and coding demonstrations. Individuals can practice simultaneously or while watching recordings or by using step-by-step instructions.

Table 3 shows the generic topics covered in the *Using Containers* webinar series.

Table 3. Core topics for container training (webinars).

Using Containers (Generic Webinars)
Introduction, Running applications in containers
Building containers, Setting up graphical applications
High performance containers (MPI, GPUs, I/O intensive)

During core training, learners can ask questions using the designated communications channel. We find that each domain has a “standard” or usual means of communicating, which we leverage. Through this reuse, we try to eliminate “noise” or learning distraction. When the group is large, we may use multiple communication channels, again to eliminate learner distraction. For example, we may use Zoom for learner-instructor communication and Slack for instructor-facilitator communication.

The webinar series is a prerequisite to the workshop series [7].

4.2 Building Expertise: Bespoke, Hands-on Container Workshops

The multi-day workshop series that follows the webinars makes up the tailored part of the training. In 2–3 three-hour sessions, learners are guided along a path from learning demonstrations to guided coding/application to extrapolated coding/application to mini-hack or BYO (Bring Your Own) code or pipelines. A forum discussion wraps up the series, providing an opportunity to summarize learning outcomes and share reflections and feedback.

While the overall workshop program has similarities across domains, in detail it is bespoke. The learners’ journey — its goals, its starting and ending points, and its focus along the way — comes out of a close collaboration between Pawsey trainers, domain experts and facilitators, and learners themselves.

For example, in the Containers pre-workshop Expression of Interest, we collected key inputs. Potential attendees specified their profiles/roles/locations/institutions, levels of (self-reported) expertise in Containers, frequency of Container use, and “hot topics” of interest. They could enter “blue sky” topics; however, we also guided topic selection through a list pre-vetted by domain experts and facilitators. This latter approach ensured that we had on-hand facilitators skilled in the topics being offered.

Table 4 shows the listing of topics that came out of the Bioinformatics Expression of Interest.

Table 4. Hot topics for container training (workshops).

Using Containers (Bespoke Workshops)
Containerizing a full workflow
Designing and building container images
Collaborative and reproducible research with containers and workflow tools
Transitioning from Docker to Singularity
Using a virtual file system to encapsulate large numbers of I/O files
Setting up and configuring your own Singularity installation
Simplifying your user experience by “hiding” container syntax
Using RStudio, JupyterHub or other web platforms via containers

To construct an appropriate set of topics and groupings for the workshop, we asked the 16 facilitators to self-assess their expertise per hot topic. We then grouped attendees with the appropriate facilitators.

Table 5 shows the complexity of the groupings per session. These sessions reflect not only the matchings but also our training partner’s goals for learning outcomes and community building.

Table 5. Virtual workshop details by grouping, focus, goals.

	Session 1	Session 2	Session 3
Grouping content	Learners and facilitators by institution (sub-grouping of skill level)	Learners by topic of interest mapped to facilitators with expertise (sub-grouping of skill level)	Learners by hot topics; facilitators pick topic per expertise and bring case study(s)
Focus	Detailed scripts and code / samples	Training content more guided, less prescriptive	Own pipelines, projects; problem solving
Goals	Significant guidance and scaffolding; early “wins” (for confidence); reinforce local community	Engage, motivate; community of shared interest; learners apply previous learning in a guided setting	Facilitator freedom to choose interaction; expand learner knowledge; wider community

We expected our biggest challenge to this live virtual training would be the complex, facilitator-attendee matching, but we were wrong. Laptop setup posed the largest challenge. Preparing attendees’ institutional laptops to enable access to our cloud or super computing resources caused significant distraction at course startup and in one instance became a barrier to participation. When in-person, trainers can troubleshoot setup by quickly glancing at attendee screens. Virtually, and especially when class sizes are 50 or more, troubleshooting becomes significantly more difficult and time consuming. Our initial attempts at real-time problem solving were chaotic. We trialed several techniques and decided on the use

of numerous operating system-specific breakout rooms, which we opened 30 minutes before the training started.

Finally, we debated the value of recording the workshop sessions for public viewing because of the individualized nature of hands-on workshop experiences. We decided to make workshop recordings available and let learners view if desired. The three-part, full workshop recordings are available on the Australian BioCommons YouTube channel [8]. Recorded, modularized workshop topics are available on the Pawsey YouTube channel [9].

5. DELIVERING INTERACTIVE HANDS-ON VISUALISATION TRAINING

Like other Pawsey training, visualization training presents challenges when moving online because of the preference to use hands-on exercises to best learn and practice new skills. The Pawsey Visualization Team found that the key to delivery of interactive, virtual hands-on Vis training was to redesign and re-plan the training from scratch. The result? Web-based remote visualization training.

By transforming the training to a web-based visualization focus, technological challenges were lessened, and training prerequisites minimized. The only requirement was for attendees to have a laptop/desktop with a web browser installed to perform web-based remote visualization.

We modularized the three-hour, in-person training into smaller topics across three days (one 90-minute session per day). We simplified the training slides with detailed screenshots explaining each step. Also, we published the training slides online for the students to access. This approach to slideware not only gave the students a copy of the original slides with high resolution images, but also enabled them to go through the hands-on parts at their own pace.

During the interactive hands-on parts, students were divided into smaller groups and moved to breakout rooms with facilitators. This gave the opportunity to enable two-way communication when needed, such as for questions and troubleshooting.

Planned breaks and scheduled large group Q&A sessions helped us to keep on track — despite the full learning agenda — and finish on time.

6. COMMUNICATING WITH AND SUPPORTING TRAINERS AND TRAINEES

Clear communication and support are some of the most important aspects of training and education activities for trainees and trainers. In contrast to the non-verbal cues so readily available in face-to-face training, communication and interaction in virtual training must be explicit and purposeful.

Good communication between instructor and student is key. Devoid of in-person cues (e.g., body language), trainers online rely on videos, which many trainees turn off. For this reason, we include frequent and varied checkpoints, such as virtual polls, voice Q&A, and hands-on activities for large and small groups.

We also rigorously practice “talk out loud” training and thinking. No on-screen activity, such as coding, is done without speaking about it, even when — especially when — we have a coding “glitch”. Glitches provide “learning moments” rich in impact; they are opportunities for students to watch and listen as an expert “unpacks” an issue. Such moments can build immediacy with the instructor and engagement with the content.

Pawsey trainers rarely train alone. Co-training or including facilitators, or helpers, allows the presenting instructor to focus on delivery and engagement, not on the technology or the “chat.” Facilitators also make breakout rooms possible. These small, intimate working groups are key to foster attendee participation in the conversation. Attendees are encouraged to ask questions. The platform’s chat function is mostly used, but Slack and Google Docs are also used. For example, in the *Using Containers in Bioinformatics* training, trainees filled 18 pages of an online document with robust discussion.

7. TALKING THE TALK, WALKING THE WALK

Pawsey is not only developing a new suite of virtual training targeted externally — at Australian teachers, students and Pawsey users — but also sourcing and co-designing virtual training for Pawsey staff.

The Pawsey Supercomputing Centre is in the process of refreshing supercomputing, data, storage, and networking equipment. New supercomputing systems bring new opportunities as well as challenges for both users and Pawsey staff. To prepare staff for the new infrastructure, Pawsey staff are working to a rigorous learning schedule to address identified skill gaps and to build requisite skills for the new systems.

As with the Pawsey user training, all internal staff training sessions are being conducted virtually, and we are learning numerous techniques and online learning practices, by being students ourselves.

8. RESULTS: TRAINING DELIVERY AND SCALABILITY

At time of publication, we have conducted one full round of our virtual core trainings (plus one-off requests), and we are designing new intermediate and advanced online trainings.

In Table 6 and 7, we compare attendance numbers (attendee reach) for in-person and online training for two sets of conducts: a single conduct of a two-day, in-person roadshow versus a single conduct of the replacement suite of online training, **and** the reach, to date, of the Container trainings.

Table 6. Sample single conduct attendance: in-person vs virtual.

Single Conduct – 2 days In-person roadshow		Single Conduct (10.5 hours) New Virtual Training Suite*	
Live	Views	Live	Views
20	NA**	66	448

*Note that at the time of writing, *Using Nimbus Research Cloud – Part 2* was not published on the Pawsey YouTube.

**Not Applicable

Table 7. Consolidated conduct attendance: Sample virtual training.

3 Conducts New Virtual Container Training	
Live	Views
473	1,890

Tables 6 and 7 present training reach only. They do not consider associated costs, such as travel costs / lost “opportunity costs” / trainer travel fatigue for in-person training or development costs for virtual training.

The virtual numbers report a level of scalability unattainable through in-person roadshows or through our previous method of in-person advanced training.

Previously Pawsey had run an advanced webinar on GPU programming. Attendance numbers were strong (85), when compared to in-person attendance figures. However, when Pawsey collaborates with one or more partners, we reach far wider and deeper than a “solo” event.

That our numbers show an increase in attendee reach is not revolutionary, when comparing in-person to online training. However, our increased bandwidth to focus on scaling and sustainability is new. We now have time to work with partners similarly incented to upskill learners in essential and advanced skills in super compute, cloud compute, data, and/or visualization, and we can continue to move our training program forward because we have “reserves,” that is, trainers are not experiencing travel-induced trainer fatigue.

The “what’s next” conversations in the online training space — too long postponed — are happening in earnest, and we are able to design training programs with best practice online experiences in mind — continually adding onto and refining our virtual toolset.

The opportunities being offered up through teaching and learning online are enormous.

9. CONCLUSIONS

We found that moving online wholesale is both challenging and rewarding. What it is not — is merely shifting materials to a new medium. Going online requires from content developers and trainers to employ a fresh perspective and a willingness to try techniques — and try yet more techniques when the earlier ones do not work as expected.

The positive outcomes to our move online are many — some intended and others unintentional. Attendance is more inclusive. Before, geographical barriers prevented individuals from attending Pawsey’s two-day, in-person training, held only in capital cities nationally. Now, location-specific barriers are removed.

Pawsey is reaching out more broadly and actively to find partners with which to collaborate. Finding domain, institutional, and other partners enables us to tailor our technical training to include partner-relevant examples and to focus on “hot topics” of interest to the group. These approaches increase the relevance and impact of Pawsey training, and build cloud, super compute, data and visualisation skills in our user base, and beyond.

We have found that there is no universal, one-size-fits-all learning solution. Rather, there are various solutions and platforms that need to be carefully selected for different groups of learners.

10. ACKNOWLEDGMENTS

We would like to acknowledge the Whadjuk people of the Noongar nation as the traditional custodians of this country, where the Pawsey Supercomputing Centre is located and where we live and work. We pay our respects to Noongar elders past, present, and emerging.

We would like to acknowledge all past and present Pawsey Supercomputing Centre staff who actively contributed to the development and running of our training and outreach activities over the years and to our new colleagues who are already actively contributing to designing and delivering Pawsey virtual training.

We would also like to thank all collaborating institutions especially Software Carpentry, Australian BioCommons, Pawsey Partner institutions as well as all participants of Pawsey's training and education programs.

11. REFERENCES

- [1] Koch, C. and Wilson, G. (Ed.). 2016. Expertise and Instruction: Expert Awareness Gap. In *Software Carpentry: Instructor Training*. Version 2016.06, May 2016. <https://carpentries.github.io/instructor-training/>, 10.5281/zenodo.57571.
- [2] Capps, M., Dunaway, S., et al. 2020. Backward Design: A Handy Tool for Remote Teaching. In *William & Mary Law School's Conference for Excellence in Online Teaching Legal Research & Writing*. 18 June 2020. https://scholarship.law.wm.edu/excellence_online_teaching/zoomsessions/june18/
- [3] Tobolowsky, B. F. (Ed.). 2014. Paths to learning: Teaching for engagement in college. National Resource Center for The First-Year Experience. University of South Carolina. 2014.
- [4] Leslie, H.J. 2020. Facilitation fundamentals: redesigning an online course using adult learning principles and trifecta of student engagement framework. *Journal of Research in Innovative Teaching & Learning*. 18 May 2020. Emerald Publishing Limited. DOI=<https://doi.org/10.1108/JRIT-09-2019-0068>
- [5] Dixon, M.D. 2010. Creating effective student engagement in online courses: What do students find engaging? *Journal of the Scholarship of Teaching and Learning*. 10, 2 (June 2010), 1–13. <https://scholarworks.iu.edu/journals/index.php/josotl/issue/view/159>
- [6] Pawsey Supercomputing Centre. 2020. Training Portal. <https://pawseysc.github.io/training.html>
- [7] Pawsey Supercomputing Centre. 2020. Containers Training Series. <https://pawseysc.github.io/containers.html>
- [8] Australian BioCommons 2020 Training YouTube Channel. Australian BioCommons YouTube channel
- [9] Pawsey Supercomputing Centre. 2020. YouTube Channel. <https://www.youtube.com/c/PawseySupercomputingCentre/>