



# DevSecOps on NeSI: A roadmap towards secure research platform

Michael Karich (michael.karich@nesi.org.nz), University of Auckland, NeSI, and UniServices  
Nathalie Giraudon (nathalie.giraudon@nesi.org.nz), NeSI

## ABSTRACT / INTRODUCTION

NeSI has started and progressed down the path of securely developing and operating bespoke research products (Flexi HPC, AGDR, Rakeiora, and others). This progress has required forward-thinking architecture while allowing for pragmatic solutions to novel challenges. Along this journey, we have progressed into a DevOps mindset of infrastructure and systems as code where possible and feasible. We are progressing towards DevSecOps to bring automated code testing and validation into our pipeline. We share some challenges and how we overcame them along with the tools, stacks, and processes we have used or want to use.

## ABOUT THE AUTHOR(S)

Michael Karich is Deputy Chief Information Security Officer – Research at the University of Auckland, NeSI, and UniServices. Across these roles, he strives to support and enable secure research operations across the entire data lifecycle. Before this, he held roles in HPC and cloud infrastructure, data management, and research administration. With Computer Science and IT Operations qualifications, Michael pulls from a broad background to build and enable secure yet efficient research. He is currently focusing on the enhancement of internationally known capability through governance, staff enablement, and visibility.

Nathalie Giraudon: Nathalie Giraudon holds an electrical engineering degree from ENSEEIHT, École nationale supérieure d'électrotechnique, d'électronique, d'informatique, d'hydraulique et des télécommunications (French engineering school based in Toulouse). After an exchange year to McGill University, Montreal, Canada, Nathalie has worked more than 30 years in telecommunications industry (Nortel Networks, Tait Communications). She is passionate about system design and business analysis. With her knowledge in product research and development, she is now the product engineering lead of the Collaboration and Integration team at NeSI.



Session 1, September 18, 2023, 11:30 AM - 12:30 PM

# Software Design and FAIR Principles: Bridging the Gap

Eirian O Perkins  
New Zealand eScience Infrastructure  
eirian.perkins@nesi.org.nz

## ABSTRACT / INTRODUCTION

Findable, accessible, interoperable, and reusable (FAIR) software is achievable when the design phase is driven by software engineering principles, namely software design patterns. Incorporating these principles is as significant as understanding researcher needs when implementing tools to support academic endeavours. The Aotearoa Genomic Data Repository (AGDR), built on Gen3 technology, allows researchers to browse and upload taonga species data, and has a reasonably comprehensive metadata schema in order to ensure the data are FAIR and CARE - have collective benefit, Authority to control, Responsibility and Ethics. One of the implementation challenges was to validate metadata for ingest and provide sensible feedback to the user. The validator was built in Python and takes advantage of structural and behavioural design patterns, and is loosely coupled with AGDR-specific implementations so that the tool can be shared back to the wider Gen3 community. While designing software to support FAIR principles seems obvious, emphasizing user experience over software design or vice versa may result in outputs that are neither interoperable nor reusable.

## ABOUT THE AUTHOR

Eirian is a devops and integration programmer at the New Zealand eScience Infrastructure, as well as a PhD candidate at the University of Auckland in the School of Biological Sciences and has over 10 years of combined industry and public sector experience.

# Advancing FAIR with CARE in Medical Imaging Research: A Federated Analysis Approach

Authors name(s): Eryn Kwon, Joshua McGeown, Maryam Tayebi, Leigh Potter, Samantha Holdsworth  
Organisation(s): Mātai Medical Research Institute

Authors Email(s): [e.kwon@matai.org.nz](mailto:e.kwon@matai.org.nz), [j.mcgeown@matai.org.nz](mailto:j.mcgeown@matai.org.nz), [m.tayebi@matai.org.nz](mailto:m.tayebi@matai.org.nz),  
[l.potter@matai.org.nz](mailto:l.potter@matai.org.nz), [s.holdsworth@matai.org.nz](mailto:s.holdsworth@matai.org.nz)

## ABSTRACT / INTRODUCTION

The FAIR data principles, established in 2016, have enhanced research data discoverability and usability. Within a few years, the CARE principles complemented FAIR to address ethical aspects of data sharing. Implementing these principles has been left to individual research groups, leading to a diverse operational implementation to uphold data sovereignty in New Zealand. Our goal is to increase the compatibility in outcomes.

Our proposal advocates a transformative approach to promote reproducibility, transparency, and collaboration using federated analysis. This approach allows individual research groups to retain raw data custody, while shared standardized protocols for acquiring and analysing data enable impactful research outcomes. Federated analysis facilitates collaborative deployment of tools, ensuring privacy and accuracy. Shared protocols reduce ambiguity while enhancing study comparability and generalisability.

A medical imaging research (MRI study) is used to showcase operationalisation example of a federated analysis, as well as a future suggestion of a searchable metadata database made. The federated analysis model will allow better control of the raw data, and prepare the data to be interoperable and reusable with the right consultations and agreements. This can ultimately optimize the value extraction of data from multiple research without compromising the integrity of research values.

## ABOUT THE AUTHOR(S)

Eryn Kwon completed a Bachelor of Engineering and Science (conjoint) from the University of Auckland, majoring in biomedical engineering, physics and chemistry. She started working on the forensic modelling project first as her Master's degree in Mechanical Engineering at the University of Auckland, which continued into her PhD. Eryn is currently working on a mild traumatic brain injury (mTBI) project, investigating early detection and modelling of the injury, which is a joint project between the Auckland Bioengineering Institute, Centre for Brain Research and Māta. Leveraging from her PhD experience in modelling large and rapid deformation to the complex biological structure (brain), her recent work focused on obtaining data to build foundations to mTBI biomarker detection platform. When combined with the advanced MR imaging capability of Mātai, the data will result in accurate computational model of mTBI providing objective assessment and prognosis specific to each individual and impact scenario.

# Mastering Essential Skills: Incorporating Software Development Principles in Statistical Practice

Olivia Angelin-Bonnet<sup>1</sup>

<sup>1</sup>The New Zealand Institute for Plant and Food Research Limited (PFR)

[olivia.angelin-bonnet@plantandfood.co.nz](mailto:olivia.angelin-bonnet@plantandfood.co.nz)

## ABSTRACT / INTRODUCTION

Modern statistician and data science roles involve working with big and messy datasets, developing complex analysis pipelines, communicating results through reports and presentations, and implementing software or packages for disseminating our work. A number of principles and tools from the software development community, such as data and code versioning, unit testing, or workflow management systems, can support statisticians through these activities, to ensure their work is sustainable and reproducible. However, these skills are often left out of statistical degrees. In this talk, I will discuss how these principles can be applied to a statistician role, and how mastering them is rapidly becoming a core competency for statistical scientists.

## ABOUT THE AUTHOR(S)

Olivia Angelin-Bonnet completed her PhD in Statistics at Massey University, where she worked on unravelling genotype-to-phenotype relationships from multi-omics data, with a focus on polyploid organisms. After a year as a lecturer in Statistics at Massey University, she is now a statistical Scientist at Plant & Food Research. Her research interests include Systems Biology, multi-omics data integration, the study of biological networks from a statistical and computational perspective, and the development of visualisation tools for omics data.

# Establishing a AU-NZ bioinformatics software accelerator program

Ove Johan Ragnar Gustafsson, Ziad Al Bkhetan, Sarah Beecroft, Matthew Downton,  
Dinindu Senanayake, Steven Manos

Australian BioCommons, University of Melbourne, Pawsey Supercomputing Centre, NCI Australia, New Zealand eScience Infrastructure

[johan@biocommons.org.au](mailto:johan@biocommons.org.au), [ziad@biocommons.org.au](mailto:ziad@biocommons.org.au), [sarah.beecroft@pawsey.org.au](mailto:sarah.beecroft@pawsey.org.au),  
[Matthew.Downton@anu.edu.au](mailto:Matthew.Downton@anu.edu.au), [dinindu.senanayake@nesi.org.nz](mailto:dinindu.senanayake@nesi.org.nz), [steven@biocommons.org.au](mailto:steven@biocommons.org.au)

## ABSTRACT / INTRODUCTION

The creation of quality and performant bioinformatics software is a growing need as the field of data-driven omics continues to rapidly grow and expand. A proposed response to this challenge is the creation of a software accelerator program that supports the development, optimisation and sharing of software for bioinformatics, by providing the authors with expertise, best practice guidelines, and access to HPC and cloud facilities.

This idea arises from the existing Australian BioCommons Leadership Share (ABLES) program, by providing environments on peak compute systems to test, scale, debug and enhance bioinformatics software (tools and workflows) in lock step with support and expertise in making software findable, understandable and citable (i.e. FAIR).

The ultimate vision is to accelerate the transition towards a culture of best practice bioinformatics at scale on peak infrastructures.

In this BoF we will collectively discuss and shape this proposed program, by exploring what the RSE community considers to be critical, and how we might share the outcomes of this effort across infrastructure partners in the region and beyond.

This concept is being actively pursued as a collaborative effort by the Australian BioCommons, Pawsey Supercomputing Centre, New Zealand eScience Infrastructure and the National Computational Infrastructure.

# A Nationally Collaborative Approach to Carpentries Training in NZ

Murray Cadzow, Tom Saunders, Hayden Thorne, Nisha Ghatak  
University of Otago, University of Auckland, Victoria University of Wellington, NeSI  
[murray.cadzow@otago.ac.nz](mailto:murray.cadzow@otago.ac.nz), [tom.saunders@auckland.ac.nz](mailto:tom.saunders@auckland.ac.nz), [hayden.thorne@vuw.ac.nz](mailto:hayden.thorne@vuw.ac.nz),  
[nisha.ghatak@nesi.org.nz](mailto:nisha.ghatak@nesi.org.nz)

## ABSTRACT / INTRODUCTION

Workshops from The Carpentries (<https://carpentries.org>) have been taught at individual institutions in New Zealand to teach researchers introductory digital literacy skills. Following the shift to online teaching during the pandemic, the traditional Carpentries workshop format of two days of in-person teaching has been adapted to offer shorter, modularised online workshops.

A pilot series of regular occurring workshops teaching the “Introduction to the UNIX Shell” lesson from the Software Carpentry curriculum was established. The change to delivering workshops as modules of the full curriculum has made it easier to share instructors, helpers, and learners across institutions. Delivery of shorter ‘modules’ online and kept within “core-hours” aimed to improve accessibility of workshop participation, due to location, personal commitments, or limited resources, as well as reducing ‘zoom fatigue’.

This talk will run through some of the experiences and lessons learned from the pilot series of workshops.

## ABOUT THE AUTHOR(S)

Murray is a Scientific Programmer within Research Teaching IT Support at the University of Otago. Prior to this he spent 11 years researching the genetic basis of gout and related diseases. Murray has been heavily involved in computational literacy and bioinformatic training at the University of Otago - organising Research Bazaar Dunedin and the Otago Bioinformatics Spring School. He is both a Carpentries instructor and instructor trainer. His teaching has focused on delivering digital literacy training to researchers, and the development and support of the local Carpentries community at Otago.

Dr Tom Saunders has a background in entomology and biosecurity but is now an Engagement Specialist in the Centre for eResearch at the University of Auckland. Tom organises, hosts, and delivers digital research skills workshops, and helps researchers to manage their data throughout the research lifecycle. He collaborates on the organisation and delivery of ResBaz Aotearoa and cross-institution Carpentries events as a certified instructor.

Dr Hayden Thorne is currently a Digital Research Consultant at Te Herenga Waka - Victoria University of Wellington. He is part of a small team delivering digital research support across the whole university. The team supports everything from survey and reference management tools through to an in-house HPC environment, and run regular in-house Carpentry events as well as being involved in these national collaborations.

Dr Nisha Ghatak is the Research Communities Advisor and Training Lead at NeSI. Currently based at the University of Auckland, her role focuses on growing the eResearch capability in New Zealand through training delivery and community building. She is also the Regional Coordinator and Executive Council member for The Carpentries in Aotearoa New Zealand.

# What's up, docs?

## Documentation, collaboration and community-building

Matt Plummer, Andre Geldenhuis  
Victoria University of Wellington

[matt.plummer@vuw.ac.nz](mailto:matt.plummer@vuw.ac.nz), [andre.geldenhuis@vuw.ac.nz](mailto:andre.geldenhuis@vuw.ac.nz)

### ABSTRACT / INTRODUCTION

Often seen as a “background task”, documentation can play a key role not just in supporting the reproducibility of research, but also in developing a resource that could help build a community of users. This paper will explore the multi-year, iterative development of user documentation for Te Herenga Waka’s High Performance Computing Cluster, Rāpoi, and explore its evolution from a few markdown files into a git-backed, version controlled, automatically deployed ecosystem to which users can contribute their own knowledge and use cases.

### ABOUT THE AUTHOR(S)

Matt Plummer is a Senior Research Partner in Victoria University of Wellington’s Research Office, where he works with researchers from different disciplines to facilitate collaborative projects, especially those which utilise technology in innovative and transformative ways. He’s assisted with the development of a range of open-source projects, augmented reality applications and research tools.

Andre Geldenhuis is a research software engineer based in Victoria University of Wellington’s Centre for Academic Development Andre works with academics and post-graduate students to deliver high-technology research and teaching outcomes and solutions.

# Deep learning and real-world problems: balancing opportunity with responsibility and dodging the hype trap

Brent Martin

Manaaki Whenua Landcare Research  
martinb@landcareresearch.co.nz

## ABSTRACT

Artificial intelligence is everywhere, solving (or poised to solve) the hard problems, and threatening to replace all our jobs. Against this background of hype, AI techniques such as deep learning are being quietly applied to a diverse set of domains, including difficult mapping problems such as wetlands, invasive plants and habitat loss.

Deep learning is powerful, but its not magic. For every success, there are multiple failures. Rampant optimism in its abilities has been replaced first with caution, and now is often met with undue scepticism. Proponents and practitioners of deep learning (and AI in general) can now be met with anything from “but company XYZ say they can do that for a fraction of the price” to “we’ve paid someone to do something similar before and it didn’t work – it’s all a con”.

Despite this somewhat bleak state of affairs, Manaaki Whenua has been quietly applying deep learning to mapping and other computer vision problems for over five years, with more than a dozen projects completed or underway, and several successes on hard problems. We will introduce some of the more contentious projects we have tackled, and discuss how we have managed expectations, and lessons learned from the experience.

## ABOUT THE AUTHOR

Brent Martin is a senior data scientist at Manaaki Whenua Landcare Research. He has more than twenty years of experience in artificial intelligence, machine learning and deep learning. Brent lectured Computer Science at Canterbury University for ten years, teaching artificial intelligence, machine learning, programming and software engineering, and conducting research and development in artificial intelligence-based educational software and games, as well as various machine learning projects. He also has experience in the application of AI and machine learning in various industry and public good contexts, including designing and developing investigative intelligence systems at Jade Software/Wynyard Group, and conducting machine learning projects for various organisations including Google, Transpower, Plant and Food Research, the Department of Conservation, Cacophony Project and Toitū Envirocare. In his current role at Manaaki Whenua, Brent leads research into deep learning for remote sensing, including automated mapping from satellite and aerial imagery and predator species detection as part of the Predator Free 2050 initiative and related activities. In between research roles, Brent has spent time as a software engineer, manager and architect in several private companies including Digital Equipment Corporation and Hewlett Packard.



# Hydra: Streamlining Deep Learning Project Configuration

Simon Planzer  
Manaaki Whenua Landcare Research  
planzers@landcareresearch.co.nz

## ABSTRACT

Deep learning projects often involve complex configurations of hyperparameters, data preprocessing, model architectures, and training procedures. Managing these configurations efficiently is critical for achieving optimal performance and reproducibility. The Hydra framework emerges as a potent solution to streamline the configuration process in deep learning projects.

This presentation will delve into Hydra's features, showcasing its ability to create hierarchical configurations, effortlessly switch between different setups, and integrate with popular deep learning frameworks. Attendees will gain insights into structuring configurations for different model architectures, data preprocessing steps, and training scenarios. Moreover, the talk will highlight Hydra's role in facilitating collaboration among researchers and improving project scalability.

By adopting Hydra for configuring deep learning projects, researchers and practitioners can significantly reduce the time spent on configuration management, enabling them to focus on innovation and achieving better model performance.

## ABOUT THE AUTHOR(S)

Simon Planzer is a member of the research team at Manaaki Whenua Landcare Research. His work primarily involves the development and implementation of deep learning workflows. Simon's focus is on applying these techniques to address remote sensing problems. Through his contributions, he aims to contribute to the broader understanding of how technology can intersect with environmental research.

# Introducing NeSI's Flexible HPC: Cloud native platform for enabling challenging solutions

Jun Huh ([jun.huh@nesi.org.nz](mailto:jun.huh@nesi.org.nz)), NeSI  
Chris Scott ([chris.scott@nesi.org.nz](mailto:chris.scott@nesi.org.nz)), NeSI  
Kahu Anderson ([kahu.anderson@nesi.org.nz](mailto:kahu.anderson@nesi.org.nz)), NeSI

## ABSTRACT

In 2021, the New Zealand eScience Infrastructure (NeSI) procured and integrated Flexible HPC (FlexiHPC), a new high-performance private cloud platform. It is envisaged that NeSI's FlexiHPC will provide the sector with a programmable platform for collaboration around scientific data and support a scalable approach to mid-tier HPC with national expertise complementing localised integration.

NeSI is launching its FlexiHPC with its essential building blocks, and aims to share how it helps RSE/DevOps users to deploy complex and challenging solutions that enable science. This talk will cover some of the use-cases from our early access tenants, key services and features, alongside a helicopter view of the core technology components, key partnerships and integrations. In particular, a recent use case of NeSI building a training environment will be looked at; with examples using technologies such as Open OnDemand and Kubernetes. We will also talk about DevOps practices applied to the system and examples of interacting with the cloud native platform via code, and how these best practices help develop solutions in ways that are sustainable and reproducible, and also enable collaboration.

## ABOUT THE AUTHOR(S)

Jun Huh is a Product Manager at NeSI. Jun brings his experience from start-up industries into the field of eResearch. He has been involved in genomic data management, and data sovereignty related projects in the recent past years, including Aotearoa Genomic Data Repository and Rakeiora Pathfinder project. Recently he has been focused on service design facilitation to help launching of NeSI's new Flexible HPC.

Chris Scott is a Research Software Engineer working at NeSI with a background in Computational Materials Modelling and High Performance Computing. He works with researchers around New Zealand through NeSI's Consultancy Service to help them develop and optimise code for their research.

Kahu Anderson is a Research DevOps Support Engineer at NeSI with a strong background in Information Technology. He is currently the 'pathfinder' when it comes to NeSI's new FlexiHPC Platform, providing the Support team with information and documentation to facilitate collaboration between researchers and NeSI's new platform.

## Can workflows be applied to massively parallel optimisation problems?

Alexander Pletzer<sup>1,2</sup> Sandy Elliott<sup>2</sup>, Thanh Dang<sup>2</sup>, Linh Hoang<sup>2</sup> and Chris Scott<sup>3</sup>  
<sup>1</sup>NeSI, <sup>2</sup>NIWA, <sup>3</sup>University of Auckland  
alexander.pletzer@nesi.org.nz

### ABSTRACT / INTRODUCTION

Parameter optimisation/calibration problems and sensitivity analyses arise in many areas of science and engineering. These problems often involve exploring a large parameter space, thus requiring many simulations. Such endeavours tend to be computationally expensive and can thus be highly suited to run on a high-performance computing platform. Here, we will discuss to what extent a workflow approach could help reduce the long wall clock time of parameter scans and calibration runs in the case of the Soil Water Assessment Tool (SWAT), a code that is used in hydrology for water quality assessment purposes.

### ABOUT THE AUTHOR(S)

Alex is a high-performance Research Software Engineer working for NeSI at NIWA. Alex helps researchers run better and faster on NeSI platforms

Sandy is a principal scientist who applies mathematical techniques to address water quality issues, especially predicting the effects of diffuse pollution and its control at catchment scale

Thanh is a water quality modeller whose work centres around hydrology, water quality and nutrient cycling

Linh is a catchment modeller interested in applying numerical models to predict water quantity and quality

Chris is head of the computational science team at NeSI. Through NeSI's consultancy service, Chris helps scientists become more productive on NeSI's platforms

# MetaAnalysis.jl: Development and deployment issues of Meta analysis package in Julia

Arindam Basu  
University of Canterbury  
arindam.basu@canterbury.ac.nz

## ABSTRACT / INTRODUCTION

Meta-analysis refers to the process of information synthesis where researchers ask focused questions, identify primary studies based on inclusion/exclusion criteria, abstract and synthesise study level or individual level data from the primary studies and arrive at summary estimates along with regression of variables on study level estimates and examination of biases. Meta-analyses of RCT and Epidemiological are de rigeur for evidence appraisal in Medicine and Public Health and also widely used in Economics, and Social Sciences. Julia does not have a meta-analysis package and this is the first such package MetaAnalysis.jl (can be downloaded from github repository: <https://github.com/arinbasu/MetaAnalysis.jl.git> and also submitted from Julia's Project Repository).

This package has implemented a basic meta-analysis package in Julia that allows fixed and random-effects meta analysis, generation of forest plots and funnel plots, and meta-regression and subgroup analysis. This package is undergoing rapid developments. The purpose of this presentation is to engage the audience on meta-analysis, julia, and this package, and seek feedback for further development.

## ABOUT THE AUTHOR(S)

Arindam Basu is an Associate Professor of Epidemiology and Environmental Health at the Faculty of Health at University of Canterbury, I have developed and maintained the MetaAnalysis.jl package.

# Developing an fMRI workflow – lessons learned on integration and reproducibility

Josh McGeown<sup>1</sup> and Maxime Rio<sup>2,3</sup>

<sup>1</sup>Mātai, <sup>2</sup>NeSI, <sup>3</sup>NIWA

[j.mcgeown@matai.org.nz](mailto:j.mcgeown@matai.org.nz), [maxime.rio@nesi.org.nz](mailto:maxime.rio@nesi.org.nz)

## ABSTRACT / INTRODUCTION (Up to 200 words)

Advanced MRI sequences provide invaluable data with high temporal and spatial resolution to understand how the brain can be affected by a particular condition. Before deriving scientific conclusions from a dataset collected on a cohort of subjects, one needs to process many individual images in a series of preprocessing steps before aggregating the results across subjects. Opensource toolboxes are available to help with part of this process. However, it can be challenging to chain them and scale the computations for large datasets.

In this talk, we will present the work we did to coordinate these tools in a coherent and reproducible workflow using Snakemake. We will discuss the delicate art of gluing tools with their own set of assumptions, praise the beauty of community owned open data formats and digress on our adventures in scaling up computations on an HPC while keeping portability.

## ABOUT THE AUTHOR(S)

Josh McGeown is a neuroscientist specializing in traumatic brain injury. Josh's current work leverages multimodal neuroimaging, wearable sensors, and AI/ML in an effort to create a complete picture of the brain's state after exposure to brain injury and how changes in this state relate to recovery patterns.

Maxime Rio is a data science engineer and data scientist at NeSI and NIWA. He enjoys helping researchers to analyse their data, from visualisation to probabilistic modelling.

# Repurposing GIS data: Training neural networks to identify land use types

Kevan Cote  
Eco-index®

kevincote@moosedesign.co.nz

## ABSTRACT / INTRODUCTION (Up to 200 words)

When creating neural networks for image segmentation, acquisition of large datasets for 1) training models and 2) ensuring models are reproducible, are significant challenges. One aim of the Eco-index® programme is to identify ecosystem types (e.g., wetlands, kahikatea forest) using temporally and qualitatively variable data. Our approach has been to combine available GIS shapefiles with accessible satellite and multispectral imagery to create training data for an object detection-focused neural network. Specifically, we have utilised open source Python packages (GDAL) and pre-existing Github repositories to split up the satellite imagery and train a model in a UNET framework. Our pipeline is open to a variety of data, thus providing flexibility in the types of land we can detect and the type of training data we can use (e.g., existing council/government land cover GIS data). Upcoming work includes the incorporation of additional data sources (e.g., soil moisture, elevation) alongside the satellite imagery, and the addition of a temporal element. The latter will allow for land change monitoring, and could be especially useful for tracking ecosystem restoration efforts. Adhering to the programme's Digital Public Good status, the final methodology will be packaged in a Docker container, thus making it reproducible in other applications.

## ABOUT THE AUTHOR(S)

Kevan Cote is an applied data scientist with the Eco-index® programme. In addition to his love of machine learning he has a unique background in manufacturing and bringing products to market which allows him to apply systematic methodology to machine learning applications. Seeing value in sharing and contributing the open source community, he strives to make code usable and easily adaptable by others.

# Optimising a complex next-generation weather and climate code

Wolfgang Hayek  
NIWA  
wolfgang.hayek@niwa.co.nz

## ABSTRACT

Despite impressive advances with AI, traditional weather and climate models still form the backbone of numerical forecasting with high skill. The UK Met Office, in collaboration with NIWA and other institutions in the Unified Model Partnership, is working on the next-generation code LFRic. The new model will improve support for current and future supercomputing architectures and enable global weather prediction at very high spatial resolution. Code optimisation is a key part of the development process to ensure a minimal resource footprint and reduce time-to-solution.

In this talk, I will present my work on improving computational performance of the LFRic code at the source code level using loop tiling. I will also discuss the complexities of a modern weather and climate model and current high-performance computing (HPC) architectures, which aim to balance performance and energy efficiency, and pose significant challenges for creating maintainable software.

## ABOUT THE AUTHOR(S)

Wolfgang Hayek is a HPC Research Software Engineer at NIWA, and group manager of NIWA's scientific programming group, with many years of experience in scientific computing and HPC.