

Powering science gateways with HPC: What are the options?

David Perry

University of Melbourne, Parkville, Australia, perry.d@unimelb.edu.au

The traditional user experience for High Performance Computing (HPC) centres around the command line, and the intricacies of the underlying hardware. At the same time, scientific software is moving towards the cloud, leveraging modern web-based frameworks, allowing rapid iteration, a renewed focus on portability and reproducibility, and greater accessibility. This software still has need for the huge scale and specialist capabilities of HPC, but leveraging these resources is hampered by variation in implementation between facilities. Differences in software stack, scheduling systems and authentication all get in the way of developers who would rather focus on the research problem at hand. This demonstration reviews tools that can help overcome these barriers, covering APIs like Agave that help with scheduler and data management, possibilities for single-sign on between web apps and HPC, and container technologies like Docker and Singularity that can improve portability.

Format: Demonstration, although this works as a traditional presentation as well (mixing demo with slides), and I could use 20min instead of 10 if your program allows.