



SAMCloud

Simplifying research driven cloud deployments for education and more

Geoffrey Squire^{1*}, Louis Moresi², Romain Beucher³, Julian Giordani³, Stuart Woodman¹, Carsten Friedrich¹, Rob Woodcock¹



1. CSIRO, 2. Australian National University, 3. University of Melbourne
* Contact geoffrey.squire@data61.csiro.au

Introduction

There is continued high demand for geoscience analytics research and the software and infrastructure to support it. Researchers and organisations creating the algorithms and products share a common need to foster the publication, sharing, use and citation of their research and outcomes. A collaboration between the AuScope GRID team at CSIRO and the Underworld team at the University of Melbourne resulted in a prototype system that can simplify and automate the configuration and deployment of research software services and applications in the cloud.

Overview

SAMCloud provides a simple web interface where researchers can select the software bundle to be deployed, configure it according to their needs, select the compute and storage resources required, specify the lifetime for the deployment and manage it once it has been deployed. The system takes care of deploying and configuring the software in the cloud, along with the related infrastructure (public URL, TLS, routing, scaling etc).

The Vision

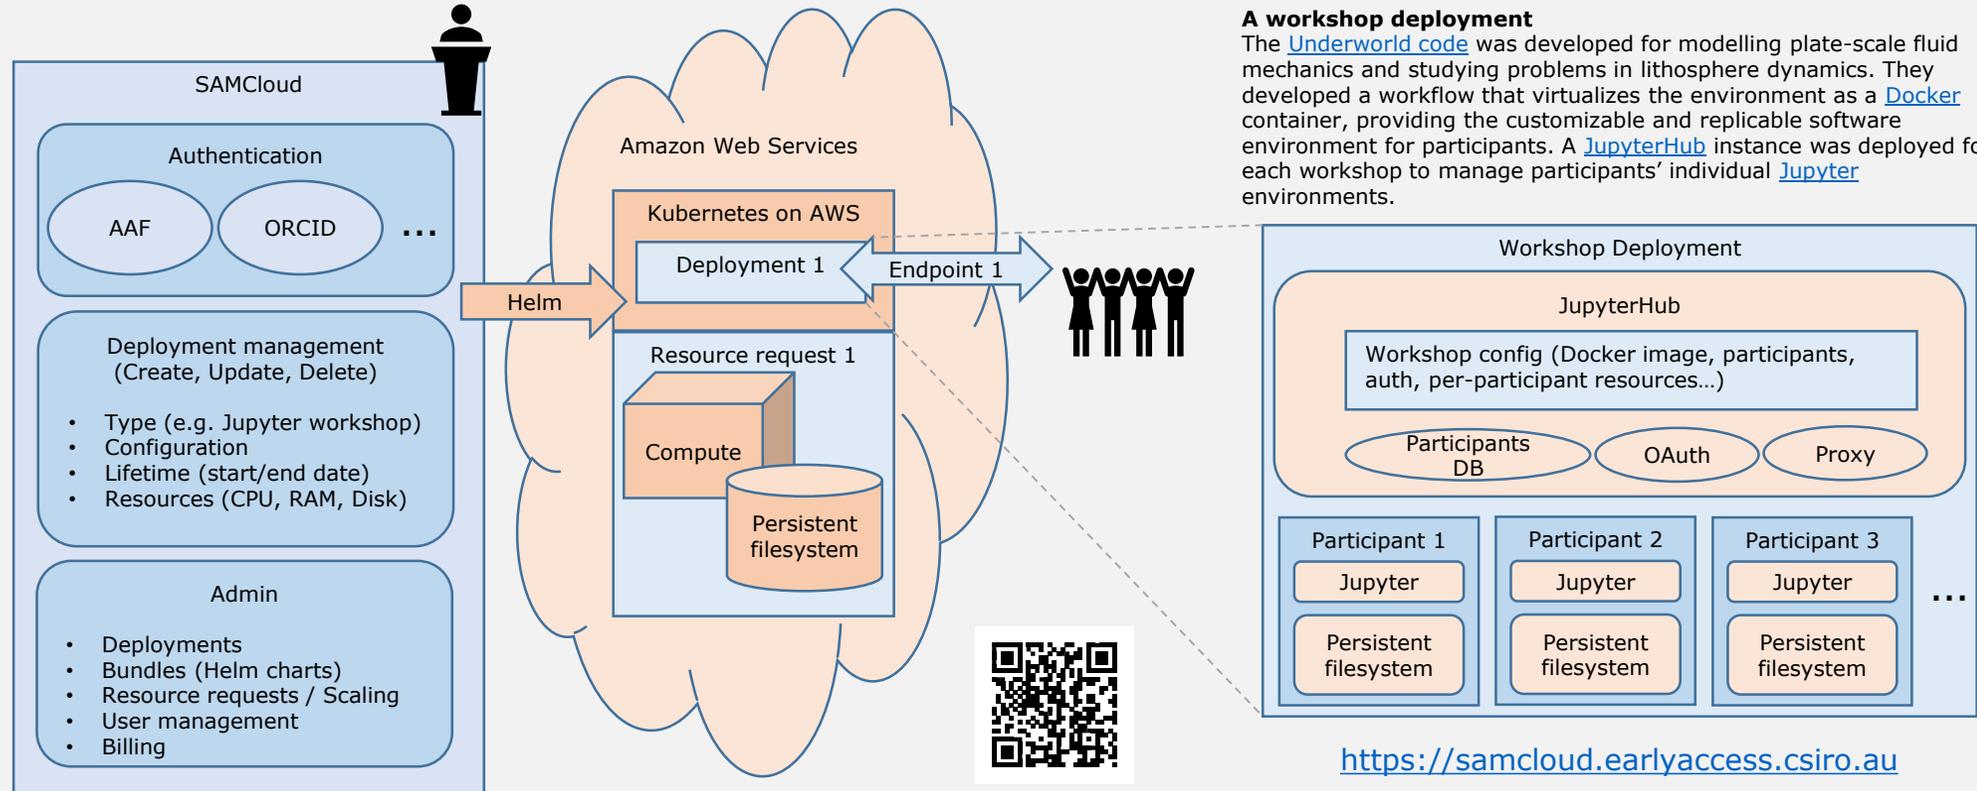
While this prototype was used for running workshops, the system as developed is more general. It can be used to deploy long-running services or applications into the cloud from a curated catalogue. Ongoing work is exploring how it can be applied to other research software and for different use-cases, including running distributed workloads and providing access to large-scale data.

Automating deployments

A deployable bundle consists of a description and a [Helm chart](#). [Kubernetes](#) and [AWS](#) provide infrastructure that is flexible, programmable, robust, secure and scalable with constraints.

Configuration options and requirements are described so that the system can provide the appropriate interface and automation for setting up the deployment. Computed configuration values can be used to simplify setup or facilitate integrations with third party systems (e.g. guided setting up a new AAF or OAuth app for the service).

A Helm chart allows us to package a complete service or other software environment for deployment into a [Kubernetes](#) cluster. Along with the configuration required for SAMCloud, we can make specific configuration options available to users.



The prototype: an Underworld workshop

Requirements

- Student environment easily customizable by the instructor
- Identical software stack for every workshop participant
- Participants can take and reuse course environment
- Cross-platform and cross-device access by users
- *Robust authentication and isolation for users*
- *Scalability, a few students in class to thousands remotely*
- *Straightforward for the instructor to set up a workshop*

A workshop deployment

The [Underworld code](#) was developed for modelling plate-scale fluid mechanics and studying problems in lithosphere dynamics. They developed a workflow that virtualizes the environment as a [Docker](#) container, providing the customizable and replicable software environment for participants. A [JupyterHub](#) instance was deployed for each workshop to manage participants' individual [Jupyter](#) environments.