

Desperately Seeking Supercomputer

The team at RSE Sheffield utilised Alces Flight on Amazon Web Services (AWS) to create a temporary supercomputer for a two-day training course in MPI for less than forty dollars.



Customer Profile



Company	RSE Sheffield
Industry	Software Engineering
Country	United Kingdom
Focus	Education Research
Website	http://rse.shef.ac.uk/

Business Need

In order to carry out a two-day MPI training course the RSE team required access to 128 cores and creation of 35 individual training accounts. This needed to be configured to mimic the The University of Sheffield’s on-premises ShARC Supercomputer installed by Alces in 2016.

Solution

The RSE team subscribed to the Alces Flight Community Edition product available via Amazon Web Services (AWS Marketplace). Using the template available within the subscription the team configured their requirements and then supplemented this with local configuration settings to create a familiar, scalable training environment.

Benefits

- Allowed the creation of a complete HPC cluster to serve a specific, short-term purpose
- Enabled scalability of research compute facilities to deal with peak workloads as required in order to reduce cost
- Exploited public cloud to deal with an immediate problem over a protracted engagement with on-premises machines
- Kept strict focus on keeping costs low due to limited initial start-up budget

“Exploiting the ability to replicate an on-premises machine, but with scalability, meant that our students were getting access to similar resources while we could save on costs. This combination is a win for public cloud and for those wanting to learn more about HPC.”

Mike Croucher, Co-founder, RSE Sheffield and Head of Research Computing at the University of Leeds

Solution at a glance

- Public Cloud Templates
- Training
- High Performance Computing

RSE Sheffield specialises in improving research software in a range of areas. One of the many services they provide is hosting training courses for external providers. One of these training sessions, focused on MPI, required supercomputer access. Problem was their on-premises HPC cluster wasn't available to host all the delegates.

The RSE team at Sheffield had set a two-day training course on parallel programming using MPI for an external group focused on high performance computing (HPC). The date, attendees, and materials were all in place - what was left out was securing access to their supercomputer, ShARC, for the duration of the course.

Since its 2012 inception at a Software Sustainability Institute Collaborations Workshop, the field of research software engineering (RSE) has exploded onto the HPC scene as a driving force in creating faster, more efficient, and highly optimised code for universities and institutions. In this day and age of research becoming increasingly more complex, the need to have an RSE engineer working alongside the subject matter expert is imperative for ensuring project success. "That's not to say that things don't go wrong in the process," explains Mike Croucher, Co-Founder of RSE Sheffield. "I often quote a law which roughly states that mistakes are inevitable and while planning for this training course I once again proved correct. In short, I forgot to ensure supercomputer provision for a supercomputing course. Thankfully public cloud is perfect for just this sort of situation.

Building a temporary supercomputer

Croucher and his team took stock of their requirements and then turned to Amazon Web Services (AWS) and Alces Flight to turn them into reality. "I wanted a cluster that would scale with the work. I knew with the course that the group would likely peak at 128 virtual cores but for the most part we'd actually need just 4 worker nodes running. My past experience with Alces and AWS showed scalability wasn't an issue and that I could configure this quickly with the Alces Flight service." Croucher's focus was also on cost. "Because our forecast initially included free access to ShARC we had no additional budget set aside. Thankfully the Community Edition of Flight Compute is free so we only had to focus on ensuring we got the best price from AWS."

Ready for flight

Thanks to Croucher's experience with using Alces Flight Compute combined with his RSE know-how the Community Edition was all he required to get his class going on AWS. "We were so pleased to find out that Mike and his team got what he needed from our free tool," said Cristin Merritt, Partnership Manager at Alces Flight. "Our aim in creating a no-cost subscription was so that anyone, anywhere with an interest in HPC can be up and running on AWS in minutes. It was a real pleasure to see that Mike did this, and more, for his training course."

The RSE Sheffield team not only succeeded in teaching their MPI class, they also managed to keep their compute costs low. "Alces Flight Compute integrates with AWS Spot pricing. We were able to determine the best price available from AWS and tailor our bid through the Alces Flight service, so the end result was less than forty dollars spent for the entire two day course. Believe it or not, this was less than I spent on coffee!" Croucher documented his experience as well as created a GitHub repository for anyone interested in replicating his work. Links can be found on his blog: www.walkingrandomly.com/?p=6392

"We were able to achieve our training goals and keep costs down. I was pleased to find out that the HPC cluster cost less than the budget set for coffee."

Mike Croucher,
Co-Founder, RSE Sheffield and
Head of Research Computing at the
University of Leeds

Products & Services

Services

Self-service documentation

Community-driven support

Products

Alces Flight Solo Community Edition

Cloud Service Provider

Amazon Web Services

View all Alces Flight success stories at alces-flight.com/success

