# You have to Break It to Make It: How On-Demand, Ephemeral Public Cloud Projects with Alces Flight Compute Resulted in the Open-Source OpenFlightHPC Project

Extended Abstract

Wil Mayers
Alces Flight Limited
Bicester, United Kingdom
wil.mayers@alces-flight.com

Mark Titerenko
Alces Flight Limited
Bicester, United Kingdom
mark.titerenko@alces-flight.com

Steve Norledge
Alces Flight Limited
Bicester, United Kingdom
steve.norledge@alces-flight.com

Cristin Merritt
Alces Flight Limited
Bicester, United Kingdom
cristin.merritt@alces-flight.com

David Marchant
Alces Flight Limited
Bicester, United Kingdom
david.marchant@alces-flight.com

Stuart Franks
Alces Flight Limited
Bicester, United Kingdom
stuart.franks@alces-flight.com

## ABSTRACT

Over three years ago the Alces Flight team made a decision to explore on-demand public cloud consumption for High Performance Computing (HPC). Our premise was simple, create a fully-featured, scalable HPC environment for research and scientific computing and provide it as a free subscription service to researchers in order to study how they would approach and consume this then new technology. This tool, Alces Flight Compute, would set out to chart how far away from the traditional bare-metal platforms our subscribers were willing to go. What we did not expect was that to get to their destination, our users would proceed to take our tool apart. This deconstruction has resulted in a new approach to HPC environment creation (the open-source OpenFlightHPC project), helped us better understand cloud adoption strategies and spending models, and handed over a set of guidelines for those looking to bring public cloud into their HPC solution.

## CCS CONCEPTS

• **Computer systems organization → Cloud computing**; • **Software and its engineering → Cloud computing**;

## KEYWORDS

Clouds and Distributed Computing, Problem-solving environments

**ACM Reference Format:**
Wil Mayers, Mark Titerenko, Steve Norledge, Cristin Merritt, David Marchant, and Stuart Franks. 2019. You have to Break It to Make It: How On-Demand, Ephemeral Public Cloud Projects with Alces Flight Compute Resulted in the Open-Source OpenFlightHPC Project: Extended Abstract. In *Proceedings of Scientific Computing (SC'19).* ACM, New York, NY, USA, Article 4, 3 pages.

## 1 INTRODUCTION

Back in 2016 the concept was simple:

- Create an HPC environment that works across public cloud platforms;
- Put it out with a solid number of open-source applications to use;
- Make the subscription free so that anyone could approach and consume HPC via public cloud;
- Record how people work through their ephemeral (temporary) projects and;
- Seek permission to deep dive with those cases willing to go public with their findings to benefit the HPC community as a whole.

The reality was different. The Alces Flight team created, distributed, and gathered adoption data via the Alces Flight Compute tool - but what we found was that the approach and use of our tool was much more nuanced and, many times, required the user to modify the Flight Compute template in order to gain the result they were after. This deconstruction resulted in our decision to move our project out of a subscription model and into the open source OpenFlightHPC project where users could readily customise the environment to their preferred platform. Along this journey we also:

- Created a sliding scale of public cloud adoption[1];
- Identified the characteristics of what entails the current ideal cloud project and;
- Opened us to the insights on where a project might balance their spend on hardware and public cloud.

[1] See: Figure 1, page 2


*SC'19, November 2019, Denver, Colorado USA*

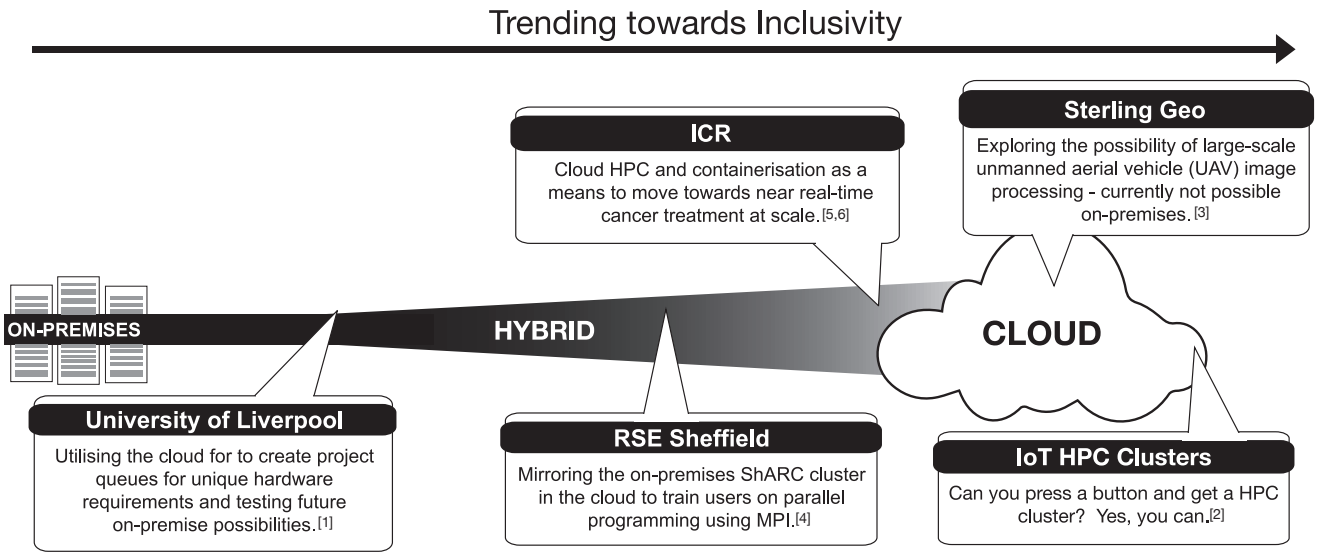


**Figure 1: CCase studies plotted in a 'sliding' format denotes how far from traditional HPC the project was / how close the project was to being cloud. [5]**

[6] [3] [1] [4] [2]

## 2 METHODOLOGY AND FINDINGS

Our research of cloud HPC adoption via Alces Flight focused in four areas:

### 2.1 How 'cloudy' was the project?

Of the 241 sites actively running Alces Flight Compute over the three years of the project we were able to publish public case studies on 5 completed projects. These studies were plotted based on the following criteria:

- Closeness to replicating an on-premises HPC cluster.
- How focused the project was on exploiting public cloud.
- If the project was aimed at using HPC in a new or novel way.

### 2.2 What are the cloud spending models?

Our original research prior to the launch of the subscription leaned to the prospective conclusion that spending patterns would shift gradually into public cloud spend. We have since learned that the community is less likely to shift slowly shift, instead spending would come in defined periods, or epochs, which can be identified as:[2]

- Significant on-premises spend, with public cloud for testing.
- Near 50/50 focus, with public cloud used to expand HPC access for training or new fields.
- Stronger cloud focus, usually for new ways of using HPC.

### 2.3 What projects are ideal for cloud HPC?

Of those sites that allowed both formal and informal survey we were able to hone in on the characteristics that denote the ideal cloud HPC project. Those characteristics are:

- Temporary/Ephemeral
- Risk is moderate or high (due to unknown level of spend)
- User-focused (no heavy Administrator requirement)
- Open-source in design of application
- Embarassingly parallel
- Platform agnostic (ideal platform not yet identified)
- Strictly time banded in their run

### 2.4 One-size fits all?

In every public case study we can report some level of customisation took place, making a subscription model unfeasible long term. The modifications completed were as follows:

- Modified to allow more shared user access;
- To mirror existing systems and arrangements;
- To make the tool platform agnostic;
- To create some amount of IP.

## 3 CONCLUSION

Our time spent with Alces Flight Compute gave us a good insight into the uptake and consumption of public cloud for HPC. It is through the interaction with public case studies that we were able to conclude that a move to open source would prove more beneficial for the community as public cloud becomes more common in HPC solutions. This project, OpenFlightHPC, is now aimed at assisting users in building a cluster environment regardless of platform and within the consideration of architecture, platform, environment, and workflow.

## ACKNOWLEDGMENTS

The authors would like to thank Brendan Bouffler from Amazon Web Services (AWS) for his collaboration efforts.

[2]Within Figure 1 it is important to note that closer to traditional HPC the project was the more conservative the cloud spend model.

## REFERENCES


[1] Cliff Addison. 2018. The Liverpool Barkla Cluster: Exploring HPC Hybridisation. https://www.newswire.com/news/dell-emc-and-alces-flight-with-aws-to-create-a-unique-hybrid-high-20089263

[2] Matthew Berryman and Adrian White. 2017. Real-world High Performance & High Throughput Computing on AWS. https://www.youtube.com/watch?time_continue=10&v=d3uz9ZxSPZA

[3] Phil Cooper, Cristin Merritt, and Wil Mayers. 2017. Breaking the boundaries, saving lives: How cloud can prevent future crisis. https://medium.com/@alcesflight/breaking-the-boundaries-saving-lives-how-cloud-can-prevent-future-crisis-c0b8aa232785

[4] Mike Croucher. 2018. Bespoke High Performance Computing Clusters in the Cloud with Alces Flight. http://www.walkingrandomly.com/?p=6392

[5] Wil Mayers. 2016. Near Real-Time: How Cloud is Shaping Cancer Treatment. https://medium.com/@alcesflight/near-real-time-how-cloud-is-shaping-cancer-treatment-e26fd2576f2d

[6] Peter Ziegenhein, Igor N. Kozin, Cornelis Ph Kamerling, and Uwe Oelfke. 2017. Towards real-time photon Monte Carlo dose calculation in the cloud. *Physics in Medicine & Biology* 62, 11 (2017), 4375. http://stacks.iop.org/0031-9155/62/i=11/a=4375