On demand HPC

(or – HPC clusters on your terms)

@CIUK18 – December 2018

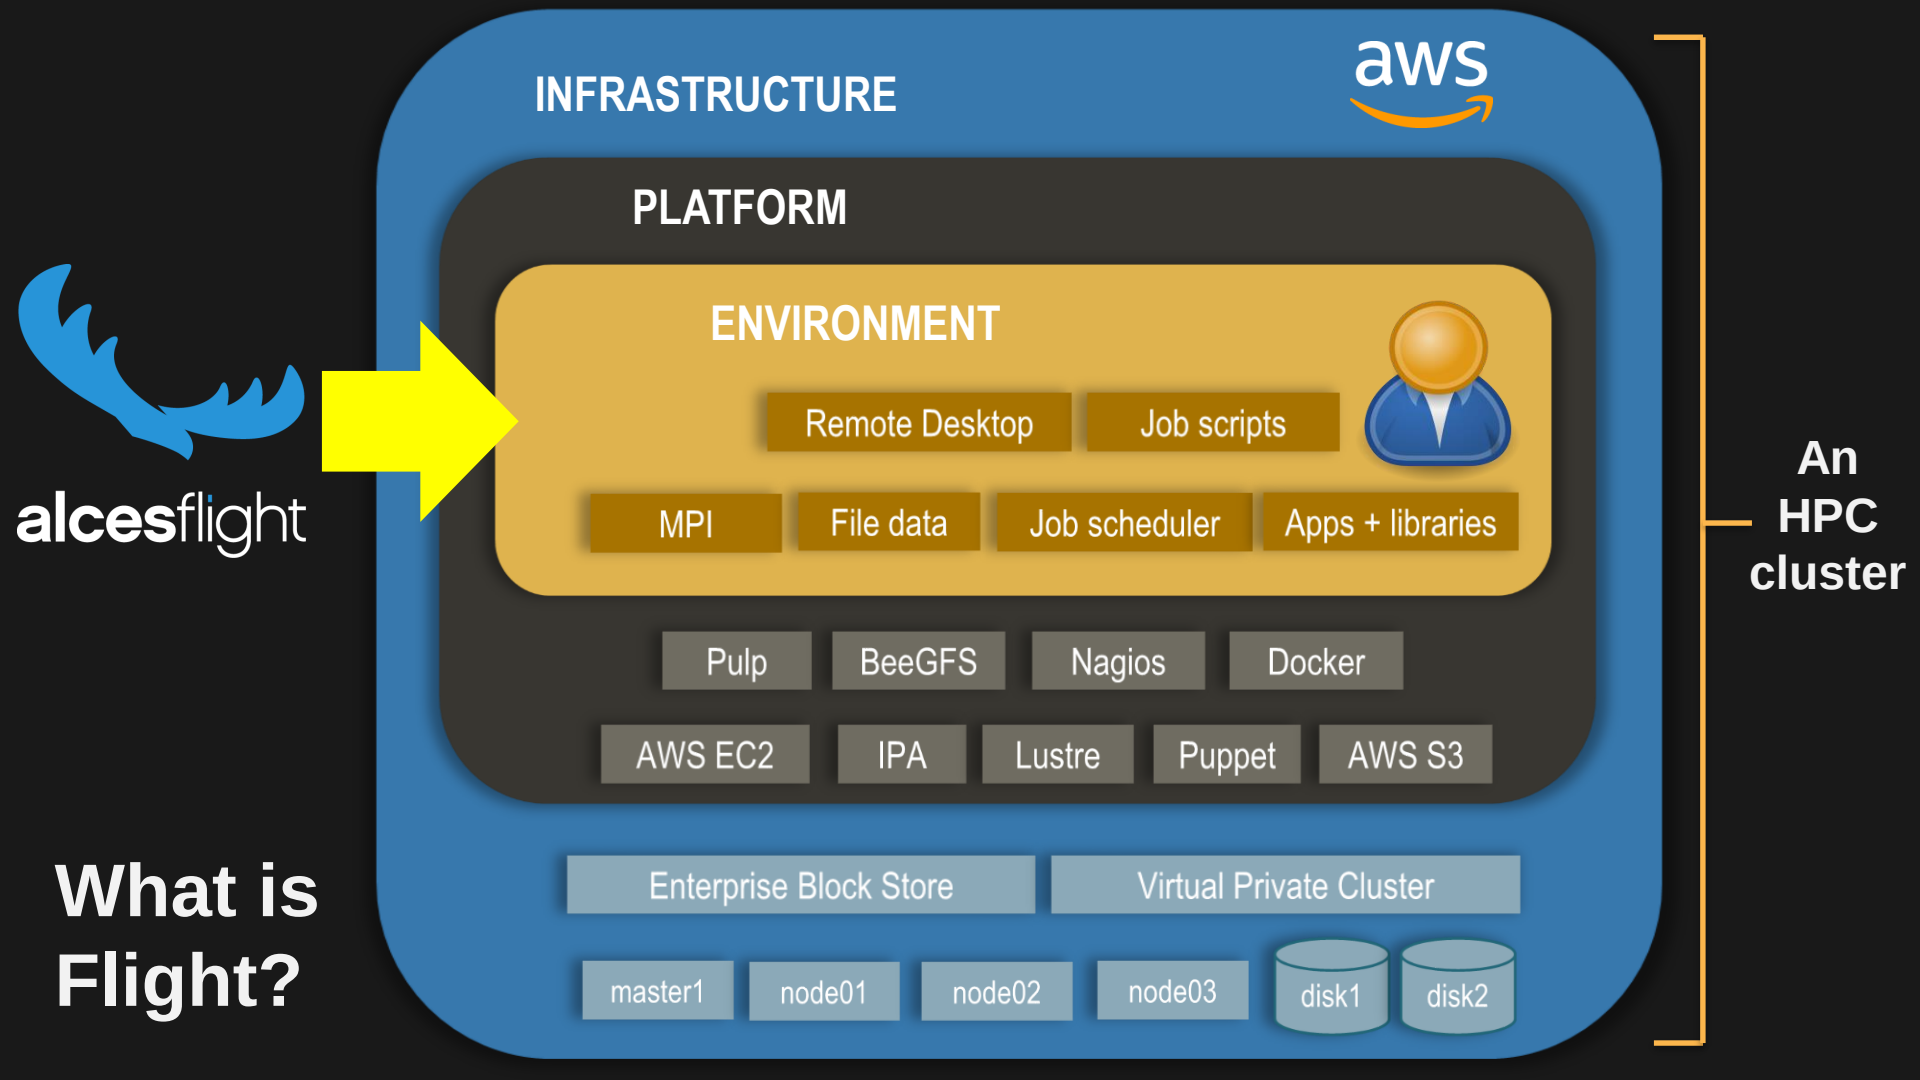
alcesflight
What is Flight?
INFRASTRUCTURE
aws
PLATFORM
ENVIRONMENT
Remote Desktop
Job scripts
MPI
File data
Job scheduler
Apps + libraries
Pulp
BeeGFS
Nagios
Docker
AWS EC2
IPA
Lustre
Puppet
AWS S3
Enterprise Block Store
Virtual Private Cluster
master1
node01
node02
node03
disk1
disk2
An HPC cluster

What does Flight provide?

- One-click HPC cluster environment
- CentOS based Linux cluster
  - *Genders* and *PDSH* utility
  - Collaborative graphical desktop
  - Storage management tools
  - HPC job-scheduler
- Curated HPC user assets
  - Job-scripts, data locations
  - Application repository
  - Libraries, dependencies, etc.
  - Data sync with persistent storage

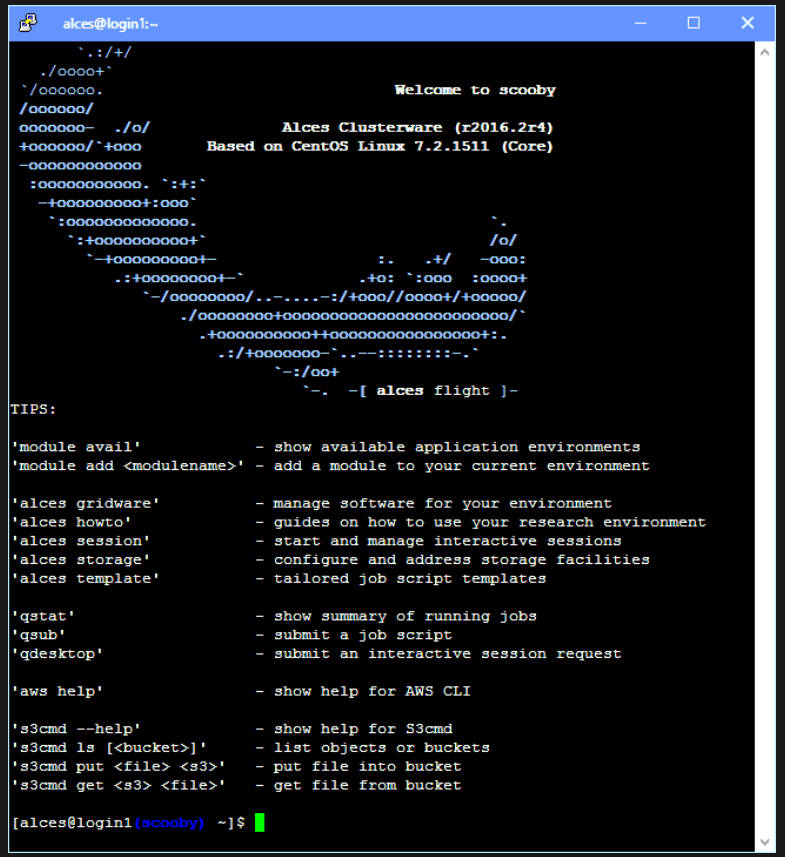

Alces Gridware Application library

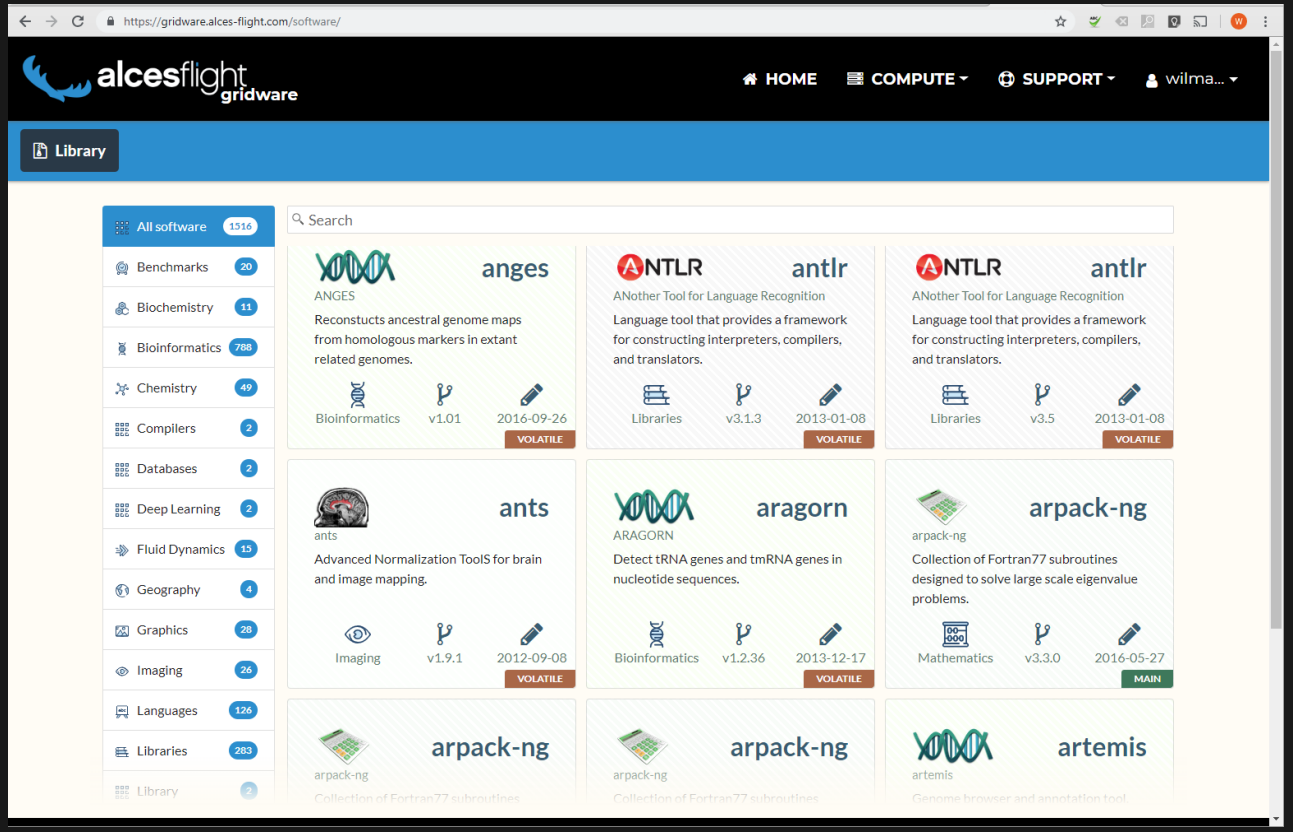

- Over 1,500 apps, libraries and MPIs
- Includes *modules* environment management
- Support for commercial and licensed applications

Flight Connections program

alcesflight

- Connect your favourite ISV applications to your clusters
- Manage subscriptions and licenses from a single point
- Includes support for:
  - Docker containers
  - Singularity containers
  - Easybuild environment

Self-service HPC on AWS and Azure Marketplace

Alces Flight - auto-scaling HPC clusters instantly ready to compute

- Run your own HPL benchmark on an Alces Flight cluster!
- Single-user, self-managed HPC clusters on-demand
- All you need is a few AWS account to launch

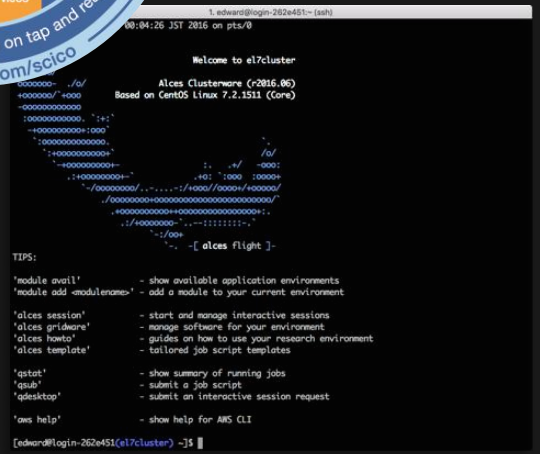

- 1500+ popular scientific applications
  - Pre-installed & ready to run
  - Multiple app versions, complete with libraries and various compiler optimizations, ready to run
- Available via AWS Marketplace

aws

alcesflight

Why use flight?

- Convenience
  - Build a cluster quickly and easily
  - Scale to any size simply at launch or automatically while running
  - Repeatable process for training and education
- Performance
  - Flight can optimize your applications automatically
  - Choose cloud platform features for HPC; e.g.
    - AWS enhanced networking for EC2 instances
    - Instance Placement groups to reduce MPI latency
    - EBS-optimized for best storage performance

Worldwide sites using Alces Flight

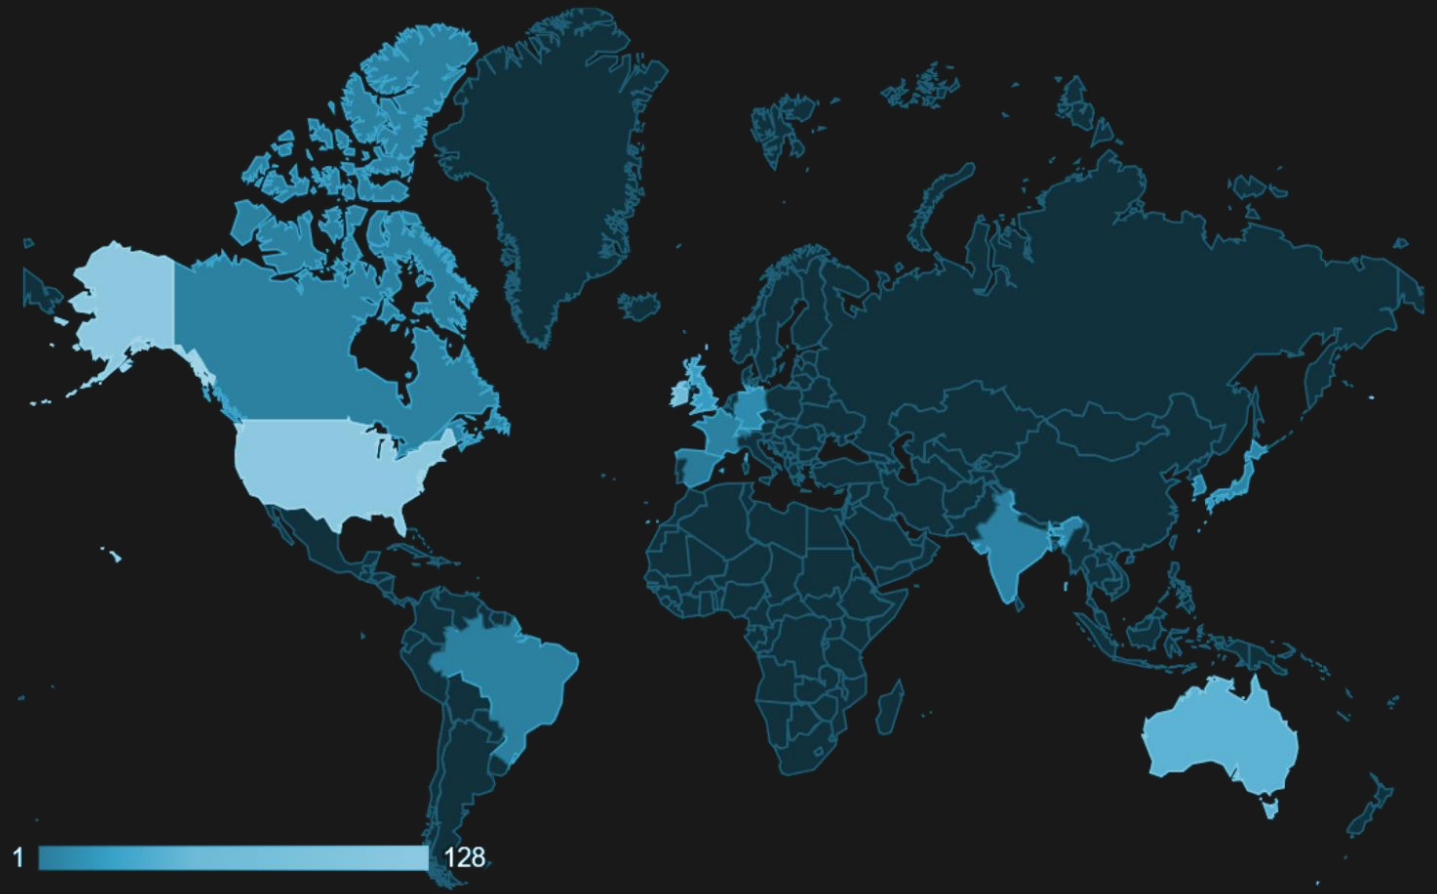

CIUK Top500 performance run

- Used 320 x SkyLake 36C 3Ghz nodes (AWS c5.18xL)
  - 25Gb interconnect between nodes
- Timescales
  - Used an existing AWS account with high EC2 region limits
  - Alces Flight cluster launched from AWS Marketplace in 19m 03s
  - 2m 40s setup time
  - 4m 15s cluster launch time (to SSH login)
  - 12 minute SPOT acquisition time for nodes

alcesflight

TOP 500 The List.

CIUK Top500 performance run

- Installed *Benchmark* depot using *Alces Gridware*
- Downloaded job-script and HPL.dat file from AWS S3
- Submitted job to Slurm job-scheduler
- Achieved 654,124 Gflops
  - Between #413 and #414 on Top500 list
- Run on AWS in spot market
  - Average node cost $1.54/hr*

* based on c5.18xlarge spot rate in EU-West2 region; excludes sales tax where applicable

* November 2017 Top500 list

Approximate costs

- 1 iteration using 11,520 cores taking 3 hours
- Total time:
  - Cost of 320 x C5 instances AWS spot market* = $1,479
  - Cost of one C4 login node for 3 hours* = $5.76
  - Cost of 100GB EBS volume for apps* = $0.02
  - Alces Flight software cost = $0.00
- Total cost per iteration = $1,485*

* based on c5.18xlarge spot rate in EU-West2 region; c4 on-demand instance; EBS st1 volume; excludes sales tax where applicable

alcesflight

If all you run is HPL…

- 24x7 usage, 100% load, all nodes
  - Public cloud is more expensive than on-prem
- Monday-Friday, 9am-5pm
  - Public cloud is 25% cheaper than buying a physical cluster
- Most HPC workloads are less demanding than HPL
  - Use Flight to auto-scale your cluster to match demand
  - Choose different node types for more flexibility and agility
  - Add-in GPU and FPGA devices when you need them
  - Launch a cluster when you need it and destroy it when done

The Smart Bit

- Use the correct platform for your work
  - No platform is good for everything
- Value your own time
  - And that of your users, researchers, scientists, accountants….
- Agility is your friend
  - Follow best software-engineering practice
  - Create and reuse repeatable solutions
  - Help your users manage change over time
- Economics are driving the cloud revolution

What about the innovative stuff?

- Flight allows you to do new things
  - Quickly install hard-to-get-running applications in a cluster
  - Try out new job-schedulers
  - Create your own development platform to try new things out
  - Try out new hardware
    - Latest GPUs, FPGA devices, latest processors
  - Give users (or groups) their own HPC clusters
    - Move workloads to different physical platforms
    - Train up new users and administrators in a safe environment
    - Maintain compatibility between many different cloud platforms

Commercial customer public cloud trial

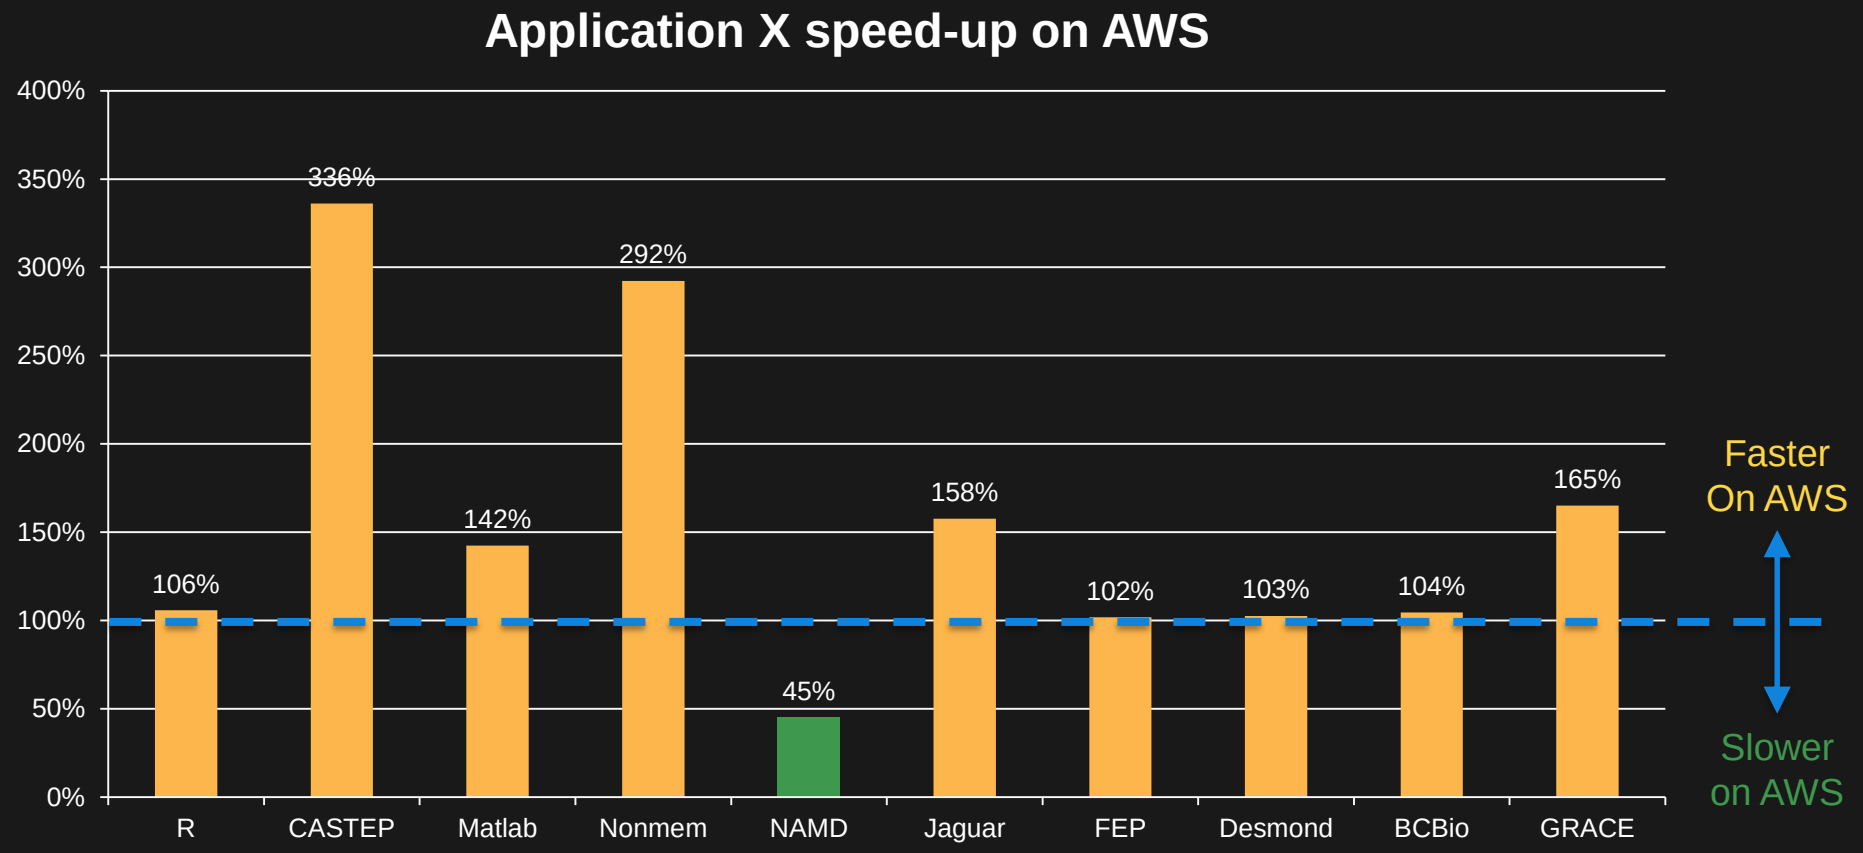

Success stories

- HPC simplicity
- Hybrid solutions
- Cloud-first agility

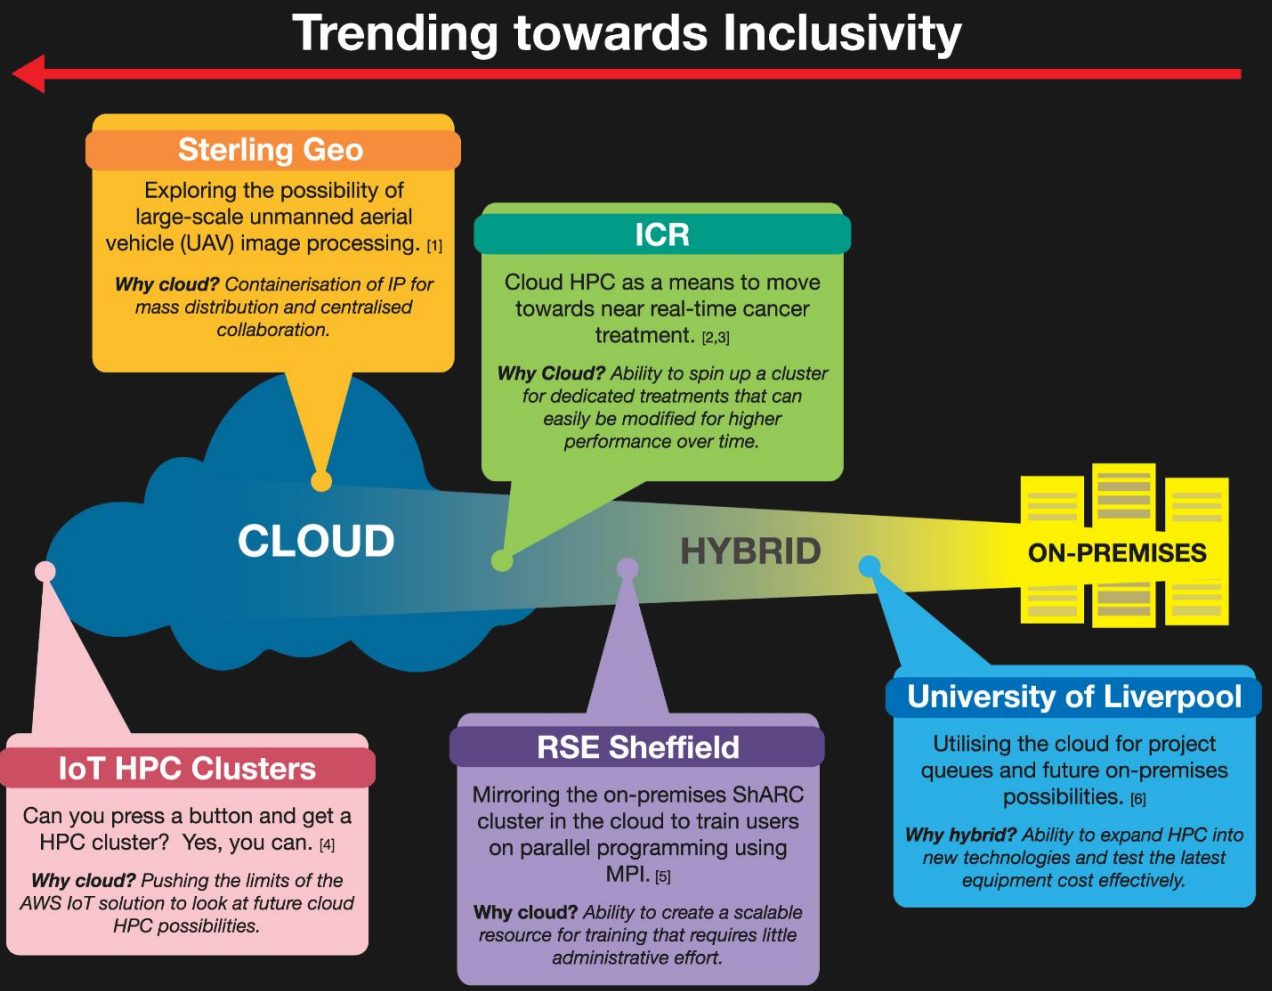

alcesflight

Your Christmas holiday project

- Try out HPC on public cloud
    1. Sign-up for a free public cloud account on Azure or AWS
    2. Subscribe to the free Alces Flight Community product
    3. Launch a cluster
    4. Install an application
    5. Run a job

Questions?

alcesflight