Supercomputing by Subscription

Wil Mayers

wil.mayers@alces-flight.com

alcesflight

Introduction to Alces

- UK based; founded in 2008
- Builds HPC and AI systems
- Works in partnership
- Subscription model

Our goals

- Simplify Supercomputing
- Reduce waste
- Inspire new users
- Make a difference

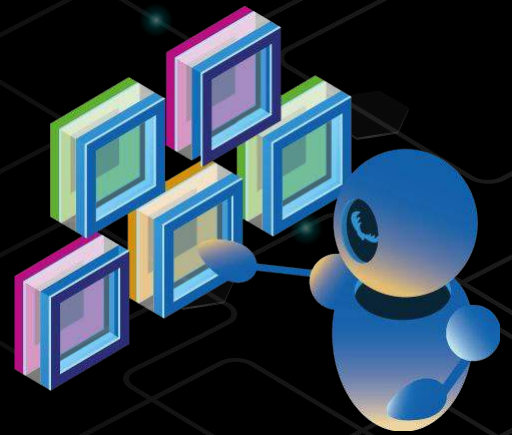

Things we worry about

(when deploying a new system)

alcesflight

Challenges deploying a new system

- Choosing the wrong hardware
- Taking too long to get into service
- Paying too much for power/cooling
- Our part in climate change
- What is it that my users want?

Cost: Spend more = better system?

- Many different hardware options to compare
  - Range of CPU families available
  - Interconnect: What do you actually need?
  - To GPU or not to GPU?
  - No amount of storage capacity/performance is ever enough
  - Resiliency costs money
- We often forget about software costs

Location: does it matter?

- Communications latency
- Carbon impact
- Data sovereignty
- Length of commitment
- System management impact

Time: Delays cost money

- Deciding what researchers need
- Running a tender process / purchasing
- Committing to multi-year agreements
- Manufacturing delays
- Data-centre infrastructure preparation
- Deployment, configuration + testing

Your team: Your most valuable asset

- End-users are your customers
  - They will need help to get the most out of any system
- The value delta
  - 100 x users needs…
  - 6 x Research software engineers, needs…
  - 2 x Infrastructure system engineers

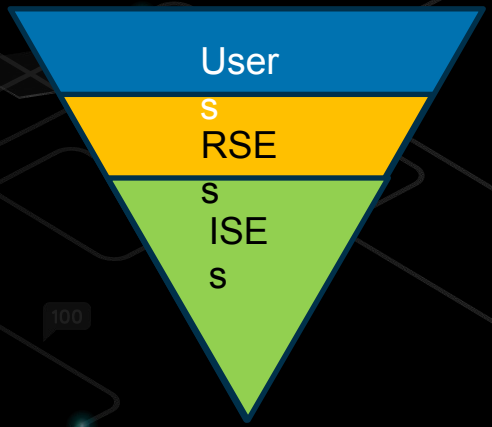

Annual Costs per GFlop

Synthetic Floating Point Benchmark

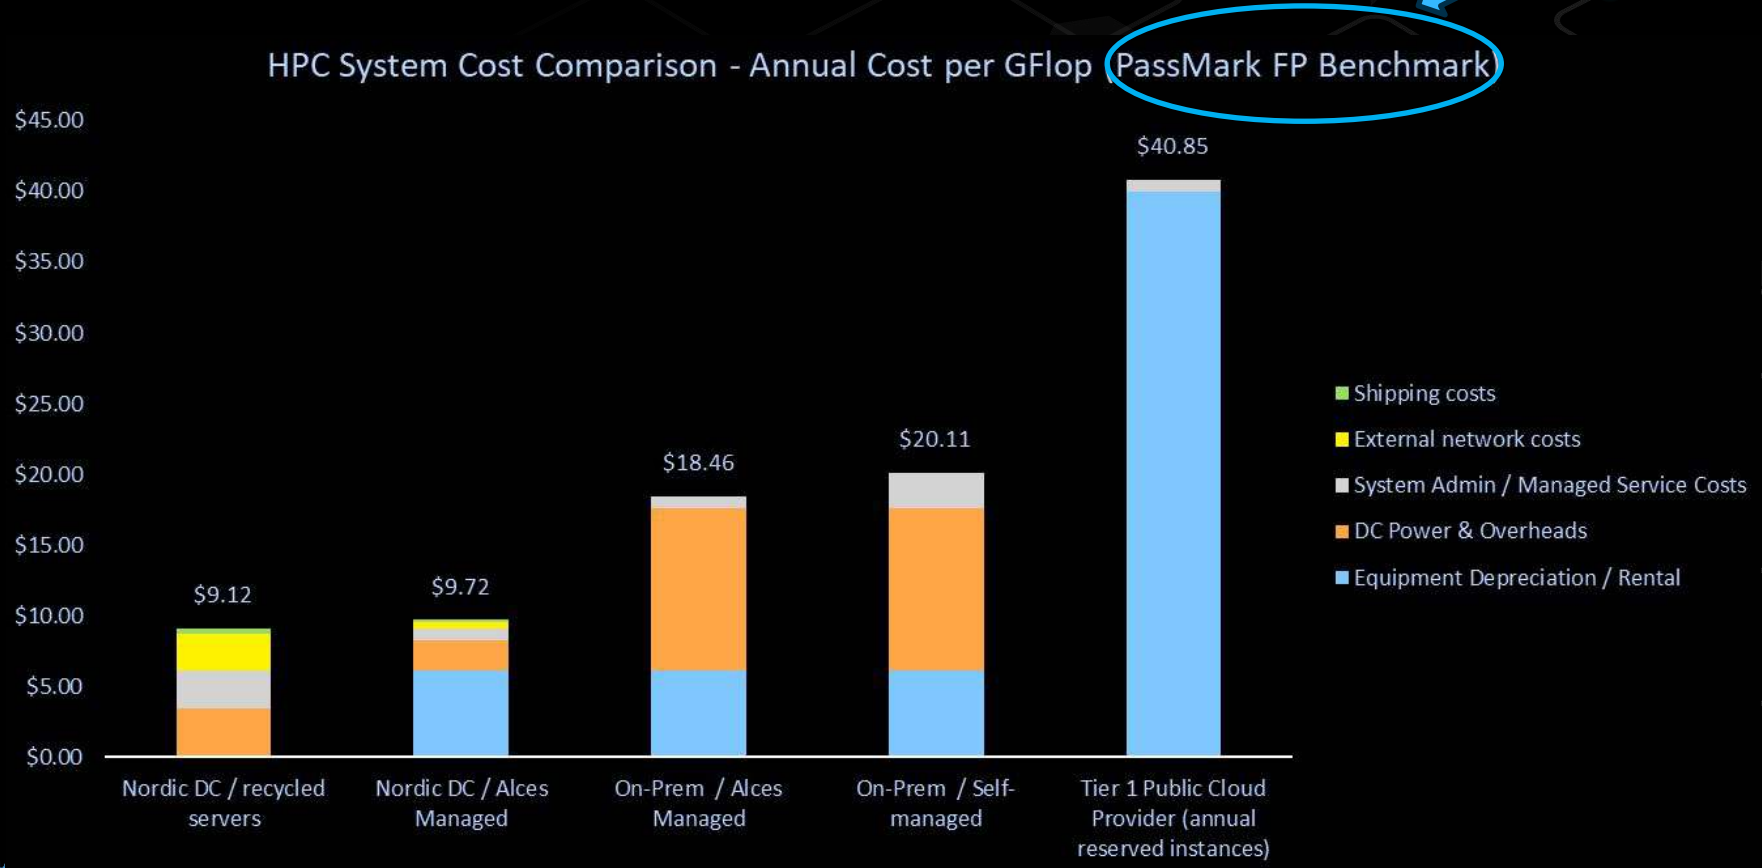

Annual Carbon Footprint per Gflop

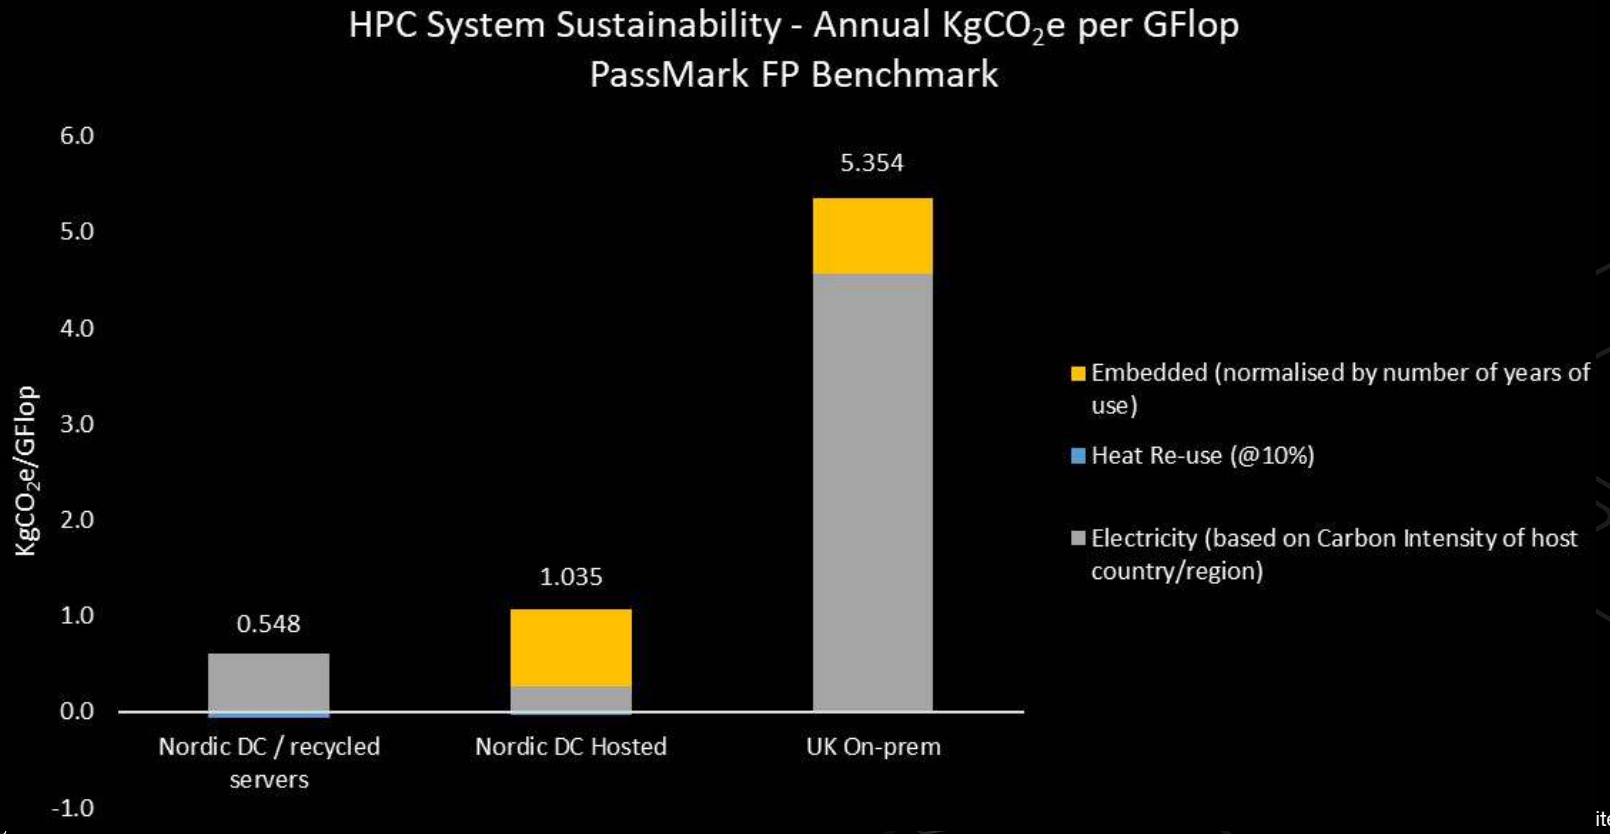

Supercomputing subscriptions

(what are they and why do I want one?)

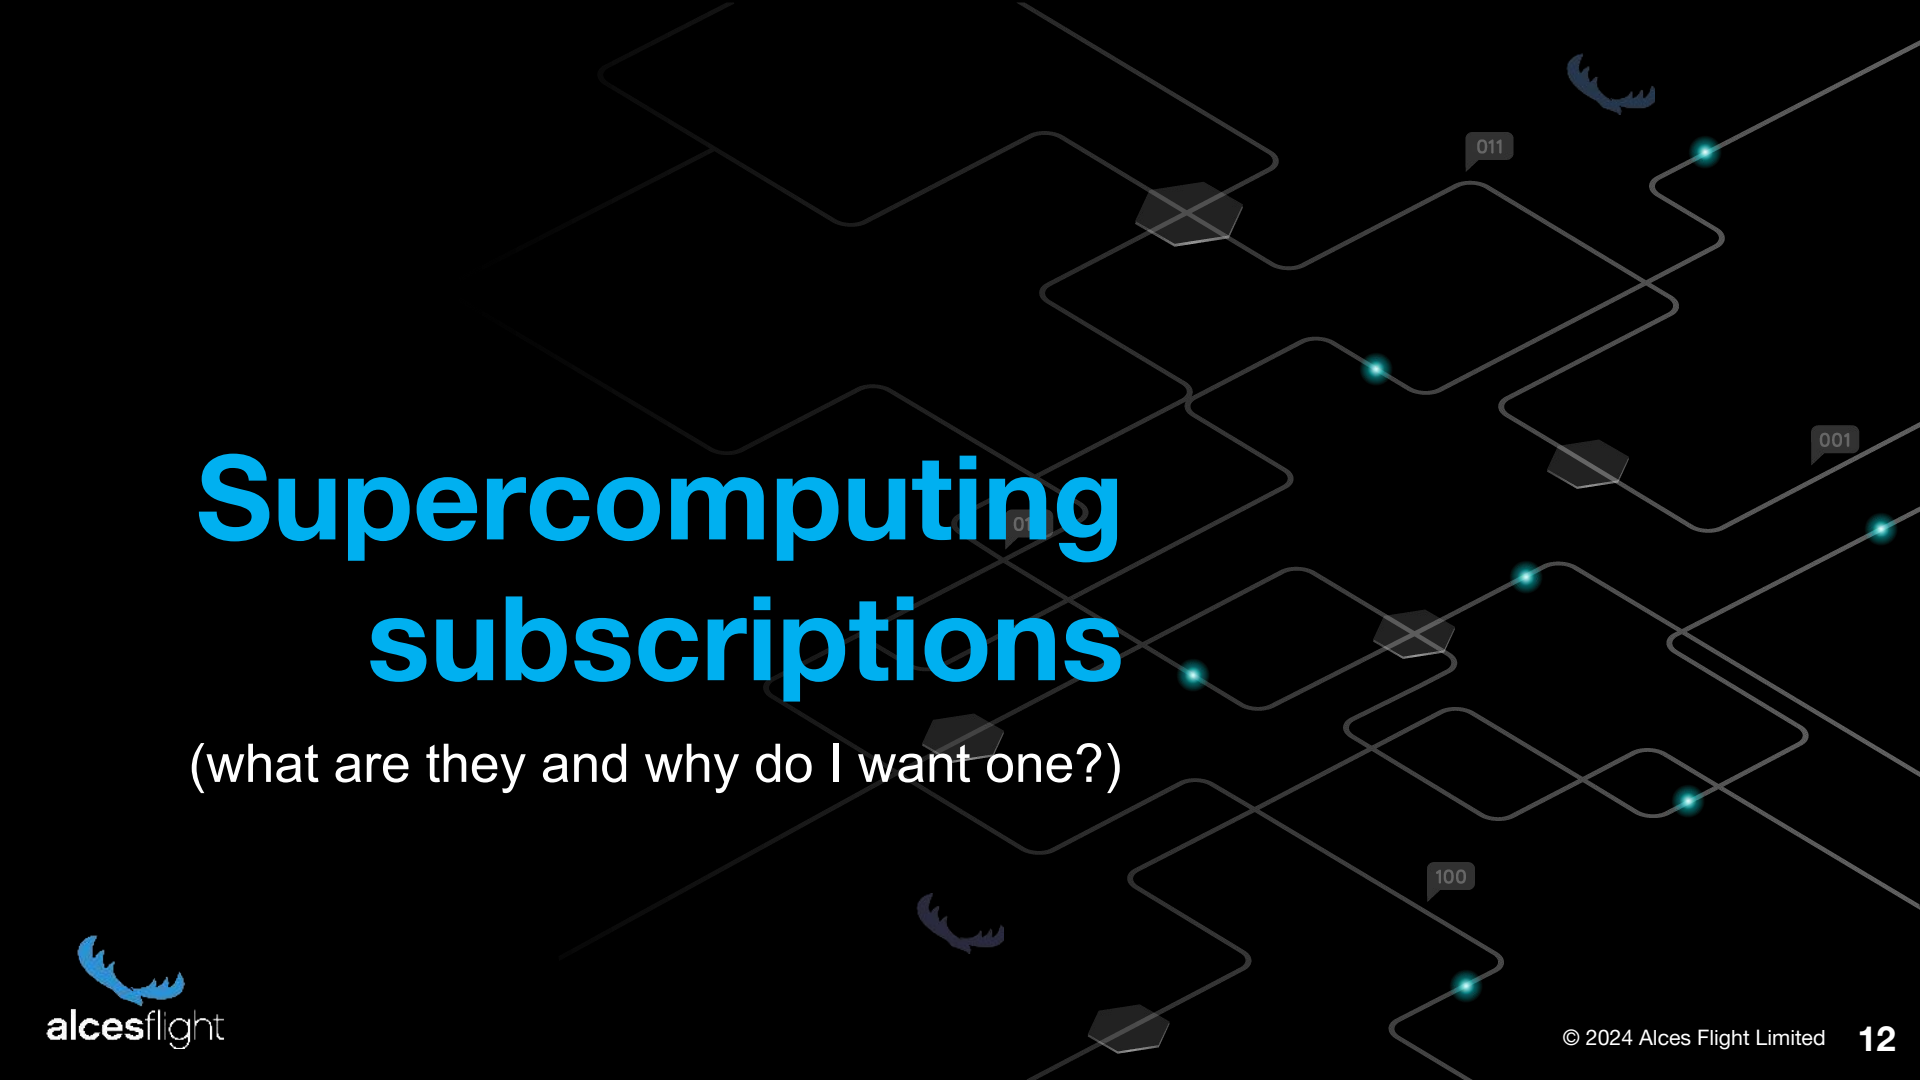

alcesflight

What can you subscribe to?

- Managed Supercomputer service
  - Fixed and pay-as-you-go options
  - Hardware / cloud instances management
  - Filesystem / storage management
  - HPC job-scheduler / AI container service
  - Apps, workflows, pipelines

> 500,000 cores
> 25,000 users
> 100PB data
> 99.9% uptime

What can you subscribe to?

- Blocks of resource
  - Compute
  - MPI fabric
  - GPU blocks
  - Storage capacity
  - Access and visualisation

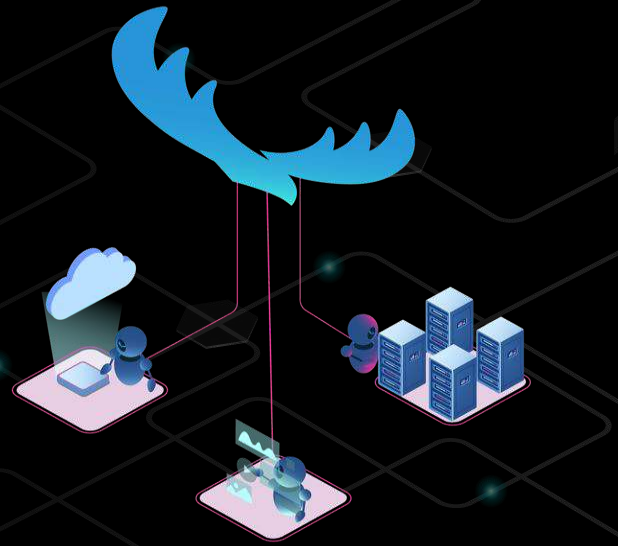

What can you subscribe to?

- Equipment hosted in a range of data-centres locations
  - UK-based, Northern Europe, Iceland
  - Range of different options covering:
    - Network bandwidth and latency
    - Power cleanliness and carbon impact
    - Commitment term length
    - Setup and monthly power costs

Why subscribe?

- Faster to deploy your workloads
- Flexibility to choose your term length
- Latest hardware technologies
- Choice of hosting locations
- Reduce your risk

Moving the needle

(why does this matter?)

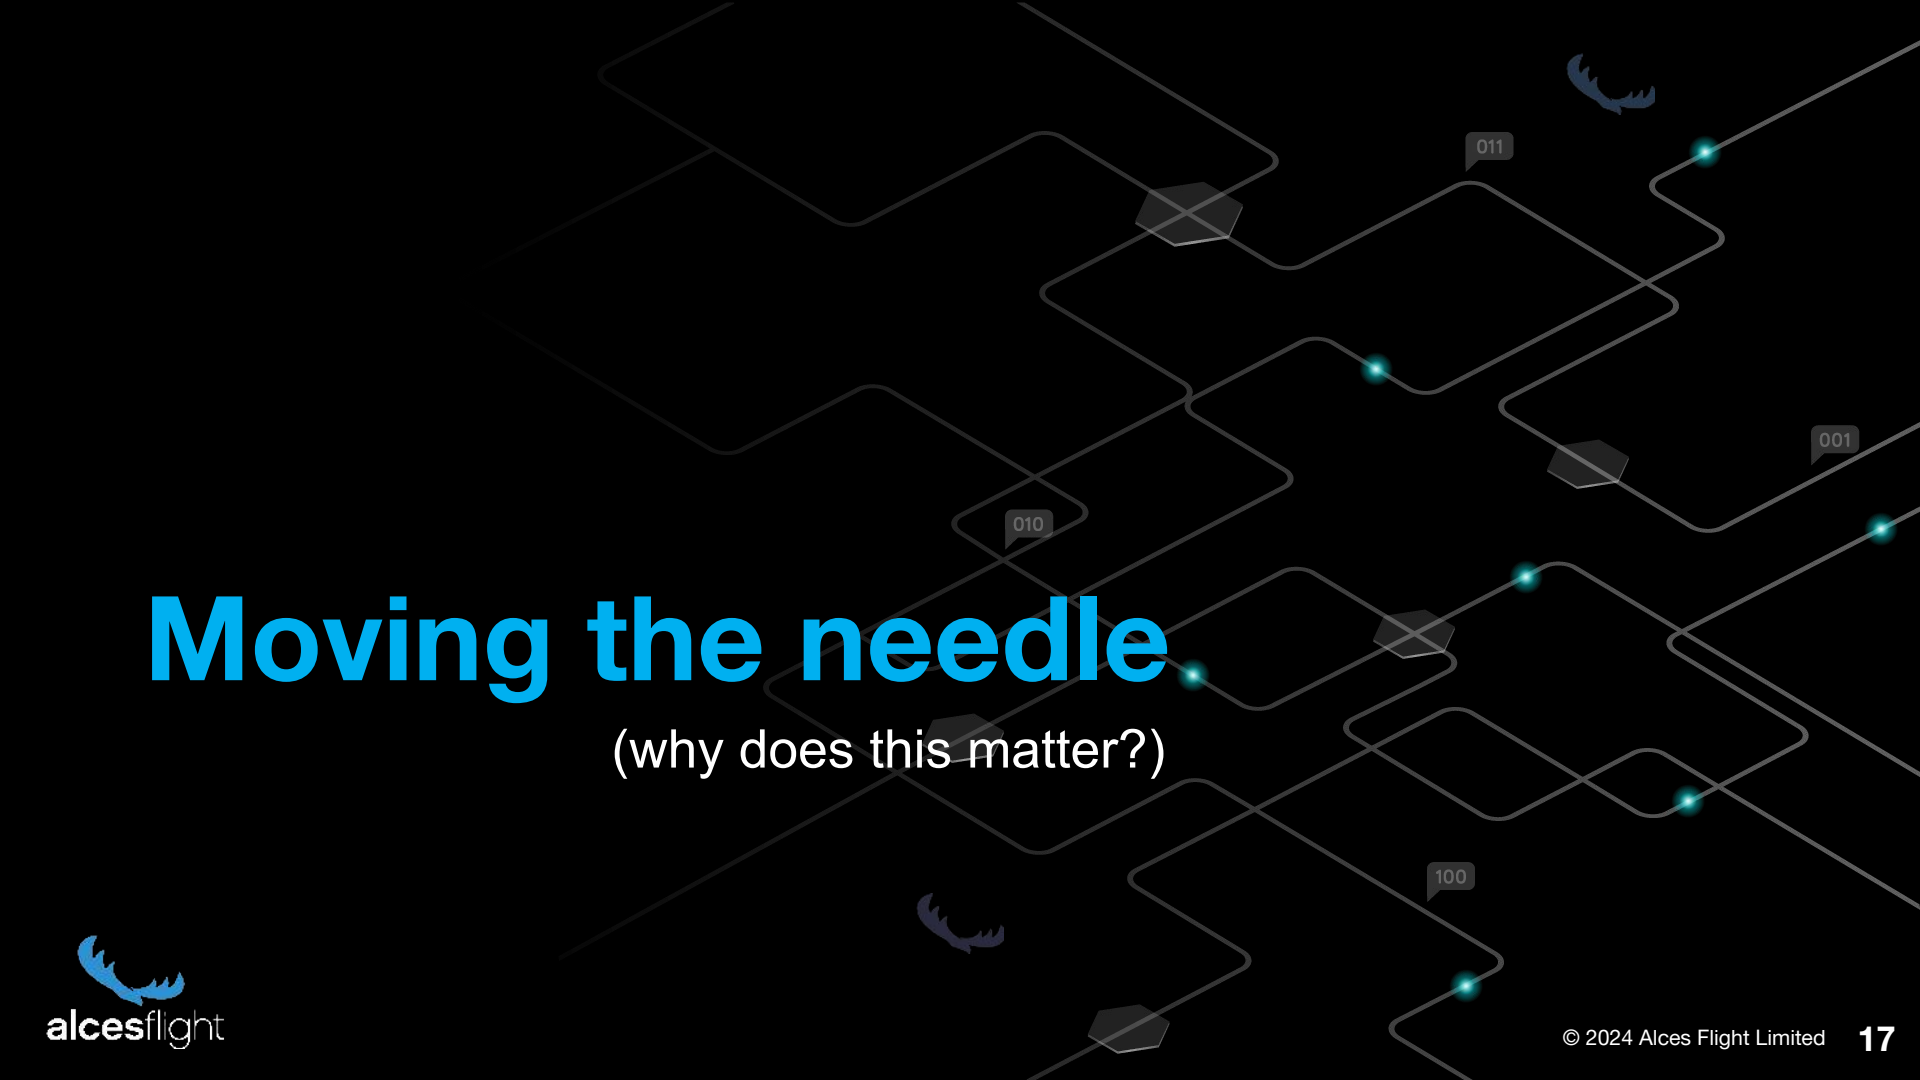

alcesflight

Case study: An actual HPC cluster

- Modelled on a mid-size HPC cluster deployment
  - 64-nodes / 10,000 CPU cores @ 4GB/core
  - 84 x interference GPUs
  - 500TB solid-state + 4PB bulk + backup storage
  - 200Gb low-latency interconnect
  - Tier1 vendor; £2M inc VAT; 110KW max power

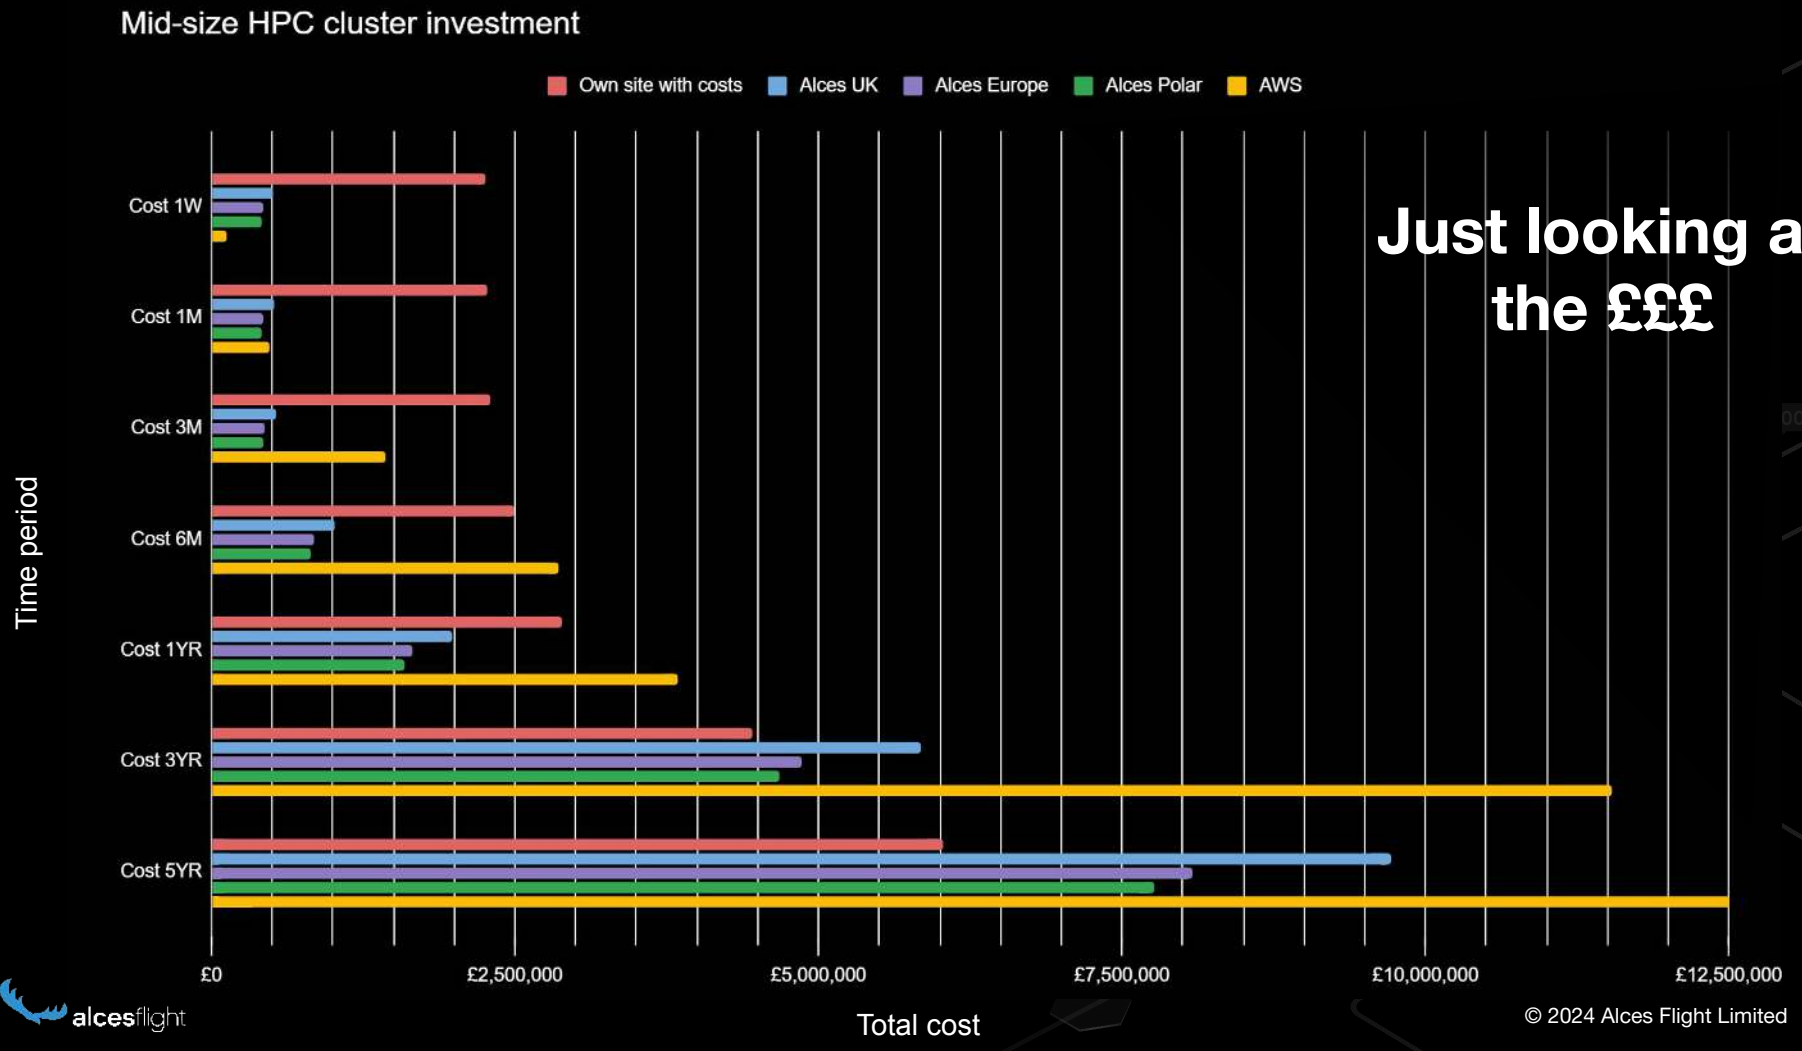

Just looking at the £££

Estimated carbon emissions by location for mid-size HPC cluster

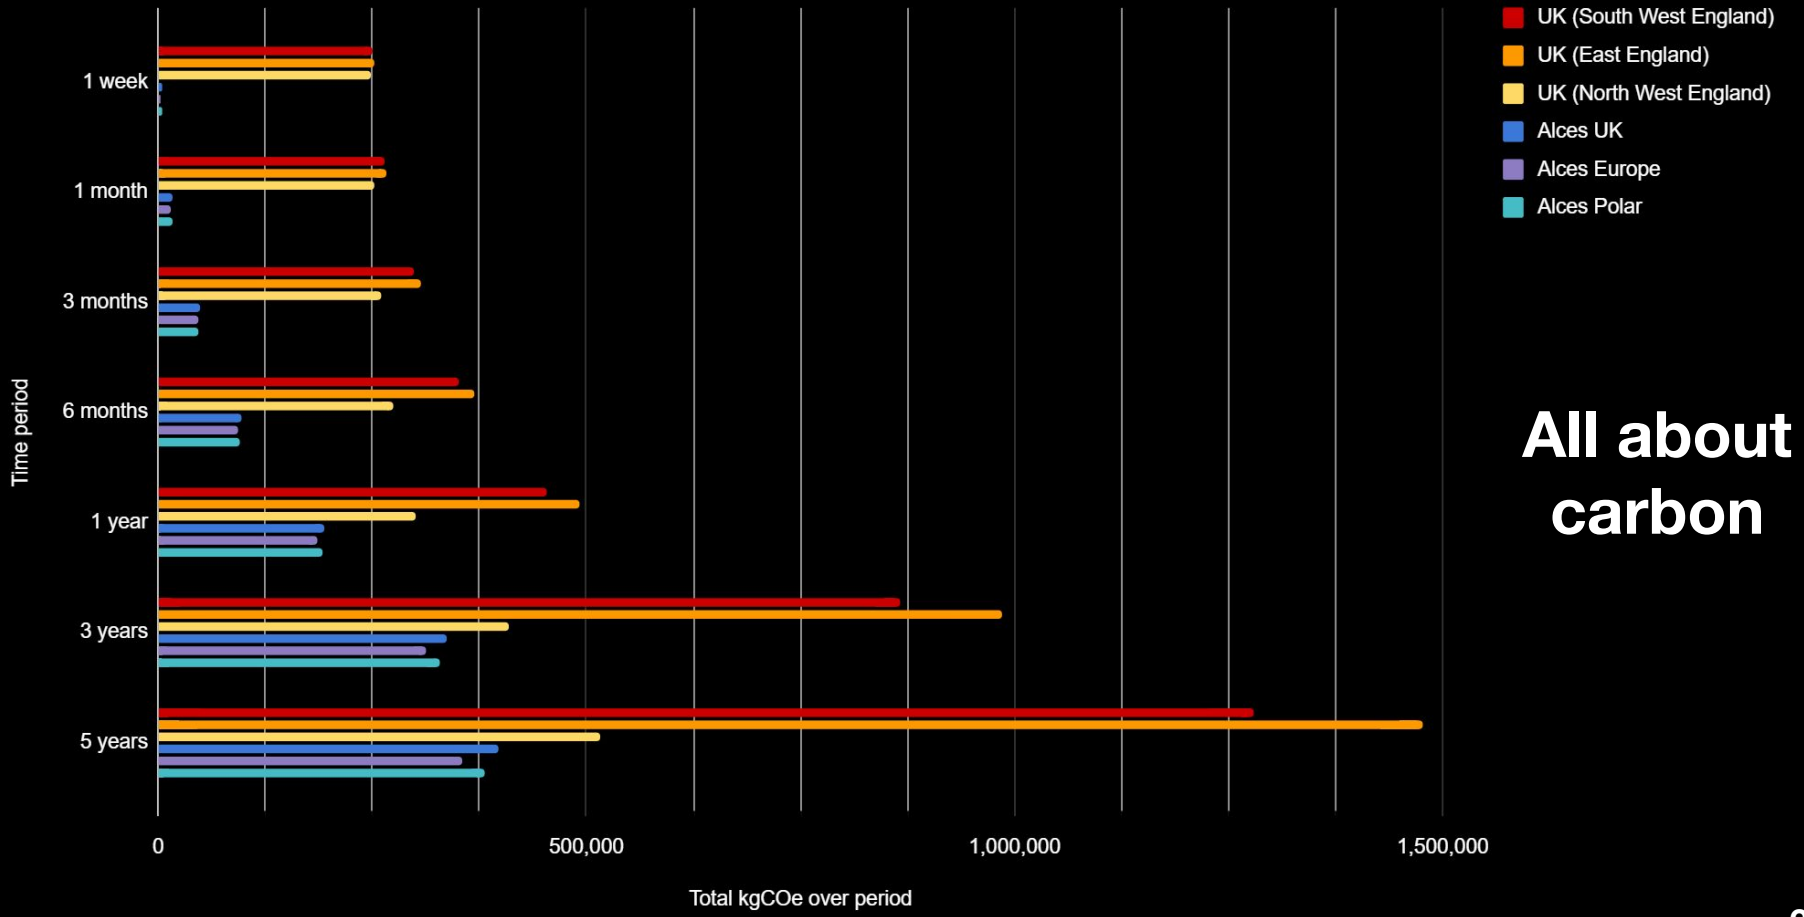

All about carbon

Weeks of delay after hardware is available

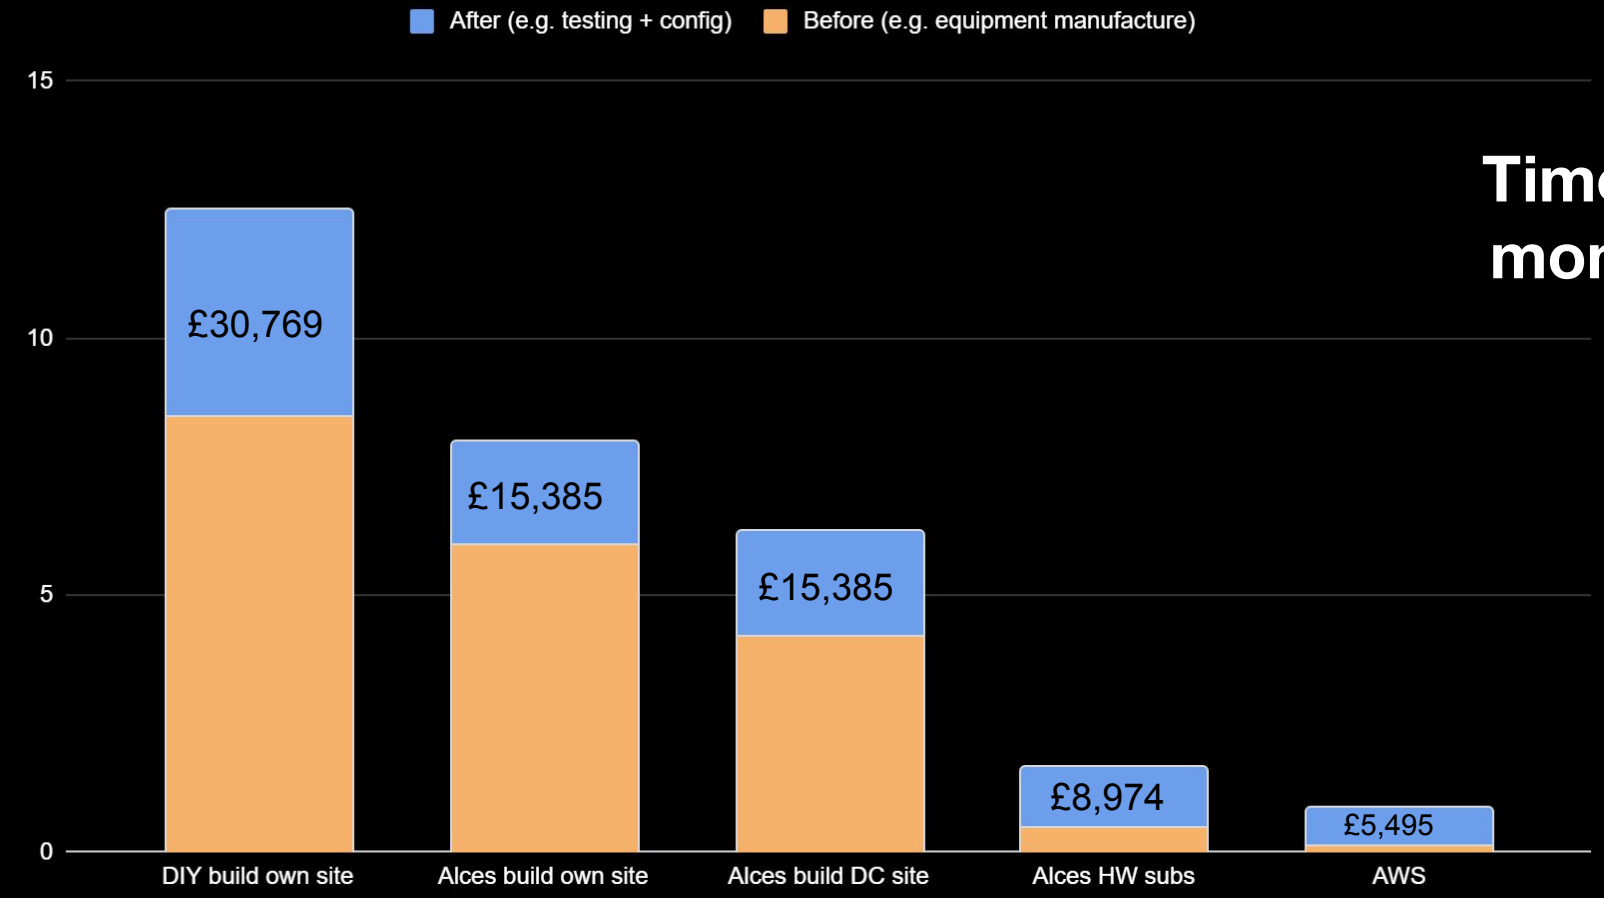

Time is money

Conclusions

(Things to think about)

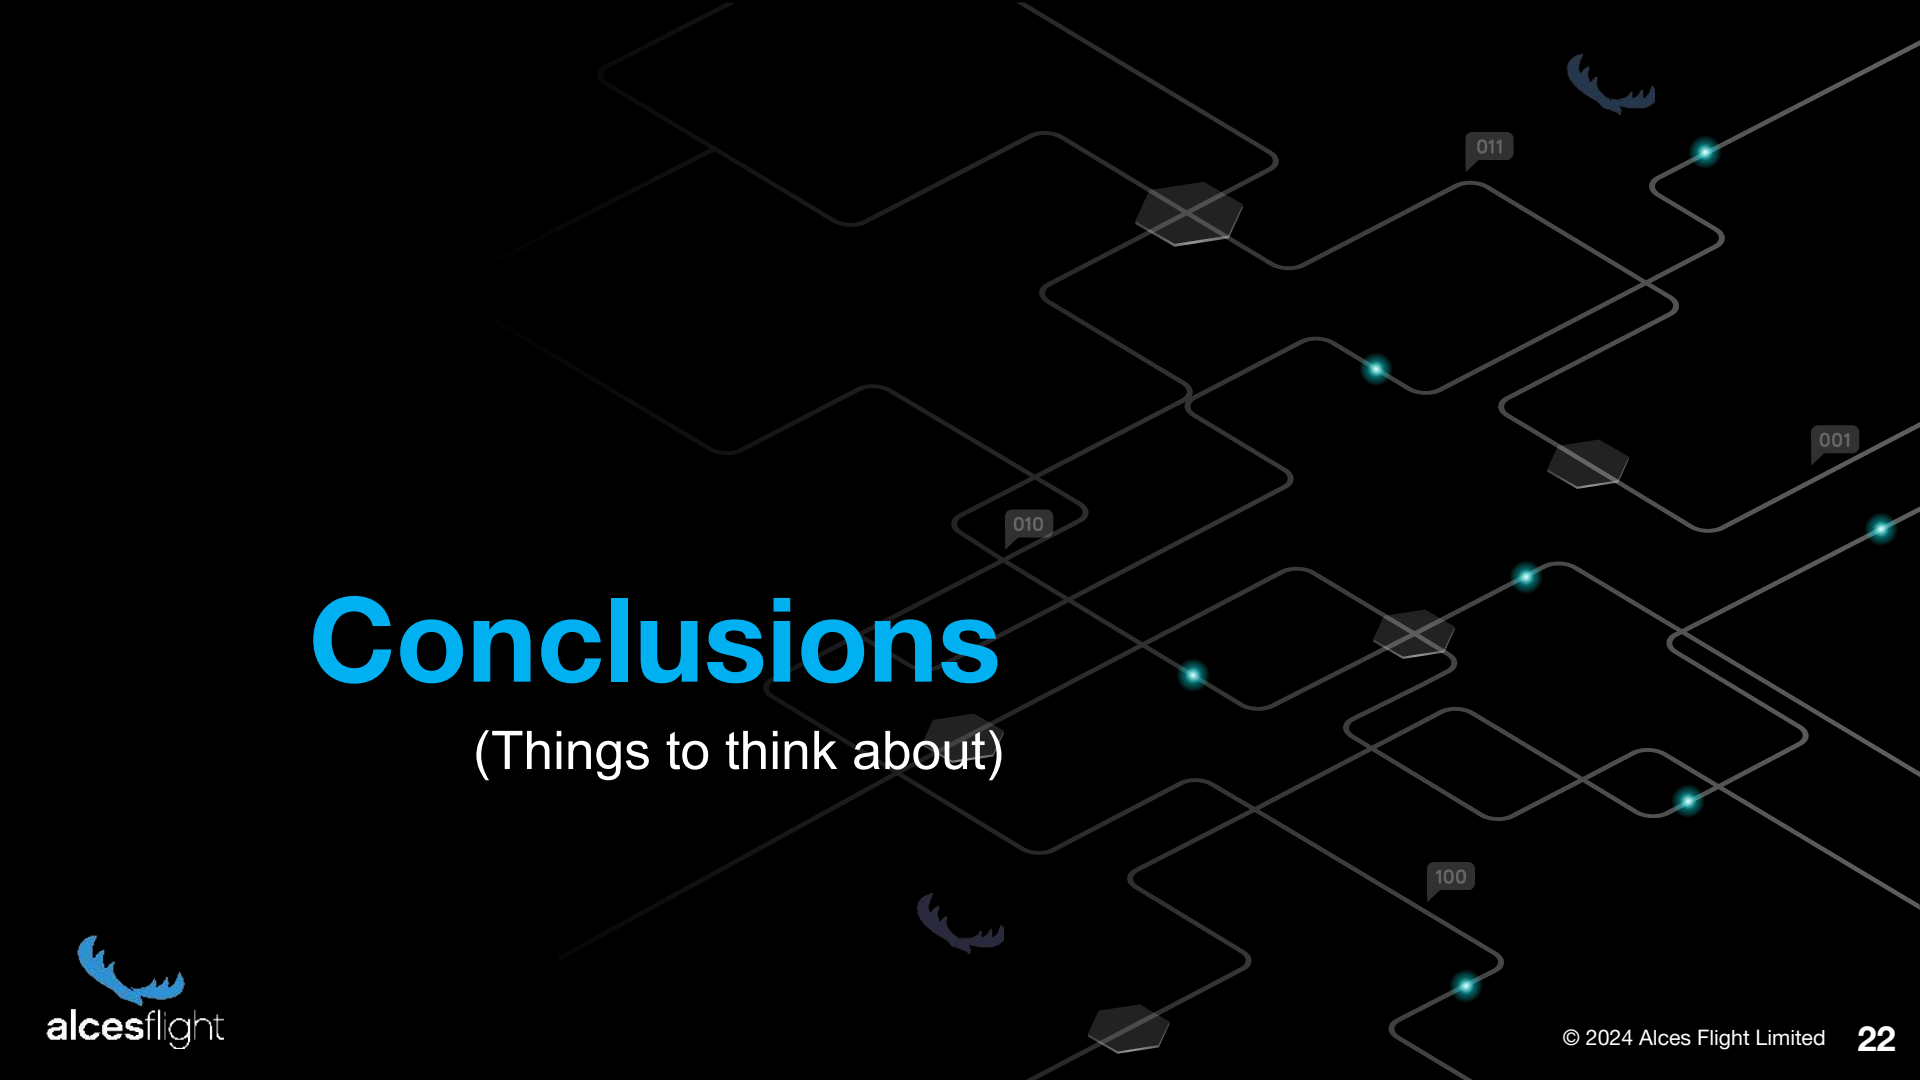

alcesflight

Conclusions

- Rinse & repeat not always the best option
- Understand all your costs including:
  - Financial
  - Time
  - Risk
  - Environmental impact
- Move the needle