



Exploring the cloud

Steven Chapman, University of Bath
Owen Thomas, Red Oak Consulting

HPC-SIG Bath, 25th February 2020



Balena HPC

- free-at-the-point-of-use and premium SLAs
- dedicated batch, shared interactive & visualisation
- 1,200 users from 16 out of 18 depts
- top4 depts: chemistry, physics, maths & chemical engineering
- approx. 200 unique applications/modules
- last 12 months:
 - 88% batch utilisation
 - 140 unique users per month
 - 20 new users per month
 - ~50% usage =<1 node



	Intel Ivybridge E5-2650v2	Intel Skylake Gold 6126
# nodes	196	17
# CPU cores	3,136	408
Memory per node	64GB, 128GB, 512GB	192GB
Accelerators	Nvidia K20x & P100, Xeon Phi	
Network	Intel TrueScale 40Gbit/s	
Storage	0.6PB BeeGFS scratch	

Also have access to...



 **JADE 2**
Tier 2 HPC

 **Isambard 2**

So what's next?

Our goals

- Continued increase of user base and domains using HPC
- Grow the number of inter-department collaborations
- Make our researchers more productive

Other influencers:

- Growing expectation within Bath to use cloud
- Enterprise infrastructure is beginning to move to Azure
- IT is refocusing around DevOps



UNIVERSITY OF
BATH

Known unknowns

- What are the costs?
- What will happen to my data?
- Cost of moving data?
- How will the HPC service model change?
- How can we manage budgets?
- How will this impact grant and non-grant funded research?
- How can negative consequences be avoided or mitigated?
- What is the best way support researchers in the transition period?
- What are the barriers to entry for the proposed system?
- What will the potential collaboration with a cloud supplier look like?
- Is it as easy to justify investment in cloud as in physical infrastructure?
- Will the new service must be at least as good as it is now?
- Will the service be available to all?
- Will the service must be secure?
- Will the cloud be integrated with research storage?
- How far can I scale my code/models within the cloud?
- How does cloud support my goals?
- Does cloud offer value for money?
- Can we trust public cloud providers?
- How can cloud help transform my research?
- What are the potential compliance issues with data rules (e.g. GDPR)?
- What are the data transfer speeds to the cloud?
- What is the performance of the cloud like?
- Is there a performance hit?
- Who else has moved to cloud?



Let's run a pilot

- Focus on end-user experience
- How can costs and budgets be managed, controlled and charged correctly
- Comparison of on-premise and cloud TCO models
- How to incorporate cloud costs and charges on to research grants
- Investigate impact on internal processes, policies, procedures and integration with enterprise environment
- Test stability, reliability, resilience, reproducibility and performance of a cloud platform
- Understand staffing and training implications
- Discover any new unknowns?



What are we doing?

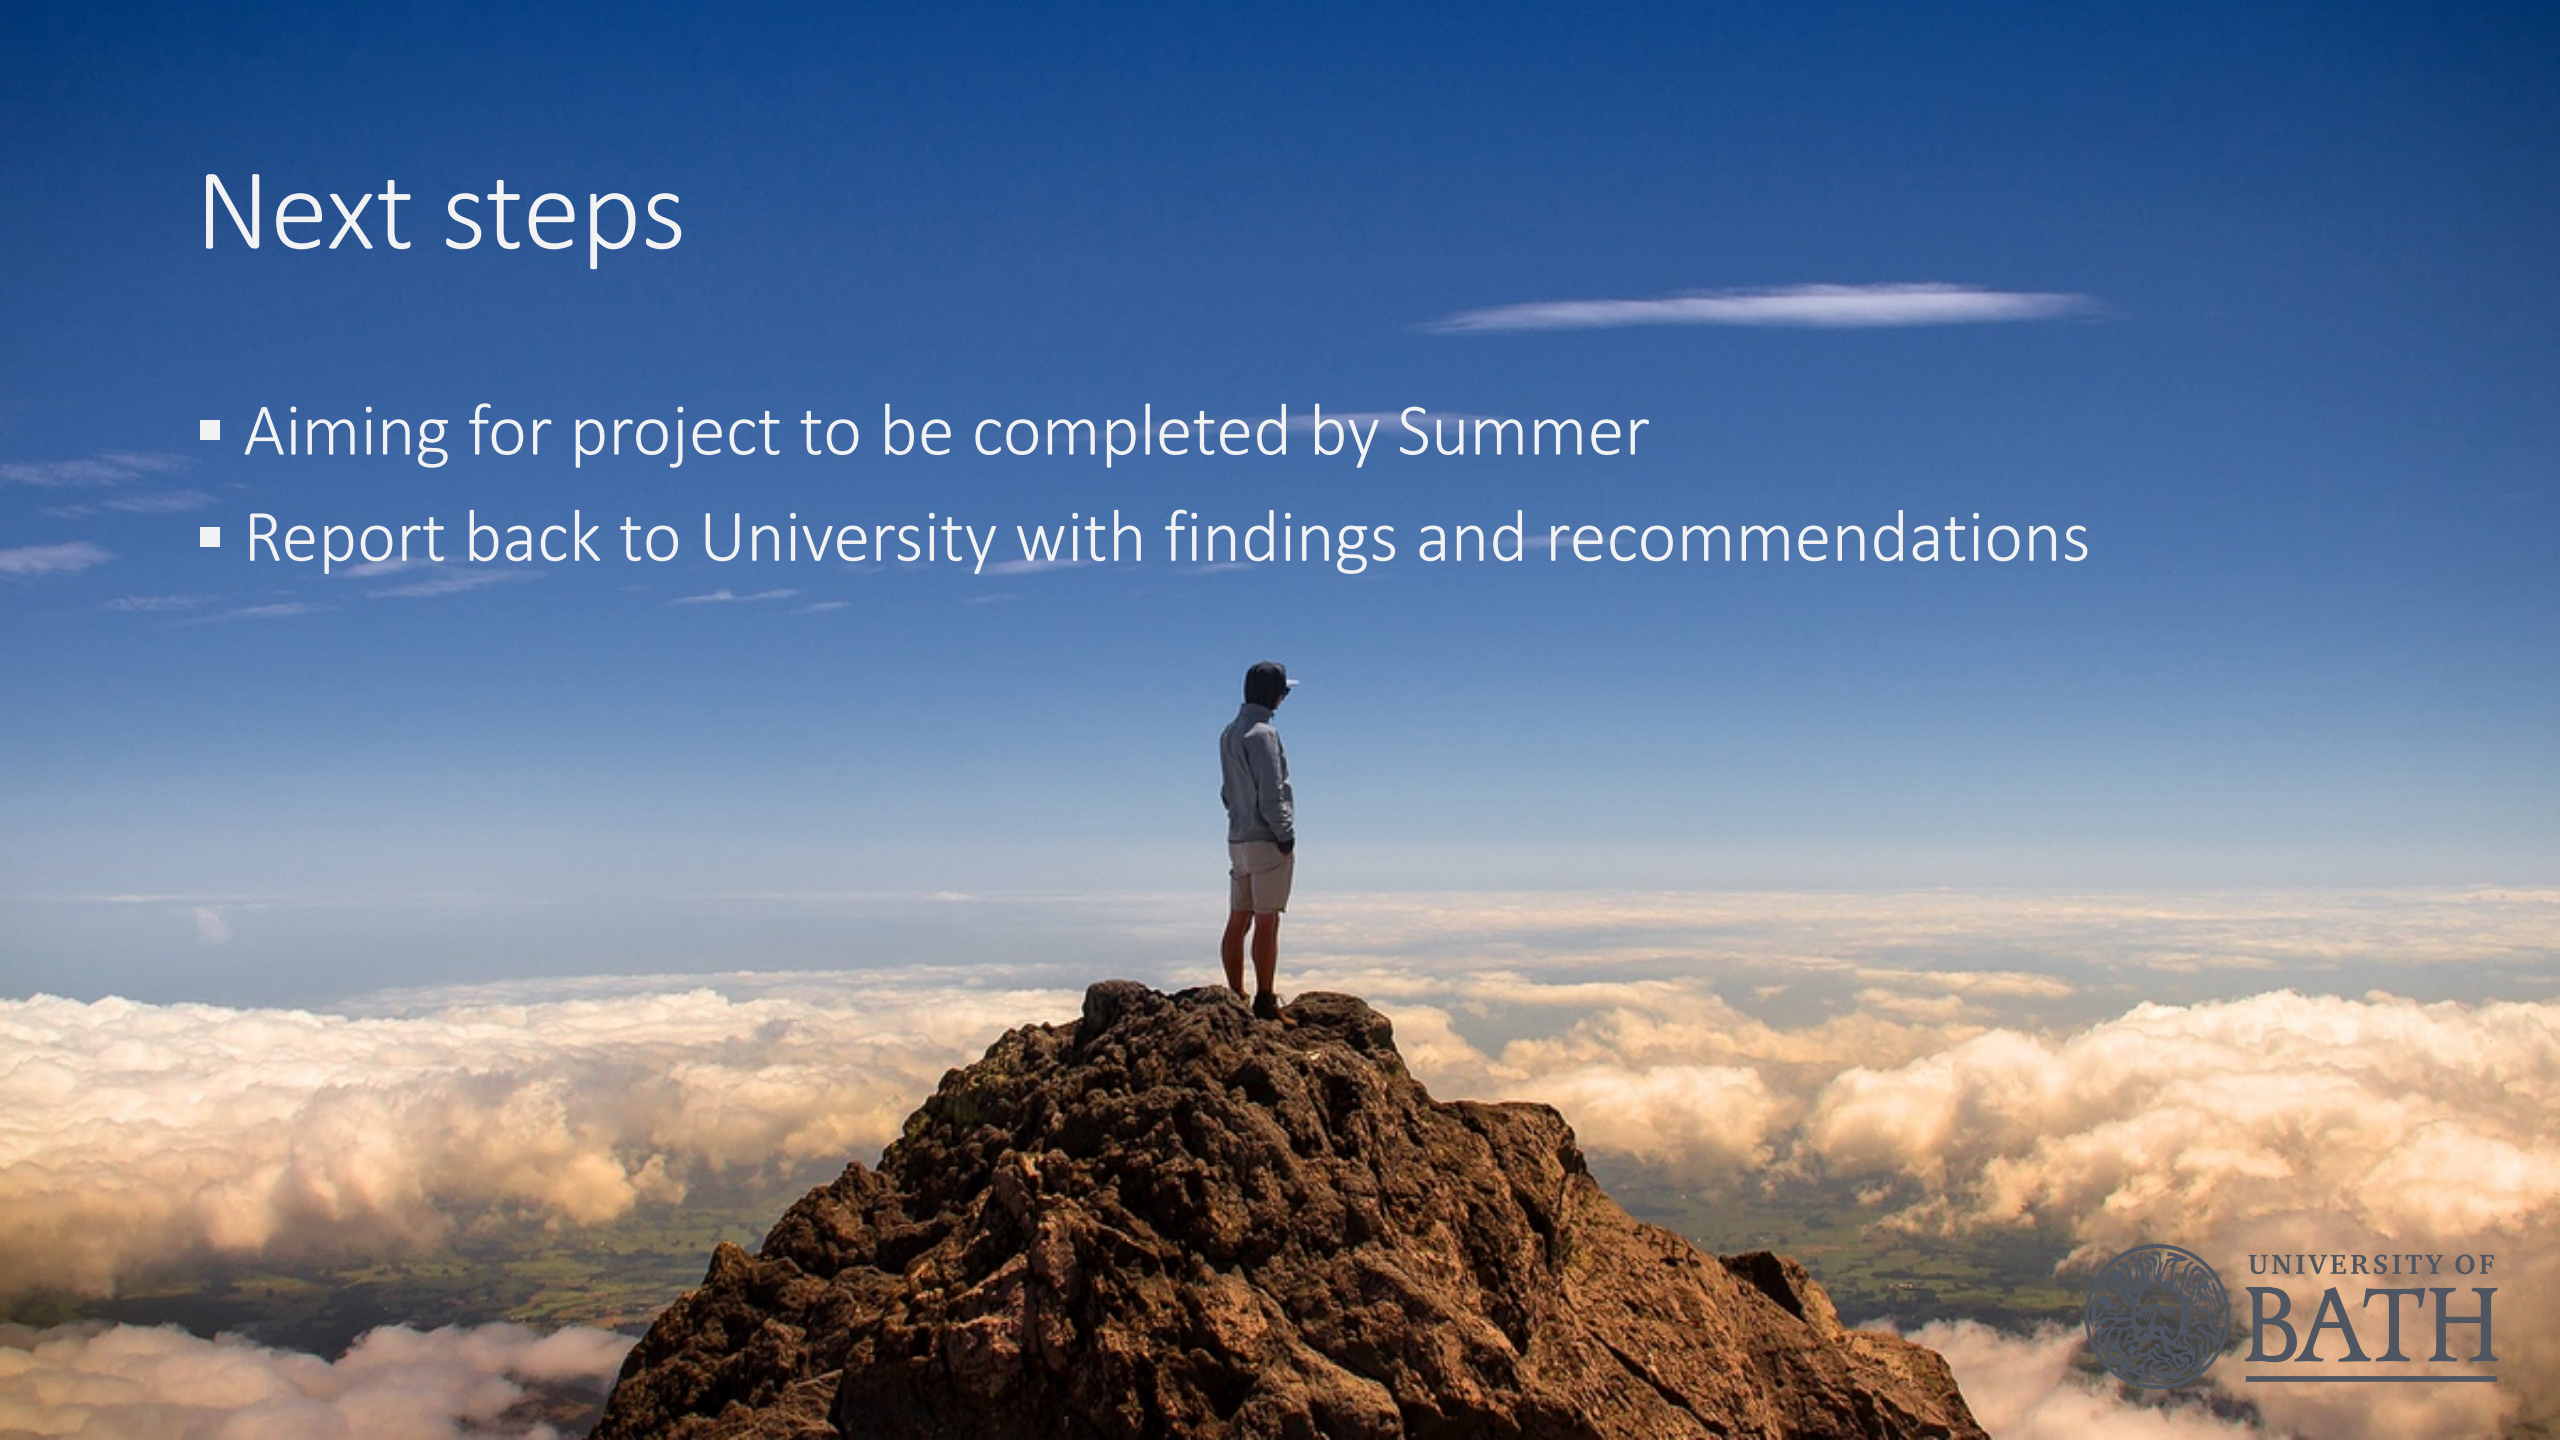
- Running a pilot project with Red Oak Consulting
- Standing up 8-10 different types of real workflows in Azure
- Based on HPC, HTC, visualisation
 - From a general fixed-capacity HPC service to a fully scalable service
- Using Microsoft Azure – Intel Skylake (Hc), AMD Rome (Hb) and AMD Milan (Hbv2)
- Understanding how we can manage costs and budgets
- Produce synthetic reports and bills for internal charging

What have we learnt so far?

- Stakeholder engagement, everyone needs to be part of this journey
- Interesting clauses in software agreements
- Delays in cost reporting and ability to map true costs to workloads
- Expecting more lessons to be learnt....

Next steps

- Aiming for project to be completed by Summer
- Report back to University with findings and recommendations



Over to Owen and Red Oak ...

