



The Taxonomy of HPC Sustainability




Dr Paul Calleja : Director Research Computing Service

Today's Talk



DAWN



Cambridge RCS – Key characteristics

- Running for last 18 years, major UK National HPC provider, one of the largest sites in UK in terms of HPC performance and users
- Providing 48% of UK Top500 performance over last 7 year
- Multiple National AI/HPC stakeholders via a pool investment shared infrastructure, rolling upgrade infrastructure, large economies of scale, agility, flexibility
- STFC DiRAC, STFC IRIS, EPSRC Tier 2, STFC SKA, UKAEA
- Currently providing the UK fastest AI supercomputer “DAWN” Along side Bristol kick Starting the DSIT / UKRI AIRR service
- Strong focus on technology transfer across Africa



DiRAC



UK Atomic
Energy
Authority

Cambridge RCS – Key characteristics

- Serving a cross domain HPC user community of over 3500 users
- Bringing together users from Physical Sciences, Biology, Medicine and Engineering on a single heterogeneous system, strong multi-domain RSE and user support team
- Back in 2006 pioneered large scale commodity / open standards based HPC systems as UK national scale HPC systems, when all large national HPC systems at that time were legacy proprietary systems with 100X lower performance per pound.
- Today driving software defined, cloud computing middleware, supporting dev-op HPC environment, with agile dynamic software defined research platforms on top of a shared heterogeneous infrastructure

This is very different to traditional monolithic HPC provision models



Cambridge RCS – Key characteristics

- Strong in-house technology team
 - *Undertaking system design, implementation and system/user support function*
 - Provides best :- fit between science need and technology deployed, strong VFM, agility, lower risk
- Focus on technology innovation via industry partnership & co-design Open ZettaScale Lab long term co-design partnership with Dell, and wider technology supply chain partners
- Our technology development is strongly directed by science need
- Strong partnerships and some commercialisation
 - Long term partners
 - Dell, Intel, StackHPC, Mellanox, Cornelis networks, SchedMD
 - Spin outs
 - Zetta Genomics, Cambridge Research Computing Ltd



Cambridge RCS Facts & figures

- Provide both **national & Local** HPC, AI and Data services (90/10)
- Undertake growing volume of **HPC / AI development** activities
- **£100M** of HPC equipment in operation
- Currently **~40 staff rising to 55** over next year. Highly capable HPC design, implementation and operations teams.
- **1.8 MW** water cooled HPC Data center 100 Racks
- **30 PF** heterogeneous converged HPC/AI system
- **3000** Dell servers X86/GPU Intel & NVIDIA
- **45 PB** storage (disk/tape)



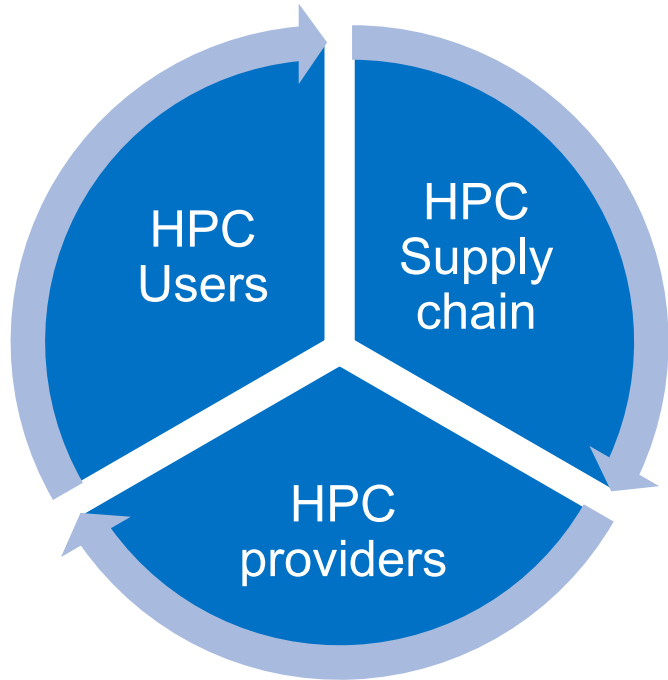
- **ISO27001** compliant secure computing for commercial & medical users



- **\$10M 4-year program** - Academic / industrial partnership for the **co-design**, development and testing of leading-edge HPC, AI and HPDA solutions
- Pushing the boundaries of performance and functionality but more importantly making large scale HPC system more accessible thereby **democratising HPC/AI and HPDA technologies for everyone at every-scale**



The co-design virtuous circle



- Fusion of science use-case, service provider and technology vendor input
- Critical mass of experience, skills and infrastructure
- Strong science led technology development
- Driving the innovation cycle requirements, develop, deploy, evaluate, iterate

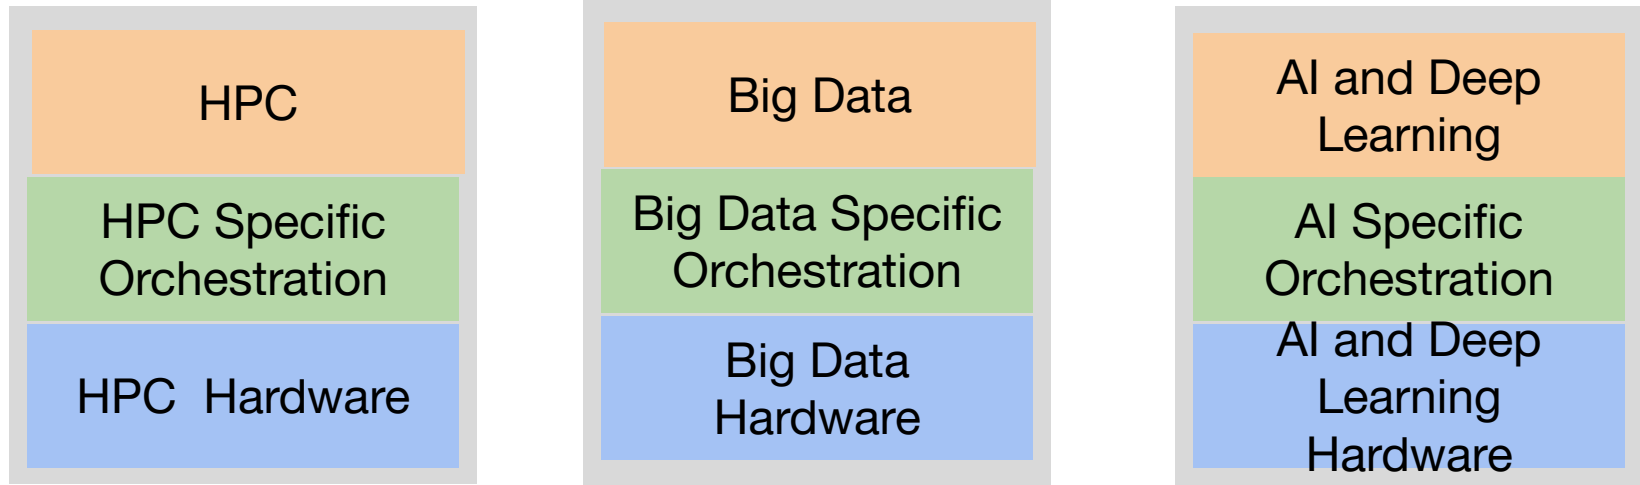


ZettaScale Lab technology development themes

- Energy efficiency
- oneAPI Centre of Excellence
- Research Computing middleware, accessibility & tools
- Large scale tiered storage solutions
- AI workflows and tools merging AI cloud into HPC infrastructure
- HPC networking technologies
- Extreme scale visualisation
- Health informatics (TRE's) HPC in the clinic

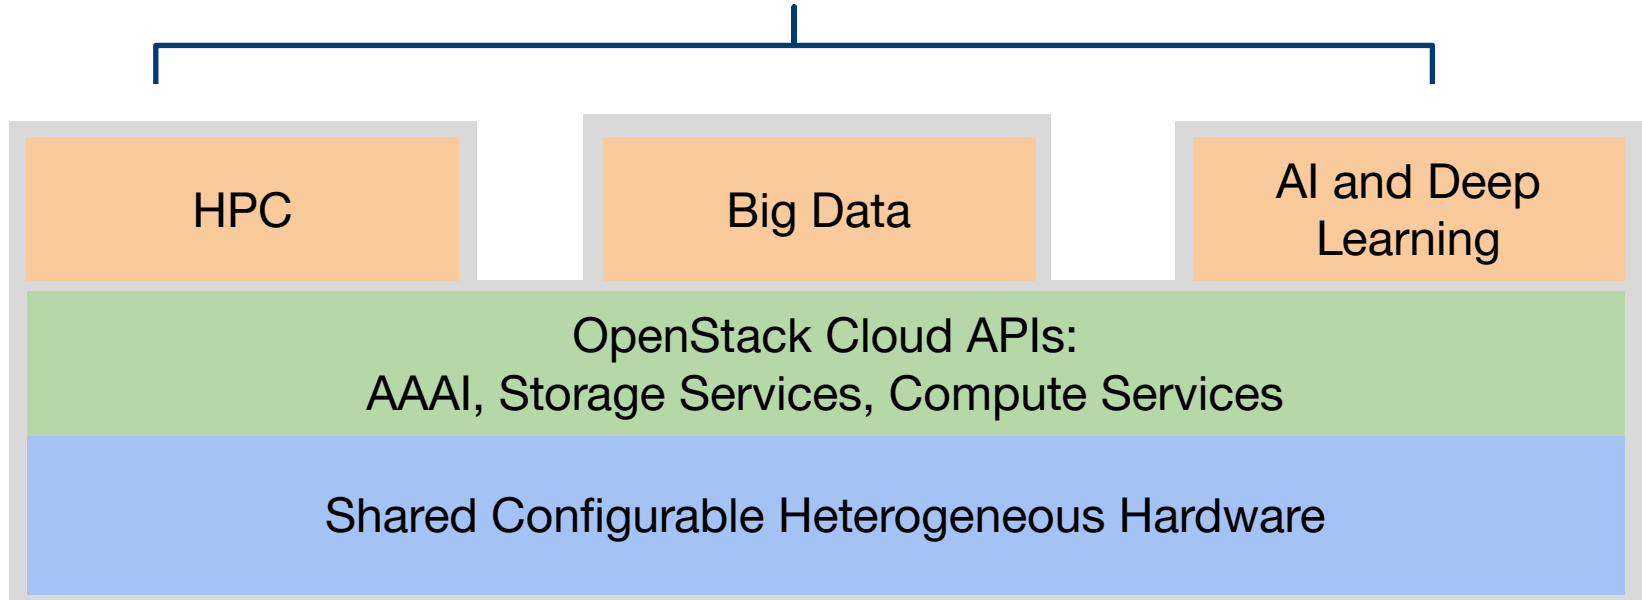


Traditional HPC Stack



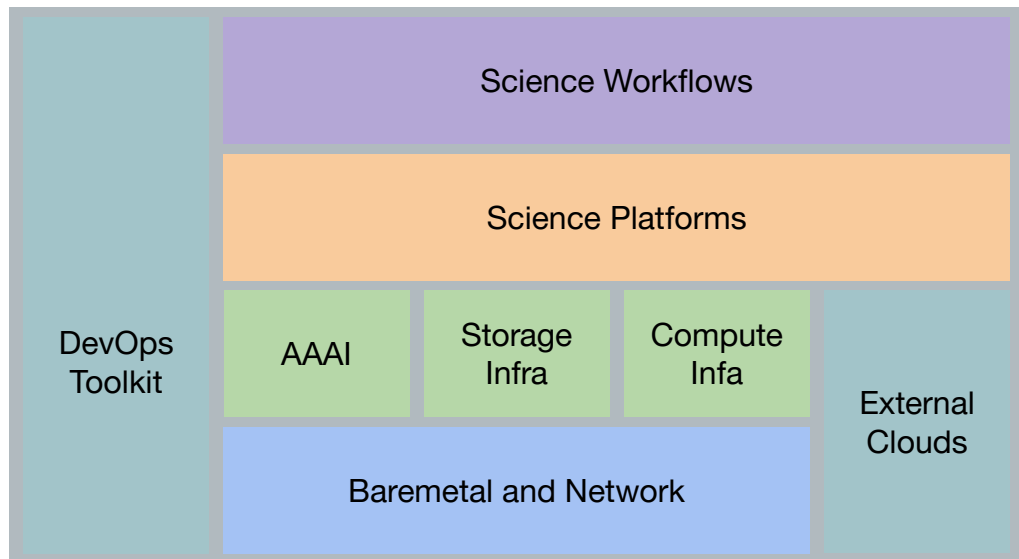
OpenStack for Research Computing

Use Openstack to provide shared heterogeneous infrastructure
Use Ansible to build software defined, reconfigurable science platform ontop of that shared Infrastructure



Scientific OpenStack – On Premise Science Cloud

Full-Stack Science platforms



- Deliver HPC & AI via cloud APIs
- Under strong active development in partnership with **StackHPC** funding UKRI, SKA, Industry
- Creating UK “Community Cloud Middleware Stack” **UKRI project**
- Revolutionises flexibility and end user functionality of HPC systems
- Makes hybrid HPC systems real
- Controlling all Cambridge infrastructure
- Gaining traction across UK and Europe



DAWN

- UK's Fastest AI supercomputer 19.45 PF, #40Top500 at launch (11/2023) #50 now
- Highly innovative co-design, co-investment partnership - Dell, Intel, Cambridge, UKAEA, UKRI, DSIT, StackHPC
- Kick starting UK's National AI Service (AIRR) making the largest AI systems easily usable by scientists, government and business
- Creating federation demonstrator across AIRR Cambridge Bristol, UCL, Durham StackHPC using Waldur access portal, which is used across Euro HPC sites.




DAWN

- 256 Dell XE9640 2U DLC cooled GPU servers
- Each with :-
 - 2 * 4th Gen Intel Xeons
 - 4 way SMP Intel Data Centre Max GPU
 - 1TB RAM, 4NVMe, 4 * HDR200 (fully non blocking)
- 1024 Data Centre Max GPU – 19.45 PF HPL
- 3000 port fiber core, 200TBs maximum throughput , 20Km cabling
- Large spinning disk Luster + additional 3PB NVMe storage with over 3000 GB/s R/W bandwidths



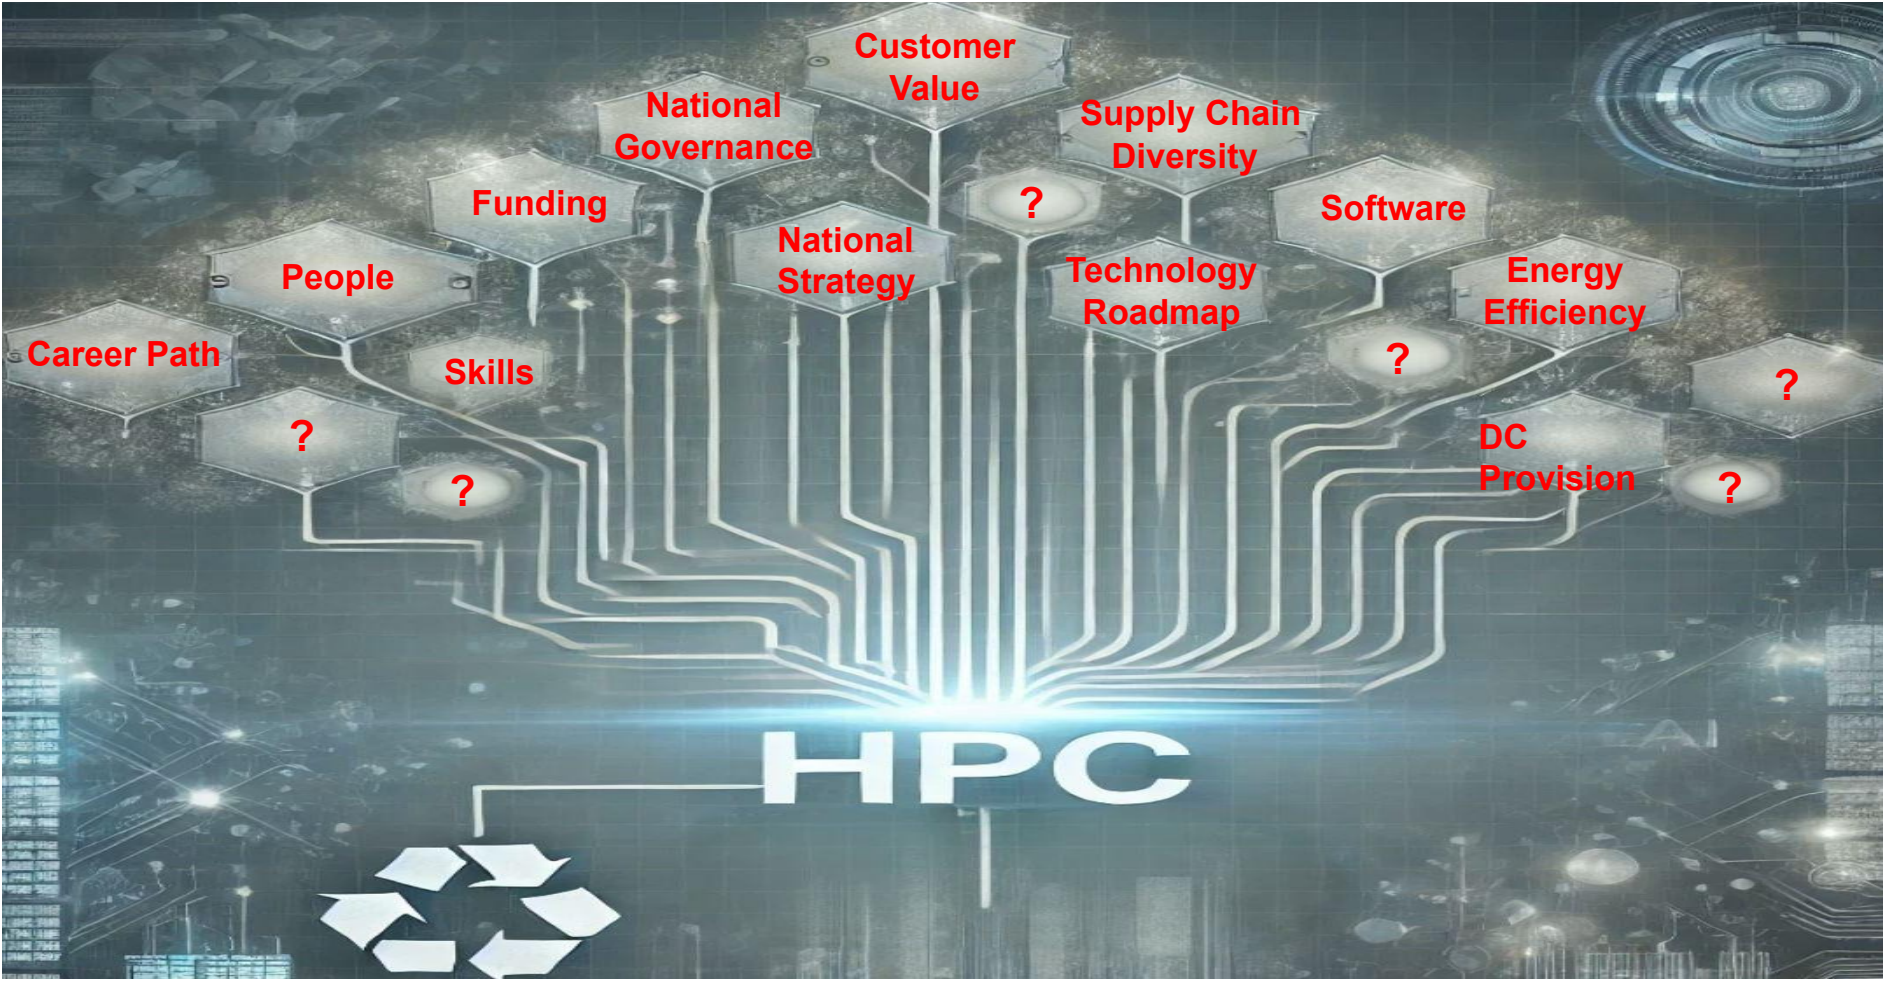
DAWN

- Consumes 1 MW of power
- DLC water cooling
- 240 L/s of flow !!
- 8 weeks P/O to delivery - 3 weeks delivery to top500 !!!
- Already up and in early science mode



DATA





Energy Sustainability

- Cooling
 - Air to crack - traditional
 - Air to water - water back doors
 - Water to water - DLC
- Compute device
 - CPU
 - GPU
 - Latest generation
- Operational efficiency
 - Frequency scaling
 - Code optimization / less mpi wait less I/O wait
- Heat reuse
- Green energy input



Energy Sustainability

- Cooling
 - Air to crack + DX traditional air con
 - Air to water back doors + Adiabatic 40%
 - Water to water – DLC 20%
- Compute device
 - CPU
 - GPU 300% but only good for 50% of workloads by volume
 - Latest generation 30%
- Operational efficiency
 - Frequency scaling 20%
 - Code optimization / less mpi wait less I/O wait 20%
- Heat reuse 50%
- Green energy input 330% !!!

That would save me £1.5M a year
or 3X performance out of same DC !!

