

# CONVERGENCE OF HPC AND DATA SCIENCE

---

AT THE EDINBURGH INTERNATIONAL DATA FACILITY

Professor Mark Parsons


EPCC Director

Dean of Research Computing




## Edinburgh & SE Scotland City Region Deal

- In 2016 EPCC helped develop a “Science and Innovation Audit”
- Identified strengths in our region for Data Driven Innovation
- City Deals are funding from UK and Scottish Governments
- Aim is to stimulate economic growth in UK regions
- £1.1 billion Edinburgh Region City Deal announced in Summer 2017

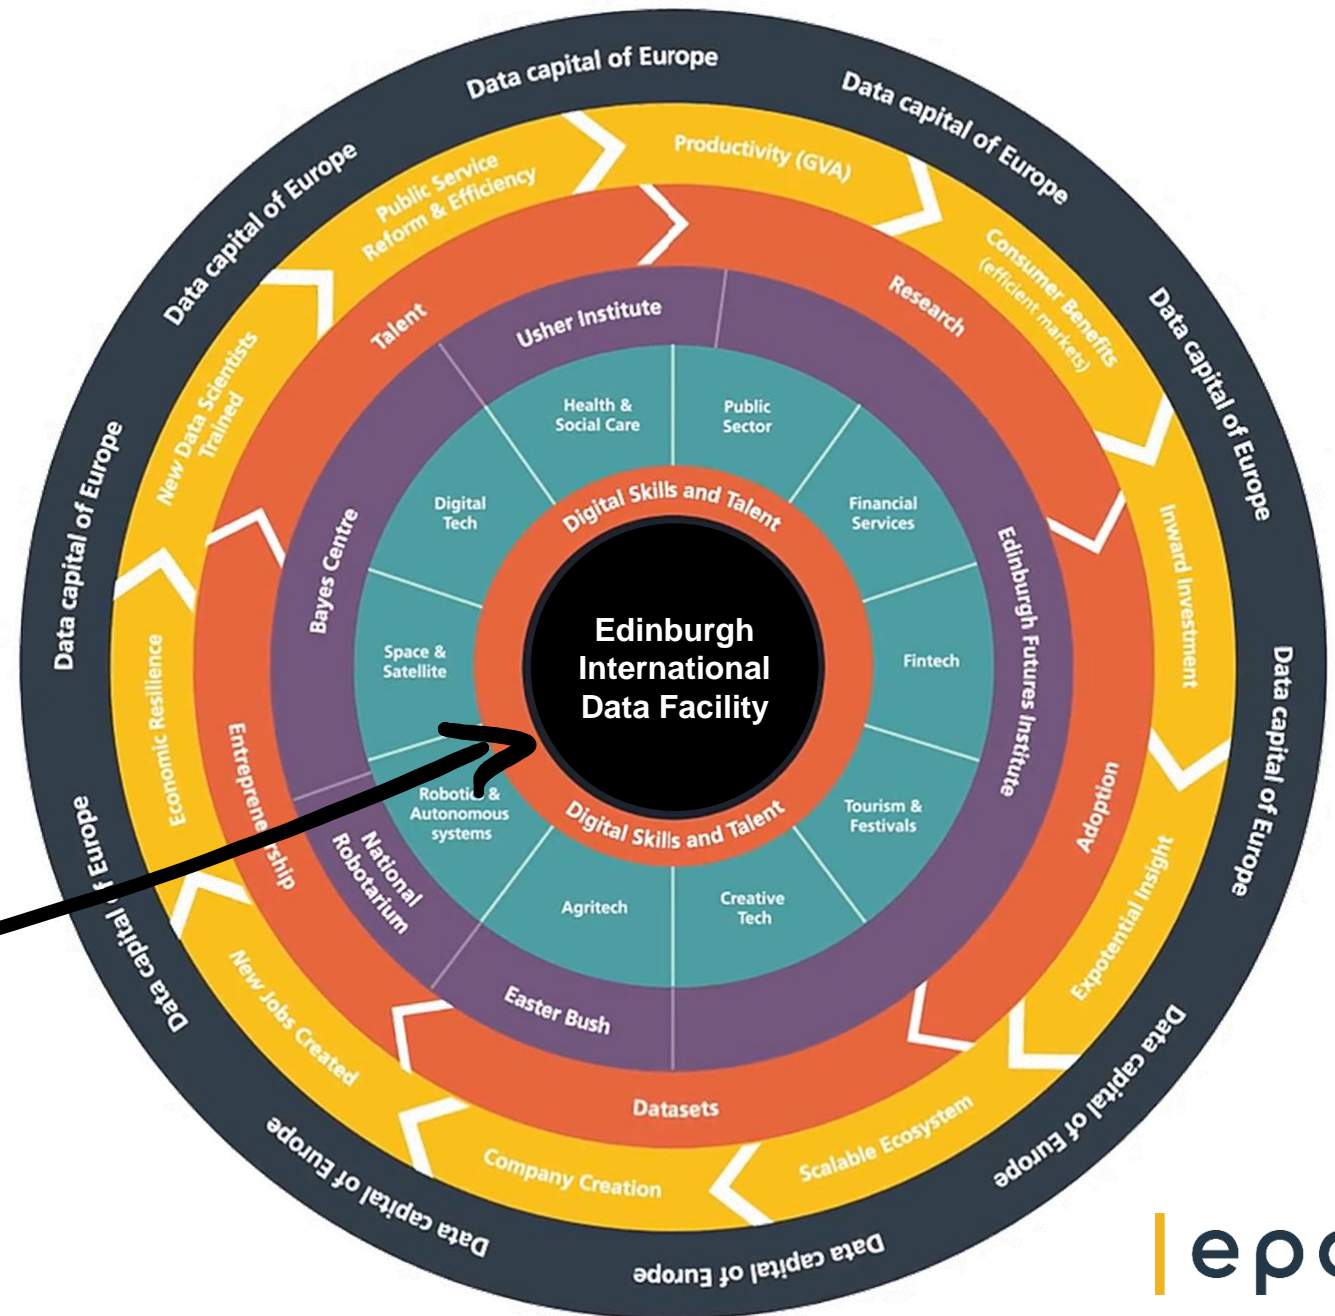


## Aims of City Deal Data Driven Innovation programme

- Capitalise on our expertise in Data Driven Innovation
- Make Edinburgh City Region the “Data Capital of Europe”
- Create a trusted public-private-third sector partnership
- Unlock economic opportunities worth £5 billion+
- Train 100,000 people in data technologies
- Develop an underpinning infrastructure – the **Edinburgh International Data Facility**



The Edinburgh International Data Facility is at the heart of the Data Driven Innovation programme








# Many projects already being supported within EIDF

## Running projects

- National Safe Haven
- Administrative Data Research – Scotland
- Global Open Finance Centre of Excellence
- Scottish Covid-19 Research Database
- ISARIC-4C research service
- Scottish Genome Partnership research service
- Data Science projects on EPCC's HPC systems
- Scottish Government SPACe analytic workbenches
- iCAIRD research service
- Data SlipStream
- Text Data Mining project

## Late stage development

- IoT data service
- National Collection of Aerial photography
- DataLoch
- ... plus many more now emerging



## EIDF IT infrastructure contract

- Keen to simplify purchase of equipment
- Didn't want to make mistake of many procurements with a “big bang” approach
- Want to build-out data infrastructure according to need
- £100m OJEU procurement – ran from November 2019 to May 2020
  - £80m City Deal + £20m EPCC/UoE – 10 year contract
  - Hewlett-Packard Enterprise won contract
- To date 16 orders have been placed
  - Trying to stay one step ahead of demand
- Building block process allows components to be ordered off a catalogue pricelist – updated ~2 times per year

# Building block approach

Large Scale Systems

(1) THE UNIVERSITY COURT OF THE UNIVERSITY OF EDINBURGH

(2) HEWLETT-PACKARD LIMITED

## MASTER CONTRACT

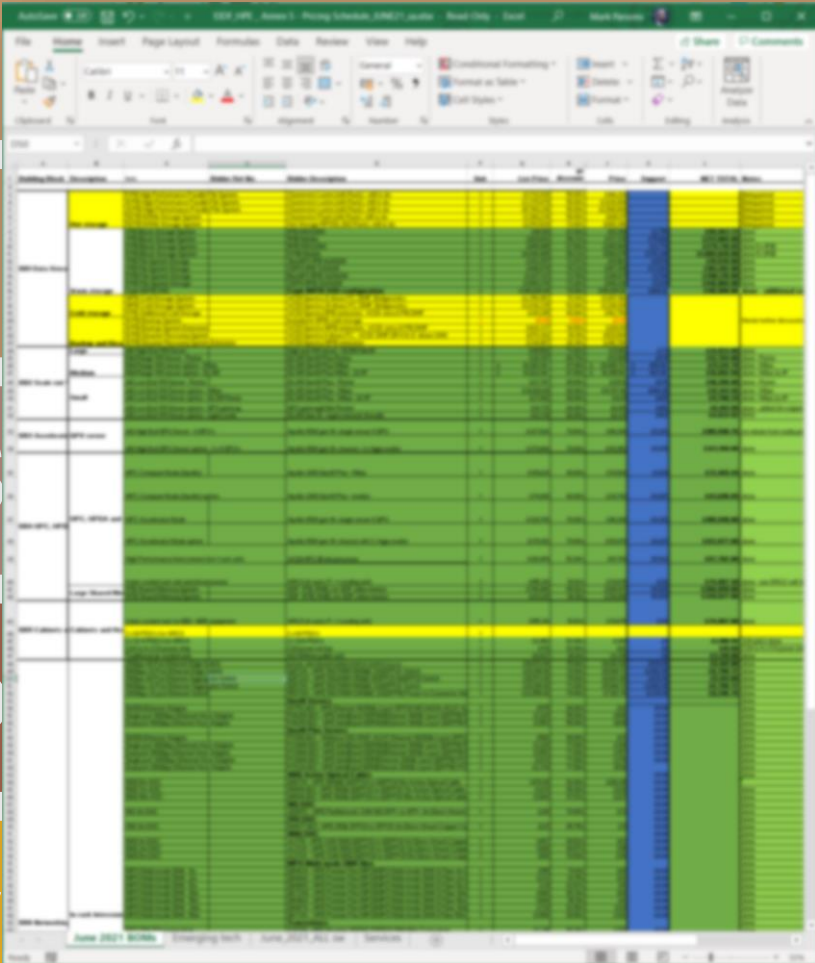
for the provision of technology products and  
related services

  
Pinsent Masons

Warm S

DR / Co

/ Secur



The screenshot shows an Excel spreadsheet with a complex data table. The table has multiple columns and rows, with some cells highlighted in yellow and others in green. The spreadsheet is titled 'Building Block Approach' and contains various data points, likely related to the 'Building block approach' mentioned in the slide title. The table is organized into several sections, with some rows highlighted in yellow and others in green, suggesting different categories or statuses. The spreadsheet is displayed in a window titled 'Building Block Approach - Excel'.





## Computer Room 4

£20m – CR 4 + PR D

£8.6m – 30MVA additional power


Space for 270 standard racks

Opened Dec 2020



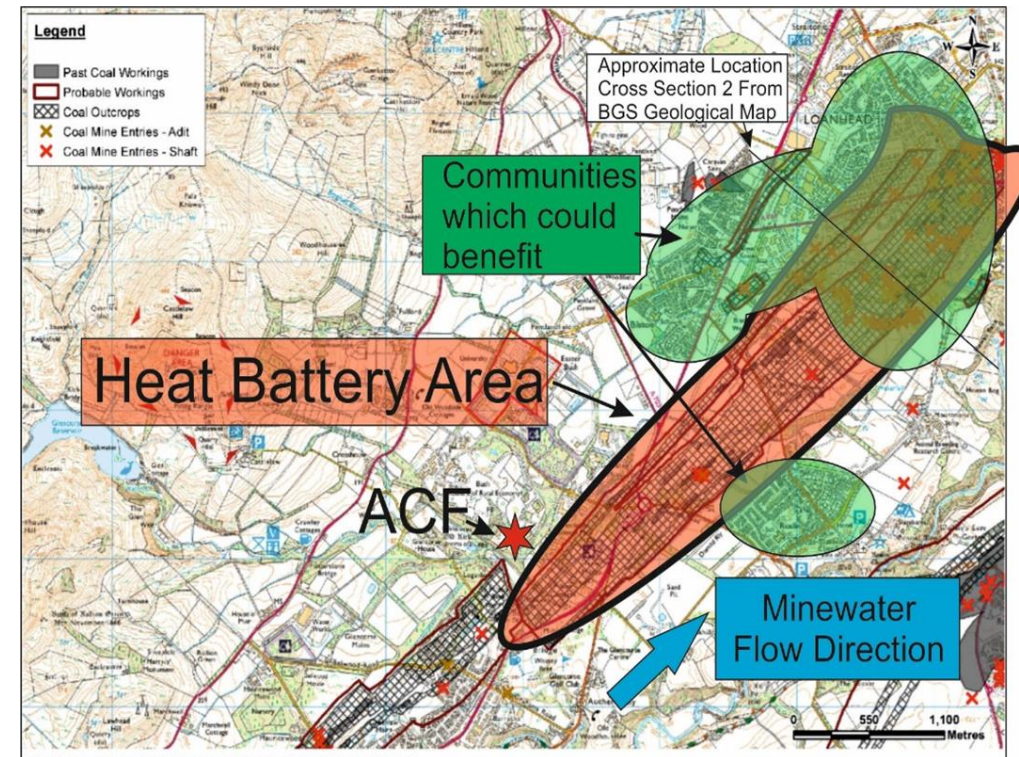


# Aiming for better than Net Zero



- Funding now secured to drill three 150m test wells
- Extraction point, Injection point and Test point
- Minewater temperature is 14C

- Detailed feasibility study completed to use hot water to heat abandoned mine workings
- Will create geothermal heat battery for us by homes, public and commercial buildings
- Battery will extend into South Edinburgh







# MACHINE HISTORY

## EPCC hardware timeline

**2019** Fulhame, HPE  
Apollo 70 (4,096 Arm  
ThunderX2 cores)



**2021** DiRAC ATOS  
Sequana XH2000  
(4,416 cores/456  
GPUs)



**2021** ARCHER2  
HPE Cray Ex  
(750,080 cores)



**2019** NEXTGenIO,  
Custom Fujitsu  
PRIMERGY design  
(1,632 cores, 102TB  
persistent memory)



**2018** DiRAC  
HPE SGI 8600  
(35,424 cores/  
32 GPUs)



**2017** Cirrus SGI  
ICE XA Cluster  
(13,248  
cores/  
152 GPUs)



**2014** ARCHER  
Phase 2: Cray XC30  
(118,080 cores)



**2013** ARCHER  
Phase 1: Cray XC30  
(72,192 cores)



**2012** UK-RDF  
(23PB)



**2011** DiRAC IBM  
BlueGene/Q (98,304  
cores)



**2011** EDIM1  
(240  
cores,  
750 TB  
disk)



**2005** IBM  
BlueGene/L  
(2,048 cores)



**2005** HPCx  
Phase 2a: IBM p5-575  
(1,536 processors)



**2006** HPCx  
Phase 3: IBM p5-575  
(2,560 processors)



**2007** FHPCA  
Maxwell (64 FPGAs)



**2007** HECToR  
Phase 1: Cray XT4  
(11,328 cores)



**2009** HECToR  
Phase 2a: Cray XT4  
(22,656 cores)



**2010** HECToR  
Phase 2b: Cray XE6  
(45,544 cores)



**2011** HECToR  
Phase 3: Cray XE6  
(90,112 cores)



**2004** QCDOC  
(14,464 processors)



**2004** HPCx  
Phase 2: IBM p690+  
(1,600 processors)



**2002** HPCx  
Phase 1: IBM p690  
(1,280 processors)



**2002** Sun Fire  
E15K (52  
processors)



**2002** Sun Fire  
6800  
Cluster (66  
processors)



**1997** Hitachi  
SR2201  
(8 processors)



**1997** Cray T3E  
(344 processors)



**1996** Cray J90  
(10 processors)



**1982** ICL DAPs  
(2 x 4,096 processors)



**1986** Meiko T800  
CS (400 processors)



**1988** AMT  
DAP608 (1,024  
1-bit processors)



**1990** Meiko i860 CS  
(64 processors)



**1991** TMC CM-200  
(16k 1-bit processors)



**1992** Meiko i860 CS  
(16 processors)



**1994** Cray T3D  
(512 processors) +  
CRAY Y-MP



**1995** Meiko CS-2  
(22 processors)





Pre-EPCC

EPCC



## ARCHER2 is finally here

- The 23 cabinet system finally opened for all users on 22nd November
- Very difficult 18 months
- No 22 in Top 500 – 19.5 Petaflop/s HPL
- Busy from Day 1 – and has remained busy





# The Supercomputing versus Data Science conundrum

## Supercomputer system

- Largely batch system based
- Bare metal optimised for numerical performance – limited memory
- Limited flexibility from back-end
- OS typically Linux
- Non-standard network
- High performance parallel storage e.g. Lustre

This environment over-constrains the software that can be run on it


## Data Science system

- Virtualised using e.g. OpenStack
- Kubernetes widely used to deploy containerised software
- Distributed storage e.g. CEPH
- Range of workloads – often tunable
- Full internet connectivity required by applications and operating system
- OS typically Ubuntu

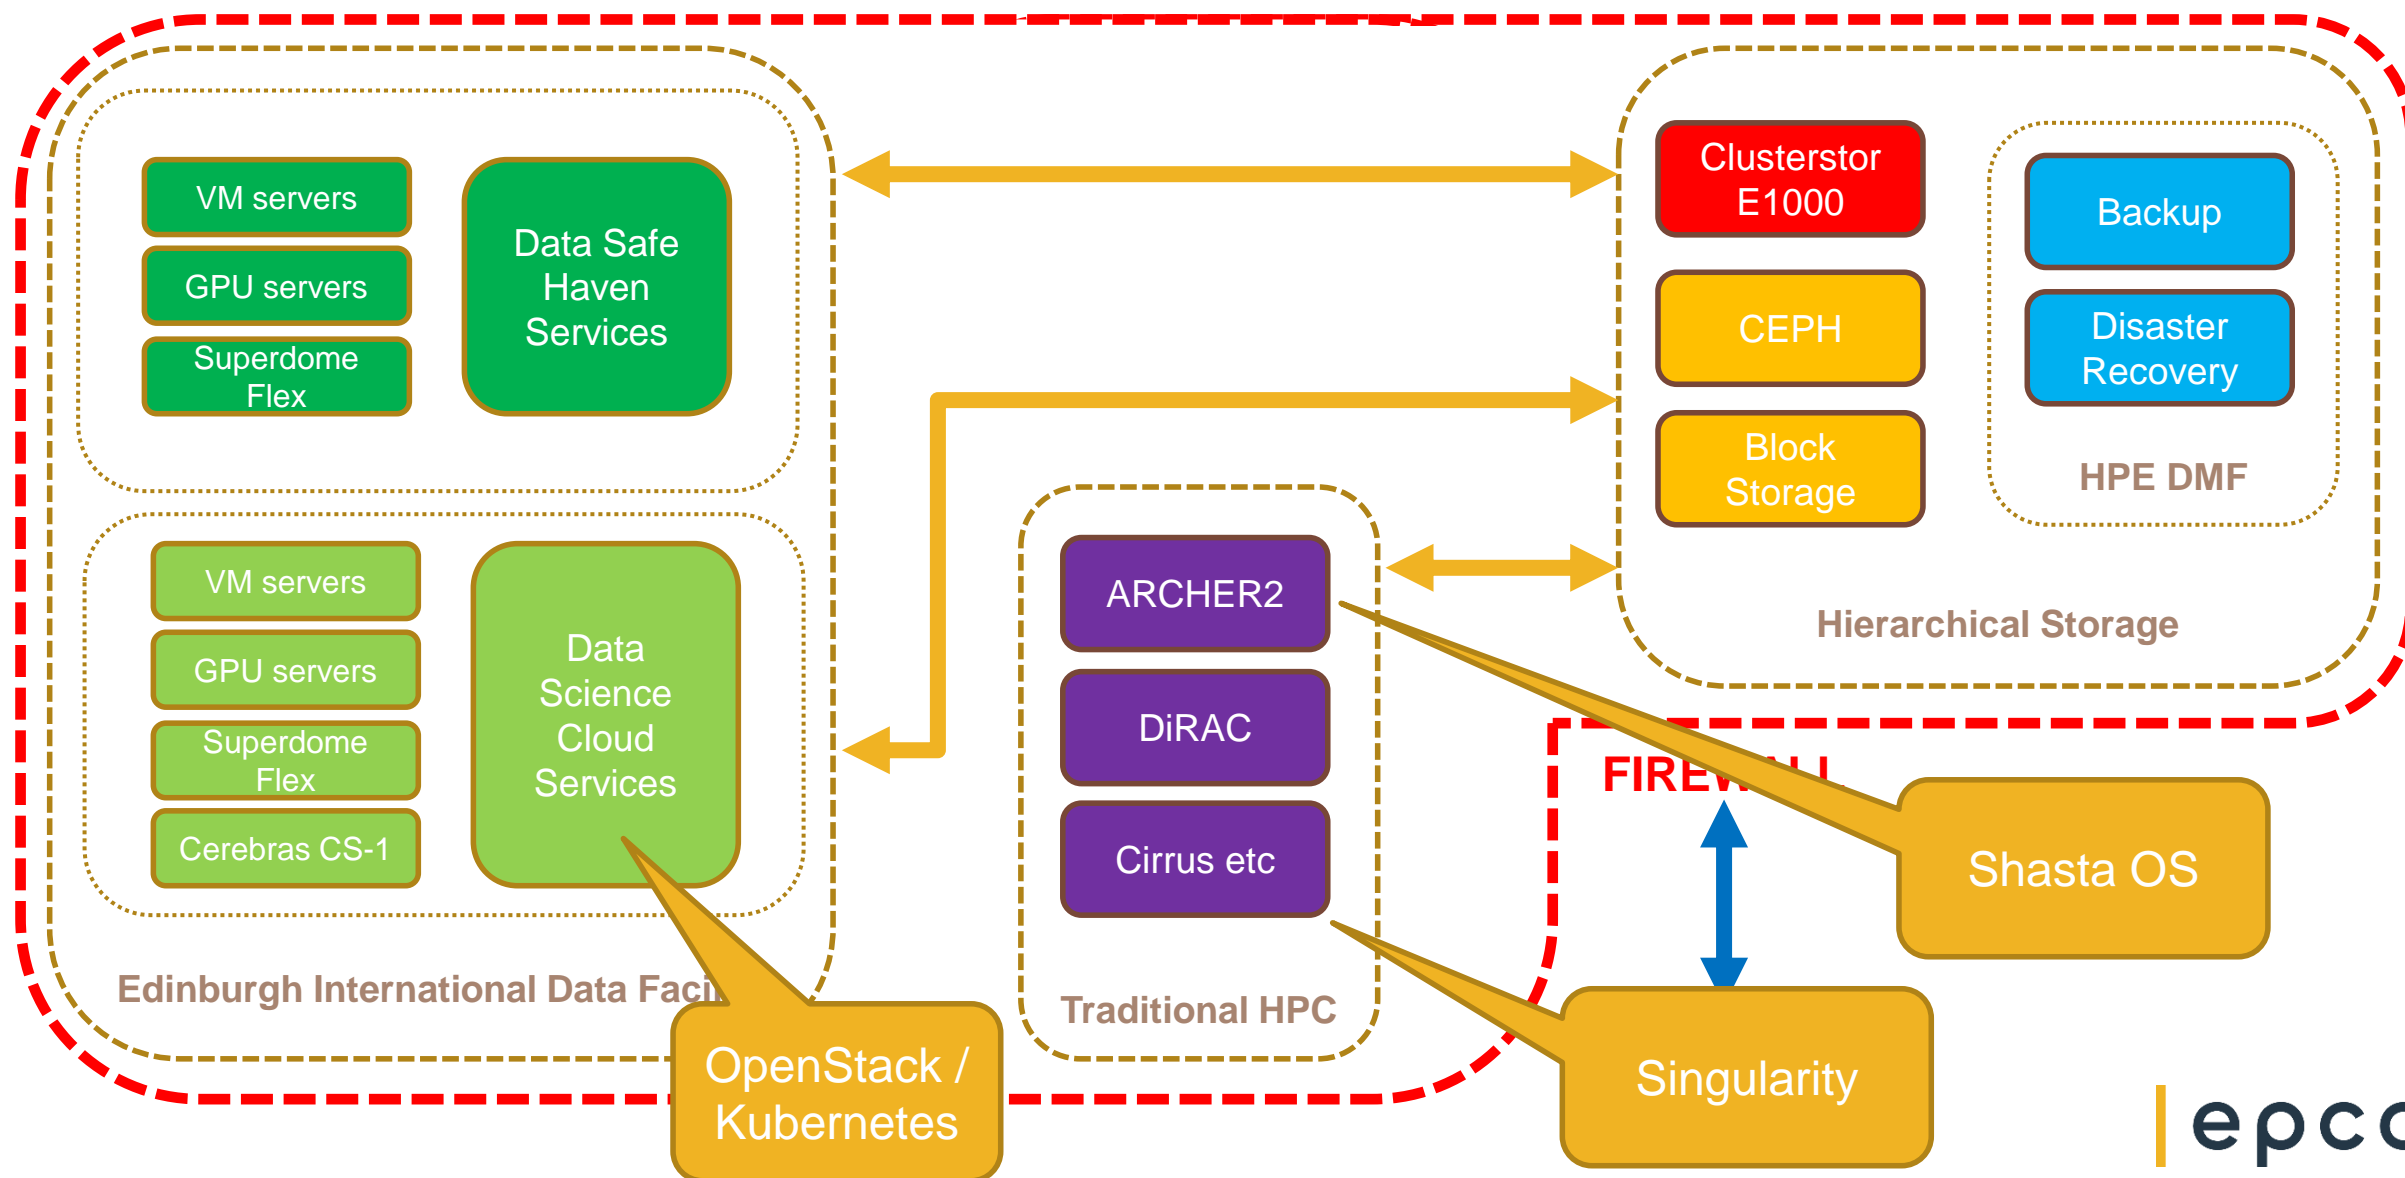
This environment often runs out of horsepower for demanding users

## Technology – recent Exascale vendor briefings

- Memory is changing
  - Many Exascale blades include High Bandwidth Memory
    - Some designs have no DRAM at all
  - But recently LPDDR5 is being mentioned more
- Four-way competition for CPUs and/or GPUs
  - **INTEL** versus **AMD** versus **ARM** versus **NVIDIA**
- GPUs market is broadening
  - AMD is strongly competing with NVIDIA
- Cabinet energy densities are rocketing
  - Today's 80-100KW cabinets will be eclipsed by cabinets at 300KW+
- Multicore CPUs are also getting AI Deep Learning features





# Edinburgh International Data Facility – Data Science meets HPC






# The data science middle ground?

- As data science challenges become more demanding VM servers run out of power
- But the jump to traditional HPC is too large
- Large shared-memory systems are proving very useful middle ground
- EIDF has three HPE Superdome Flex systems
  - Each has 576 cores and 18TB shared memory
- Three use cases
  1. In the Data Science Cloud
  2. Inside the Data Safe Haven
  3. Hosting our Cerebras CS-1 AI supercomputer

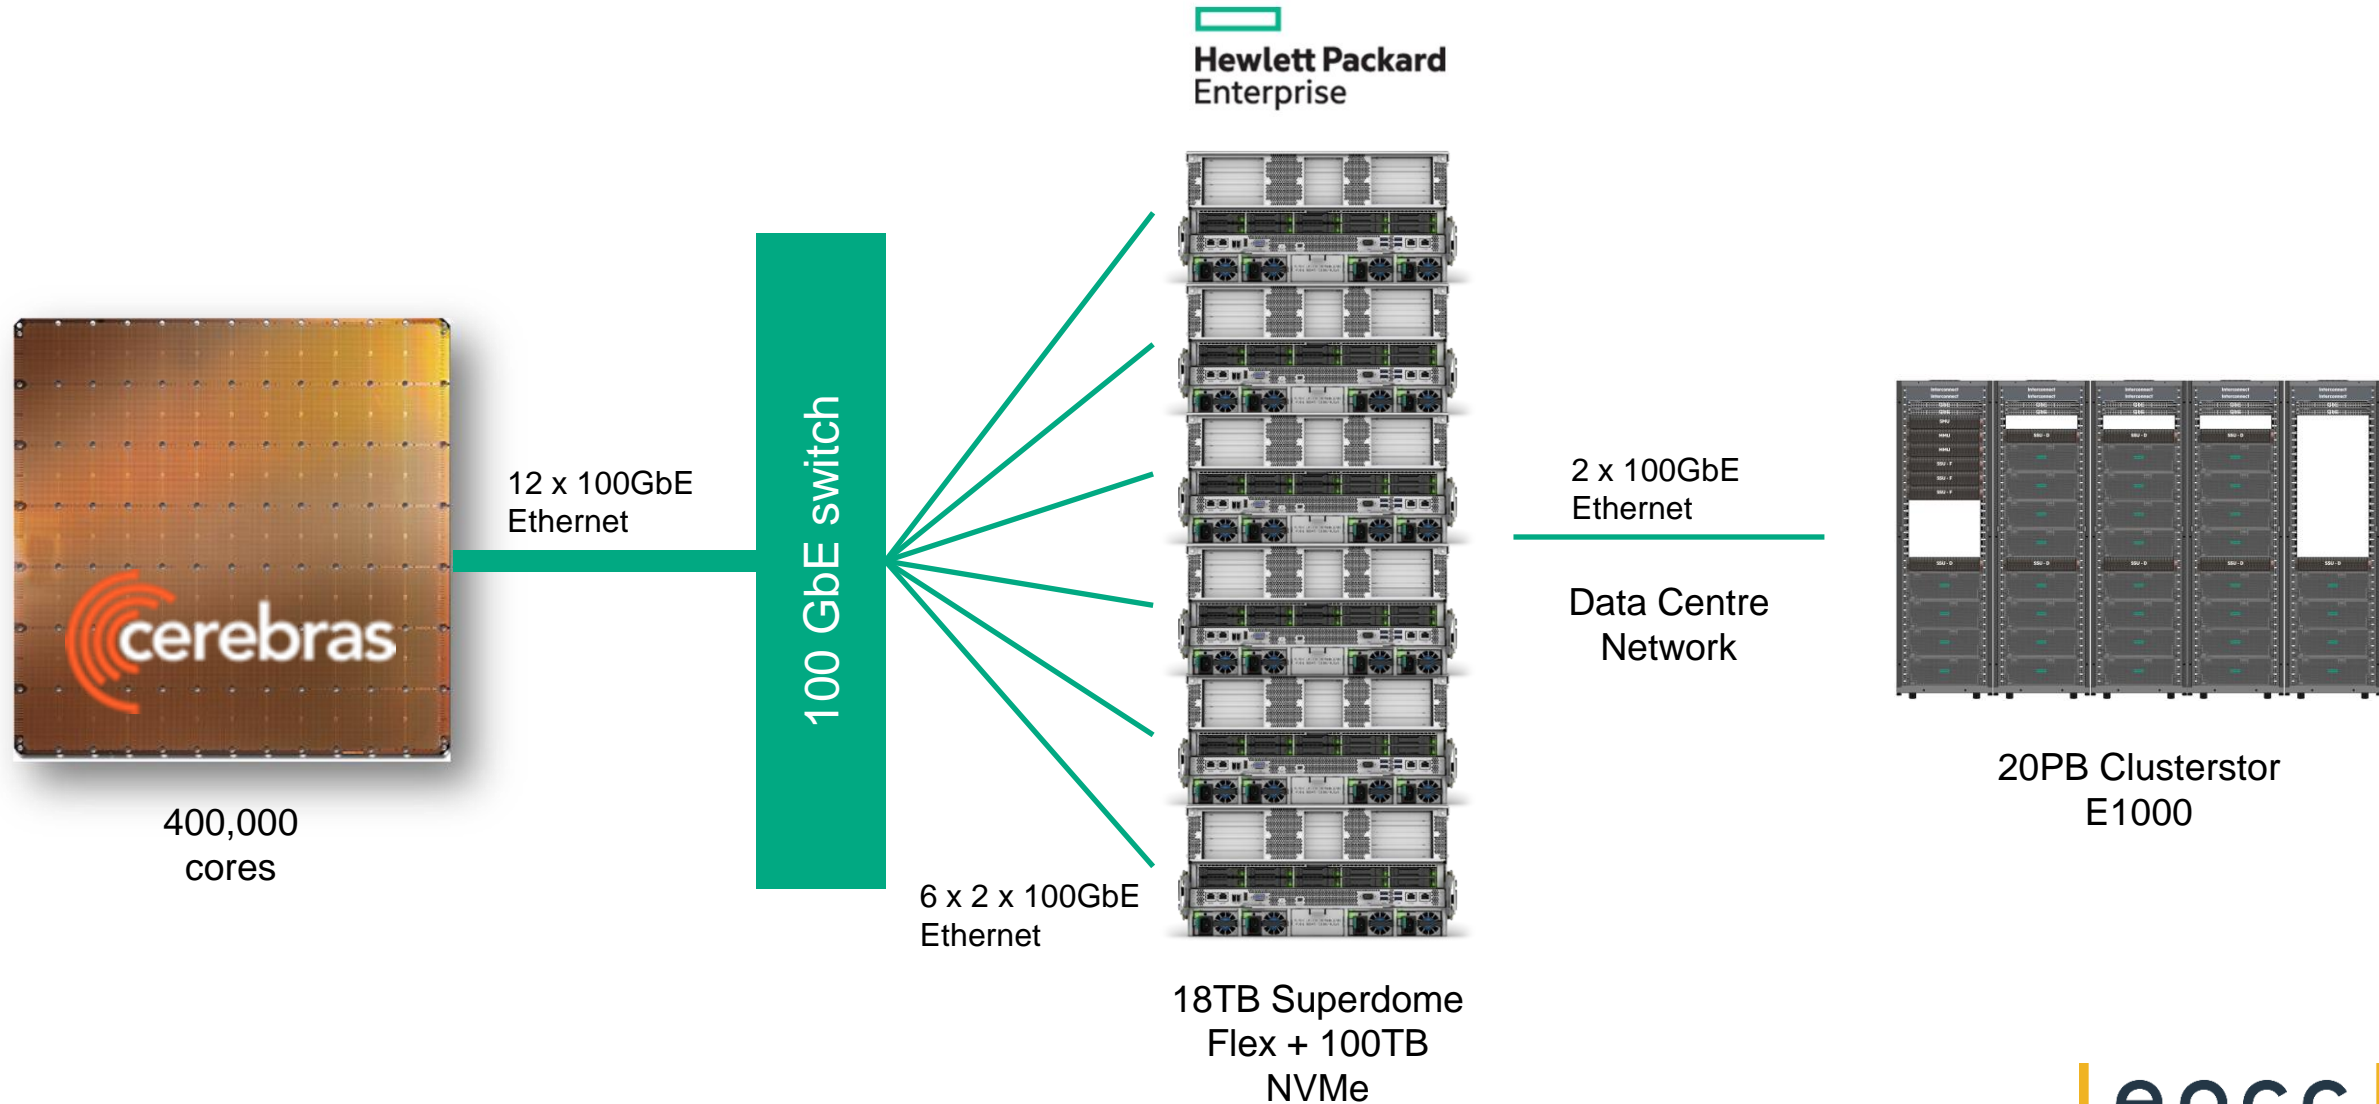


## Cerebras CS-1 arrives at EPCC

- EPCC has installed Europe's first Cerebras CS-1 – arrived March 2021
  - 400,000 AI cores optimised for sparse linear algebra – 1.2 trillion transistors
  - 18GB on-chip SRAM
  - 100 Pb/s internal interconnect with 1.2Tb/s ethernet connection to host
  - Focussed on TensorFlow and PyTorch etc
- Hosted within Edinburgh International Data Facility using 18TB SuperDome Flex
- Focussed on largest Deep Learning problems




# Cerebras CS-1 system setup





# 2D Mesh of 400,000 fully programmable processors

A cluster on a wafer




## First projects:

- Text data mining
- GWAS
- Cybersecurity
- Biomedical AI
- Natural Language Processing
- ML benchmark
- Robotics

No need for cache. No OS on the processors.  
Single-cycle message overhead.

## First projects on CS-1 underway

- The system went live in early May
- Initial projects include:
  - Experiments with running **BERT models for NLP** using Tensorflow
  - **Text data mining** in the social sciences
  - Investigation of **GCN/LSTM/Conv1D networks** using Tensorflow
  - **Bio-medical** AI PhD projects
  - **Cyber-security** AI project (commercial company)
  - **Robotics** AI project (commercial company)
  - Initial investigation of **programming models**



PyTorch now  
supported and  
numerical  
programming SDK

## Conclusion

- Building a converged Supercomputing and AI infrastructure is a fascinating challenge
- This is the future of supercomputing but there are many complex systems and software configuration challenges
- The EIDF is becoming one of the largest data science infrastructures in Europe
- Delivering Data Driven Innovation for Edinburgh & SE Scotland in partnership with HPE

