Service Specification

Please find specific information about each service below. This will help guide your application and dertermine which service to apply to.

ARCHER2

Service details

Service Contact Details	support@archer2.ac.uk
Service Webpage	https://www.archer2.ac.uk/
Service Reference	ARCHER2 PR17125

Hardware and Technical specifications

System name	ARCHER2
Compute nodes & Processors	5,848 compute nodes, each with dual AMD Rome 64 core CPUs at 2.2GHz, for 748,544 cores in total and 1.57 PBytes of total system memory
Interconnect	Cray Slingshot
Storage	14.5 PBytes of Lustre work storage in 4 file systems
Software available	https://www.archer2.ac.uk/about/hardware.html
Additional information on the hardware available	https://www.archer2.ac.uk/about/hardware.html

Resources available through this call

Unit of Allocation	ARCHER2 allocates its compute resource in ARCHER2 Compute Units (CU). Please note: • 1 node hour on ARCHER2 costs 1 CU, unless jobs are submitted in low priority queues where a discount applies. • 1 CU on ARCHER2 should (at a minimum) provide at least as much scientific throughput as 1.5156 kAU on ARCHER for most codes. This is based on conservative estimates of the performance of ARCHER2 relative to ARCHER, and thus is subject to variability based on the code used.

Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	Up to 840000 CU, 2.5% of EPSRC's ARCHER2 compute.
% compute allocated to EPSRC mechanisms (including but not limit to this call)	~77-83%, this is the total % of ARCHER2 EPSRC can utilise each year i.e. EPSRC's ARCHER2 compute.
Storage available	N/A

Requirements on applications for the service

Eligible EPSRC research areas	All
Project length restrictions over and above those in the call	1 year
Maximum and Minimum requests	Users must request more than 4000 CU and less than the call budget, users who want less can use the Pump-priming access route to ARCHER2 from the start of the service, or the instant access route to ARCHER beforehand http://www.archer.ac.uk/access/instant-access/.

Cirrus

Service details

Service Contact Details	support@cirrus.ac.uk
Service Webpage	http://www.cirrus.ac.uk/
Service Grant Reference	EP/P020267/1 (Phase I) EP/T02206X/1 (Phase II)

Hardware and Technical specifications

System name Cirrus HPE/SGI ICE XA Cluster

Compute nodes:

280 dual CPU compute nodes and 2 quad GPU nodes (Phase I)
144 NVIDIA V100 GPUs and an accompanying fast storage layer (Phase II)

Processor:

Cirrus standard compute nodes each contain two 2.1 GHz, 18-core Intel Xeon E5-2695 (Broadwell) series processors. Each of the cores in these processors support 2 hardware threads (Hyperthreads), which are enabled by default. The standard compute nodes on Cirrus have 256 GB of memory shared between the two processors. The Cirrus GPU compute nodes each contain two 2.4 GHz, 20-core Intel Xeon Gold 6148 (Skylake) series processers. Each of the cores in these processors support 2 hardware threads (Hyperthreads), which are enabled by default. The nodes also each contain four NVIDIA Tesla V100-PCIE-16GB (Volta) GPU accelerators connected to the host processors and each other via PCIe.

Cirrus Phase II adds a GPU node upgrade with a further 36 'Plainfield' blades (single GPU node with two Intel processors and four GPU's) into the empty 4th rack of the system along with the necessary power supplies, EDR IB switches and cables. These blades are similar to the two in the existing system, except that they will have Intel 'CascadeLake' processors (6248), 2933 MHz memory and will use EDR IB mezzanine cards and EDR IB switches. Each GPU node will have four NVIDIA V100's (16GB) for a total of 144 GPU's.

Storage:

A single filesystem Lustre file system has a total of 406 TiB available.
Cirrus Phase II will include fast storage to the new GPU nodes using HPE XFS/RPOOL with NVMe devices.

Interconnect	FDR Infiniband Hypercube
Software available	See Service Catalogue: http://www.cirrus.ac.uk/about/Cirrus_Service_Component_Catalogue.pdf
Additional information on hardware	See: http://www.cirrus.ac.uk/about/hardware.html

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	Projects awarded for autumn 2019 RAP ranged from 1,000,000 CPUhs to 4,000,000CPUhs
Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	Approximately 70,000,000 CPUhs available per year for Cirrus Phase I. Will have Cirrus Phase II resource available.
% compute allocated to EPSRC mechanisms (including but not limit to this call)	Cirrus Phase I 70% Cirrus Phase II 70%
Storage available	Default 250GiB per project, can allocate more space if justified

Requirements on applications for the service

Eligible EPSRC research areas	All
Project length restrictions over and above those in the call	1 year
Maximum and Minimum requests	Flexible with justification

CSD3

Service details

Service Contact Details	resources@hpc.cam.ac.uk
Service Webpage	www.hpc.cam.ac.uk
Service Grant Reference	EP/P020259/1

Hardware and Technical specifications

System name	peta4-skylake
Total compute nodes	1152x Dell PowerEdge C6420
EPSRC funded nodes	427 nodes: 65% via RAP, 15% Cambridge EPSRC users internal call, 20% industrial usage
Processor	Intel Xeon Gold 6142 CPU @ 2.60GHz (2 sockets, 32 cores)
Memory	192GB and 384GB
Interconnect	Intel Omni-Path

System name	peta4-cascadelake
Total compute nodes	672 x Dell PowerEdge C6420
EPSRC funded nodes	276 nodes: 65% via RAP, 15% Cambridge EPSRC users internal call, 20% industrial usage
Processor	Intel Xeon Platinum 8276 @2.2GHz (2 sockets, 56 cores)
Memory	192GB and 384GB
Interconnect	Mellanox HDR Infiniband

System name	peta4-knl
Total compute nodes	342x Dell PowerEdge C6320p
EPSRC funded nodes	190 nodes: 65% via RAP, 15% Cambridge EPSRC users internal call, 20% industrial usage
Processor	Intel Xeon Phi CPU 7210 @ 1.30GHz (single socket, 64 cores)
Memory	96GB
Interconnect	Intel Omni-Path

System name	wilkes2-gpu
Total compute nodes	90x Dell PowerEdge C4130
EPSRC funded nodes	200 GPUs: 65% via RAP, 15% Cambridge EPSRC users internal call, 20% industrial usage
Processor	Intel Xeon CPU E5-2650 v4 @ 2.20GHz (single socket, 12 cores)
GPUs	4x NVIDIA Tesla P100-PCIE-16GB per node
Memory	96GB
Interconnect	Mellanox EDR Infiniband

Storage	3120 TB lustre storage available to Tier2 Dell ME4 Series

Software available	A large range of software packages are preinstalled. Licenced packages (e.g. VASP) are available - please contact support@hpc.cam.ac.uk to determine access.
Additional information on the hardware available	https://www.hpc.cam.ac.uk/

Resources available through this call

Please note that allocations on CSD3 must start from February 2021

Indicative sizes of previously successful applications (not a restriction)	1M-10M Skylake CPU core hours 50,000-200,000 KNL node hours 50,000-200,000 GPU hours
Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	33M cpu hours on Skylake 38M cpu hours on Cascadelake 464K KNL node hours 488K P100 GPU hours
% compute allocated to EPSRC mechanisms (including but not limited to this call)	80% Skylake 80% Cascadelake 80% KNL 80% P100 GPU
Storage available	3120 TB lustre storage (available to all of Tier2)

Requirements on applications for the service

Eligible EPSRC research areas	Any
Project length restrictions over and above those in the call	1 year
Maximum and Minimum requests	None

Isambard GW4 Tier-2

Service details

Service Contact Details	Prof Simon McIntosh-Smith S.McIntosh-Smith@bristol.ac.uk +44 117 3315324
Service Webpage	https://gw4.ac.uk/isambard/
Service Grant Reference	EP/T022078/1

Hardware and Technical specifications

System name	Isambard 2 phase 1, a Cray XC50 Arm-based system
Compute nodes	332, each dual socket, 21,248 cores in total
Processor	Arm-based Marvell ThunderX2 32 core 2.1 GHz (2.5GHz turbo)
Interconnect	Cray Aries (same as ARCHER)
Storage	1 PByte
Software available	Full Cray software stack (Cray compiler, MPI, debugger, profiler, performance tools, math library) Full GNU software stack (compilers et at). Full Arm software stack (Clang/LLVM based compiler, math library, Allinea tools etc). Many widely used applications pre-installed on the system ready to use.
Additional information on the hardware available	Most codes should just compile and run on the Arm-based system just like they do on any other supercomputer. Most users won't even be able to tell they are on an Arm system.

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	1-5M core hours over 6 months.

Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	~516,000 node hours, 33M core hours. Projects requesting anywhere up to but not exceeding 10M core hours across the 6 month RAP period could therefore be considered (projects in the 1-10M range are most likely to succeed).
% compute allocated to EPSRC mechanisms (including but not limit to this call)	A total of 40% of the system is available during the 6 month period.
Storage available	Up to tens of TeraBytes per RAP project.

Requirements on applications for the service

Eligible EPSRC research areas	All areas.
Project length restrictions over and above those in the call	6 months typically for Access to HPC. Maximum 1 year by exception.
Maximum and Minimum requests	Minimum - 0.5M core hours. Maximum – 10M core hours.

Baskerville

Service details

Service Contact Details	baskerville-tier2@contacts.bham.ac.uk
Service Webpage	https://www.baskerville.ac.uk/
Service Grant Reference	EP/T022221/1

Hardware and Technical specifications

System name	Baskerville Accelerated Compute Facility
Compute nodes **Processor**	46 Lenovo ThinkSystem SD650-N V2 compute nodes, each with two 36-core Intel CPUs, 960GB SSD, 512GB DDR4 RAM, and 4 Nvidia Ampere A100-40 GPUs interconnected with NVLink and connected to the CPUs with PCIe4. GPU bandwidth is 6.2TB/s per node.
Interconnect	HDR Infiniband
Storage	Lenovo DSS-G providing 5.2PB of useable HDD and 0.5PB of useable SSD for high-throughput workloads.
Software available	We support the machine learning frameworks for GPU-accelerated computing (Tensorflow, PyTorch, Rapids.ai), and the major packages for molecular dynamics and materials modelling (GROMACS, LAMMPS). We encourage early enquiries from prospective users to help us develop the software base to match community needs.
Additional information on the hardware available	Baskerville is hosted at the University of Birmingham on behalf of EPSRC and the project partners Diamond Light Source, the Rosalind Franklin Institute, and the Alan Turing institute. It is designed for GPU-accelerated computing and aims to serve both machine learning and simulation communities. It is especially well suited for very large data workloads, featuring 186 Nvidia A100-40 GPUs and a large amount of high-speed storage to maximize throughput.

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	As Baskerville is a new facility, we are not yet able to provide data on previous applications.
Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	N/A
% compute allocated to EPSRC mechanisms (including but not limit to this call)	Up to 40% of Baskerville will be available through this call
Storage available	N/A

Requirements on applications for the service

Eligible EPSRC research areas	It is designed for GPU-accelerated computing and aims to serve both machine learning and simulation communities
Project length restrictions over and above those in the call	6 months typically for Access to HPC. Maximum 1 year by exception.
Maximum and Minimum requests	N/A

Sulis

Service details

Service Contact Details	sulis@warwick.ac.uk
Service Webpage	www.sulis.ac.uk
Service Grant Reference	EP/T022108/1

Hardware and Technical specifications

System name	Sulis
Compute nodes	25,216 CPU compute cores configured as 167 dual processor CPU compute nodes plus 30 nodes equipped with 3x Nvidia A100 GPUs
Processor	CPU compute nodes are equipped with dual AMD Epyc 7742 2.25GHz processors with 64 cores per socket and 512GB of DDR4-3200 RAM. GPU nodes additionally contain three Nvidia Ampere PCIe A100 GPUs with 40GB RAM. Each A100 can be partitioned into multiple virtual GPUs allowing for replication of several GPU-equipped workstation-scale configurations per node.
Interconnect	Mellanox ConnectX-6 HDR100 (100 Gbit/s) InfiniBand
Storage	2PB storage array implemented as a 200TB SSD tier and 1.8PB of underlying hard disk storage.
Software available	Core software will consist of standard compilers, MPI, numerical libraries etc, including Python frameworks (joblib, DASK) to support ensemble computing workloads, DMTCP for checkpoint/resume etc.
Additional information on the hardware available	N/A

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	As Sulis is a new facility, we are not yet able to provide data on previous applications.

Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	N/A
% compute allocated to EPSRC mechanisms (including but not limit to this call)	25% of CPU and GPU resource is available to use via this mechanism.
Storage available	Projects will be allocated 1TB of storage. Additional capacity can be provided with reasonable justification.

Requirements on applications for the service

Eligible EPSRC research areas	Projects suitable for this service will consist of workflows that concurrently execute calculations over ensembles of inputs via replication of many workstation-scale calculations, likely several per node. These should be arranged as a single job submission per calculation via GNU parallel , an appropriate high-level framework such as Python joblib, parallel tasks in Julia or similar.
Project length restrictions over and above those in the call	6 months typically for Access to HPC. Maximum 1 year by exception.
Maximum and Minimum requests	N/A

JADE

Service details

Service Contact Details	wes.armour@oerc.ox.ac.uk ResearchComputePlatforms@turing.ac.uk
Service Webpage	https://www.jade.ac.uk/
Service Grant Reference	EP/P020275/1

Hardware and Technical specifications

System name	JADE
Compute nodes	22x NVIDIA DGX-1V
Processor	Per node: 8x V100 16GB 2x 20 core Xeon E5-2698
Interconnect	InfiniBand ERD to filestore (so not really designed for heavy node-to-node communication). Ideal problems are those that fit within a node (so across 8x GPUs), codes that use NVLink (nccl) will also benefit.
Storage	512 GB DDR4 per node, 4x 2TB SSD (RAID0) per node. 1TB spinning disk filestore.
Software available	Anything in Nvidia NGC: https://ngc.nvidia.com/catalog/all?orderBy=modifiedDESC&pageNumber=1&query=&quickFilter=&filters=
Additional information on the hardware available	Designed for AI/Machine Learning. We also support some Molecular Dynamics (MD) work and have the standard set of MD codes available.

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	Varies significantly, we consider any application.

Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	Due to JADE's model of open access for AI/Machine Learning research, applicants who wish to conduct AI/Machine Learning projects on JADE should contact the service directly (at the above addresses) to discuss suitable levels of resource. Applicants in other research areas should refer to the eligibility section below.
% compute allocated to EPSRC mechanisms (including but not limit to this call)	80% of JADE's total capacity, split between AI and Molecular Dynamics research.
Storage available	Without asking users to remove data we have around ~200TB free at the moment.

Requirements on applications for the service

Eligible EPSRC research areas	AI & Machine Learning projects can get access through this call. Projects in the area of Molecular Dynamics should apply through HEC BioSim at http://www.hecbiosim.ac.uk/jade/application-form. Projects in other research areas are not eligible for access to JADE.
Project length restrictions over and above those in the call	6 months with the possibility of renewing for up to another 6 months
Maximum and Minimum requests	N/A

MMM Hub

Service details

Service Contact Details	rc-support@ucl.ac.uk
Service Webpage	https://mmmhub.ac.uk
Service Grant Reference	EP/T022213/1

Hardware and Technical specifications

System name	Thomas 2 (provisionally)
Compute nodes	576 HPE 40 core compute nodes with 192 GB of RAM, 3 with 3TB of RAM and 3 with 6TB of RAM
Processor	2x 20 core Intel Xeon Cascade Lake
Interconnect	Omnipath in 36 node 1:1 blocks
Storage	1PB Lustre
Software available	Standard UCL application stack: 700+ software modules supporting development tools (compilers from Intel, Python etc) and user applications (e.g. VASP, GROMACS, CP2K...)

Resources available

The MMM Hub operates a different allocation mechanism to the other Tier 2 Centres. Access is only available via membership of two HEC consortia and not via this call. For more information see: https://mmmhub.ac.uk/thomas

% compute allocated to EPSRC mechanisms (including but not limit to this call)	30% of the Thomas 2 facility

Requirements on applications for the service

Eligible EPSRC research areas	As noted above, research within the broad area of materials and molecular modelling.

NI-HPC (Kelvin-2)

Service details

Service Contact Details	v.purnell@qub.ac.uk
Service Webpage	www.ni-hpc.ac.uk
Service Grant Reference	EP/T022175

Hardware and Technical specifications

System name	Kelvin-2
Compute nodes	Standard: 60x Dell PowerEdge R6525 with 768GB RAM Hi-memory: 4x Dell PowerEdge R6525 with 2TB RAM GPU: 8 x Dell DSS8440 (each with 2x Intel Xeon Platinum 8168 24 Core CPU). Provides 32x NVIDIA Tesla v100 32GB
Processor	AMD Rome 2x64core 7702
Interconnect	Mellanox EDR infiniband
Storage	2PB usable lustre for scratch storage Metadata Servers: Dell R640 Metadata Targets: Dell Powervault ME2024 with 1TB SSD Object Storage servers: Dell Powervault ME4084
Software available	Centos 7.7 Lustre file system Alces flight cluster manager Applications – see attachment.

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	n/a

Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	2688 standard compute cores 2 hi-memory compute nodes 12 GPUs
% compute allocated to EPSRC mechanisms (including but not limit to this call)	35
Storage available	2PB shared scratch (no project quota planned in 2020-21)

Requirements on applications for the service

Eligible EPSRC research areas	Priority areas: computational neuroscience, advanced chemistry, innovative drug delivery, precision medicine, food fingerprinting and hydrogen deflagration Then: any EPSRC related area
Project length restrictions over and above those in the call	1 year
Maximum and Minimum requests	Max number of cores per job: 1344 Min number of cores per job: 600 GPUs per job max: 4 GPUs per job min: 1

Northern Intensive Computing Environment (NICE)

Service details

Service Contact Details	rebecca.appleby@durham.ac.uk +44 (0) 191 33 42520
Service Webpage	https://n8cir.org.uk/supporting-research/facilities/nice
Service Grant Reference	EP/T022167/1

Hardware and Technical specifications

System name	bede.dur.ac.uk
Compute nodes	32x IBM AC922 with 0.5TB and 4x32GB V100 GPU, 4x IBM IC922 with 256GB and 4xT4 GPU, 2x IBM IC922 with 256GB memory and FPGA
Processor	AC922: 2x16core 2.7Ghz Power 9. IC922: 2x20core 2.9Ghz Power 9.
Interconnect	Mellanox EDR
Storage	2Pb, 10GB/s Lustre filesystem for running jobs.
Software available	https://n8cir.org.uk/supporting-research/facilities/nice/software
Additional information on the hardware available	https://n8cir.org.uk/supporting-research/facilities/nice/hardware

Resources available through this call

Indicative sizes of previously successful applications (not a restriction)	N/A – but say 5000 node hours over 12 months for guidance
Indicative level of computational resource available through this call (subject to fluctuations in overall demand)	85000 node hours over 12 months/42,500 node hours per 6-month call.

% compute allocated to EPSRC mechanisms (including but not limit to this call)	38%
Storage available	N/A

Requirements on applications for the service

Eligible EPSRC research areas	All
Project length restrictions over and above those in the call	1 year
Maximum and Minimum requests	Requirement for requests to utilise unique features of Power 9 architecture, such as GPU/CPU memory coherence and/or multi-GPU. Whole nodes will be allocated to individual jobs.

Full Application Checklist

When submitting your application via the smart survey (see 'Submitting an Application') at https://www.smartsurvey.co.uk/s/DU5WMR/. Please ensure the following are attached in the specified sections of the survey.

Text/Document	**Maximum Page length**
Document 1 - Completed Application Form including:	
Objectives	½
Description of the proposed research and its context	2 ½
Importance	1
Expertise and track record of the team	1
Other associated resources	½
Resource Management	1 ½
Document 2: Diagrammatic Work Plan	1
Document 3: Completed Technical Assessment	N/A

Details on the expected content for each of these sections can be found in the 'Guidance on 'Writing an Application'' section.

In addition to the above, applicants also have the opportunity to add an optional cover letter as a separate attachment to the smart survey. This will only be seen by EPSRC. See 'Guidance on 'Writing an Application'' for details.

Related Content

Please find links to the relevant EPSRC policy below:

- Use of animals
- Responsible research and innovation
- Ethical considerations
- Equality, Diversity and Inclusion
- Conflicts of interest

Change log

Name	**Date**	**Version**	**Change**
Joseph Marriott	10/12/2020	1	N/A