



High Performance Computing: EPSRC Service Specification

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

ARCHER2	1
Cirrus.....	2
CSD3.....	4
Isambard	4
Baskerville	4
Sulis.....	5
JADE	7
MMM Hub.....	7
NI-HPC (Kelvin-2).....	7
Northern Intensive Computing Environment (Bede).....	9
Change log	10

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	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
Processors	
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 hardware available	https://www.archer2.ac.uk/about/hardware.html
Use cases particularly suited to this Service	Large capacity jobs

Resources available through this call

Unit(s) of Allocation	ARCHER2 allocates its compute resource in ARCHER2 Compute Units (CU). Please note: <ul style="list-style-type: none">• 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	Up to 3.4 MCUs, 10% of EPSRC's ARCHER2 compute
Indicative sizes of previously successful applications (not a restriction)	Access to HPC Call: Autumn 2022 Ranged from 12 kCUs – 611 kCUs
% compute allocated to EPSRC mechanisms (including but not limited to this call)	~77- 83%, this is the total % of ARCHER2 EPSRC can utilise each year i.e. EPSRC's ARCHER2 compute.
Storage available	Flexible with justification

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. Users who want less can use the Pump-priming access route to ARCHER2, see https://www.archer2.ac.uk/support-access/access.html .

Cirrus

Service details

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

Hardware and Technical specifications

System name	Cirrus
Compute nodes	280 dual CPU compute nodes and 2 quad GPU nodes 144 NVIDIA V100 GPUs and an accompanying fast storage layer.

Processors 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 processors. 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 added a GPU node upgrade with a further 36 'Plainfield' blades (single GPU node with two Intel processors and four GPU's)

Along with the necessary power supplies, EDR IB switches and cables. These blades are similar to the two in Cirrus Phase I, except that they have Intel 'CascadeLake' processors (6248), 2933 MHz memory and use EDR IB mezzanine cards and EDR IB switches.

Each GPU node have four NVIDIA V100's (16GB) for a total of 144 GPU's.

Interconnect FDR Infiniband Hypercube

Storage /home: CEPH filesystem for /home
 /work: Lustre file system (406 TB) (will shortly be replaced with 1PB E1000)
 HPE XFS/RPOOL with NVMe devices for fast storage

Software available See Service Catalogue:
http://www.cirrus.ac.uk/about/Cirrus_Service_Component_Catalogue.pdf

Additional information on hardware available See: <http://www.cirrus.ac.uk/about/hardware.html>

Use cases particularly suited to this Service Good mixture of use cases including CPU and GPU

Resources available through this call

Unit(s) of Allocation	CPUhs and GPUhs
Indicative level of computational resource available through this call	25 MCPUh 800 kGPUh
Indicative sizes of previously successful applications (not a restriction)	Autumn 2022 Call: CPU requests ranged from 2 KCPUh – 6.9 MCPUhs GPU requests ranged from 31 KGPUhs – 96 KGPUhs
% compute allocated to EPSRC mechanisms	70%

(including but not limited to this call)

Storage available Defaults are /home 100GiB and /work 250GiB, 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	None
Maximum and Minimum requests	Flexible with justification

CSD3

The CSD3 Service will be decommissioned during the next quarter so no HPC resources are available for this Access 2 HPC call.

Isambard

The Isambard2 Service will be decommissioned during the next quarter so no HPC resources are available for this Access 2 HPC call.

The Isambard3 Service will be fully operational for the next Access 2 HPC call, so HPC resources will be available for projects starting around July 2024.

Baskerville

Service details

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

Hardware and Technical specifications

System name	Baskerville
Compute nodes	52 Lenovo SD650-N V2 liquid cooled compute nodes
Processors	<ul style="list-style-type: none">All nodes host 2x Intel® Xeon® Platinum 8360Y CPUs, each with 36 cores at 2.4GHz (with boost to 3.5GHz)46 nodes host 4x NVIDIA A100 40GB GPUs 6 nodes host 4x NVIDIA A100 80GB GPGPUs
Interconnect	All nodes are connected by <ul style="list-style-type: none">1x 25GbE NVIDIA® Mellanox® (on-planar ConnectX-4 port)1x HDR (200Gbps) NVIDIA Mellanox Infiniband port (ConnectX-6 PCIe gen4 adapter)

Storage	5PB of HDD and 0.5PB of SSD are available. The storage systems are Lenovo DSS-G running IBM® Spectrum Scale
Software available	A full list of installed software is available at https://apps.baskerville.ac.uk/applications/
Additional information on hardware available	N/A
Use cases particularly suited to this Service	Baskerville is designed primarily for GPU-accelerated workloads and is suitable for both single-GPU and multi-GPU jobs, with parallelisation over up to 8 nodes (32 GPUs) supported as standard. Workloads that are predominantly CPU-focussed will not be accepted.

Resources available through this call

Unit(s) of Allocation	GPU hours
Indicative level of computational resource available through this call	337900 GPU hours available for the first half of 2024 allocation period. 360000 GPU hours available for the second half of 2024 allocation period.
Indicative sizes of previously successful applications (not a restriction)	Previous successful applications have been awarded between 5000 and 150,000 GPU hours in a six-month period.
% compute allocated to EPSRC mechanisms (including but not limited to this call)	40%
Storage available	Up to 500TB of storage can be requested. Globus is available to support parallel data transfer.

Requirements on applications for the service

Eligible EPSRC research areas	All areas are supported.
Project length restrictions over and above those in the call	Applications for 6 or 12 month periods will be accepted.
Maximum and Minimum requests	None

Sulis

Service details

Service Contact Details	sulis@warwick.ac.uk
Service Webpage	https://sulis.ac.uk
Service Reference	https://sulis-hpc.github.io

Hardware and Technical specifications

System name	Sulis
Compute nodes	Dell PowerEdge R6525 compute nodes each with 128 cores per node; 512 GB DDR4-3200 RAM per node
Processors	2 x AMD EPYC 7742 2.25 GHz 64-core processors per node GPU nodes contain 3x Nvidia A100 40 GB (PCIe variant)
Interconnect	Mellanox ConnectX-6 HDR100 (100 Gbit/s) InfiniBand
Storage	2 PB IBM Spectrum Scale (GPFS)
Software available	Large suite of compilers/libraries available via environment modules, Python and R packages and end-user applications. Others installed on request. Containers are supported via Singularity.
Additional information on hardware available	Sulis has 4 high memory (1 TB RAM) and 3 very high memory (4 TB of RAM) nodes available for pre- and post-processing of data, memory intensive analytics etc.
Use cases particularly suited to this Service	Sulis is particularly suited for high throughput and ensemble computing workloads consisting of many small (workstation scale) calculations that can be bundled up into a smaller number of job submissions. See sulis-hpc.github.io/advanced/ensemble for examples. Sulis is less suited for large calculations that require high speed low-latency communication between GPUs or CPU nodes, e.g. high fidelity CFD simulations or data processing that relies on tightly coupled GPUs.

Resources available through this call

Unit(s) of Allocation	CPU core hour, GPU hour
Indicative level of computational resource available through this call	24M CPU core hours and 95k GPU hours.
Indicative sizes of previously successful applications (not a restriction)	Typically 1-5 million CPU core hours or 15-25k GPU hours over 12 months.
% compute allocated to EPSRC mechanisms (including but not limited to this call)	25%
Storage available	Default 2TB per user. Larger allocations can be accommodated on request. An SSD scratch partition is available for workloads performing intensive and non-sequential IO operations.

Requirements on applications for the service

Eligible EPSRC research areas	Any
--------------------------------------	-----

Project length restrictions over and above those in the call

Projects up to 12 months in duration are welcome

Maximum and Minimum requests

No explicit limits. Projects needing less than 100,000 CPU core hours or 1,000 GPU hours should request exploratory access via sulis.ac.uk rather than applying via this call.

JADE

The JADE2 Service will be decommissioned during the next quarter so no HPC resources are available for this Access 2 HPC call.

MMM Hub

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/young/>.

NI-HPC (Kelvin-2)

Service details

Service Contact Details v.purnell@qub.ac.uk

Service Webpage www.ni-hpc.ac.uk

Service Reference EP/T022175

Hardware and Technical specifications

System name Kelvin-2

Compute nodes Standard

- 60x Dell PowerEdge R6525 with 768GB RAM
- 21x Dell PowerEdge R6525 with 1TB 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
- 4 x XE8545 (each with 4x NVIDIA A100 80GB) AMD Rome 2x64core 7702

Processors 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
Additional information on hardware available	New infrastructure planning to come online in March 2023: Standard: 7x Dell PowerEdge R6525 with 1TB RAM Hi-memory: 2x Dell PowerEdge R6525 with 2TB RAM X-series compute: 2x Dell PowerEdge R6525 with 1TB RAM with dual 7773X CPU series – to review CFD optimisation possibilities New high-bandwidth scratch filesystem servers to optimise GPU throughput workloads.
Use cases particularly suited to this Service	Priority areas from original bid focussed on: computational neuroscience, advanced chemistry, innovative drug delivery, precision medicine, food fingerprinting and hydrogen deflagration Service is general purpose machine so can deliver on a wide range of use cases.

Resources available through this call

Unit(s) of Allocation	Core hours
Indicative level of computational resource available through this call	Average 7% of Kelvin2 resource. Based on 5 equal projects allocated sharing the 35%.
Indicative sizes of previously successful applications (not a restriction)	N/A – current users from calls have shown a very wide usage profile on the service.
% compute allocated to EPSRC mechanisms (including but not limited to this call)	35
Storage available	2PB shared scratch (no project quota planned in 2023-24)

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 (Bede)

Service details

Service Contact Details	Rebecca Appleby (Rebecca.appleby@durham.ac.uk), +44(0) 191 33 4250
Service Webpage	https://n8cir.org.uk/bede/
Service 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,
Processors	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://bede- documentation.readthedocs.io/en/latest/software/index.html
Additional information on hardware available	https://bede- documentation.readthedocs.io/en/latest/hardware/index.html
Use cases particularly suited to this Service	Accelerated computing that requires more than what can be provided by a single accelerator: Extending accelerator memory into CPU (memory coherence); using multiple accelerators per node & using multiple accelerators across nodes. Both simulation and AI/ML workloads.

Resources available through this call

Unit(s) of Allocation	Node hour / GPU hour
Indicative level of computational resource available through this call	40,000 node hours (160,000 GPUh) per year.
Indicative sizes of previously successful applications (not a restriction)	3,500 node hours over 12 months. (14,000 GPUh)
% compute allocated to EPSRC mechanisms (including but not limited 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 10,000 node hours (40,000 GPUh) / year maximum

Change log

Name	Date	Version	Change
Richard Bailey	01/03/2023	1	Initial version
Richard Bailey	27//09/2023	2	Updated for Autumn 2023 call