Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing Service Specification

Jade 2.5

Service details

Service Contact Details	jade-support@arc.ox.ac.uk
Service Webpage	https://www.jade.ac.uk/

Hardware and Technical specifications

System name	Jade 2.5
Compute nodes	3× Lenovo ThinkSystem SR685a V3
Processors	2× AMD EPYC 9534 64-Core CPUs per node 8x AMD MI300 GPUs (PCIe) per node
Interconnect	Mellanox NDR
Storage	32TB of local SSD scratch network bulk storage WEKA scratch
Software available	Apptainer, GCC, ROCm
Additional information on hardware available	2.2 TB DDR5 Registered ECC RAM 8× AMD MI300X GPUs with 192GB RAM per node System supplied by Lenovo System integration, hosting and management by ARC, University of Oxford
Use cases particularly suited to this Service	

Resources available through this call

Unit(s) of Allocation	GPU/h
Indicative level of computational resource available through this call	42,000 GPU/h over a six-month period
Indicative sizes of previously successful applications (not a restriction)	4,000 GPU/h
% compute allocated to UKRI mechanisms (including but not	40%

limited to this call)

Storage available	10TiB per project

Requirements on applications for the service

Project length restrictions over and above those in the call	6 months
Maximum and Minimum requests	20,000 GPU/h (Maximum), 4,000 GPU/h (Minimum)

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

ARCHER2

Service details	
Service Contact Details	support@archer2.ac.uk
Service Webpage	https://www.archer2.ac.uk/
Hardware and Technical specifications	
System name	ARCHER2
Compute nodes	5,860 compute nodes
Processors	Each node has 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	Information on the HPE-provided software is on https://www.archer2.ac.uk/about/hardware.html Information on software packages and libraries installed by the ARCHER2 team at EPCC can be found at https://docs.archer2.ac.uk
Additional information on hardware available	https://www.archer2.ac.uk/about/hardware.html
Use cases particularly suited to this Service	Large scale jobs
Resources available through this call	
Unit(s) 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.
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)	From around 10,000 CUs to more than 1,000,000 CUs

% compute allocated to UKRI mechanisms (including but not limited to this call)	Around 95%
Storage available	Flexible with justification
Requirements on applications for the service	
Project length restrictions over and above those in the call	1 year The current planned end date for ARCHER2 is November 2026.
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

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

Baskerville

Service details

Service Contact Details	baskerville-tier2@contacts.bham.ac.uk
Service Webpage	https://www.baskerville.ac.uk/

Hardware and Technical specifications

System name	Baskerville
Compute nodes	50 Lenovo liquid cooled compute nodes, each providing four NVIDIA GPUs.
Processors	37 nodes with 2x Intel® Xeon® Platinum 8360Y, 512GB, and 4x NVIDIA A100 40G GPUs 11 nodes with 2x Intel® Xeon® Platinum 8360Y, 512GB, and 4x NVIDIA A100 80G GPUs 2 nodes with 2x AMD EPYC® 9554, 768GB, and 4x NVIDIA H100 80G GPUs
Interconnect	All nodes are connected by • 1x HDR (200Gbps) NVIDIA Mellanox InfiniBand port (ConnectX-6 PCIe gen4 adapter) • 1x 25GbE NVIDIA® Mellanox® (on-planar ConnectX-4 port) It is worth noting that the use of InfiniBand represents a valuable USP of Baskerville compared to other Tier 2 (and indeed Tier 1) providers, allowing it to excel at tasks where rapid communication across multiple nodes is required.
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 for GPU-accelerated workloads and is suitable for both single-GPU and multi-GPU jobs, with parallelisation 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 core hours (GPUh)
Indicative level of computational resource available through this call	175,000 GPUh per quarter
Indicative sizes of previously successful applications (not a restriction)	Previous successful applications have been awarded between 5,000 and 100,000 GPUh per quarter.
% compute allocated to UKRI mechanisms (including but not limited to this call)	40%
Storage available	Up to 100TB of storage can be requested. Globus is available to support parallel data transfer.
Requirements on applications for the service	
Project length restrictions over and above those in the call	3 months
Maximum and Minimum requests	No formal limits

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

Cirrus

Service details

Service Contact Details	support@cirrus.ac.uk
Service Webpage	https://www.cirrus.ac.uk

Hardware and Technical specifications

System name	Cirrus
Compute nodes	256 nodes
Processors	73,728 cores
Interconnect	Dual-rail HPE Cray Slingshot 11
Storage	/home: Ceph file system (1PB) /work: E1000 Lustre file system (1 PB)
Software available	https://docs.cirrus.ac.uk/
Additional information on hardware available	http://www.cirrus.ac.uk/about/hardware.html
Use cases particularly suited to this Service	Good mix of use cases for CPU

Resources available through this call

Unit(s) of Allocation	CPU core hours (CPUh)
Indicative level of computational resource available through this call	258 M CPUh per year 65 M CPUh per quarter
Indicative sizes of previously successful applications (not a restriction)	Previous call on previous hardware: CPU requests ranged from 2 kCPUh – 2 MCPUh
% compute allocated to UKRI mechanisms (including but not limited to this call)	40% of compute allocated to UKRI
Storage available	Defaults are /home 100GiB and /work 250GiB, can allocate more space if justified

Requirements on applications for the service	
Project length restrictions over and above those in the call	Allocations will be made on a quarterly basis up to a maximum of 12 months.
Maximum and Minimum requests	Flexible with justification

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

Isambard 3

Service details

Service Contact Details	brics-enquiries@bristol.ac.uk
Service Webpage	https://docs.isambard.ac.uk

Hardware and Technical specifications

System name	Isambard 3
Compute nodes	384
Processors	NVIDIA Grace-Grace CPUs, 72 cores per socket, 144 cores per node at 3.1 GHz. 240 GBytes LPDDR5X memory providing ~1 TByte per second of memory bandwidth per node. NVIDIA's Grace CPUs implement the Arm instruction set (aarch64), rather than the x86 instruction set used by Intel and AMD.
Interconnect	HPE Slingshot 11 200 Gbps
Storage	2 PetaBytes HPE ClusterStor Lustre
Software available	Cray Programming Environment (CPE), GNU compilers and libraries, Clang/LLVM compilers and libraries, NVIDIA compilers and libraries. Further software described on service webpage.
Additional information on hardware available	See service webpage for detailed information for "Isambard 3 Grace"
Use cases particularly suited to this Service	General-purpose HPC codes with a focus on memory bandwidth.

Resources available through this call

Unit(s) of Allocation	Node hours (NH)
Indicative level of computational resource available through this call	Up to 60% of Isambard 3's compute resource is available for allocation through this call. For the 12 months of this call, this amounts to about 1.2 million node hours (1.2M NH)
Indicative sizes of previously successful applications (not a restriction)	We anticipate projects will apply for allocations in the range of tens of thousands to low hundreds of thousands of node hours.

% compute allocated to UKRI mechanisms (including but not limited to this call)	60% for UKRI in total.
Storage available	The 2 PetaByte storage system in total can support projects needing up to tens of TeraBytes each, with appropriate justification. Note that Isambard 3's storage is scratch space only, and project data that needs to be stored safely and securely beyond the project will need to be stored elsewhere, as per the Isambard 3 terms and conditions.

Requirements on applications for the service

Project length restrictions over and above those in the call	12 months
Maximum and Minimum requests	Applications will preferably demonstrate that the codes have already successfully run on an Arm-based platform, such as a previous incarnation of Isambard, or on Fugaku or AWS Graviton etc. If this is not possible, reasonable evidence that the codes are not x86-specific should be provided. Additionally, evidence that the codes have already been shown to scale well to node sizes appropriate for Isambard 3 and Tier 2 more generally should be provided.

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

Northern Ireland HPC

Service details	
Service Contact Details	v.purnell@qub.ac.uk
Service Webpage	www.ni-hpc.ac.uk
Hardware and Technical specifications	
System name	Kelvin-2
Compute nodes	Standard • 60x Dell PowerEdge R6525 with 768GB RAM • 28x Dell PowerEdge R6525 with 1TB RAM X-series compute: • 2x Dell PowerEdge R6525 with 1TB RAM with dual 7773X CPU series Hi-memory • 6x Dell PowerEdge R6525 with 2TB RAM GPU • 32 x NVIDIA Tesla v100 GPUs in 8 nodes. • 16 x NVIDIA Tesla A100 GPUs in 4 nodes. • 4 x AMD MI300X GPU's in 1 node. • 4 x Nvidia H100 GPUs in 1 node. • 4 x Intel Max 1100 GPUs in 1 node.
Processors	AMD Rome 2x64core 7702
Interconnect	Mellanox EDR infiniband
Storage	4PB 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	Rocky8.10 operating system Lustre scratch file system Alces flight cluster manager
Additional information on hardware available	

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 Kelvin-2 is a general purpose cluster so can deliver on a wide range of use cases.
Resources available through this call	
Unit(s) of Allocation	**CPU core hours (CPUh), GPU core hours (GPUh)**
Indicative level of computational resource available through this call	**8.5M CPUh per quarter** **33,000 GPUh per quarter**
Indicative sizes of previously successful applications (not a restriction)	**2M CPUh per quarter** **10,000 GPUh per quarter**
% compute allocated to UKRI mechanisms (including but not limited to this call)	**35**
Storage available	**4PB shared scratch**
Requirements on applications for the service	
Project length restrictions over and above those in the call	**None**
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**

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – Autumn 2025 Service Specification

Sulis

Service details	
Service Contact Details	sulis@warwick.ac.uk
Service Webpage	https://sulis.ac.uk
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) or 3 x NVIDIA L40 48 GB RAM Details at https://sulis-hpc.github.io/techspecs/
Interconnect	Mellanox ConnectX-6 HDR100 (100 Gbit/s) InfiniBand
Storage	800TB (SSD) 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	15M CPU core hours and 40k 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 UKRI mechanisms (including but not limited to this call)	25%
Storage available	Default 2TB per user. Larger allocations might be accommodated on request depending on available space.
Requirements on applications for the service	
Project length restrictions over and above those in the call	Projects are limited to finish by end of March 2026 after which machine status is uncertain. Exploratory access will be granted in advance of project start to facilitate preparation to exploit allocated time.
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.

Engineering and Physical Sciences Research Council

UK Research and Innovation

Access to High Performance Computing – spring 2025
Service Specification

Northern Intensive Computing Environment (Bede)

Service details

Service Contact Details	Rebecca Appleby (arc.admin@durham.ac.uk), +44(0) 191 33 4250
Service Webpage	https://n8cir.org.uk/bede/

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, 6/7*x Grace Hopper (96GB GPU RAM) 480GB LPDDR5X RAM 1/0*x Grace-Grace 480GB LPDDR5X RAM 1 DUAL 2xGrace Hopper (144GB GPU RAM) 480GB LPDDR5X RAM *Grace-Grace node normally configured as login environment, but can be swapped in as a compute node with a Grace-Hopper node configured for login if requested.
Processors	AC922: 2x16core 2.7Ghz Power 9. IC922: 2x20core 2.9Ghz Power 9. Grace Hopper: Grace 72-core aarch64 CPU 3.483Ghz, Hopper H100 96/144GB HBM3e
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	GPU core hours (GPUh)
Indicative level of computational resource available through this call	23,500 GPUh on V100 1st Jan - 31st March 2026. Access to Grace-Hopper seed units for evaluation
Indicative sizes of previously successful applications (not a restriction)	10,000 GPUh over 9 months
% compute allocated to UKRI mechanisms (including but not limited to this call)	38%
Storage available	N/A

Requirements on applications for the service

Project length restrictions over and above those in the call	9 months
Maximum and Minimum requests	40,000 GPUh / year maximum