

Guidelines for peer reviewed research project access to NeSI High Performance Computing facilities (v 1.1)

If you are applying for or have already received a peer reviewed grant for your research project and require High Performance Computing, NeSI can help you complete your research at very low cost, right now. Funded to support and enhance research that depends on computation, the New Zealand eScience Infrastructure (NeSI) consortium hosts specialised HPC facilities across NZ, can provide access to equipment that fits your needs, and is keen to work with you to realise your research goals.

Only research projects that contribute to the objectives of a research grant that has been awarded through a peer review process conducted across an institution or nationally/internationally will be eligible for peer reviewed research project access.

Researchers already awarded a peer reviewed research grant can apply for access, as covered in the section **Access for already peer reviewed projects**. If you do not have a peer reviewed research grant, there are other options, so please read on.

These guidelines outline **eligibility** and other **requirements** to researchers applying for / who have already received peer reviewed research grants. Find out how to plan for **access to the High Performance Computing (HPC)** facilities designed to meet the needs of New Zealand research communities, including:

- HPC facilities available to support researchers in their projects
- classes of allocation for a project
- support available for each class of allocation
- application proposal and review process for access to the HPC facilities
- role of NeSI's Technical Qualification Panel during proposal technical review
- costs and payment structure
- acknowledgement in publications
- security and data ownership

The benefits of including NeSI facilities within a research grant include higher priority access to our facilities, a significant in-kind contribution from NeSI to your project and advice on how best to successfully use your allocation.

Application Process overview

It's easy to include an initial estimate of NeSI facilities in a peer reviewed research proposal. If you already understand your computing needs, all you need do is:

1. Estimate the number of CPU core hours required to complete your project. To do this, use the **Costs** guidelines below and include costs as a project operating expense in your proposal.

Then, when requested by NeSI or the grant provider, or if you already have a peer reviewed research grant:

2. Submit a technical proposal for an allocation on NeSI facilities at www.nesi.org.nz. This proposal should justify your need to use NeSI's HPC facilities, and identify all resources needed including: processor core hours, data storage and external (i.e. NeSI) human resources required to complete the project, which will then undergo a technical review by our Technical Qualification Panel.

note: NeSI's Technical Qualification Panel doesn't duplicate scientific research peer review processes. Further details of the review process, including evaluation criteria for the technical review, may be found in NeSI's Access Policy at www.nesi.org.nz/access-policy

If you are a researcher new to HPC then you may apply for a Proposal Development allocation – to test the HPC component of your research proposal without a usage charge. The information gained through that allocation enables you to complete a technical proposal for access to NeSI facilities.

Read these full guidelines for background including details of the complete process and requirements.

Facilities available to researchers

Account Management

All aspects of your NeSI account can be managed through our website: www.nesi.org.nz

- submit proposals for allocations on HPC facilities
- manage your allocations and check your usage
- HPC facilities access, information and tools to run jobs and manage files
- contact NeSI experts to support your access to our HPC facilities

High Performance Computing

Our HPC facilities host a range of hardware architectures and operating environments.

Review www.nesi.org.nz/facilities for further information on facilities' technical specifications, access mechanisms and related services.

Scientific applications: Facilities install and maintain a wide variety of computational science applications and packages to support your research, though only for applications suitable to be run on our HPC facilities.

Support

NeSI's national team performs the following roles in support of successful proposals:

- Advisor: An Advisor is assigned to all Research and Private Industry allocations, as the primary contact point for your project within our team. Our team advises the PI on how to plan and monitor project progress, from getting started through to project completion, for the best possible use of their allocation.
- Self Service: Manage your account, including allocations, jobs, and support requests, through our self-service website. Review guidelines and frequently asked questions and other related support materials.
- Incident: Request support for a specific technical issue with using our systems, and our team will work with you to resolve it.

Additionally, project proposals can request the support of an Expert, in which case a NeSI scientific applications expert may be assigned to successful projects as an associate on your project. Such experts may, for example, contribute to optimizing your applications and software to better scale and perform on our HPC facilities. Please note – this resource is limited and will be allocated on a merit basis.

Allocations, Allocation Classes, Review Cycles

To gain access to HPC facilities, Principal Investigators (PI) can apply for an allocation of CPU core hours, specifying an appropriate Allocation Class for each proposal. The Technical Qualification Panel reviews proposals on a periodic cycle with allocations and access then granted. All projects receiving successful allocations will have a short project summary published on the NeSI public website: www.nesi.org.nz/projects

Allocations are expired, on timeframes that vary by allocation class, so that resource demand doesn't build up and can be managed over time. This approach to capacity management on HPC facilities ensures the most demanding requirements can be met by having the ability to schedule resources allocated.

Requests may be submitted in any period before the one in which they are needed, during designated submission windows in advance of proposal deadlines.

Full information is available at www.nesi.org.nz/allocations

Costs

HPC cost is based on the total CPU core hours required by the research. PIs should estimate the amount of computational time (in CPU core hours) necessary to accomplish project goals. As system performance constraints require allocating CPU cores by the node or partition, each costing must recognize the appropriate Allocation Unit, being the number of CPU cores allocated on a specific platform to meet your needs. The specific costs by HPC platform are:

Platform	Cost per CPU core hour (GST excl.)	CPU cores per Allocation Unit
P575/POWER6 P755/POWER7	\$0.20	32 CPU cores / node
Intel Cluster	\$0.20	12 CPU cores / node
BlueGene/P	\$0.05	256 CPU cores / partition

There is an exception, addressed below, wherein it is viable to share a node between multiple jobs without interference. Both cases take into consideration the Job Size, being the number of CPU cores used for a typical job. For example, if your code typically runs on up to 4 CPU cores but consumes as much memory as is available on the node, and the Allocation Unit has 12 CPU cores, you need to cost the full 12 CPU cores, even though you're not able to use 8 of them, as your job will consume all the available memory and hence make the remaining 8 CPU cores unusable by another job.

Projects are billed for 20% of the total cost of HPC resources while NeSI contributes the remaining 80% by way of a subsidy to the research project.

Cost calculation

To calculate the cost, follow the steps below:

1. There is a general formula to calculate the number of CPU core hours required to be scheduled, taking into consideration your Job Size (being the number of CPU cores your code typically runs on). Calculate the Job Size multiplied by the upper limit to the time the job needs to run for (commonly referred to as Wall Clock time) to arrive at CPU core hours. The general formula is:

$$\text{Job Size} \times \text{Wall Clock hours} = \text{CPU core hours}$$

Another formula takes into consideration how your job fits within and across the Allocation Unit – see the Scaled calculation below.

2. Multiply the result by the Cost per CPU core hour.

Two versions of the calculation required to complete this step are below, with notes on when to use each one.

Full or Shared:

There are two cases for using the below calculation:

Full: If Job Size is a multiple of the CPU cores per Allocation Unit

Shared: If Job Size is not a multiple of the CPU cores per Allocation Unit, yet can be configured to run in a way that it will not interfere with other jobs running on the same Allocation Unit

$$\text{CPU core hours} \times \text{Cost per CPU core hour} = \text{HPC cost}$$

In the Shared case, researchers must provide evidence to show clearly that the job will not reduce the performance of the shared resource, for example by oversubscribing memory or other system resources. Where a job impacts another job or overall system performance, an increased charge will apply.

Scaled: Jobs that are not a multiple of the CPU cores per Allocation Unit yet require exclusive access will be charged by the full CPU cores of the Allocation Unit. An example of this is when a job requires exclusive access to the full memory or other system resources accessible on the Allocation Unit. Such jobs will require exclusive access to a larger allocation of resources than stated by their Job Size. Use the formula below to calculate the correct HPC cost when jobs do not precisely fill an integer number of Allocation Units.

$$(\text{CPU core hours} \times \text{Cost per CPU core hour}) \div (\text{CPU cores consumed per Allocation Unit} \div \text{CPU cores per Allocation Unit}) = \text{HPC cost}$$

Note #1: In the case of those wishing to use the BlueGene/P the Scaled cost formula is used throughout since Allocation Units can only be utilised exclusively for any single user.

To provide a budget of billable charges for the project multiply the total HPC cost by the 20% cost to the researcher.

HPC cost x 0.20 = project cost

To note the contribution by NeSI, multiply the total cost by 80%.

HPC cost x 0.80 = contribution by NeSI

Cost calculation examples

Example 1: Using the Full or Shared calculation:

A researcher has a Job Size of 32 CPU cores and needs 100,000 CPU core hours on the NIWA P575/POWER6 system, meaning the 32 CPU cores will concurrently execute their job for 3125 hours. They use the Full or Shared calculation and enter the following operating expense into their budget:

100,000 CPU core hours X \$0.20 Cost per CPU core hour = **\$20,000** HPC cost

\$20,000 HPC cost X **0.20** = **\$4,000** project cost

They note the following contribution from NeSI to the project:

\$20,000 HPC cost X **0.80** = **\$16,000** contribution by NeSI

Example 2: Using the Scaled calculation:

A researcher has a Job Size of 128 CPU cores and needs 86,016 CPU core hours on the Intel Cluster system, meaning the 128 CPU cores will concurrently execute their job for a total Wall Clock time of 672 hours (28 days). Their job requires more memory per CPU core than is available on the Allocation Unit, so the job is set to use only 10 of the 12 available CPU cores per Allocation Unit.

They use the Scaled calculation and enter the following project cost into their budget:

(86,016 CPU core hours X \$0.20 Cost per CPU core hour) ÷ **(10** CPU cores consumed per Allocation Unit ÷ **12** CPU cores per Allocation Unit) = **\$20,644** HPC cost

\$20,644 HPC cost X **0.20** = **\$4,129** project cost

They note the following contribution from NeSI to the project:

\$20,644 HPC cost X **0.80** = **\$16,515** contribution by NeSI

Note 1: If the researcher wanted to **run the above job 10 times**, they would multiply the HPC cost by 10.

Note 2: As the above CPU cores consumed per Allocation Unit is 10 while the total CPU cores per Allocation Unit is 12, the researcher covers the cost for their exclusive use of all 12 CPU cores per Allocation Unit.

Note 3: If the researcher can clearly show (through previous experience or a Proposal Development allocation or credibly defining system resources required such as memory) that their job can share an Allocation Unit without affecting the performance to another job running on the shared system, they can use the Full or Shared calculation. In the above example, this would involve the researcher being able to provide evidence that their use of 10 CPU cores of 12 on an Allocation Unit still leaves 2 CPU cores available for another researcher to use, without any performance degradation.

note: The Cost Calculator on the NeSI website assists in calculating your required HPC CPU core hours and costs.

Acknowledgement, Publication

Our suggested approach to acknowledgement is to include the following statement:

The author(s) wish to acknowledge the contribution of NeSI high-performance computing facilities to the results of this research. New Zealand's national facilities are provided by the New Zealand eScience Infrastructure (NeSI) and funded jointly by NeSI's collaborator institutions and through the Ministry of Science & Innovation's Research Infrastructure programme. URL <http://www.nesi.org.nz>

Notify NeSI (pubs@nesi.org.nz) of any publications in which an acknowledgment to these facilities is made.

Applying for Access

Allocation requests

Proposals for all allocation classes are supported through www.nesi.org.nz.

Research allocations require a related peer reviewed research grant – guidelines specific to the Research allocation class are below.

NeSI's Access Policy defines review criteria and guidelines, including the principles guiding the Technical Qualification Panel in reviewing any proposal. More information is available in our Access Policy: www.nesi.org.nz/access-policy

Research Project Proposal Process

The research proposal submitted to the grant provider should include a budget for access to HPC facilities as a project cost, as defined in the Costs section above.

If you are just getting started and not certain of your computational needs then you should apply for a Proposal Development allocation, which you then use to understand the magnitude of your computational needs and how your application code scales.

During the grant provider's proposal review process the following steps occur for each proposal once the provider accepts a research proposal is inside their current call:

1. The grant provider provides advice to the PI that a NeSI technical project proposal is required through the NeSI technical proposal submission process.
2. Researcher completes and submits the related NeSI technical proposal.
3. NeSI's Technical Qualification Panel (Panel) receives proposals.
4. Panel conducts a review to assess technical HPC requirements and feasibility, seeking additional information from the proposer as necessary. Panel discusses and agrees response for each proposal.
5. Response from Panel sent to associated grant provider.
6. Grant provider completes review of proposal and contacts applicant on overall outcome.

On notice from the grant provider reviewing your research proposal, or if you have an existing successful peer

reviewed research grant that requires such resources and can cover the costs, complete a proposal for a Research allocation.

note: The PI is responsible for meeting the requirements of both the grant provider and NeSI as defined in the NeSI Access Policy. Proposal assessment will be based solely on information contained in the original research proposal and the technical proposal made to NeSI and any subsequently requested additional information.

On notice of success: Successful applicants will be advised of their approval to access HPC facilities at the same time they are advised of their award.

On notice of failure: Researchers may apply for access to other allocation classes at lower priority. Further information on allocation classes is available at www.nesi.org.nz/allocations

On notice of further information required: Researchers may be contacted to provide further information if requested by the Technical Qualification Panel. Any such requests will be at the discretion of the Panel.

On notice of appeal: The appeals process is confidential, and information may be provided to the researcher is at NeSI's discretion.

Deadlines for Proposal Submissions

For deadlines for proposal rounds, see NeSI's website: <http://www.nesi.org.nz>

Access for already peer reviewed projects

Researchers already awarded a peer reviewed research grant can apply for access via submission of a proposal directly to the Technical Qualification Panel for review. In this case, the above proposal process applies, excepting that all coordination and communication is directly with NeSI, rather than through the grant provider. All other criteria apply, including the requirement that a Proposal Development allocation is successfully completed first or similar experience demonstrated.

Eligibility

Refer to our website for conditions of Eligibility: <http://www.nesi.org.nz/eligibility>

Policies

Access Policy: www.nesi.org.nz/access-policy

Acceptable Use Policy: www.nesi.org.nz/acceptable-use

Payment

Advance payment is required for resource allocation and in most cases will be made directly by the grant provider upon approval of a proposal. Payment information will be provided in cases where direct payment is necessary.

Security and Data Ownership

NeSI has stringent privacy and storage policies that conform to international standards and best practices.

Security requirements specific to a proposal should be listed during the application process. This includes instructions for how to deal with processing proprietary or sensitive data.

Note that NeSI does not provide a long-term archive service.

Terms of use

NeSI may publicly publish project summaries. Maximum duration of 12 months from resource allocation, except in exceptional circumstances as agreed during the proposal review. The project is allocated storage resources for the duration, for both high performance data access and for archiving of output data during computing. Computational resources are allocated according to the NeSI Access Policy, including requirements for ageing and expiration of allocations over time. Agreement may be granted for special execution requirements to allow for the full exploitation of the awarded resources.