

Guidelines for Resource Allocation on the National E-infrastructure

(Allocation period 2023.2)

Abbreviations:

Administration: Sigma2

Committee: Resource Allocation Committee (Ressursfordelingskomitéen, RFK)

Partners: Sigma2, NTNU, UiB, UiO, UiT

NRIS: Norwegian Research Infrastructure Services

Definitions:

Resources:

The national e-infrastructure includes a variety of computing and storage facilities. Advanced user support is also considered a resource.

Allocations:

- *Computing facility allocations* are in the form of computing hours (processor core hours).
- *Storage facility allocations* are in the form of Tebibytes (TiB).
- *Advanced user support* takes the form of person-hours.

Allocation Period:

The computing and storage resources are allocated per period. Unless stated otherwise, an allocation period for project data storage and HPC resources is six months, starting 1 April or 1 October.

NRIS:

A collaboration of five organizations to pool competencies, resources and service (UiB, UiT, UiO, NTNU and Sigma2)

Background

The term e-infrastructure in this document denotes the infrastructure provided by the resources financed in whole or in part by Sigma2, in particular:

- HPC:
 - Norwegian infrastructure for high-performance computing
 - HPC resources available through international agreements, such as
 - EuroHPC (LUMI)
- Data Storage: National Infrastructure for Research Data (NIRD)
- HPC and Storage for sensitive data: Services for Sensitive Data (TSD)
- Platform for Data Analytics:
 - NIRD Service Platform, a Kubernetes-based cloud infrastructure providing platform for data analytics, pre/post-processing, visualization, and AI/ML workflows.

The infrastructure is available for research and education, including individuals and groups in all scientific disciplines at Norwegian universities, university colleges, Norwegian hospitals and research organizations. The services are provided for education and research projects financed by the Resource Council of Norway or ministries. In addition, the infrastructure may be available for commercial research, contributing to the funding of the infrastructure.

Sigma2 AS has the strategic and operational responsibility of the e-infrastructure.

The Board of Sigma2 appoints the Resource Allocation Committee¹. The Committee evaluates proposals based on scientific merit and the evaluation of the science of the research activity and documented needs. The resources are allocated by the Committee to individuals, groups and projects.

The Committee reports each year to the board of Sigma2. The report shall include statistics on allocations, usage and approval rates for proposals, and an assessment of resource access. The report shall also include a high-level evaluation of the quality of the proposals, project results and future needs.

Access is granted based on submitted applications. Large proposals might be subjected to application performance evaluation.

Organization

The Committee meets before the start of each allocation period to decide on allocating available resources to projects for the upcoming period.

A minimum of five Committee Representatives are required to make the meeting quorate.

The Administration takes care of the secretarial tasks for the Committee:

- Maintains contact with projects that have (or would like to have) access to the resources in the e-infrastructure.
- Provide the Committee with required material and information to evaluate the proposals and allocate the resources.

Working Group

Before the Committee-meeting, a Resource Allocation technical working group (hereafter referred to as RFKwg) whose members have good knowledge of the characteristics and status of the resources and the (existing) projects that use the e-infrastructure meets. The working group's purpose is to support the Committee in the technical evaluation of the proposals. For each proposal, the working group gathers information about the usage of the facilities by the projects, particularly concerning efficiency, scalability and relevance of the applications.

The group comments on technical aspects that could be relevant information for the evaluation:

- Requested resources
- Software needs
- The feasibility of using the resources
- The efficiency of the software

¹ Mandate for Research Allocation Committee for e-infrastructure:
<https://www.sigma2.no/sites/default/files/media/Dokumenter/824083-eng-mal-mandat-korr-scan.pdf>

- Patterns, structure and usage of the storage systems

For consumable resources, such as Tebibyte of storage, the RFKwg suggests the amount of TiB available for allocation in the given period and might provide recommendations to the administration regarding the need for future expansion.

Principles for proposal-based resource allocation

A primary task of the Committee is to evaluate the proposals and decide on allocating the available resources. This includes the distribution of research projects over the available facilities to ensure optimal utilization of the resources.

The following principles form the basis for the evaluations and allocations by the Committee:

- **Transparency:** The criteria for assessing proposals and details of the evaluation process are published before applicants submit proposals. The criteria define how the Committee operates and manages the evaluation process.
- **Assessment:** The Committee may use external experts to (i) assess the scientific merit of proposals against the published criteria and to (ii) assess the technical feasibility of proposals.
- **Confidentiality:** The Committee and the Administration treat proposals in confidence. Those who advise the Committee are required to do the same.
- **Prioritization:** Proposals may be prioritized for access by assessing the merit of each proposal against general or published criteria for assessment and against other proposals.
- **Managing Interests:** All participants in the evaluation and allocation process (including the Committee members) are required to declare interests so that any conflicts are identified and managed.
- **Reporting:** Applications might experience reduced allocations due to missing publication reporting. This is further described in the application form and allocation letters.
- **Fairness to Science:** The evaluations and allocations must be fair to the science proposed rather than to an individual applicant, group or institution.

Proposals

The Committee issues calls for proposals twice a year:

- A call in January/February, with allocations effective 1 April
- A call in July/August, with allocations effective 1 October

Researchers may apply for allocations on one or more facilities. Researchers working in the same unit or group (or on the same project) are encouraged to submit a single proposal and share the granted resources. Multiple research projects with overlapping activity may submit a single proposal, except in cases where projects handle sensitive data. In the latter case an application must be per authorized research project. Special needs for resources or support must be stated in the proposal. If additional applications for the same period are received, the last one will always override the previous ones.

The Administration maintains communication with the applicants about the status of their proposals and may request additional information to ensure the completeness and correctness of the submissions. The tasks of the Administration include:

- Preparation and publication of the calls for proposals
- Maintenance of web pages with information for users (and potential users) on the e-infrastructure
- Communication of the results of the evaluation and allocations to the applicants

Proposals must include sufficient information for the Committee to judge the scientific quality of the proposed activity and the resources requested fit for the purpose. Essential elements of each proposal are:

- Details regarding the applicant
- Details on how the (research) activity is funded
- The type and amount of resources requested
- Software needs
- Data Management Plan (if storage is requested)
- Backup needs (if storage is requested)
- Core count and core memory needs
- A brief project description
- Expected measurable output (PhD degrees, master's degrees, scientific publications, patents, others).

The applicant must have a permanent position, be a postdoctoral fellow, or be a temporarily employed researcher at their home institution. Research proposals are granted resources only if they are financed fully or in parts by public funding, i.e., money from governmental sources (e.g., the Research Council of Norway, ministries, the European Union) to individuals, organizations or entities (e.g., universities, university colleges, research institutes).

Applications for quotas can span multiple allocation periods. However, for multi-period applications, the applicant shall confirm the requested quotas for the next period by the deadline for the specific call. Applications must be confirmed within the deadline to be case-handled.

Proposals submitted outside the regular calls

If a project exhausts its allocations before the end of the ongoing allocation period, the project manager can submit a request for an extra allocation for the remainder of the period. There are no deadlines for such requests. Requests for extra allocations are subject to available capacity on the facilities. This also applies to proposals submitted outside regular calls. HPC applications submitted after the deadline and extra applications granted in part or full will receive a lower scheduling priority than ordinary applications. This does not apply to applications for new projects.

Proposals for small-scale exploratory work

To explore the feasibility of the e-infrastructure, it is possible to apply for a small-scale allocation through a free-text email to contact@sigma2.no in which purpose, anticipated resource needs, software needs, and duration are described. The maximum small-scale allocations for computing and storage resources are 20,000 CPU hours and 500 GiB storage on HPC or up to 5 TiB storage on NIRD. Small-scale proposals may be submitted outside regular calls.

Access to cloud resources

Researchers that have a storage project can utilise resources on the NIRD Service Platform / NIRD Toolkit. In order to use these resources a mail shall be sent to contact@sigma2.no.

Evaluation and allocation

The Committee evaluates proposals for access to the e-infrastructure. The evaluation of proposals can include:

- The scientific merit of the proposed activity
- The need for access to the requested resources
- The feasibility of using the (requested) resources for the proposed activity
- Previous usage and publication

When deemed necessary, the Committee may require and use additional information in the evaluation process, e.g., from experts in the scientific and technical area of a proposal, or experts in techniques for efficient resource utilization. Proposals for large allocations may be sent to external experts.

The Committee assigns allocations based on the evaluation of the proposals:

- For continuation of a project that spans multiple allocation periods, resource usage in preceding allocation periods is considered.
- For the continuation of a project that spans multiple allocation periods, recent research output is considered.

After each allocation meeting, the Committee produces a list of all proposals, including allocations applied for, approved allocations and a brief motivation for the decisions.

Allocations are granted per allocation period, i.e., a project may apply for, but is not granted allocations for multiple periods. This limitation is needed to be able to take into account (i) changes in the total amount of allocations requested (by all proposals) for each new allocation period and (ii) regular changes in the available resources/facilities. Limiting the allocations per period, allows that the projects are treated equally (or projects are prioritized) in the allocation decisions for each new period, independent of the time when the proposals for these projects were submitted.

The Committee has the right to modify the requested resources by:

- Modifying the total allocation for the project in each allocation period
- Granting allocations to other resources than those applied for

Considerations in modifying the requested allocations in a proposal are:

- The quality of the proposal and outcome of the evaluation
- The total amount of allocations requested (for all proposals) versus the total amount of resources available.
- The history of the applicant/project, e.g., in terms of requested allocations and actual usage in previous allocation periods, and scientific output
- How well justified is the request for access to specific platforms compared to more generic capacity resources.

Conditions on allocations

The following allocation principles are in effect for the allocation of storage and compute resources:

- The Committee may, at any time, revise the allocations for a project during an allocation period. In such case, the project responsible shall be informed, and the reason stated.
- The allocations may only be used for the purposes described in the proposal. Abuse may result in a revision of allocations.
- The Committee and the Administration cannot be held responsible for any damage that results from usage of allocations granted by the Committee.
- Projects granted access to the computing and storage facilities must provide a list of all publications resulting from using these resources. Such reporting must be performed through the Norwegian research publication registry, Cristin (www.cristin.no).

Registering publications in Cristin is compulsory for all projects granted access to the national e-infrastructure. Scientists accessing both storage and compute resources must register publications connected to both projects, separately. It is the duty of each project responsible to ensure that relevant publications are reported to the Cristin system and tagged with the correct funding source.

Projects are required to acknowledge the use of the national e-infrastructure in their publications.

Projects will be decommissioned as per the data decommissioning policies².

Specific conditions for the allocation of computing resources

- a. Projects that have not registered publications in Cristin in the last three years before the period applied for and have received and used allocations before these three years will not be granted resources.
- b. Unused compute allocations at the end of an allocation period cannot be transferred to a future allocation period.
- c. The Committee might reduce applied computing time due to resource limitations, generally or for specific project applications.
- d. Applications submitted after the deadline will be granted non-priority CPU allocation only. The Committee may deviate from this general rule in specific cases. New projects always receive priority time.
- e. The Committee tacitly assumes resources are used evenly (approximate linear usage) throughout the allocation period. For projects using significantly less than this projected usage (underusing), part of the allocation might be transferred from priority to non-priority time.
- f. The Administration might decide on allocations and extra allocations below 1% of a system's available RFK capacity. Extra allocations are cumulative during the period, i.e., the sum of extra allocations within an allocation period exceeding 1% will be handled by the Committee.
- g. Regarding access to LUMI:

² <https://www.sigma2.no/data-decommissioning-policies>

- a. A project may apply for LUMI resources through the regular resource application in the Metacenter administration system (MAS). If resources are granted, the users will be granted access through the Puhuri portal linked to LUMI.

Specific guidelines for allocation of storage

- a. Allocation of storage resources below 25 TiB is delegated to the Administration (does not require the Committee's authorisation).
- b. New projects usually are not granted more than 10 TiB in the first allocation period. However, the Committee can allocate higher quotas in extraordinary cases (for example when the data already exist in other storage media).
- c. Extra allocations between calls for applications can be processed and authorised by the Administration for requests up to 25 TiB. Extra allocations are cumulative during the period, i.e., the sum of extra allocations within an allocation period exceeding 25 TiB will be handled by the Committee.
- d. Applications for storage must be complemented with an up-to-date Data Management Plan, either created with easyDMP or with any other tool. Received DMPs are stored in Sigma2's internal database. It is the project owner's responsibility to submit updated versions of the DMP. The Committee reserves the right not to grant new resources and deny access to the old quota if a valid DMP has not been provided within the deadline for the application. Should there be a need for exceptions to the DMP requirement, these should be documented by the deadline and will be subjected to evaluation of the Committee on a case-by-case level.
- e. Projects that have not started using their allocation by the end of two consecutive periods will be discontinued.
- f. Projects discontinued for a period or more but not decommissioned, will be evaluated as the quota has been allocated without discontinuity. Invoicing will be done accordingly.
- g. Under-usage (below 75%) of the quota allocated in the previous period might result in adjustment of the allocation for the following period below the requested quota. Users are invited to make accurate evaluation of the needed resources, and report to the administration the foreseeable under-usage as soon as a deviation from the planned utilization is observed. This is to avoid reduction of the quota at the next period.
- h. Projects that have received a quota in the last three years but have not registered publications in Cristin, might undergo a reduction of the requested resources. It is the project owner's responsibility to record the publications.
- i. Provided that the usage from the previous allocation period is within the guidelines for allocation (no under-usage, <75%) and there is no shortage of resources, the resources awarded for one allocation period are guaranteed for subsequent periods. In case of shortage of resources, the resources awarded might be below the quota allocated in the previous period, but not below the current usage.

Specific guidelines for allocation of sensitive data services (TSD)

- a. TSD resources (HPC hours and/or TiB storage) are granted only to research projects that have obtained approval from the data protection authority of reference (Regional Ethical Committee, Norwegian Center for Research Data, Datatilsynet or similar). Checking permits and formal approval of projects is the responsibility of TSD operations at UiO/USIT and not the responsibility of the Committee.

- b. A TSD project (indicated by the alphanumeric code *pXXXX*) must exist in order to submit a request for Sigma2 granted resources. It is the responsibility of the applicant to get a TSD project before sending an application for resources to Sigma2.
- c. If the application is granted, the allocated resources will be assigned to the *pXXXX* project indicated in the application form.
- d. Allocated quotas apply to the actual period regardless of when the TSD project has been activated.
- e. Allocation of 10 TiB storage and/or up to 10 000 CPU-hours is delegated to the administration (do not require the authorisation of the Committee).
- f. New projects are normally not granted more than 10 TiB in the first allocation period.
- g. Under-usage (below 75%) of the quota allocated in the previous period might result in adjustment of the allocation for the following period below the requested quota. Users are invited to make accurate evaluation of the needed resources, and report to the administration the foreseeable under-usage as soon as a deviation from the planned utilization is observed. This is to avoid reduction of the quota at the next period.
- h. Extra-allocations between calls for applications can be processed and authorised by the Administration for requests up to 10 TiB and 10 000 CPU hours. Requests above this limit should be forwarded to the Committee.
- i. Provided the usage from the previous allocation period is within the guidelines for allocation (no under-usage, below 75%) and there is no shortage of resources, the awarded resources for one allocation period are guaranteed for subsequent periods. In case of shortage of resources, the resources awarded might be below the quota allocated in the previous period, but not below the current usage.