

Resource Allocation on Berzelius

1 Introduction

Resource allocation on Berzelius, that is, the allocation of compute resources and data storage for projects, is carried out by an allocation committee of four people. The committee consists of three representatives from the National Supercomputer Center (NSC) at Linköping University (LiU) and one from Wallenberg AI, Autonomous Systems and Software Program (WASP). In addition, two NSC personnel perform part of the technical review.

Berzelius is owned, operated and allocated by NSC. Application for, and review of projects using Berzelius is made in the Swedish User and Project Repository (SUPR). The process described in this document concerns academic research projects exclusively, see [Qualifications for Applicants](#) for those who are eligible to apply for projects on Berzelius. The application process itself is described in detail in the section [Project applications on Berzelius](#)

Berzelius is a research infrastructure intended for Swedish academic research and industrial doctoral students / industrial postdocs affiliated strategic initiatives by the Knut and Alice Wallenberg Foundation (KAW), but can also accommodate other projects, including commercial ones, under conditions described separately.

This document defines the Principal Investigator (PI) as the person applying for a project. A project member is a person who has been granted membership in a project by its PI via SUPR and the project group are the members of the project, the PI included. The term “resource” is used in the document as a collective term for compute time and storage space.

1.1 Overview of Resource Allocations on Berzelius

Resource allocation takes place after a review of the project application carried out by the allocation committee. The review must in principle treat three questions about each project application as a basis for allocating resources:

1. How scientifically significant is the project?
2. How likely is the project to be successful?
3. How much resources are required by the project?

Based on the answer to these questions, an allocation of resources on Berzelius is made. The answers consist of assessments in the form of grades in five categories from which a total score is formed. For details see [Grading](#). The categories are:

Possibility How feasible is the project for the project group? How do the research group’s qualifications match the problems treated and the methods used? How qualified is the research group?

Efficiency How efficient is it to carry out the project on Berzelius? How well does the described research and the tools used match against Berzelius’ features and capabilities?

Track Record How have previous allocations on Berzelius been used? Have any previous allocations on Berzelius resulted in published results? Is there a detailed activity report from any previous projects?

Data Management Plan How have any previous storage allocation been used? Is there a clear plan for how the results will be taken care of after the project is completed? Is there a clear and reasonable calculation for the storage resources requested?

Affiliation and Methods Is the project linked to a strategic initiative by KAW? Are AI/ML methods the subject or the main tool for conducting research in the project?

In the grading of the first four categories, a balance is made between the scientific goals and the estimated probability that they will be achieved for the project, such that a high grade indicates an estimated high scientific value in the project and a high probability that the goals will be achieved.

The latter category reflects the priority to be given to specific research areas and research initiatives. The background to the category lies in the purpose and conditions for the donation of Berzelius to NSC, which states that this should be prioritized when allocating resources. Together, the grades in the five categories are considered to answer the questions above.

The size of the resource allocation is based on the overall grade for the application. The resource allocations are awarded for 6 months maximum at a time for projects and there is a maximum allocation of 28800 GPU hours per running 30 day period. The projects are evaluated monthly as they are received.

2 Project Applications on Berzelius

The section deals with information that concerns PIs applying for projects on Berzelius, that is, how the application is made and what information needs to be included in the application to provide a basis for review.

2.1 Apply for Projects in SUPR

Applications for projects on Berzelius are made in [SUPR](#) by **qualified PIs**. The practical procedure for applying for a project is:

1. [Register in SUPR](#) if you have not already registered.
2. Log in to SUPR via the method you chose when you registered, [SWAMID](#) , “E-mail and Password” or with a [certificate](#). You may need to enter a Time-based One-Time Password (TOTP). Clients for this are available on most platforms (smartphones have such), for example, Google Authenticator is available for both Android and iOS.
3. Apply by selecting “Round” -> “AI/ML” -> “LiU Berzelius”
 - For a new project, select option 1 “Create Proposal from Scratch”
 - To create a continuation project, select option 2 “Create Proposal Based on Earlier Proposal”. Continuation project clones an existing project which is then edited to correspond to current conditions. Continuation projects are not discussed further in this guide.

When a project is created, it is in edit mode. Once the editing is complete, the application must be submitted. The editing of the application is described [below](#).

NSC may, in the event of low occupancy of Berzelius and in certain circumstances, grant projects that fall outside the framework described under [Qualifications for Applicants](#). Such projects can contact berzelius-support@nsc.liu.se directly in the matter.

2.2 Parts of the Application

The application has several categories with items to be filled in, which are expanded from the main page of the application, see figure 1 below. The categories are accessed by clicking in the grey boxes and mandatory points are marked with a red asterisk. The categories with entries to fill in before submission are,

- PI information
- Basic information. The project is described here.
- Continuation. Only to be filled in if the project is a continuation of a previous project on Berzelius.
- Co-investigators. Non-mandatory, only fill in if a deputy in the project is wanted.
- Resources. The resource request is made here.
- Directory name. Name the work data directory used in the project here.

2.2.1 Personal Data for the PI

Important parts to fill in in this category consist of

- Contact information
- Approval of the Berzelius User Agreement

2.2.2 Basic Information

In this category, the project is described in terms of goals and purpose, specifically how the size of the request for resources has been computed, the period for which the request applies and the funding the PI has for the project.

Abstract Describe the project briefly with regards to, for example

- Goals
- The importance of the project.
- Expected goal fulfilment at the end of the project period.
- Software and methods to be used.

Resource Usage Specify how the requested compute resources are to be used here, i.e., a justification for the size of the compute time request. For instance, specify:

- Number of active users
- Number of jobs per active user
- Amount of resources per job

Abridged Data Management Plan Specify how the requested data storage is to be used, i.e., a justification for the size of the request, but also the possibilities the project has to take care of results and transfer data to and from the system.

- Number of files
- Total size
- Plan for handling data after the end of the project
- Data transfer

Classification Scientific classification code for the PI's research in accordance with the Standard for Swedish classification of research topics ([Standard för svensk indelning av forskningsämnen](#)).

Requested Duration and Start Date Duration up to 6 months can be requested and the start date is "ASAP" unless otherwise stated. Special requests for the allocation/distribution of requested time during the application period can be stated under "Resource Usage" or during the project through email to support.

Financial Support References to all grants the applicant researcher has are listed here. Especially, list grants linked to [strategic initiatives from KAW](#), as this gives a higher grade in the review and thus priority when allocating resources.

Demonstration Project (Berzelius-2023-1)

Instructions: Get started by editing the basic information for your proposal below. Change the continuation information if needed. Add any co-investigators you have for your proposal. Finally, submit the proposal (it will not be handled unless you do that).

Information about this round is available at <https://www.nsc.liu.se/systems/berzelius/>.

Proposal

Dnr: Berzelius-2023-1
State: Editing (not yet submitted)
Round: LiU Berzelius 2023
Principal Investigator: Peter Munger

[Check Principal Investigator Information](#)

Basic Information

Project Title: Demonstration Project
Abstract:
Resource Usage:
Data Management Plan:
Affiliation: Linkoping University
Primary Classification Code:
Requested Duration: 6 months

[Edit Basic Information](#)

Continuation

This proposal is not registered as a continuation of an earlier project.

[Change](#)

Members

When the project is created:

- The PI and co-investigators of this proposal will become members of the project.

Co-Investigators

There are no co-investigators for this proposal.

If you want to designate a proxy, you need to add that person as a co-investigator first.

[Add Co-Investigator](#)

Resources

Click on the linked resource name below to show more information about the resource in this proposal.

Resource	Centre	Requested Unit	Requested Unit
Berzelius Compute	NSC	240 GPU-h/month	
Berzelius Storage	NSC	5 000 GiB	5 000 000 files

Directory Name

You have not yet selected a directory name for project storage.

[Select Directory Name](#)

Cancellation

If you do not intend to submit the proposal for consideration, you can cancel it using the button below.

[Cancel the Proposal](#)

Submission

You need to submit the proposal if it is to be taken into consideration for allocation.

The submission deadline is 2024-01-01 00:00.

[Submit the Proposal](#)

Figure 1: Start page for a new project application

2.2.3 Resources

On Berzelius, there are two types of resources to request, compute and storage. Compute resources must be requested specifically and explicitly while storage has a standard allocation and only deviations from this need to be specified.

1. **Berzelius Compute.** The format is GPU-h/month format for compute resources and has a maximum request size of 28800 GPU-h/month, which corresponds to 40 GPUs full time per running 30 day period or 5 DGX-A100.

A special category of projects are the “small” projects. “Small” projects are those that request up to 240 GPU-h/month (i.e., 1/3 GPU full time per month on average), which undergo only an initial review and a technical review, thus providing a fast track to get started on the system.

2. **Berzelius Storage.** Requested storage space in GiB, where the default is 5000 GiB. The number of files can be specified if other than the default is wanted, which is 5M.

2.3 Especially Important Information in the Project Application

To facilitate a fair review of the project application, the motivation for the requested resources is particularly important. It is also important that an activity report from any previous projects on Berzelius is submitted and that all published research resulting from use of Berzelius acknowledges this use.

Overall, the degree of scrutiny of project applications follows the amount of resources requested in the application, and consequently large requests require more in-depth motivation than small requests.

3 Project Application Reviews on Berzelius

The section describes the process for reviewing and allocating resources to project applications. It is primarily aimed for project application reviewers, but also contains information that may be of interest to applicants.

3.1 Qualifications for Applicants

The PI and the project described must meet certain criteria in order to be granted resources on Berzelius. These are:

1. Level of education for the PI from doctoral student and higher. The affiliation of the PI must sort under one of the following alternatives:
 - a. Academic researcher active in Sweden.
 - b. Industry doctoral students / industrial postdocs affiliated a strategic initiative by KAW.
 - c. Exceptions may be granted for 1.a and 1.b in exceptional cases, see item 2.
2. The project must conduct research aimed at publication unless specific reasons for exceptions can be stated. As special reasons, for example, great public benefit or societal interest can be stated or that the project pays for resources on Berzelius. Exceptions can only be granted when Berzelius is undersubscribed at the time of the award.
3. The project must not process sensitive data (e.g., identifying personal data), as Berzelius does not currently meet the legal requirements in this regard. Projects processing sensitive data is expected to be catered to at a later stage, but not initially.

In cases where the criteria for the application are not met, it is rejected, and the applicant PI is notified.

The qualifications for projects and PIs covered in this section only apply to research projects. Any other projects¹ are only affected by the requirement that sensitive data may not be processed on the resource and that the project must be mainly conducted in Sweden.

3.2 Process

The following process is applied to project applications on Berzelius.

1. The application is received via SUPR.
2. The application is assessed.
 - a. Preparatory review
 - i. Performed by the **review manager**
 - ii. Reconcile with the **checklist** and proceed to **technical review** if all points are met.
 - b. Technical review
 - i. Performed by one or more technical reviewers.
 - ii. The answer goes back to the review manager.
 - c. Grading review
 - i. Performed by a reviewer appointed by the review manager.
 - ii. The review must contain grades in the categories (found in the review form in SUPR)
 - Possibility
 - Efficiency
 - Track record
 - Data management plan
 - Affiliation and Methods
 - d. Allocation is decided.
3. The applicant is notified of the allocation decision.

3.3 Distribution of Responsibilities in the Review Process

The review of projects is carried out by the allocation committee and two technical reviewers. The technical reviewers assess whether the projects match Berzelius' capabilities, and the allocation committee then rates the proposal and decides on the award. The review process is conducted by the review manager of the allocation committee.

3.3.1 Allocation Committee

The allocation committee consists of the review manager and two other reviewers from NSC as well as a scientific expert in the AI / ML field from WASP. All members of the allocation committee also serve the role of reviewer. The Allocation Committee is jointly responsible for

1. Ensuring that applications are reviewed within a reasonable time.
2. Review applications and allocate resources depending on the quality of the application.
3. Inform the applicant PI of the decision and handle communication about this.

The review of project applications and the communication around them can be done by individual reviewers or by the committee together. However, the responsibility for implementation is always shared.

¹For example, socially beneficial projects or projects on commercial terms. For projects on commercial terms, the purchase of calculation time on Berzelius is not expected to be possible during the earliest phase of the resource but will be implemented at a later stage.

3.3.2 Review Manager

The person responsible for the review process, the review manager, is a designated member of the allocation committee. The review manager carries out an **initial review** of the application ensuring that the applicant and the application meet all the criteria for further processing of the application, and forwards the application for **technical review** if the application meets the requirements. If the application does not meet the requirements, the review manager notifies the applicant PI.

If the application passes the technical review, the review manager distributes the project application for further review, otherwise the PI is notified of the rejection.

3.3.3 Reviewer

The task of a Reviewer in the allocation committee is to rate the project application and allocate resources based on the grade.

Reviewers in the committee receives the project application from the Assessment manager for review after it has met the criteria in the **checklist** and passed the **technical review**. The reviewer then **grades the application** in the various categories and based on the average grade determines an allocation of resources.

3.3.4 Scientific Expert

The Scientific Expert is a representative from WASP with expert knowledge in AI / ML that complements the knowledge of the reviewers from NSC.

The role includes:

- Review of applications
- Consultation for other reviewers in AI / ML issues that may arise in the review of applications.

3.3.5 Technical Reviewer

A technical reviewer performs a comprehensive **technical review** at the request of the **Review manager**. The technical reviewers consist of staff from NSC.

3.4 Assessment of the Project Application

The review of each application has three elements,

1. Initial review
2. Technical review
3. Grading review

After completing these steps, resources are allocated based on the grade in step 3 weighed against what is requested in the application.

3.4.1 Initial Review

The initial review is performed by the Review manager and consists of checking the application against a checklist (see below) of criteria a proposal must meet to proceed in the process. The step aims to discover as early as possible any major errors in the application and reject it or take other measures that may be warranted (e.g., request additional information).

Checklist The checklist defines the criteria for proceeding in the review of the application.

1. Is this a scientific project, i.e., aimed at scientific publication?

- a. If yes: Assess whether the application meets the criteria in **Qualifications for Applicants**
 - b. If no and Berzelius is fully subscribed, the application will be rejected, otherwise the allocation committee will convene to discuss the proposal. An assessment must be made as to whether the application meets the conditions for non-academic projects on Berzelius.
2. Is the documentation in the application from the PI sufficiently detailed?

3.4.2 Technical Review

The technical review is performed by technical reviewers who are to assess:

1. Can the described software and workflow perform the task described?
2. Can the described software run efficiently on the system, i.e., use GPUs efficiently?
3. Is the group judged to be able to handle its software themselves, both running and installing it, i.e., is only a low amount of support expected to be needed in this regard?
4. Can the way in which the software is to be run achieve the result described in the application?

The purpose of the technical review is to allow the review process to proceed unless direct errors are discovered. The answer from the technical review is either, “yes, the process can continue” if the answer to all the questions is yes, otherwise the answer is “no”. The answers should be judged to be “yes” to the questions above unless there are strong indications to the contrary.

If “no,” the application is rejected, and the PI is notified.

3.4.3 Grading

Grading of the project application is performed by a reviewer and comprises assigning scores between 1 and 5 in five categories in which rating 5 is the highest. The five categories are; “Possibility”, “Efficiency”, “Track Record”, “Data Management” and “Affiliation and Methods”, which are described below.

The arithmetic mean of the grades in the five categories forms the summary grade (or, total score) for the application and constitutes the main basis for allocating resources.

The summary rating can be considered an estimate of the ratio “reward” over “risk” for the project where a high value means high reward and low risk and conversely for low values. “Reward” consists of the assessed value of the expected scientific results and “risk” the estimated probability the project has of not achieving results with its methodology and available knowledge.

3.4.4 Grading Categories

In general, the reviews in the various categories are about the relationship between the strengths that exist in the project weighed against its weaknesses. The scale should be as far as possible linear, i.e., if the total strengths are judged to be equal to the total weaknesses, the grade 3 is awarded, if the strengths outweigh the weaknesses, the grade 4 or 5 is awarded depending on how strong the overweight for the strengths is.

The reviews should strive to be (in some sense) absolute and not relative other projects on Berzelius. Other projects on Berzelius are considered however when allocating resources to the project, where projects seeking similar amount of resources are compared with each other, among other things regarding grades.

An exception applies to the category “Affiliation and Methods” - where three grades can be awarded; 1, 3 and 5 - for which the grade levels have clear criteria. For the grade “5”, the PI is connected to a strategic initiative by KAW, there are no requirements for method or research field. For the grade “3”, the PI is not connected to a strategic KAW initiative, but

uses or develops methods for Artificial Intelligence and Machine Learning (AI/ML) in the project. For the grade “1”, the PI is neither affiliated with a KAW strategic initiative nor uses or develops methods for AI/ML in the project.

The reason for the special treatment of projects based on the PI's affiliation and the project's purpose / methods lies in the background to KAW's donation of Berzelius to NSC, where NSC undertakes to make priorities of this kind in the allocation of resources. The method described above makes this prioritization transparent and predictable for everyone involved.

Possibility The review in the category only refers to the requested computing time (i.e. not storage), and shall assess the project's scientific merits weighed against it's possibility to carry it out. Guiding questions for the reviewer

- How scientifically advanced is the project's goal?
- How qualified is the applicant/research group overall?
- Can the applicant carry out the computations described in the application?
- Has the applicant previously used the resources in the proposal?
- Has the applicant previously used the resources in the proposal in the proposed way?
- How viable is it to use the described software on the resource?
- Do they have required licenses for the software described?

Efficiency The review in the category only refers to the requested computing time (i.e. not storage), and shall assess how well the project utilizes different aspects of Berzelius' performance. Guiding questions for the reviewer

- Is this an innovative use of methods considering the problem studied?
- Is the proposed way to use Berzelius efficient, considering Berzelius' technical profile?
- How well is the hardware performance utilised compared to the most efficient way to use it?
- Is the planned use of the resource in line with “best practice” for the scientific field?

Track Record The review in the category only refers to the requested computing time (i.e. not storage), and assesses the scientific outcome of previous allocations on the resource and to what degree they have been used. Guiding questions for the reviewer

- Which scientific results have previous allocations produced? Examples
 - Publications
 - Conference submissions
 - Dissertations/Theses
- Has the applicant used any previous allocations on the resource as they have described?

The evaluation shall weigh in explanations from the applicant on any discrepancies between the planned and actual use of the resource in previous projects.

Data Management This is a compound category for data storage composed of “Possibility”, “Efficiency” and “Track Record” corresponding to the categories for calculation above, and the grade given should be seen as an overall grade regarding these aspects. The guiding questions in the compute categories can to a some extent also apply here, but regarding data management.

Affiliation and Methods A five-point scale with three levels is applied here: 1, 3 and 5. The grade is not so much a review as it is a check of the degree of fulfilment for two criteria. Ratings and criteria are:

- Grade 5: KAW affiliated.
- Grade 3: Non-KAW affiliated. AI/ML used as the main tool or research subject.

- Grade 1: Non-KAW affiliated. Research unrelated AI/ML.

3.5 Resource Allocation

Resource allocation for received applications takes place monthly and can be done by an individual reviewer or at a meeting of the allocation committee, depending on the circumstances. The general procedure for resource allocation is:

- The resource allocation to the project applications received is made in order of priority, from applications with high to low average grades.
- A calculation is made of what resources are available, i.e., not already allocated.
- An assessment is made for each application if the requested resources are well justified. If not, the request is scaled down to what is deemed justified.
- A calculation of what the total requested resources amount to in the applications is compiled.
- The allocations are then made in order of priority until the available applications or resources are exhausted.

4 Other

4.1 Conflicts of Interest

All reviewers in the process, committee members and technical reviewers, must follow the policy regarding conflicts of interest. The policy is based on *Förvaltningslagen* (2017:900), which states that:

- Den som känner till en omständighet som kan antas göra honom eller henne jävlig omedelbart ska anmäla detta till myndigheten.
- Den som är jävig inte får ta del i handläggningen av ärendet och inte heller närvara när ärendet avgörs.

Translated to the current context, the policy regarding conflicts of interest for reviewers of project applications on Berzelius becomes:

- Any circumstance making a reviewer having a conflict of interest in the processing of a project application shall immediately be reported to the other reviewers.
- Reviewers having a conflict of interest may not perform any part of the review of the project application or have any other influence on the allocation of resources to the project.

In the event of a conflict of interest, all measures taken to comply with the policy must be documented. The reason for the conflict of interest does not have to be stated by the person reporting it.

4.2 Confidentiality

Project applications and the review process are covered in all parts by confidentiality and professional secrecy for the reviewers vis-à-vis outsiders. The reviewers consist of the allocation committee and the technical reviewers. Documents provided during the case may not be distributed and must be removed after processing by the reviewer. No outside party may take any part of the case. All communication in the case between the reviewers and applicants is done via SUPR.

4.3 Document Information

Revision 1.2

Diary number LiU-2023-01789

Date 2023-04-12

4.3.1 Previous versions

Revision number	Date
1.1	2021-06-17
1.0	2021-05-03