

**REGULATIONS GOVERNING
APPLICATIONS
FOR USE OF THE FLEMISH TIER-1
SUPERCOMPUTING PLATFORM**

Effective as of 01 January 2021

PART I. PRELIMINARY PROVISIONS

§ 1 – Definitions

For the purposes of these regulations, the following terms shall have the following meanings:

1. **Application:** an application for computing time and disk storage capacity to perform computing tasks on the Tier-1, other than a Starting Grant application or a Collaborative Grant application;
2. **Applicant:** a Researcher who submits an Application;
3. **Core-hour:** unit of computing time, corresponding to one hour on one CPU core; kcore-hour and Mcore-hour represent a thousandfold and a millionfold respectively;
4. **GPU-hour:** unit of computing time, corresponding to one hour on one graphics processing unit; kGPU-hour and MGPU-hour represent a thousandfold and a millionfold respectively; with every GPU-hour allocated, the applicant automatically receives an amount of core-hours that can be used exclusively on the CPU cores of the node(s) on which the graphics processing units are installed;
5. **Researcher:** a natural person who, within the framework of a statutory appointment or assignment, an employment contract, a contracting contract, a research grant, a voluntary collaboration or a collaboration as a student, performs research at a Public Research Institution;
6. **Public Research Institution:**
 - a. a university within the Flemish Community;
 - b. a college of higher education within the Flemish Community;
 - c. a research institution under the authority or supervision of a university or a college of higher education, such as a special university institution, as referred to in article III.73 et seq. of the Higher Education Code of 11 October 2013;
 - d. a university hospital, as referred to in article 4 of the law on hospitals and other care facilities, coordinated on 10 July 2008;
 - e. the Flemish strategic research institutions;
 1. Flemish Institute for Biotechnology (VIB);
 2. Imec;
 3. Flemish Institute for Technological Research (VITO);
 4. Flanders Make;
 - f. an institution for post-initial education;
 - g. a centre for policy-relevant research as referred to in the Decision of the Flemish Government of 15 September 2006 concerning centres for policy-relevant research;
 - h. The Expertise Centre for Research and Development Monitoring as referred to in Title V, Chapter VI, of the Decree of 30 April 2009 on the organisation of science and innovation policy;
 - i. research institutions whose activities are fully or partially financed or subsidised by the Flemish Community or the Flemish Region;
 - j. a research institution under the authority or supervision of the Flemish government, as referred to in article III.73 et seq. of the Higher Education Code of 11 October 2013;

- k. research institutions whose activities are fully or partially financed or subsidised by the Federal Science Policy Office;
 - l. a research institution or an institution under the authority or supervision of the Federal government
7. **Starting Grant:** allows the holder to try out Tier-1 compute to perform a benchmark or software tests, by way of preparation for a future Application;
 8. **Tier-1 compute:** all of the partitions of the compute component of the Flemish Tier-1 supercomputing platform;
 9. **Tier-1 compute partition:** a portion of the Tier-1 compute infrastructure, logically divided by installation location and hardware specifications; at present, this involves two logical BrENIAC partitions (Broadwell and Skylake) at KU Leuven and one logical Hortense partition at UGent (doduo);
 10. **Tier-1 data:** the data component of the Flemish Tier-1 supercomputing platform;
 11. **Tier-1 cloud:** the cloud component of the Flemish Tier-1 supercomputing platform;
 12. **Queue:** the list of tasks that are waiting to be carried out;
 13. **Wall clock time:** the time that elapses according to the regular clock between the start and the end of a (sub-)task.
 14. **SCRATCH:** temporary, fast storage used to store input data for calculations or output data from calculations. This storage is not backed up Data that has been deleted cannot be recovered.
 15. **Standard quota:** the maximum amount of "SCRATCH" disk storage available to one application on the BrENIAC or Hortense partitions of Tier-1 compute, and the maximum total number of files that one application can create.

PART II. REGULAR APPLICATIONS

Chapter 1. Scope and Institutional Provision

§ 2 – Scope

(1) As of 01 January 2021, this part governs the submission and processing of Applications for computing time and disk storage on Tier-1 compute, as well as the implementation of granted Applications.

(2) Applications submitted and approved before 01 January 2021 will be processed and implemented in accordance with the regulations effective at the time of submission of the Application.

§ 3 – Evaluation Committee

(1) The Board of Trustees of the FWO delegates the decision-making competence regarding Applications to an Evaluation Committee, which operates under the authority of the Board of Trustees.

(2) The Evaluation Committee evaluates Applications against the admissibility criteria referred to in § 4, and, insofar as they are admissible, against the evaluation criteria referred to in § 5.

(3) The Evaluation Committee consists of:

1. at least three experts appointed by the Board of Trustees who are not active within the Flemish Community and have wide experience in the field of using large computing capacity;
2. a representative of the FWO.

(4) The senior official of the FWO and representatives or delegates of the VSC operational team may be invited to take part in the meetings of the Evaluation Committee as observers.

(5) The secretariat of the Evaluation Committee is held by the FWO, referred to in (3), 2.

Chapter 2. Grounds for evaluation

§ 4 – Admissibility Criteria

(1) An Applicant electronically submits an Application to the FWO using a standard application form, attached as enclosure 1 and also available at <https://www.vscentrum.be/tier-1>. The standard form must be completed and submitted through the online EasyChair platform <https://www.easychair.org/conferences/?conf=tier1>, in accordance with the instructions given at <https://www.vscentrum.be/tier-1>.

(2) An Application shall cover at least the equivalent of 500 kcore-hours or 1 kGPU-hour and no more than the equivalent of 5 Mcore-hours and/or 25 kGPU-hours. These upper limits may be exceeded if justified.

(3) The scientific quality of the relevant research project which is assessed:

- a. if the research project is part of an ongoing research project, by the Evaluation Committee established by the university, the research institution or the strategic research centre, or by an external funding agency;
- b. if the research project is not part of an ongoing research project, by a declaration from the relevant institution in which the proposed research is approved.

(4) The Wall clock time of the different sub-tasks described in an Application, shall amount to maximum 72 hours; where appropriate, the Applicant shall specify the available restart options for sub-tasks that require more than 72 hours of Wall clock time.

(5) The Standard Quota for an Application shall be 2 TiB for "SCRATCH" disk storage and 200,000 for the total number of files, unless justification for exceeding these limits is attached to the Application. "SCRATCH" disk storage must always be specified:

1. in multiples of one TiB for storage,
2. in multiples of 100,000 for the number of files, and
3. for the total duration of the Application.

(6) A standard Linux HPC software stack is available on Tier-1 compute. Where necessary, additional (Linux) software can be installed in accordance with the provisions in § 16.

(7) The Application clarifies:

1. which software will be used;
2. that the required software, which is at the Applicant's expense:
 - a. may be validly installed on the Tier-1 compute partition(s) involved, and
 - b. may be validly used by the Applicant and by the persons mandated by the Applicant;
3. that any licences that may be required are valid for use on the Tier-1 compute partition involved and are sufficient to cover the computing capacity applied for.

(8) Applications shall be drawn up in English.

(9) Only Applications that comply with the formalities and requirements set out in (1) through (8) shall be admissible.

§ 5 – Evaluation criteria

Admissible Applications are evaluated on the basis of the criteria referred to in the application form, such as:

1. the Applicant's substantiated experience regarding the use of Tier-0, Tier-1 and/or Tier-2 computing infrastructure;
2. the technical feasibility of the computing task, which is assessed on the basis of:
 - a. the justification of the reasons why the Tier-1 compute or a specified partition is appropriate to perform the proposed computing task;
 - b. the number of nodes/cores/GPUs applied for per computing task, with division of the computing time into sub-tasks, indicating the sequence of the sub-tasks and the estimated wall clock time;
 - c. where appropriate, the use of parallel processing (Open MP, MPI, OpenMP + MPI (hybrid), worker framework, atools, etc.) and the parallel efficiency of the executed tasks;
 - d. the estimated memory use of a computing task;
 - e. the requirements for disk storage on the Tier-2 "HOME" and "DATA" volumes and the Tier-1 compute "SCRATCH" volume (estimated volume in TiB and total number of files) during execution of the computing tasks.

Chapter 3. Assessment Procedure

§ 6 – Possibility of Remediation

The Evaluation Committee may request an Applicant to clarify or amend the Application. This possibility of remediation shall not serve to fill in the gaps in an incomplete or vague Application.

§ 7 – Cut-Off Dates

(1) Before the start of each calendar year the Evaluation Committee shall draw up an evaluation calendar with a number of cut-off dates. The cut-off dates for 2021 are: February 1, June 7 and October 4.

(2) In exceptional cases, the Applicant may include in the Application a reasoned request clarifying why the Application cannot wait until the next cut-off date. The chairman or deputy chairman of the Evaluation Committee shall within five working days decide on the urgent character of the Application and notify the Applicant of their decision through the most appropriate communication channel. If the urgent character of the Application is recognised, the Application shall be further evaluated in accordance with this chapter, it being understood that the decision periods for application of § 8 shall commence from the date of recognition of the urgent character.

§ 8 – Decision Periods

The Evaluation Committee shall decide on Applications within an indicative period of 30 days from the cut-off date following the date of submission of the Application.

§ 9 – Decision modalities

(1) Within the framework of the overall management of the Tier-1 compute, the Evaluation Committee can take a reasoned decision to grant only a part of the requested computing time and disk storage capacity or to limit, per cut-off date, the number of approved Applications submitted by Applicants active in the same scientific sub-disciplines.

(2) The Evaluation Committee shall define the period within which the allocated core-hours and GPU-hours have to be used for each granted Application.

(3) The Evaluation Committee shall grant each Applicant storage capacity in accordance with the Standard Quota. The Standard Quota can only be deviated from on the basis of a special justification in accordance with § 4, (5).

(4) After three months of the period (allocated in accordance with §9 (2)) have passed, the Applicant loses 20% of the initially granted number of core-hours and GPU-hours, if that 20% has not yet been used.

(5) Combined applications from institutions belonging to categories k and l (in accordance with §1 (6)), can be granted a maximum of 10% of the total available computing time per call round.

§ 10 – Notification of the decision

- (1) The FWO shall send an e-mail message regarding the decision of the Evaluation Committee on the Application within a period of 10 days from the decision of the Evaluation Committee.
- (2) The granting of an Application shall be published in the annual report of the Flemish Supercomputer Centre (VSC).

Chapter 4. Appeal procedure

§ 11 – Appeal Modalities

(1) The Applicant who is of the opinion that the decision of the Evaluation Committee is contrary to a written or unwritten legal rule or general legal principle, can lodge an appeal against the decision with the Board of Trustees of the FWO.

(2) The appeal shall be lodged within an expiry term of five days, which commences on the day following the day of sending, stated in § 10 (1).

(3) The appeal shall be lodged by means of a signed and dated notice of appeal, to be submitted to the Board of Trustees by registered mail, failing which it shall be inadmissible. It shall contain at least the identity of the Applicant concerned, the challenged decision(s) and a factual description of the objections invoked. At the same time, the Applicant shall, by way of information, send an electronic version of the notice of appeal by e-mail to the FWO (fwo@fwo.be). The date of the postmark of the registered mail shall count as the date of the appeal.

§ 12 – Hearing of the appeal

(1) The Board of Trustees shall process the appeal on the exhibits. However, it can invite anyone whose presence it considers helpful to facilitate the hearing of the appeal.

(2) The Board of Trustees has the right to demand or gather any exhibits and information it considers useful to facilitate the hearing of the appeal.

(3) The Board of Trustees can obtain the advice it considers useful for the hearing of the appeal.

§ 13 – Relativity rule

The Board of Trustees can uphold a decision against which an appeal has been lodged, despite infringement of a written or unwritten legal rule or general legal principle, if it is plausible that the interested parties have not been disadvantaged by the decision.

§ 14 – Decision

If the appeal is admissible, the challenged decision shall be reconsidered on the grounds presented in the appeal. Insofar as the reconsideration shows that there is cause to do so, the Board of Trustees shall revoke the challenged decision and, if need be, make a new decision to replace the original decision. The decision shall be announced by sending or delivering it to those to whom the decision is addressed.

Chapter 5. Provisions regarding the implementation of granted applications

§ 15 – Technical problems

The KU Leuven or UGent staff members entrusted with the Technical Operation of the respective Tier-1 compute partition(s) may at any time abort the implementation of a granted Application if it gives rise to technical problems. The Applicant shall be informed of this decision by e-mail, including a short accompanying explanation.

§ 16 – Assistance

If a granted Application requires the use of software that is not yet available on a Tier-1 compute partition, the Applicant may call on the assistance of KU Leuven or UGent staff members entrusted with the technical operation of the respective Tier-1 compute partition(s), without prejudice to the application of § 4, (5 and 6). The Applicant shall send this request via e-mail to compute@vscentrum.be, which acts as the sole contact for both parties. KU Leuven or UGent shall make reasonable efforts to respond to the request for assistance, which can by no means bring about any obligation to produce a certain result. The KU Leuven or UGent staff members entrusted with the technical operation of the respective Tier-1 compute partition(s) shall never intervene on the code of the actual software.

PART III. STARTING GRANT APPLICATIONS

§ 17 – Scope

(1) This part governs the submission, processing and follow-up of Starting Grant applications as of 01 January 2021.

(2) Starting Grant applications submitted and approved before 01 January 2021 shall be processed and implemented in accordance with the procedure effective at the time of submission of the Application.

§ 18 – General principles

(1) A Starting Grant is awarded for a maximum period of 500 kcore-hours and 1 kGPU-hour. The maximum period for use of the allocated computing time is four months.

(2) A Starting Grant is personal and cannot be shared with researchers other than those to whom the Starting Grant is awarded.

(3) Starting Grant applications can be submitted on a rolling basis.

§ 19 – Procedure and follow-up

(1) An Applicant electronically submits a Starting Grant application via e-mail to compute@vscentrum.be, using a standard application form, attached as enclosure 2 and also available at <https://www.vscentrum.be/tier-1>.

(2) The application is assessed for validity by the operational managers of the Tier-1 compute partitions, who report their findings to the VSC operational team.

(3) After approval, the operational manager of the respective Tier-1 compute partition grants access and computing time to the Researcher. If the application is not approved, an answer with justification of the decision is sent to the Applicant.

(4) The award of a Starting Grant is published in the annual report of the Flemish Supercomputer Centre (VSC).

PART IV. COLLABORATIVE GRANT APPLICATIONS

§ 20 – Scope

(1) This part governs the submission, processing and follow-up of Collaborative Grant applications as of 01 January 2021.

(2) Collaborative Grant applications submitted and approved before 01 January 2021 are processed and implemented in accordance with the procedure effective at the time of submission of the Application.

§ 21 – General principles

(1) A Collaborative Grant is awarded for a maximum period of 10 Mcore-hours and/or 75 kGPU-hours. The maximum period for use of the allocated computing time is twelve months.

(2) A Collaborative Grant can only be applied for by a consortium of at least three research groups from different Public Research Institutions (cf. §1 (6)) with a clearly defined common research topic.

(3) Collaborative Grant applications can be submitted on a rolling basis.

§ 22 – Procedure and follow-up

(1) An Applicant electronically submits a Collaborative Grant application via e-mail to compute@vscentrum.be, using a standard application form, attached as enclosure 3, and also available at <https://www.vscentrum.be/tier-1>.

(2) This Application shall in particular clarify why it is crucial for the consortium to apply for a Collaborative Grant and the proposed common research work cannot be carried out via one or more regular Applications.

(3) The application is assessed for validity by the VSC operational team.

(4) After approval, the operational manager of the respective Tier-1 compute partition grants access and computing time to the Researcher. If the application is not approved, an answer with justification of the decision is sent to the Applicant.

(5) The award of a Collaborative Grant is published in the annual report of the Flemish Supercomputer Centre (VSC).

PART V. COMMON PROVISIONS

§ 23 – Procedure and follow-up

This part contains provisions that are common to the processing and follow-up of (regular) Applications and to Starting Grant and Collaborative Grant applications.

§ 24 – Exoneration Clauses

(1) The FWO, KU Leuven and UGent can in no way be held liable for personal injury or property damage arising directly or indirectly from the implementation of a granted Application or a Starting Grant or Collaborative Grant application. The Applicant shall hold harmless the FWO, KU Leuven and UGent from and against any claim for compensation by persons suffering loss or damage in this regard.

(2) The FWO, KU Leuven and UGent can in no way be held liable if the use of Tier-1 compute were to cause any type of loss or damage to a Researcher.

(3) The FWO, KU Leuven and UGent can in no way be held liable for the non-implementation or the partial non-implementation of a granted Application or Starting Grant or Collaborative Grant application for any technical reason or any technical failure specific to Tier-1 compute. The Applicant shall hold harmless the FWO, KU Leuven and UGent from and against any claim for compensation by persons suffering loss or damage in this regard.

(4) If KU Leuven or UGent grants a request for assistance within the meaning of § 16, the Applicant can in no way hold the FWO, KU Leuven or UGent liable if the software cannot be installed or if the implementation of the Application has thus been delayed. Neither will it be possible to claim any type of compensation (validity term, amount of allocated computing time, etc.).

(5) The Researcher is the sole responsible for the timely planning and sending of computing tasks to the Queue. Waiting time as a result of the computing tasks of other users can never be invoked as a justification for total or partial non-implementation of a granted Application or Starting Grant or Collaborative Grant application. Neither will it be possible to claim any type of compensation (validity term, amount of allocated computing time, etc.).

§ 25 – Confidentiality

(1) Except for the abstract that is attached to the Application (unless that too has been identified by the Applicant as being subject to a confidentiality clause), the FWO undertakes to process the Applications with the strictest confidentiality and not to disclose, transfer or divulge them to third parties, either in full or in part, in whatever form, directly or indirectly, except within the framework of the evaluation of the Applications. In that case, the FWO undertakes to have that third party endorse the same confidentiality obligation as that contained in this article.

(2) An Applicant can, however, expressly indicate in the Application that he/she is willing to waive the confidentiality clause §25 (1), thereby authorising the FWO and representatives of the VSC operational team to use the Application in its entirety as an example for other researchers.

(3) The FWO, KU Leuven and UGent shall, in accordance with the due diligence principle, take all reasonable measures to ensure that the data stored or generated on the "SCRATCH" storage volume of the respective Tier-1 compute partition(s) within the framework of the Application are not disclosed, transferred or divulged to third parties, either in full or in part, in whatever form, directly or indirectly.

(4) The same confidentiality provisions shall apply to Starting Grant and Collaborative Grant applications.

§ 26 – Data Retention

The FWO, KU Leuven and UGent reserve the right to delete the Researcher's "SCRATCH" storage volume on the respective Tier-1 compute partition without back-up, once the period within which the allocated core-hours and GPU-hours are to be used, has expired.

The FWO, KU Leuven and UGent can in no case be held liable for data loss on the "SCRATCH" storage volume for any technical reason or any technical failure specific to the Tier-1 compute.

§ 27 – Acknowledgement

In the event of publication of the results of research for which the Tier-1 supercomputing platform has been used, the Flemish Supercomputer Centre (VSC) and the FWO shall be mentioned as follows:

De infrastructuur en dienstverlening gebruikt in dit werk werd voorzien door het VSC (Vlaams Supercomputer Centrum), gefinancierd door het FWO en de Vlaamse overheid.

The resources and services used in this work were provided by the VSC (Flemish Supercomputer Centre), funded by the Research Foundation - Flanders (FWO) and the Flemish Government.

PART VI. PUBLICATION, EVALUATION AND AMENDMENT

§ 28 – Publication

These regulations shall be published simultaneously on the websites of the Flemish Supercomputer Centre and the FWO.

§ 29 – Evaluation

Each year, no later than after the decision on the applications submitted in connection with the third cut-off date, the Evaluation Committee shall evaluate the application of the regulations and report its findings to the Board of Trustees of the FWO, formulating a number of amendment proposals where necessary.

§ 30 – Amendment

The Board of Trustees of the FWO can amend the regulations, the amended regulations shall take effect at least 4 weeks before the next cut-off date.

Attachments:

- Enclosure 1 – Standard application form
- Enclosure 2 – Starting Grant application form
- Enclosure 3 – Collaborative Grant application form

Application form: Compute component of the Flemish Tier-1 supercomputing platform

Title of the application:

EasyChair code in case of resubmission:

EasyChair code in case of continuation:

Applicant name, first name:

Institution:

Research group / department:

E-mail address:

VSC id of all mandated persons, separated by commas:

Core hours applied for:

GPU hours applied for:

Largest amount of scratch disk required (in TiB):

Largest associated number of files:

List of simulation codes and their version numbers:

This application should not exceed 18 pages, excluding possible appendices (confirmation letter of financing institution, software license, etc.) which may be considered by the Tier-1 Allocation Board.

1. Research project within the framework of which computing time is applied for.
 - Title
 - Supervisor(s) and their e-mail address
 - If available, IWETO or FRIS link
 - Financing institution or channel (FWO, BOF, VLAIO, EU, etc.).
 - Attach the confirmation letter as enclosure.
 - Attach a letter of approval of your own institution in case the project has not gone through a scientific approval process.

2. Include a short description of your research project, in layman's terms wherever possible, with a view to dissemination. Explicitly mention the scientific questions that you are planning to address and the overall scientific goals of the project. (max. 1 A4 in Arial 12)

3. Persons mandated by the Applicant to compute on the Tier-1 within the framework of the present project. Please provide for every person:
 - Name, first name
 - VSC id
 - Institution
 - Research group / department:
 - Experience with using particular HPC resources (i.e. Tier-0/Tier-1/Tier-2 infrastructure) in Belgium and abroad. Specify both the name of infrastructure and number of years it was used.
 - List of computing time allocations received during the past two years, on the Flemish Tier-1 systems, as well as other Tier-1 and Tier-0 systems.

4. Why does this project need to run on a Tier-1 system? Select appropriate checkbox(es):
 - It requires resources (#nodes, #cores, #GPUs, memory, interconnect, storage) that are not available on Tier-2 systems.
 - The total runtime for the project (e.g. due to number of jobs, job turnaround time) would be prohibitively large on Tier-2 systems.
 - Other reason, please specify.

5. Provide information for each software package that will be used.
 - If centrally installed on Tier-1 compute or a Tier-2 system within VSC, state the module name.
 - If not open source software, state that the associated license can be validly used by all mandated users on the desired Tier-1 compute partition (BrENIAC/Hortense). Add a copy of the signed license to this application.

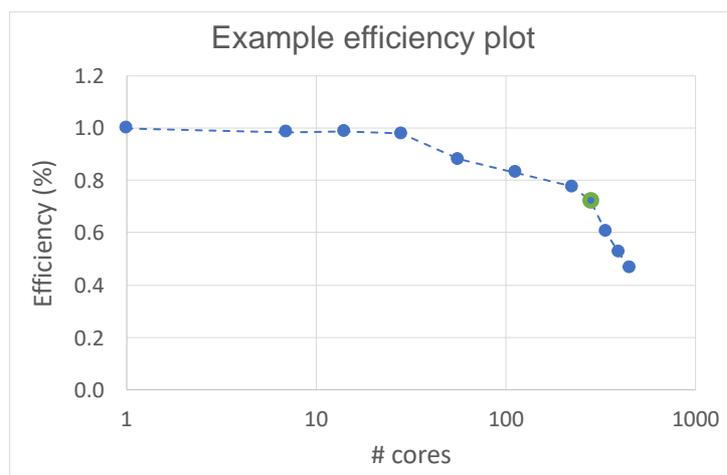
6. Provide the results of parallel efficiency tests for each software package that will be used.
 - Perform these benchmark tests on the Tier-1 compute infrastructure (using, e.g., a Starting Grant), on the desired partition (Hortense/BrENIAC). Make sure to mention this particular partition.
 - Use system/problem sizes that closely reflect those of the intended computational tasks (e.g., same mesh size, actual molecular system, similar I/O pattern, same communications patterns, etc.). If a different system/problem size is used in the tests, describe how it relates to the problem size in the application. State if I/O has been included in the tests. For example, simply run your application tasks for a limited number of iterations.
 - List the results in a table and plot efficiency versus number of cores or number of GPUs using a log scale x-axis (see example Table 1 and Plot 1).
 - Start the scaling tests of your code using the smallest number of cores or GPUs possible. If possible, the baseline is using 1 core or 1 GPU on a dedicated node. If not possible, explicitly state why.
 - Wall clock times are preferably obtained by averaging the timing results of several similar simulations for each node/core/GPU configuration. This is required when task farming jobs that vary significantly in run time. In that case, give an indication of variation (e.g. standard deviation). Otherwise, report the minimum time of at least three separate runs.
 - Task loads that don't use the maximum number of cores/GPUs per node are preferentially packed together, using the worker framework, atools, ...
 - Explain anomalies in plot and table.
 - Clarify, on the basis of the parallel efficiency plot and table, which number of nodes and cores/GPUs you plan to use for your computational tasks (cf. Section 7) and explain why.

Example Table 1

| Number of nodes | Total number of cores | Wall clock time (s) | Speed-up (w.r.t. baseline) | Efficiency |
|--|-----------------------|---------------------|----------------------------|---|
| $A_{baseline}$ | $B_{baseline}$ | $C_{baseline}$ | 1.00 | 1.00 |
| A1 | B1 | C1 | $C_{baseline}/C1$ | $(B_{baseline} * C_{baseline}) / (B1 * C1)$ |
| A2 | B2 | C2 | $C_{baseline}/C2$ | $(B_{baseline} * C_{baseline}) / (B2 * C2)$ |
| <i>Baseline = minimal configuration with which your computational task can be carried out on Tier-1.</i> | | | | |
| <i>Wall clock time is difference between start/end of the computational task, including any I/O operations as part of that task.</i> | | | | |

| Tier-1 partition on which benchmark was performed: ... | | | | |
|--|-----------------------|---------------------|----------------------------|------------|
| Number of nodes | Total number of cores | Wall clock time (s) | Speed-up (w.r.t. baseline) | Efficiency |
| 1 | 1 | 4000 | 1.00 | 1.00 |
| 1 | 7 | 580 | 6.90 | 0.99 |
| 1 | 14 | 289 | 13.84 | 0.99 |
| 1 | 28 | 146 | 27.40 | 0.98 |
| 2 | 56 | 81 | 49.38 | 0.88 |
| 4 | 112 | 43 | 93.02 | 0.83 |
| 8 | 224 | 23 | 173.91 | 0.78 |
| 10 | 280 | 19.8 | 202.02 | 0.72 |
| 12 | 336 | 19.7 | 203.05 | 0.60 |
| 14 | 392 | 19.4 | 206.19 | 0.53 |
| 16 | 448 | 19.2 | 208.33 | 0.47 |

Example Plot 1



The optimal number of cores in this example is 280, as parallel efficiency quickly drops below 70% when more cores are used.

7. Justify the number of core-hours and GPU-hours, and storage volume applied for.

Describe your planned computational tasks and the sequence in which these tasks will be performed. Start from the examples in Table 2 and Table 3 and adjust them to your project.

Note that per requested GPU-hour on Hortense, you will automatically receive 12 core-hours on the CPU cores of the node containing that GPU unit. These core-hours do not need to be specified explicitly in Table 3.

Provide additional descriptions for the computational tasks listed in the table. Resource estimates (wall clock time, number of nodes/cores/GPUs, memory, storage) should be based on the results of actual calculations on BrENIAC or Hortense (via, e.g., a Starting Grant) for system/problem sizes that match closely those of the intended computing tasks (e.g., same mesh sizes, actual molecular system, same I/O pattern, same amount of communications, etc.). If you plan to run the tasks concurrently, mention this in the description, so you can specify the correct total amount of scratch space required.

NB: After 3 months of the allocation time have passed, you will lose 20% of the initially granted core-hours and GPU-hours, if that 20% has not been used. (cf. regulations, § 9)

Example Table 2

| Computational task | Core-hour calculation | | | | | Memory usage (GiB) per node per job | OpenMP / MPI / OpenMP + MPI (hybrid) / worker framework / atools / etc. | Storage volume estimate | |
|---|-----------------------|---|--------------------------------|--------------------------------|-------------------------------------|--|--|---|--|
| | Number of such jobs | Wall clock time (in hours) per job | Number of Tier-1 nodes per job | Number of Tier-1 cores per job | Total core-hours per task | | | Tier-2 DATA/HOME volume (TiB) + number of files | Tier-1 SCRATCH volume (TiB) + number of files |
| Task • software X • parameters/conditions • system/mesh size • ... | A | B | C | D | = A x B x C x D | | | | |
| Task example CP2K • CP2K – MD • 100 ns runs • PBE functional • 1 -> 5 water molecules | 5 | 48 | 10 | 28 | 67200 | 64 | MPI | 0 TiB 0 files | 0.1 TiB 5000 files |
| Task example worker • MDTraj postprocessing • 5000 files | 10000 | 0.5 | 1 | 1 | 5000 | 3 (84 GiB for 28 jobs in one node) | These single-core jobs will be packed within 1 node using worker framework | 1 TiB 10000 files | 0.1 TiB 5000 files |
| | | | | | Sum of core-hours applied for = ... | | | | Largest amount of scratch disk required + number of associated files = ... |
| <i>Important information:</i> | | <i>3 days is the maximum wall clock time for any job.</i> | | | | <i>Memory limits (GiB/CPU node)</i> <i>BrENIAC:</i> <i>128->256</i> <i>Hortense:</i> <i>256->512</i> | | | |

The example CP2K task needs to run 5 times, for a molecular system containing 1 to 5 water molecules. Based on timing runs on Tier-1 compute partition BrENIAC, we found that one such job runs for 48 hours on 10 nodes, using all the cores (28) in the node. The job needs 64 GiB RAM in each node and produces 20 GiB of SCRATCH storage (1000 files). Since the 5 jobs (for the 5 listed molecular systems) will be run concurrently, $5 \times 20 \text{ GiB} = 100 \text{ GiB}$ of scratch disk space is required (and $5 \times 1000 = 5000$ files) for the entire task.

File postprocessing with the MDTraj tool is done in the example worker task, where 10000 jobs need to run on 5000 files on the SCRATCH volume to generate 10000 files on the Tier-2 DATA volume. Each job runs on 1 core. Based on 5 timing runs on Tier-1 compute partition BrENIAC, we found that the job duration varies between 25 and 28 minutes, and memory usage is 3 GiB at most. To be on the safe side, we foresee 30 minutes per job (0.5 hours). Wherever possible, 28 jobs will be packed on a single node using the worker framework, so 28 jobs require 1 full node for 30 minutes. For the entire task, 5000 core-hours are required.

Example Table 3

| | GPU-hour calculation | | | | | | | Storage volume estimate | |
|--|----------------------|---|--------------------------------|-------------------------------|------------------------------------|---|---|---|--|
| Computational task | Number of such jobs | Wall clock time (in hours) per job | Number of Tier-1 nodes per job | Number of Tier-1 GPUs per job | Total GPU-hours per task | Memory usage (GiB) per node per job | OpenMP / MPI / OpenMP + MPI (hybrid) / worker framework / atools / etc. | Tier-2 DATA/HOME volume (TiB) + number of files | Tier-1 SCRATCH volume (TiB) + number of files |
| Task <ul style="list-style-type: none"> • software X • parameters/conditions • system/mesh size • ... | A | B | C | D | = A x B x C x D | | | | |
| Task example QE <ul style="list-style-type: none"> • Quantum Espresso • 1,500 compounds • SCF calculation | 1500 | 8 | 1 | 2 | 24000 | 106 | MPI & OpenMP | 0.4 TiB 2500 files | 1.2 TiB 7500 files |
| | | | | | Sum of GPU-hours applied for = ... | | | | Largest amount of scratch disk required + number of associated files = ... |
| <i>Important information:</i> | | <i>3 days is the maximum wall clock time for any job.</i> | | | | <i>Memory limits (GiB/GPU node) Hortense: 256</i> | | | |

The example QuantumEspresso task needs to run 1500 times, to perform an SCF calculation on 1500 different compounds. All tasks can be executed independently of each other. Based on timing runs on the GPU nodes of Tier-1 compute partition Hortense, we found that one such job runs for 8 hours on 1 node, using 2 GPUs along with 24 CPU-cores in the GPU node. Each job requires 106 GiB of RAM, therefore two jobs can run simultaneously on a Hortense GPU node (256 GiB). The worker framework will be used to pack 2 tasks in one job that will make sure both end up on one GPU node, optimally using all GPUs of that node. Each job generates 5 files that total 0.8 GiB. For all tasks, this amounts to $1500 \times 0.8 \text{ GiB} = 1.2 \text{ TiB}$ of SCRATCH storage (7500 files). These will be regularly offloaded to the Tier-2 DATA storage in a compressed format.

8. Describe how you will manage the workflow and the resources requested in the period during which the task is to be performed.

In case you will launch a large number of computational tasks, describe how you will manage your jobs and provide details regarding job management, automation and dataflow. Will you make use of a task/workflow manager, such as the worker framework, atools or something similar? On which infrastructure or node will this manager run?

Please present how you will manage your data. Describe how the transfer of files to/from the Tier-1 compute partition will be managed and automated, if data reduction and/or compression of files will be performed. If available, provide information about IOPS.

9. Can this proposal in its entirety be made public by FWO or VSC e.g. as an example or inspiration for other researchers?

- Yes
 No

Don't hesitate to consult your local support when you are preparing a Tier-1 application:
<https://www.vscentrum.be/getintouch>

Starting Grant Application form: Compute component of the Flemish Tier-1 supercomputing platform

Title of the application:

Applicant name, first name:

Institution:

Research group / department:

E-mail address:

VSC id of the applicant:

Disclaimer

Allocated computing time on Tier-1 for a Starting Grant is 500 kcore-hours and 1 kGPU-hours at most, allocated for a period of 4 months. Total disk storage should not surpass 1 TiB. Allocated time is granted only to the applicant and can never be transferred to anyone else.

1. Motivate your application: why specifically do you want a Starting Grant?
Which Tier-1 compute partition do you want to run on (Hortense/BrENIAC)?
Would you require CPU or GPU resources?
2. Provide information for each software package that will be used.
 - If centrally installed on Tier-1 compute or a Tier-2 system within VSC, state the module name.

- If not open source software, state that the associated license can be validly used by all mandated users on the desired Tier-1 compute partition (BrENIAC/Hortense). Add a copy of the signed license to this application.
 - If not centrally installed, provide compilation instructions: software website, toolchain used, compilers, versions, easyconfig file for EasyBuild, ...
3. Short description of the computing tasks, the software tools required, expected disk storage and memory usage. Please specify if applicable:
- whether these computing tasks use diversification (OpenMP, MPI, OpenMP + MPI (hybrid), ...)
 - the estimated memory use of a computing task (maximum 256 GiB/node on BrENIAC, on Hortense maximum 512 GiB/CPU node and 256 GiB/GPU node)
 - the requirements for disk storage (estimated volume in GiB and the total number of files)

Please make sure to gather all data required for the efficiency table and description of the computational tasks when applying for a regular Tier-1 compute allocation later on.

| |
|---|
| Send the completed Starting Grant application form to: compute@vscentrum.be |
|---|

| |
|---|
| Don't hesitate to consult your local support when you are preparing a Tier-1 application: https://www.vscentrum.be/getintouch |
|---|

Collaborative Grant Application form: Compute component of the Flemish Tier-1 supercomputing platform

Title of the application:

Applicant name, first name:

E-mail address:

Consortium information (research groups, departments, institutions):

VSC id of all mandated persons, separated by commas:

Core hours applied for:

GPU hours applied for:

Largest amount of scratch disk required (in TiB):

Largest associated number of files:

List of simulation codes and their version numbers:

1. Include a short description of your research project, in layman's terms wherever possible, with a view to dissemination. Explicitly mention the scientific questions that you are planning to address and the overall scientific goals of the project. (max. 1 A4 in Arial 12)
2. Persons mandated to compute on the Tier-1 within the framework of the present project. Please provide for every person:
 - Name, first name
 - VSC id
 - Institution
 - Research group / department:
 - Experience with using particular HPC resources (i.e. Tier-0/Tier-1/Tier-2 infrastructure) in Belgium and abroad. Specify both the name of infrastructure and number of years it was used.
 - List of computing time allocations received during the past two years, on the Flemish Tier-1 systems, as well as other Tier-1 and Tier-0 systems.
3. Why does this project need to run on a Tier-1 system? What is the added value of a collaborative grant for your consortium over one or more regular proposals?
4. Provide information for each software package that will be used.
 - If centrally installed on Tier-1 compute or a Tier-2 system within VSC, state the module name.
 - If not open source software, state that the associated license can be validly used by all mandated users on the desired Tier-1 compute partition (BrENIAC/Hortense). Add a copy of the signed license to this application.
5. Provide the results of parallel efficiency tests for each software package that will be used.
 - Perform these benchmark tests on the Tier-1 compute infrastructure (using, e.g., a Starting Grant), on the desired partition (Hortense/BrENIAC). Make sure to mention this particular partition.
 - Use system/problem sizes that closely reflect those of the intended computational tasks (e.g., same mesh size, actual molecular system, similar I/O pattern, same communications patterns, etc.). If a different system/problem size is used in the tests, describe how it relates to the problem size in the application. State if I/O has been included in the tests. For example, simply run your application tasks for a limited number of iterations.
 - List the results in a table and plot efficiency versus number of cores or number of GPUs using a log scale x-axis (see example Table 1 and Plot 1).

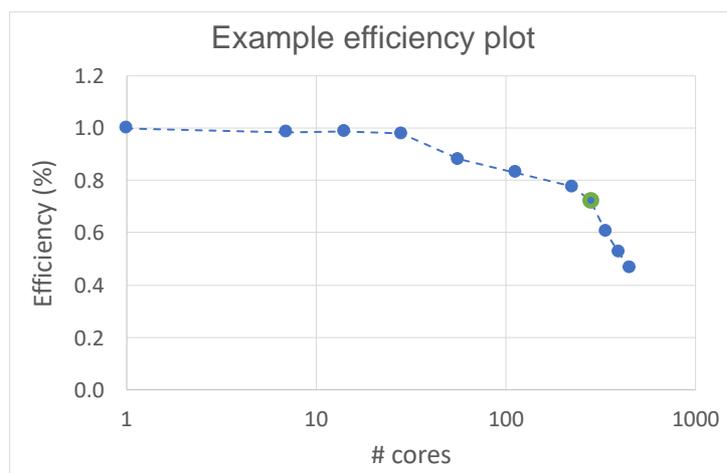
- Start the scaling tests of your code using the smallest number of cores or GPUs possible. If possible, the baseline is using 1 core or 1 GPU on a dedicated node. If not possible, explicitly state why.
- Wall clock times are preferably obtained by averaging the timing results of several similar simulations for each node/core/GPU configuration. This is required when task farming jobs that vary significantly in run time. In that case, give an indication of variation (e.g. standard deviation). Otherwise, report the minimum time of at least three separate runs.
- Task loads that don't use the maximum number of cores/GPUs per node are preferentially packed together, using the worker framework, atools, ...
- Explain anomalies in plot and table.
- Clarify, on the basis of the parallel efficiency plot and table, which number of nodes and cores/GPUs you plan to use for your computational tasks (cf. Section 7) and explain why.

Example Table 1

| Number of nodes | Total number of cores | Wall clock time (s) | Speed-up (w.r.t. baseline) | Efficiency |
|--|-----------------------|---------------------|----------------------------|---|
| $A_{baseline}$ | $B_{baseline}$ | $C_{baseline}$ | 1.00 | 1.00 |
| A1 | B1 | C1 | $C_{baseline}/C1$ | $(B_{baseline} * C_{baseline}) / (B1 * C1)$ |
| A2 | B2 | C2 | $C_{baseline}/C2$ | $(B_{baseline} * C_{baseline}) / (B2 * C2)$ |
| <i>Baseline = minimal configuration with which your computational task can be carried out on Tier-1.</i> | | | | |
| <i>Wall clock time is difference between start/end of the computational task, including any I/O operations as part of that task.</i> | | | | |

| Tier-1 partition on which benchmark was performed: ... | | | | |
|--|-----------------------|---------------------|----------------------------|------------|
| Number of nodes | Total number of cores | Wall clock time (s) | Speed-up (w.r.t. baseline) | Efficiency |
| 1 | 1 | 4000 | 1.00 | 1.00 |
| 1 | 7 | 580 | 6.90 | 0.99 |
| 1 | 14 | 289 | 13.84 | 0.99 |
| 1 | 28 | 146 | 27.40 | 0.98 |
| 2 | 56 | 81 | 49.38 | 0.88 |
| 4 | 112 | 43 | 93.02 | 0.83 |
| 8 | 224 | 23 | 173.91 | 0.78 |
| 10 | 280 | 19.8 | 202.02 | 0.72 |
| 12 | 336 | 19.7 | 203.05 | 0.60 |
| 14 | 392 | 19.4 | 206.19 | 0.53 |
| 16 | 448 | 19.2 | 208.33 | 0.47 |

Example Plot 1



The optimal number of cores in this example is 280, as parallel efficiency quickly drops below 70% when more cores are used.

6. Justify the number of core-hours and GPU-hours, and storage volume applied for.

Describe your planned computational tasks and the sequence in which these tasks will be performed. Start from the examples in Table 2 and Table 3 and adjust them to your project.

Note that per requested GPU-hour on Hortense, you will automatically receive 12 core-hours on the CPU cores of the node containing that GPU unit. These core-hours do not need to be specified explicitly in Table 3.

Provide additional descriptions for the computational tasks listed in the table. Resource estimates (wall clock time, number of nodes/cores/GPUs, memory, storage) should be based on the results of actual calculations on BrENIAC or Hortense (via, e.g., a Starting Grant) for system/problem sizes that match closely those of the intended computing tasks (e.g., same mesh sizes, actual molecular system, same I/O pattern, same amount of communications, etc.). If you plan to run the tasks concurrently, mention this in the description, so you can specify the correct total amount of scratch space required.

Example Table 2

| Computational task | Core-hour calculation | | | | | Memory usage (GiB) per node per job | OpenMP / MPI / OpenMP + MPI (hybrid) / worker framework / atools / etc. | Storage volume estimate | |
|--|-----------------------|---|--------------------------------|--------------------------------|-------------------------------------|--|--|---|--|
| | Number of such jobs | Wall clock time (in hours) per job | Number of Tier-1 nodes per job | Number of Tier-1 cores per job | Total core-hours per task | | | Tier-2 DATA/HOME volume (TiB) + number of files | Tier-1 SCRATCH volume (TiB) + number of files |
| Task <ul style="list-style-type: none"> • software X • parameters/conditions • system/mesh size • ... | A | B | C | D | = A x B x C x D | | | | |
| Task example CP2K <ul style="list-style-type: none"> • CP2K – MD • 100 ns runs • PBE functional • 1 -> 5 water molecules | 5 | 48 | 10 | 28 | 67200 | 64 | MPI | 0 TiB 0 files | 0.1 TiB 5000 files |
| Task example worker <ul style="list-style-type: none"> • MDTraj postprocessing • 5000 files | 10000 | 0.5 | 1 | 1 | 5000 | 3 (84 GiB for 28 jobs in one node) | These single-core jobs will be packed within 1 node using worker framework | 1 TiB 10000 files | 0.1 TiB 5000 files |
| | | | | | Sum of core-hours applied for = ... | | | | Largest amount of scratch disk required + number of associated files = ... |
| <i>Important information:</i> | | <i>3 days is the maximum wall clock time for any job.</i> | | | | <i>Memory limits (GiB/CPU node)</i> <i>BrENIAC:</i> <i>128->256</i> <i>Hortense:</i> <i>256->512</i> | | | |

The example CP2K task needs to run 5 times, for a molecular system containing 1 to 5 water molecules. Based on timing runs on Tier-1 compute partition BrENIAC, we found that one such job runs for 48 hours on 10 nodes, using all the cores (28) in the node. The job needs 64 GiB RAM in each node and produces 20 GiB of SCRATCH storage (1000 files). Since the 5 jobs (for the 5 listed molecular systems) will be run concurrently, $5 \times 20 \text{ GiB} = 100 \text{ GiB}$ of scratch disk space is required (and $5 \times 1000 = 5000$ files) for the entire task.

File postprocessing with the MDTraj tool is done in the example worker task, where 10000 jobs need to run on 5000 files on the SCRATCH volume to generate 10000 files on the Tier-2 DATA volume. Each job runs on 1 core. Based on 5 timing runs on Tier-1 compute partition BrENIAC, we found that the job duration varies between 25 and 28 minutes, and memory usage is 3 GiB at most. To be on the safe side, we foresee 30 minutes per job (0.5 hours). Wherever possible, 28 jobs will be packed on a single node using the worker framework, so 28 jobs require 1 full node for 30 minutes. For the entire task, 5000 core-hours are required.

Example Table 3

| | GPU-hour calculation | | | | | | | Storage volume estimate | |
|--|----------------------|---|--------------------------------|-------------------------------|------------------------------------|---|---|---|--|
| Computational task | Number of such jobs | Wall clock time (in hours) per job | Number of Tier-1 nodes per job | Number of Tier-1 GPUs per job | Total GPU-hours per task | Memory usage (GiB) per node per job | OpenMP / MPI / OpenMP + MPI (hybrid) / worker framework / atools / etc. | Tier-2 DATA/HOME volume (TiB) + number of files | Tier-1 SCRATCH volume (TiB) + number of files |
| Task <ul style="list-style-type: none"> • software X • parameters/conditions • system/mesh size • ... | A | B | C | D | = A x B x C x D | | | | |
| Task example QE <ul style="list-style-type: none"> • Quantum Espresso • 1,500 compounds • SCF calculation | 1500 | 8 | 1 | 2 | 24000 | 106 | MPI & OpenMP | 0.4 TiB 2500 files | 1.2 TiB 7500 files |
| | | | | | Sum of GPU-hours applied for = ... | | | | Largest amount of scratch disk required + number of associated files = ... |
| <i>Important information:</i> | | <i>3 days is the maximum wall clock time for any job.</i> | | | | <i>Memory limits (GiB/GPU node) Hortense: 256</i> | | | |

The example QuantumEspresso task needs to run 1500 times, to perform an SCF calculation on 1500 different compounds. All tasks can be executed independently of each other. Based on timing runs on the GPU nodes of Tier-1 compute partition Hortense, we found that one such job runs for 8 hours on 1 node, using 2 GPUs along with 24 CPU-cores in the GPU node. Each job requires 106 GiB of RAM, therefore two jobs can run simultaneously on a Hortense GPU node (256 GiB). The worker framework will be used to pack 2 tasks in one job that will make sure both end up on one GPU node, optimally using all GPUs of that node. Each job generates 5 files that total 0.8 GiB. For all tasks, this amounts to $1500 \times 0.8 \text{ GiB} = 1.2 \text{ TiB}$ of SCRATCH storage (7500 files). These will be regularly offloaded to the Tier-2 DATA storage in a compressed format.

7. Describe how you will manage the workflow and the resources requested in the period during which the task is to be performed.

In case you will launch a large number of computational tasks, describe how you will manage your jobs and provide details regarding job management, automation and dataflow. Will you make use of a task/workflow manager, such as the worker framework, atools or something similar? On which infrastructure or node will this manager run?

Please present how you will manage your data. Describe how the transfer of files to/from the Tier-1 compute partition will be managed and automated, if data reduction and/or compression of files will be performed. If available, provide information about IOPS.