

DiRAC Resource Allocation Committee

Guidance Notes for Applications

1	<i>Introduction</i>	2
2	<i>DiRAC Resource Allocation Committee</i>	2
3	<i>DiRAC RAC Calls for Proposals</i>	2
4	<i>Enquiries</i>	2
5	<i>Proposal Types</i>	2
5.1	Short Projects	2
5.2	Thematic Projects	3
5.3	Discretionary/Seedcorn Proposals	3
5.4	Research Software Engineering Support	3
6	<i>How to apply</i>	4
7	<i>General guidance for applicants</i>	4
8	<i>Assessment Criteria</i>	5
9	<i>Assessment Process</i>	6
9.1	Discretionary/Seedcorn Proposals	6
9.2	Thematic Projects and Short Projects	6
10	<i>How to Submit a Proposal</i>	6
10.1	Submission process	6
10.2	Principal Investigator	7
10.3	Discretionary/Seedcorn Proposals	7
10.4	Thematic Projects and Short Projects	7
11	<i>Code Efficiency</i>	8
12	<i>Data Management Plan</i>	9
13	<i>Project Reporting</i>	9
14	<i>Storage Policy</i>	9
Annex 1	<i>: RAC Scoring Matrix</i>	11

1 Introduction

DiRAC (Distributed Research utilizing Advanced Computing) is the STFC national HPC resource for the UK astronomy, cosmology, particle physics and nuclear physics communities. Information on the HPC services that DiRAC offers can be found at www.dirac.ac.uk. This document explains the process for the allocation of time on the DiRAC facilities.

2 DiRAC Resource Allocation Committee

The DiRAC Resource Allocation Committee (RAC) was established to oversee the time allocation for DiRAC project proposals. The membership of the RAC is available at www.dirac.ac.uk. To maximise the quality of the scientific output of DiRAC, the allocation of time for Thematic Projects and Short Projects will be determined via robust, transparent peer review. The RAC has two sub-panels, one for Particle Physics & Nuclear Physics and one for Astronomy & Cosmology, which will consider the proposals within their respective remits. A RAC meeting with representatives from both sub-panels will then determine the overall allocation of DiRAC time across the whole portfolio of proposals.

3 DiRAC RAC Calls for Proposals

The RAC will issue one call to apply for time on the DiRAC facility per year.

Closing Date	Allocations Start	Proposal Types
1 st October 2018	1 st April 2019	All proposal types
Anytime	As soon as possible	Discretionary and Seedcorn proposals

4 Enquiries

All Call forms and documentation can be found on the DiRAC website: www.dirac.ac.uk

Further enquiries should be directed as follows:

- RAC process and remit: STFC Swindon Office (DiRACRAC@stfc.ac.uk)
- Technical questions: Technical Working Group (Lydia.Heck@durham.ac.uk)
- Direct allocations or discretionary requests: DiRAC Director, Dr Mark Wilkinson (miw6@leicester.ac.uk)

Enquiries received within 7 days of a closing date are not guaranteed to be responded prior to the closing date.

5 Proposal Types

The categories of proposals are:

1. Short Projects
2. Thematic Projects
3. Discretionary/Seedcorn
4. Research Software Engineer support

5.1 Short Projects

A Short Project is a self-contained research problem typically lasting 3-6 months, up to a maximum of 12 months. The category will include proposals intended to develop exploratory study by users new to HPC or to DiRAC. Investigators on Thematic Projects may use the Short Project route to apply for resources for new sub-projects that fall outside the scope of the original, peer-reviewed Thematic Project proposal, provided this will not impact on their ability to deliver the research programme associated with the Thematic Project.

5.2 Thematic Projects

It is anticipated that the majority of DiRAC compute time will be allocated to a relatively small number of Thematic Projects.

A Thematic Project is a clearly defined research programme of outstanding scientific merit which requires significant HPC resources over a period longer than 12 months and up to 36 months duration. The proposed research should be world-leading, with the expectation of making step changes in knowledge through the use of DiRAC resources. Applicants must demonstrate a track record of the productive use of HPC. Thematic projects must be centered on a singular scientific theme, but can contain a small number of sub-project activities, and can span multiple Research Organisations/Institutes. They can consist of a number of institutions or groups working on a large project activity.

Large projects with significantly varying scientific themes are advised to submit separate applications. Applicants should consider the range of material held within submissions as very large proposals can have a detrimental effect to the peer review system due to effective justification and the need for reviewer diversity. No single application may request more than 80% of the RAC available time on any individual machine (up to date figures can be found on the DiRAC website). Requests above this will not be considered by the RAC.

5.3 Discretionary/Seedcorn Proposals

These are very small allocations of DiRAC resource (up to 50k resource hours) for projects that fall into the following categories:

- Scientifically outstanding projects where DiRAC resources could enable a breakthrough to be made but where the impact of the research would be lost if the project were submitted according to the scheduled calls for proposals.
- Very small projects where the researcher is not already a member of an existing Short Project or Thematic Project.

Seedcorn applications cannot be used to uplift existing project activities. Proposals may be submitted at any time. Decisions are normally made within two weeks.

5.4 Research Software Engineering Support

All time allocated projects will be given support for code profiling and simple optimisation on a best efforts basis. Existing and new projects, can apply for additional resources in order to undertake more extensive research software engineering support.

- All Projects will be required to complete the simple optimisation step within 2 months of the award.
- This service will provide a profile of your code during runtime and simple optimisations will be applied to the workflow. We envisage that very few lines of code will be changed during this process.

- We will ensure that you are using an optimised version of the appropriate community code if this is applicable.

All proposals are able to ask for further assistance from a research software engineer to carry out the tasks listed below. Please note:

- Research Software Engineer Support can only be requested for software which is relevant to existing projects which are currently running on DiRAC facilities, or for projects where computing time is requested in parallel with the research software engineer support.
- This effort is technical in nature and is NOT research support. In particular:
 - These resources are not meant to be a replacement for PDRA/Post-grad student activity. The construction of a piece of scientifically valid code is the project's responsibility and is NOT the role of research software engineers.
 - RSE's role to help to produce fast optimised code based upon work that you have already done and demonstrated is scientifically valid.
- In total 2.0 FTE of effort per year will be available. We expect a typical award to be 1-2 months of staff effort, but up to 12 months per year can be awarded if the RAC assess this to be well justified.
- Note that resource requests and science goals in DiRAC computing time applications should not depend on progress or goals you aim to achieve with the support of DiRAC Research Software Engineer Resources.
- Research Software Engineering Support will be distributed according to ranking and availability by the RAC.
- Any general or non-project specific code optimisation to community codes are not awarded through this mechanism; these requests should go directly through the DiRAC Director and is outside the 2 FTE discussed here.

6 How to apply

Applications for Short Projects, Thematic Projects and RSE support all require a previously submitted technical case. Any received proposals will not be considered by the RAC unless this has been completed. **Technical cases can be submitted direct to the Technical Working Group at any time up to 1st September 2018.** If technical cases are submitted after this date it cannot be guaranteed that there will be sufficient time for assessment and subsequent consideration by the RAC.

Following the submission of the technical case, proposals will have to complete the following documentation on or prior to the closing date of 1st October 2018:

1. Completed application form
2. Scientific case
3. Management case
4. RSE application form (if relevant)
5. Any letters of support (non-mandatory, maximum of 3)

7 General guidance for applicants

New application rules:

No single application may request for more than 80% of the available time on any individual machine.

Existing Thematic Projects wishing to apply for more computing time due to additional resources becoming available at DiRAC can submit using the following methods:

- Applications with scientific themes distinct from the existing award can be submitted as a separate proposal.
- Applications building on the same scientific theme as an existing award should apply as a new project, and this new award would then replace any existing compute award. **PI's requesting for a revised or updated long project must justify this request fully; the RAC will take into account all currently active projects which hold a comparable science case.**

Proposals should include the following:

- Proposals should be focused on scientifically coherent themes, but should hold sufficient technical and scientific detail. It should be noted that proposals with greater numbers of themes will generally result in poor coverage and potentially weaker reviews. In these cases it may be beneficial to submit multiple thematic requests.
- Proposals should include adequate detail to justify the requested allocation and should be written in a way that is accessible to the RAC panel, who may not be an expert in the given field (this especially applies to Thematic Projects proposals).

The DiRAC resources are divided into four allocation periods, starting 1st April, 1st July, 1st October, and 1st January. Successful applicants will be advised of the total amount of resources they have been allocated and the periods within which the allocations must be used. Resources must be used in the allocation period to which they were assigned; they cannot be carried over to the next allocation period. For Thematic awards the first year allocation will be fixed and subsequent years are subject to change by DiRAC. Significant deviations from uniform profiles may be requested in Thematic Project proposals with adequate justification, but it is not guaranteed that they can be accommodated. For Thematic Projects a uniform resource usage profile will be assumed.

It is possible to request a delay to the start date of a project, this should be specified on the application form (start dates must be on the first day of a month).

Up to date available resources which can be applied for can be found on the DiRAC website. Please check these regularly as they are updated.

8 Assessment Criteria

Proposals will be assessed according to the following criteria:

Scientific Excellence (Total weighting of 60%)

- Significance of the proposed research goals with reference to the STFC Roadmap
- Strategic value within the STFC programme
- National and international competitiveness and leadership
- Suitability and national/international standing of the investigator(s) for the proposed research

Project Management (weighting 10%) and Data Management (weighting 10%)

- Feasibility of project timeline given DiRAC resources requested and size of investigator team
- Project risk and mitigation analysis (include compute/storage usage, staff and technical risks)

- Publication plans
- Availability of sufficient researcher effort to carry out the proposed research
- Direct response to the [STFC Data Management policy](#) guidelines

Technical Assessment by TWG (Total weighting 20%)

- Appropriateness of the proposed architecture/machine selection
- Justification for any research software engineering support requested
- Efficiency of resource usage and how well code(s) vectorise

RAC Additional Scoring considerations

- Alignment of the proposal with the call guidelines.
- Full and effective usage of previous DiRAC allocations
- Timely submission of interim and final reports of any previous DiRAC allocation (see section 13)

A score between 0 – 10 will be applied to the RAC additional scoring considerations. This will then form a percentage decrease in the overall score gained from the main criteria with a maximum of 20% decrease in score. For example, a score of 9 would produce a 2% reduction in the final score.

9 Assessment Process

9.1 Discretionary/Seedcorn Proposals

Discretionary and Seedcorn proposals can be submitted at any time. They will be reviewed by the chairs of the sub-panels; they will not be sent out to referees. Applicants will be notified of the outcome as soon as possible, usually within two weeks.

9.2 Thematic Projects and Short Projects

The relevant RAC Sub-Panel Chair will assign a RAC Sub-Panel member to act as primary introducer for each project proposal received. For Short Projects the primary introducer will initially contact four expert reviewers. For Thematic Projects the primary introducer will initially contact six expert reviewers, at least one of which should be from a non-UK based reviewer. For both Short and Thematic Projects a member of the Technical Working Group (TWG) will be asked to provide a technical assessment of the proposal (as detailed above) and a score on previously submitted technical case. Reviewers will only see the submitted documentation to the RAC, the technical assessment is solely completed by the TWG.

Project PIs will be given an opportunity to respond in writing to referees' comments. It is strongly recommended that PIs make use of this opportunity. Each proposal will be considered at a meeting of the relevant Sub-Panel. At the meeting, the Sub-Panel will grade the proposal and the PI's response to the referee reports. At this meeting an overall ranked list of the proposals will be produced.

Following the Sub-Panel meetings a meeting will be convened of a combined RAC and the resources to be assigned to each proposal will be agreed. Applicants will be notified of the outcome for their proposal and provided with feedback as soon as possible after the RAC meeting.

10 How to Submit a Proposal

10.1 Submission process

All proposals, to include all of the documentation listed above, must be sent to STFC Swindon Office to the following email address: DiRACRAC@stfc.ac.uk. Proposals will then be processed and forwarded to DiRAC. Applicants may wish to send their proposal via encrypted/secure email.

10.2 Principal Investigator

Each proposal must identify a Principal Investigator (PI) who has overall responsibility for the delivery of the proposed research and will act as the point of contact for all STFC, DiRAC and RAC communications.

In line with STFC's research grant conditions the PI must be either a) resident in the UK, or b) be employed by an overseas Research Organisation approved by STFC as eligible to apply for research grant funding. The STFC eligibility criteria can be found via this [link](#).

We welcome proposals that would be the UK's contribution to an International Research Programme. However it is expected the proposed research programme that will use the DiRAC facility will enhance the UK's research outputs and that the users fulfil the formal STFC research grant eligibility criteria.

10.3 Discretionary/Seedcorn Proposals

Applicants should submit the application form together with a science case (maximum one page) describing the work to be undertaken and explaining the reasons for wishing to use DiRAC. For proposals which are being submitted via this route due to their being time critical, the proposal should include a clear statement of why the impact of the work would be lost by delaying until the following RAC submission deadline. Proposals should be sent to STFC Programmes Directorate, DiRACRAC@stfc.ac.uk. Applicants may wish to send their proposal via encrypted/secure email.

10.4 Thematic Projects and Short Projects

Applicants should submit the application form together with a case for support. Proposals should be sent to STFC Programmes Directorate, DiRACRAC@stfc.ac.uk. Applicants may wish to send their proposal via encrypted/secure email.

The case for support must be a maximum of 12 pages for Thematic Projects, and a maximum of 8 pages for Short Projects, font size 11 point, and contain the following sections:

1. Science justification and proposed research (maximum 9 pages Thematic Projects, 6 pages Short Projects)

This should describe the proposed research programme and explain why the applicants wish to use DiRAC. It should address the following assessment criteria:

- Significance of the proposed research goals with reference to the STFC Roadmap
- Appropriateness of the proposed methods/codes
- Justification of the requested resources
- Justification for any research software engineering support requested
- Suitability of the investigator(s) for the proposed research
- Justification of any periods of machine use in exclusive mode

2. Project and Data Management (maximum 3 pages Thematic Projects, 2 pages Short Projects):

- The project management structure
- Data management plan (refer to section 13)
- Availability of sufficient researcher effort to carry out the proposed research (For Thematic Projects a table should be included showing the estimated effort, as a percentage, that the PI and each Co-I expect to contribute to the project)

- A description of the internal allocation process for the allocation of time to sub-projects (if applicable)
- A work plan, with milestones against which the progress of the project will be measured
- Publication plans
- Project risk and mitigation analysis (include compute/storage usage, staff and technical risks)

3. Relevant Publications produced via DiRAC resources over the last 3 years (short annex)

4. References (short annex)

11 Code Efficiency

The aim of the RAC process is to maximise the output of high quality research by the DiRAC facility. Scientific excellence will be the primary driver for allocation decisions and the RAC will balance ‘time to science’ against reasonable requirements on the operational efficiency of approved projects and simulation codes. It is recognised that the cutting-edge and novel nature of research across the DiRAC community means that many DiRAC codes are under active development and may not be as efficient or scalable as more mature codes. Further, results obtained in a timely manner with a sub-optimal code will often have greater impact than results delayed by extended periods of code optimisation work. However, while it is recognised that new HPC users, or users of new codes, may not have sufficient resources or experience to provide full details of code efficiency, applicants will be required to demonstrate that their operational plan is as efficient as possible (e.g. running multiple jobs sequentially on small numbers of cores versus running jobs concurrently on larger numbers of cores) and that the architecture requested is the most appropriate for the work.

In cases of similarly ranked proposals, applicants who demonstrate more efficient use of DiRAC resources either in terms of actual code efficiency or more efficient operational strategies will be given preference. Projects may request allocations of research software engineer effort to assist with code optimisation. The RAC may also allocate research software engineer effort to particular projects where concerns about efficiency have been raised.

Where scientific projects intend to migrate between architectures, RAC particularly encourage codes that have been well-optimised and scaled to larger project sizes than would have otherwise been attainable. A development queue will be available for those preparing proposals who want to demonstrate code efficiency. This queue will have strict limits on resource usage. Requests for access to this queue for new users will be dealt with by the Technical Working Group.

Requesting Exclusive Use of a DiRAC Machine:

Some Thematic Projects may include (or may consist entirely of) sub-projects, which require usage of an entire DiRAC machine, or significant fraction of a machine, for a period of longer than two days. This mode of use must be explicitly justified in the proposal, and a detailed timeline for the sub-project must be included. The technical assessment of such sub-projects will include an assessment of the efficiency of machine use.

If the request is approved, a fixed start date for the sub-project exclusive usage will be agreed to enable re-scheduling of other users to other machines during the period of unavailability. Time lost due to failure to meet the approved start deadline will not be compensated.

No more than two DiRAC machines will be operating in this mode at any one time. Where this mode of operation directly impacts on another project (e.g. due to technical requirements which mean it cannot be moved to another machine) it may be necessary to sub-divide periods of exclusive use or reserve a fraction of the cores for other projects. In cases where multiple Project proposals with overlapping scientific goals are received, if appropriate the RAC may invite the applicants to consider merging their proposals.

12 Data Management Plan

Projects are required to complete a short data management plan for all types of proposals. The completion of this section is mandatory and will be marked against in the review of the project by the DiRAC RAC.

PIs are requested to address all points which are found in the STFC Data Management Plan (DMP) guidelines ([Link](#)). This includes: which communities may have potential interest in the data being produced by their DiRAC project; how and where the data from the project will be stored; how long the data will be stored and what metadata processes will be used.

Applicants are encouraged to consult the DMPs of any collaborations they are a part of and to discuss with their universities data management units.

13 Project Reporting

All projects (Thematic, Short and Seedcorn) will be required to submit reports at the end of the project describing facility usage, progress against objectives, achievements and publications. In addition Thematic Projects will be required to complete annual progress reports. Report templates will be provided.

The annual progress report for Thematic Projects will be assessed by the appropriate RAC sub-panel against the original peer-reviewed milestones. Confirmation of resource allocations to a Thematic Project in subsequent allocation periods will be conditional on the approval of their progress report. The report should include explicit justification for any significant deviations from the science proposed in the original work plan, and any delays should be accounted for. Thematic Projects which are deemed not to be progressing satisfactorily may have their allocations in subsequent allocation periods reduced and/or may be required to submit interim reports at six month intervals. The RAC may solicit expert referees' reports in cases where scientifically significant changes to the original milestones for a Thematic Project are proposed in the annual report.

Failure to submit a satisfactory project annual report or project final report will result in subsequent proposals from the group being marked down (see section 9).

14 Storage Policy

Quota - an amount of disk that you cannot exceed.

Allocation - an amount of disk that you are guaranteed to have access to.

/home - this will be small and have quotas applied of, say, 10GB. This is for storing code, key input files, etc. but is not a working space.

"/scratch or /work - this is the main working area available to all compute nodes and is quota'd by

having separate volumes/file systems, one per project, thin provisioned on the storage. The amount of space allocated to scratch can be over-allocated, hence it is a quota rather than an allocation. It is left to the users to manage their data within the limits of the set quota. However each DIRAC site reserve the right to sweep of files which have not been accessed for a given amount of time. This is at the discretion of each site."

/data (medium term) - this is for storing results awaiting final post processing or on which more work may be required prior to publication and transfer back to the user's own institution. These will be quota-ed to reduce the risk of significant overutilisation. Sweeping will be used here as well, as it is not a permanent archive, but will be on a much longer access time, say 9 months.

Archive: This is tape storage. Applicants must specify what data products they wish to have backed up to tape.

Annex 1 : RAC Scoring Matrix

Category	N/A	Unfundable	Not Competitive	Good	Excellent	Exceptional
Score	0	1-2	3-4	5-6	7-8	9-10
Scientific Case	No Science Case submitted, or at an insufficient quality to be considered.	Science that is not novel, is of unsatisfactory quality, unrealistic objectives, is not timely and is unlikely to advance the field.	The science lacks novelty, is of low quality, not well thought through objectives and is not timely but could result in some useful knowledge.	The proposal is good quality science, which is of very good scientific merit in terms of novelty, quality, objectives, and timeliness and addresses important scientific questions.	The proposal is high quality science, which is of excellent scientific merit in terms of novelty, quality, objectives, and timeliness and addresses highly important scientific questions.	The proposed work meets exceptional scientific standards in terms of novelty, quality, objectives, and timeliness and addresses extremely important scientific questions.
	No Science Case submitted, or at an insufficient quality to be considered.	The research is unlikely to contribute to the understanding of the subject. Not a good match of research to the host group's scientific infrastructure and environment with no benefit to the group and no opportunity to develop new skills.	The research will make a marginal contribution to the understanding of the subject. Not a strong match of research to the host group's scientific infrastructure and environment with little benefit to the group and little opportunity to develop new skills.	The research is not likely to make a significant contribution to the understanding of the subject. Good match of research to the host group's scientific infrastructure and environment with some benefit to the group and the opportunity to develop some new skills .	The research is likely to make a significant contribution to the understanding of the subject. Excellent match of research to the host group's scientific infrastructure and environment with excellent benefit to the group and the opportunity to develop new skills.	The research is highly likely to make a significant contribution to the understanding of the subject. Exceptional match of research to the host group's scientific infrastructure and environment with exceptional benefit to the group and the opportunity to develop new skills.
	No resource requirements/requests or justification submitted, or at an insufficient quality to be considered.	Resource requirements/requests and justification are poorly defined, and cannot be assessed against the guidance criteria.	Resource requirements/requests and justification are not well defined, and cannot accurately be assessed on the guidance criteria.	Resource requirements/requests and justification are defined against the guidance criteria.	Resource requirements/requests and justification are well defined, and can accurately be assessed against the guidance criteria.	Resource requirements/requests and justification are clearly defined against the guidance criteria, and can easily be evaluated.
Management Case	No Project Management or Data Management Case submitted, or at an insufficient quality to be considered.	Plan does not address any of the STFC Project or Data Management frameworks.	Plans covers the minimal percentage the relevant aspects of the Project and Data Management Frameworks outlined by STFC.	Plans cover a large percentage the relevant aspects of the Project and Data Management Frameworks outlined by STFC.	Plans cover all the relevant aspects of the Project and Data Management Frameworks outlined by STFC.	Plans have exceptional coverage of the Project and Data Management frameworks outlined by STFC.
Overall	This research should not be considered	This research should not be supported.	This research should not be supported.	This research could be supported if funds are available.	This research should be supported.	This research programme should definitely be supported and it would be a loss not to do so.