

The Resource Request Document

Updated: J. Hays, 13th October 2020

Contents

Abbreviations	2
Resource Request Preparation & Submission Process.....	2
Completing the Resource Request.....	3
Completing the Spreadsheet.....	3
Assessment Criteria.....	4
1 Your Project – The Administrative Details	5
2 Science Programme to be supported by IRIS	5
3 Report on the use made of IRIS resources in the previous year.....	5
4 The Computing Model and Resource Request.....	6
4.1 Justification of Resource Request.....	6
4.2 Resource Model.....	6
4.3 Summary of Request	7
5 Computing Environment.....	7
5.1 Basic compute information.....	7
5.2 Access requirements	8
5.3 GPUs.....	8
5.4 Storage.....	8
5.5 Networking.....	8
5.6 Software	9
6 References	9

Abbreviations

- STFC: Science and Technology Research Council (<http://stfc.ukri.org>)
- IRIS: e-Infrastructure for Research in STFC (www.iris.ac.uk)
- RSAP: IRIS Resource Scrutiny and Allocation Panel
- TWG: IRIS Technical Working Group
- DB: IRIS Delivery Board
- PPAN: Particle physics, Astrophysics and Nuclear Physics
- PPRP: Projects Peer Review Panel
- PPGP: Particle Physics Grants Panel

Resource Request Preparation & Submission Process

This document is for the Activity leads on the IRIS (www.iris.ac.uk) Delivery Board to guide the preparation of resource requests for submission to the FY21 Resource Allocation round.

This resource request process is intended only for previously “peer reviewed/approved science” (see below what this means)¹. When completing the resource request document PLEASE do NOT re-justify it again as if you were in front of a Panel, just summarise it for information.

The following table highlights the timeline for the process and the key dates and deadlines.

13th October 2020	Request for Resources from IRIS Opened
13th October to 9th November	Partners prepare first draft of Resource Request Documents. Requests for clarification of the process should be directed to the technical scrutiny team
9th November	Deadline for first draft. Applications received after 4pm on 9 th November 2020 will be considered at the discretion of the DB
9th November to 11th December	Formative iteration with technical scrutiny team
11th December	Deadline for final revised submission at 4pm on 11 th December. Applications received after this deadline may be considered at the discretion of the DB
29th January 2020	Preliminary allocations are agreed by RSAP and submitted to TWG and DB. Feedback on preliminary allocation decisions sent to partners.
1st April	Formal responses to partners issued. Grants issued ASAP after this date.

¹If your science to be supported has not been peer reviewed, then it will go to a separate process using the DiRAC RAC. Please contact IRIS management about this.

Completing the Resource Request

Fill out the Resource Request Template completing each section according to the guidance given below. There is no word limit but please be sensible in how much you write. Support for this process is provided by the Technical Scrutiny Team.

The technical scrutiny team are:

- Daniela Bauer, daniela.bauer@imperial.ac.uk
- Hannah Griffin, hannah.griffin@stfc.ac.uk

They will be examining the technical details of the requests and working with applicants during the formative step after the submission of the draft applications. They will also provide support where requested in completing the draft resource request.

Applications are made in a two-step process:

- i) the initial submission of the draft documents followed by a formative iteration with the technical scrutiny team
- ii) the final submission followed by review and allocation by the panel.

The formative step is a critical part of the process to ensure that applications are well aligned with the resources that IRIS can provide and ensure that all relevant information has been supplied in time for the panel to conduct its review of the applications.

Submit your draft documents to the technical scrutiny team by 4pm on 6th November 2020

Submit two final documents to the technical scrutiny team by 4pm on 11th December 2020

- A completed Resource Request Document
- A completed Resource Request Document Spreadsheet

Failure to engage in a timely fashion with the technical scrutiny team may result in a resource request being denied if it fails to provide the required information.

You are strongly advised to contact the team prior to the initial submission.

Completing the Spreadsheet

Fill in the Resource Request Document Spreadsheet. You should make two types of line entries in the spreadsheet.

The first line entry is mandatory. **Enter the resource you are requesting from IRIS** (i.e. over and above any you may have elsewhere). Please provide:

- A firm estimate for 2021
- Preliminary estimates for 2022-2025

The entries should represent the total resource required from IRIS for that particular year – not just the “new” request for that year. For example, if in 2021 you request 1 PB of disk and in 2022 you need another 1 PB in addition the entry for 2022 should be 2 PB representing the total request for that year.

The second line entry is optional. Enter the resource you already have and expect to have in the future **from other non-IRIS sources** in each of the years 2021-2025. This will be used for reporting purposes only.

CPU units:

Please fill in the spreadsheet in approximate **CPU-Cores**. If there is a specific architecture or generation of CPU you have used to estimate your requirements, please indicate this in the CPU Notes box.

GPU units:

Please indicate the number of **GPU cards** that you request. In order to match your request to the available hardware, please indicate in the **Notes section on the request line** the type of workload you expect to run on the cards and, if known, an indicative architecture. Put any additional information in the GPU Notes box.

Storage units:

Disk and Tape are in units of **Disk-TB** and **Tape-TB** respectively.

Assessment Criteria

Requests for resource will be assessed according to:

- Validity of the request:
 - is the science activity a supported activity within the IRIS remit and has appropriate peer review of the science taken place?
- Justification of the resource request:
 - has evidence been presented of how the resource requested will support the activity
- Credibility of the resource request:
 - Are details provided of how the resource requirement was determined and how this relates to the activity
- Effective use of resource from previous allocation
 - Where previous allocation has been made, the effective use of this allocation will be used to moderate the current request. While it is recognised that events out of the control of the applicants can result in under usage of resource, significant under-utilisation of previously allocated resources must be addressed in the application.

Resource Request Guidance

Each section below corresponds to a section in the Resource Request Template document. Complete each section as appropriate following the guidance. Questions about completing the document should be addressed to the technical scrutiny team.

1 Your Project – The Administrative Details

The table should be completed with the details of the contact person associated with the request. For the organisation field provide either the institute of the applicant or the project with which the request is associated, as appropriate, and include the URL of the organization home page if relevant. The information will only be used for the purposes of carrying out official IRIS activities.

2 Science Programme to be supported by IRIS

This information will be used in order to:

- confirm the science programme lies within the IRIS remit and determine the peer-review status of the science activity to be supported with IRIS resources
- provide contextual information for your justification of the resource requested
- produce reports for STFC and/or UKRI.

In the case of the Facilities: Please summarise the science workflows that follow from instrument operation in general terms as appropriate to the Facility.

In the case of PPAN activities: Please summarise your approved science or construction programme, with particular reference to those aspects for which the resources will be used.

It is a principle of IRIS that activities are not subject to double jeopardy, thus the RSAP will not be judging your science programme where this has been peer reviewed elsewhere. In the case of Facilities, it is taken as read that those Facilities are approved to run for an annual schedule and user scientists are given time by relevant peer review bodies. In the case of PPAN activities please note the relevant peer review panel (eg. PPRP, PPGP etc) and cite the relevant grant.

3 Report on the use made of IRIS resources in the previous year

Please report on the use made of IRIS resources in the previous year with respect to your allocation. This will be used by the panel to establish that previously allocated resources were used effectively.

If no allocation was made in the previous year then please put “N/A”. You should include details of your previous allocation and what use was made of the resources requested. In the case of significant deviation between allocation and usage, you should explain the reasons for this. This should usually be given in terms of overall usage under the headings – CPU, Disk, Tape, and GPU. There is no need to break down usage into that corresponding to individual activities that

may have made use of the previous allocation except where this helps to clarify variation from expected usage.

4 The Computing Model and Resource Request

This is the main section of the resource request document. Information provided here will be used by the panel to assess the justification of the resources requested and the credibility of the request.

You should provide the following three pieces of information:

4.1 Justification of Resource Request

Include here a description of the specific activity or objectives that will be supported by the resource being requested, referencing Section 2 above as appropriate. This will be used by the panel to assess the justification of the resource requested.

You should briefly but clearly explain how the resource requested will be used to support the relevant science activities or achieve the relevant science objectives.

4.2 Resource Model

Include here an explanation of how your resource request volumes were determined. This will be used the panel to assess the credibility of the resource requested.

Whether or not you have a formal computing model, you must explain the primary inputs, operational assumptions and policies (e.g. replication policy) that determine your resource request volumes.

This may include but is not limited to:

- Fundamental types of workload (e.g. primary reconstruction, XYZ simulation programme, autocorrelation jobs, instrument data reduction, modelling, fitting, data analysis.....);
- CPU-hours and data I/O volumes per unit of data for the above (event, file, instrument run period);
- Total annual running time of machines or beamline;
- Total volume of raw data recorded per year;
- Simulation volume requirements, ML training volume requirements...;
- Production scheme, e.g. Number of times workloads need to be run per week, year etc. This might be policy driven e.g. "we reconstruct the full data set twice per year;
- Data retention and replication requirements or policies. Note: if you are subject to a DMP then separately estimate the volumes required for compliance;

Please use common sense when completing this section. The idea is that all these factors can be used to produce bottom line requirements in a comprehensible and deterministic way.

The more deterministic and numerical you make this section the better. Please feel free to put details and/or supporting information into appendices or include links to other external documentation.

4.3 Summary of Request

This should be a summary of the resource requested and should be the same as entered into the spreadsheet. You may also use this section to provide additional details, if appropriate, that do not fit onto the spreadsheet.

5 Computing Environment

In this section you should provide details of the computing environment required to carry out the relevant science activities. This will be used by the Technical Scrutiny Team and the Technical Working group for two key purposes:

- i) Help align the request with the resources available from IRIS during the formative step of the request process.
- ii) Provide information to the technical working group to guide the procurement and provision of hardware and related technologies and how we place digital asset funds.

While it is not used directly in the assessment of the resource request it is important to allow us to match requests to available resources and help align requests to what can feasibly be provided.

There are practical restrictions on the services that can be provided by IRIS and it may not be possible to fulfill a request on technical grounds.

Information will also be aggregated across all requests and used to provide evidence of community computing needs to STFC and UKRI.

Please provide what you can that makes sense for your context. If you have nothing to report in a given subsection then it can be left blank. Please feel free to include anything else that you feel is relevant that we haven't explicitly asked for. If you are unsure what would be appropriate to include please discuss this via email with the technical scrutiny team.

The bullet points in the sections below are intended as examples to guide the kind of information that is helpful for the two key purposes listed above. You do not need to answer every question and should feel free to add additional information that you feel is relevant.

You are strongly encouraged to engage with the Technical Scrutiny Team if you have any questions when completing this section.

5.1 Basic compute information

In this section provide basic information about the typical needs for the jobs that you expect to run:

- How many cores do your jobs need?
- How many cores can your jobs use on a multi-core machine?
- How much memory per core,job or node do you need?

- Can your jobs take advantage of hyperthreading?
 - And if so, did you factor this into your CPU core request?
- How much scratch space do you require?
- Do you expect multiple jobs to be able to access shared scratch space?
- IRIS is not an HPC facility, but comment on any special backplane needs or requirements for MPI etc?

5.2 Access requirements

In this section provide details of how you would prefer to access the computing resources. If you are unsure what to put in this section or what might work for your particular workflows, please contact the Technical Scrutiny Team for support.

Examples include:

- Simple ssh login with batch system
- Grid based distributed computing
 - Submission through for example DIRAC or glideinWMS
- Cloud based computing (OpenStack)
 - This supports multiple different functionalities – list the ones you wish to make use of if possible

5.3 GPUs

Please include any extra information relevant to your GPU needs. Examples might include:

- Details of existing usage
- Compatibility requirements with your software
- Type of card and expected usage – eg computation, visualisation.

Please note that the specific architectures available will be limited. Please give enough information about the types of workloads you will be running to match your request onto the available hardware.

5.4 Storage

Please include information about your storage requirements. This may include details that refer to hardware requirements as well as how you expect to collect the output of jobs running on IRIS hardware. For example:

- Do you have special disk storage requirements (e.g. high performance)
- Do you need a simple filesystem to be available?
- What access methods do you require for your storage
 - WebDAV
 - Xrootd
 - rucio
 - scp
 - Other

5.5 Networking

In this section please detail any additional information that may be relevant for network provision. This is particularly important if you have high network bandwidth requirements (for

example anything above 1Gb/s) or if you need sustained connectivity over a long period of time (more than an hour or so) – so that is if you need to transfer a lot of data very quickly or if you are streaming data for a long period of time. If you are unsure what to put in this section please contact the Technical Scrutiny Team.

5.6 Software

In this section please detail information with regards to how you would like to distribute the software to be used in running your jobs.

- How do you distribute your software (e.g. cvmfs, nfs, tarball, container, VM image)?
- Do you have any specific software requirements which you need to let IRIS know about at this stage as it may effect hardware specification?
- Where relevant, which operating system and version does your software use (e.g. CentOS7)?
- Do you need access to central databases/catalogues? If so, how do you currently access them?

6 References

Please add any references you may have used above here.