DiRAC
High Performance
Computing Facility

DiRAC USER LIFECYCLE MANAGEMENT

Authors: The DiRAC Technical Working Group

Contents

1 Introduction

This document describes the DiRAC user management policy. In particular it covers the managing of accounts and associated data (both within the SAFE and on specific DiRAC systems). The document is intended to cover the entire lifecycle of both user accounts and DiRAC projects. The policies given here are general and intended to cover all DiRAC access, but please see the appendix for site specific policies. The document was approved by the DiRAC Technical Directorate in February 2024 and circulated to the Service Management Board Chairs, DiRAC Board and DiRAC Oversight Committee.

2 SAFE enrolment and application for a DiRAC account

2.1 User accounts

For DiRAC users to be able to use any of the DiRAC compute services, they are required to enrol for a SAFE account (https://safe.epcc.ed.ac.uk/dirac). Using this SAFE login, they then apply for a DiRAC account under one or more DiRAC projects (RAC, TWG or directorate approved), on one or more DiRAC systems. The RAC DiRAC project naming is dp*lnm,* where *l, n* and *m* are numbers between 0-9. Benchmark, testing and teaching projects have a different naming conventions such as ds*lmn*, dt*lmn* or dr*lmn*.

When the user applies for a SAFE account, the user is required to enter his/her full name, email address, nationality, institute and department, and the HPC experience they have. They must also enter their title and position (termed “career stage” by SAFE, e.g. UG student, PhD student, PDRA, academic etc). They are also asked to optionally provide their telephone number and the full address of their research institution. The user's email address must be linked to the research institution. DiRAC will usually not accept private email addresses such as gmail or hotmail or other such accounts.

DiRAC users are requested to respond to an annual contact by SAFE to verify their email addresses. If they do not respond within a set period (30 days), their DiRAC accounts will be disabled on the relevant DiRAC services. The user may request the re-enabling of their account for a period of up to a year via SAFE. If an account remains disabled on a DiRAC service for a period of more than 1 year, this account may be retired (see below).

2.2 Users’ actual whereabouts

DiRAC users are from all over the U.K. and quite a few of them are based in research institutions outside of the U.K. A high percentage of the users are on temporary contracts or are PhD students with a fixed term contract with the research institution at the time of their DiRAC enrolment. When they move, users are likely to forget to update their SAFE account or inform the DiRAC site(s) which they are using as part of a project, and we lose track of their whereabouts. Emails at a previous institution can be extended and they do not necessarily form a true image of the user's current affiliation.

The individual DiRAC sites might maintain their own user email mailing lists and usually the first indication that a user has moved is the fact that emails to that mailing list will produce failed deliveries. If we use the DiRAC mailing lists, we – at the individual sites - will not see failed deliveries.

2.3 Do the PIs of projects know each user who participates in their project?

For small projects, the PI of the project will be able to identify each individual user. For the large consortia projects, however, the PI will have to trust their collaborators and CoIs to steer the potential users in the right direction and also inform the PI. This means there is some element of risk and room for potential interlopers so PIs and CoIs should take care to check who is being accepted.

2.4 Ageing of accounts and updates of account information

In universities and other research institutions, user identifiers have expiry dates. For academic members of staff, the expiry dates are far into the future, but for contract staff, the expiry date is set by the end of their contract. For DiRAC SAFE accounts we have no ageing process for user accounts, and so send an annual reminder asking them to verify their email address and check their personal details.

For SAFE (and hence DiRAC) to have better management of user accounts, we propose to apply the following or equivalent rules:

- At the time of the application to be added to a project with DiRAC resources, the applicant has already provided details such as their name, institution and email in SAFE. The request for resources will then be approved by the PI or a delegated project manager. It is assumed that the PI or project managers will have appropriate knowledge of the person applying before accepting their application. A PI or project manager MUST NOT endorse a request for an account on the DiRAC facility as part of their project if they are not fully sure of the origin or entitlement of the applicant.
- SAFE has no concept of user expiry dates: the primary driver for expiry is project

membership. While a user is a member of a valid project, it will not expire. However, on the DiRAC services, accounts should be set with an expiry date for security purposes. Therefore the expiry date of the PI's and project managers user accounts on the DiRAC services may be linked in the first instance to the end of their DiRAC project with a 3 month grace period. For all other users a default expiry date of one year may be set from the date of enrolment. Alternatively DiRAC services may choose to only disable users based on instruction from SAFE.

- When an account is close to expiry on a DiRAC service, an email will be sent to the user warning them of this, with instructions on how to extend their account (usually by raising a help desk ticket). Extension will usually be for an additional year.
- For any DiRAC user, the policy is that only their current institutional email is allowed. There may be temporary mitigations, for example for users moving between institutions, which should be approved by the DiRAC project PI (e.g. RAC-awarded project) or nominated representatives.
- DiRAC project PIs should review their project memberships on an annual basis. They will receive an annual request from SAFE to do this.
- It is the DiRAC users' responsibility to update their SAFE information. To help with that, SAFE sends an automatic reminder once a year, one month before the annual review date. The automatic email should contain the user's information as available on SAFE for review. If a user fails to confirm their SAFE information by the review date, their SAFE account is automatically suspended until their information has been confirmed. SAFE will then send tickets to DiRAC services with their associated user accounts requesting the disablement of their accounts.
- If an email to a user becomes invalid, the SAFE account is automatically suspended. The DiRAC facilities on which the user has an account will be notified, where in turn the user account will be suspended/disabled. An email will be sent to the PI of any project of which the user is a member to make them aware of the e-mail delivery failure.
- If a DiRAC service identify a failed email address, they should suspend the account immediately and inform SAFE who in turn will inform any other DiRAC sites to which the user has access. The user account associated with the failed email address will be suspended.
- If a user account becomes inactive and be suspended, this user may be removed from day-to-day system mailing lists (system downtime, etc), though it will be maintained for

purposes including data deletion (e.g. retirement of old file systems, etc). For DiRAC systems which use the SAFE mailing functionality, users use SAFE to manage their own email preferences, and they can remove themselves from the email list, and there is no manual management of the list membership.

- If a DiRAC service user account has been dormant (no login) for a period of three months (or some other period as determined by the site, which must be at least 3 months and cannot be shorter than the DiRAC default), the account will be suspended. The user will be given a warning of at least a week before hand.
- After a DiRAC account has remained suspended/disabled for a year, the project PI will be informed that the account is scheduled for closure and deletion within 3 months.

3 Projects

3.1 When a project ends or is no longer required

When a project ends

- The PI is responsible for removing the data from the system. There will be a 3 month grace period after the end of the project during which this should be done.
- If accounts are not part of any other project, they will be disabled on the DiRAC facility (see policy below on users leaving a project).

This is actioned by SAFE.

In some cases it may be possible for DiRAC to retain data after a project ends. This is not guaranteed and in general DiRAC does not have long term data curation facilities so data should be considered “at risk” after the end of a project and PIs should discuss their requirements with DiRAC to determine what is feasible. This is an interim policy. Eventually a DiRAC Data Curation Service should manage longer term storage.

3.2 When a user or PI leaves a project

If a PI leaves a project, and an alternative PI cannot be found, the project is deemed ended.

A user leaving a project is most likely to occur in one of the following scenarios:

1. Retirement of a project (typically 3 months after a project has ended, though can be extended).

2. A PI requesting removal of a user no longer working on their project.

3. A user wanting their account removed.

4. A user becoming inactive and having their account disabled, and then not re-enabling it for a 12 month period.

5. Failure by a user to respond to the annual request by SAFE to verify their email address.

Note, for PIs of large projects, who do not have an active role on DiRAC facilities, their account may well be considered moribund, and hence disabled. However, while they remain a project PI, they are still considered a user, even if they do not have a DiRAC account on any DiRAC system.

Once a user has left, the following actions are taken:

1. The DiRAC system account is suspended. Login is no longer possible
2. After a period of at least 12 months of suspension, the account is then retired:
    1. The PI is contacted
    2. Ownership of the project data is transferred to the PI. If available at the site, the PI can request that this data is transferred to a DiRAC archive with a 3-year retention period, removed from disk, and automatically removed at the end of the 3-year period.
    3. If data is added to the DiRAC archive, it is then removed from disk.
    4. After the 3-year archive retention period (at sites with this facility), data will be automatically removed.
    5. The corresponding SAFE account may be retired if it is no longer linked to active user accounts associated with active projects.

This is actioned by SAFE. Please see appendix for site-specific information regarding personal and private data.

It may be possible to restore an account up until personal data has been removed. This is not guaranteed.

A user’s right to use DiRAC is granted by a project PI, nominated approver or data manager, which in turn is allocated by the RAC, technical managers or technical directorate. Our policy is to treat the account and associated data as the property of the PI as the owner of the project and its resources. It is the user’s responsibility to ensure that any data they store on DiRAC is handled appropriately and to copy off anything that they wish to keep to an appropriate location.

A PI or project manager can revoke your access accounts within their project at any time; locking, removing or re-owning the account as appropriate.

A user may voluntarily close their account and yield control of their remaining data to the PI of the project.

When a project is due to end, the PI will receive notification of the closure of the project and its accounts 3 months before the accounts are closed and cleaned. This should be done automatically by SAFE eventually.

If the user has abandoned archive data (at sites with this facility), this will be expired after 3 years, and automatically deleted. It should be noted that hardware failure may mean this data is expired earlier.

If the user has archived data, this will be maintained as long as practical. However, it should be noted that hardware failure may still occur, and that data is not guaranteed.

If a user account is re-owned, any archived data is not re-owned. However, it can be re-owned upon restoration, if requested to the appropriate service desk where the data is held.

3.3 Key PI and User Responsibilities

PIs and individual users have a number of responsibilities in terms of the management and security of accounts and associated data. The following is not an exhaustive list but outlines the key responsibilities in each case.

PI

- The PI's key responsibility is for the overall project and its users. This entails regularly (at least annually) checking membership of the project (including who has managerial roles) and ensuring that new account requests are approved or rejected as appropriate. PIs should also regularly check allocation usage, disk usage, and be aware of any end dates of projects, allocations and services. Some of these responsibilities may be delegated to project managers, but the overall responsibility and accountability resides with the PI.

User

- Individual users' key responsibilities are for their own accounts and associated data, appropriate usage of system resources and ensuring that appropriate security measures are observed with regard to passwords, multi-factor authentication, etc. Users should also ensure that their details are kept up-to-date in the SAFE, especially contact details such as email addresses.

Appendix A: Site specific policies

Site specific policies regarding personal and private data and login authentication is given here. Usually any individual will only have one SAFE account. However, there may be reasons (e.g. training or development) that there are multiple SAFE accounts for a single individual. These accounts are hereafter labelled a SAFE person.

Cambridge

Account creation and ownership

On DiRAC (CSD3) only one user account per SAFE person (SAFE account) is created, which the user is expected to use while working on behalf of one or more projects (not all of which will necessarily be DiRAC projects). The user home directory is considered to be private filespace owned by the user. Each project additionally has an associated project storage area, private to that project. The PI of each project is considered to be the owner of all data in the associated project storage area and they (or a designated data manager) may access all files there on demand. Requests by PIs to access user home directories are not normally granted so it is important that project-related files are not stored in home directories.

Login authentication

User authentication on DiAC (CSD3) is two factor - i.e. one of password + TOTP, or SSH key + TOTP. For more information please see https://docs.hpc.cam.ac.uk/hpc/user-guide/mfa.html. Please note that all SSH private keys should be protected by passphrases.

Durham

Account creation and ownership

On DiRAC (MI) Cosma 7/8 only one user account per SAFE person (SAFE account) is created on COSMA. Users are authorised to use cosma7/8 depending on project.
A home directory is not considered private to the user - upon request, it can be reowned to the PI - after checking with the user if they are contactable. The .ssh contents would be removed first.
Data areas are created for each user and project and can be reowned to the PI upon request.

Login authentication

Authentication is via password and ssh key.

Edinburgh

Account creation and ownership

On DiRAC ES systems only one user account per SAFE Person (SAFE account) is created, which the user is expected to use while working on behalf of one or more projects. All data on the service owned by user accounts is considered property of the project and the PI named by DiRAC has ownership of all project data. Project ownership of data will be determined by their group ownership (for file system based storage) or by their archive ownership (for tape based storage). The PI may request access to any data on the system owned by their project.

Login authentication

User authentication on DiRAC ES systems is two factor and requires both a password and SSH key. All SSH private keys should be protected by passphrases.

Leicester

Account creation and ownership

On DiRAC (DIaL 2/3) systems only one user account per SAFE Person (SAFE account) is created. Users are authorised to user DIaL 2 or 3 by being a member of a project authorised to use that machine (so users can be authorized on either DIaL 2, DIaL 3 or both systems, depending on their project membership).

Home directory data is considered private to the user; scratch/data areas are created for each project and the responsibility of the project PI.

Login authentication

Authentication is via password and OTP, either via an authenticator app or email OTP. More detail can be found here: https://dial3-docs.dirac.ac.uk/Getting_started/connecting_dial3/#multifactor-authentication-mfa

Appendix B: Risks

The following table gives a risk list associated with the DiRAC project and user account lifecycle. The list here is not exhaustive and PIs and users should consider carefully any risks associated with their own projects and accounts.

ID	Risk description	Impact (1-5)	Likelihood (1-5)	Score (Impact * Likelihood)	Treat / Accept	Treatment plan	Notes
1	DiRAC (and sites) delegate approval of users to PIs who may fail to correctly police this	4	5	20	Treat	Annual monitoring of project membership by PI and/or summary emails sent to PIs annually.	
2	PIs may not be able to appropriately confirm ID/status of each user in large groups	2	5	10	Treat	Annual monitoring of project membership by PI and/or summary emails sent to PIs annually and/or possible auditing of membership.	
3	Account deletion for inactive users may lead to loss of valuable data	4	2	8	Treat	Possible archiving of data for longer period and warning emails sent by SAFE before deletion.	

4	Unauthorised access	5	1	5	Accept	MFA is in place and deployed on all systems	
5	Lack of user training leading to misuse of resources or unauthorised access	4	3	12	Treat	Continue to provide appropriate training	