

CHPC Accounts Policy Document

Note: the CHPC reserves the right to amend and/or extend this document. Notice will be given to all users of any proposed amendments to allow for feedback. The CHPC reserves the right to make final determination on the amendment. Any user who is unable to accept the amendments may voluntarily withdraw as a user.

Version: 2.3

Updated on: 13 December 2018.

1 Introduction

1.1 Scope of the document

User policies (UP) at the CHPC are divided into the following areas:

1. User account policies, including accountability, restrictions and responsibilities;
2. Resource limits;
3. User job scheduling policies, which covers the scheduling and priority of users' job on CHPC's computing systems

1.2 Summary

This user policy document (UPD) covers the following items:

1. duration of accounts: every account expires after 365 days and has to be renewed for the new annual cycle, on a per account timeline
2. renewal is conditional on progress (measured in publications, student graduations, dissertations, etc.) and appropriate usage (measured in efficiency of resources used)
3. principal investigators (PI) lead a Research Programme ("project")
4. resources are allocated to Research Programmes and not users
5. a user must be a member of a Research Programme before a user account is created or renewed
6. access to the priority queue is based on contract (payment)
7. the user has responsibilities in terms of the security, privacy and proper use of their account
8. misuse, use for work other than in the Research Programme, sharing one account, and other abuses will result in suspension of account
9. storage (quotas) and CPU limits (allocations)

10. users must have the necessary competence in using their application codes
11. attendance of CHPC training courses or groups are responsible for own training
12. supervisors are the PI for student and postdoctoral accounts; postdocs cannot be PIs
13. users are to provide essential information for the CHPC user database (UDB): pertinent personal information, contact details, status (staff/student), publication history, workshop attendance, and other relevant information.
14. student users are expected to present a poster at the annual CHPC conference (subject to funding and space)

2 User Accounts

2.1 Research Programmes and Principal Investigators

Computational and storage resources at the CHPC are only provided on a per Research Programme basis. Each Research Programme is lead and managed by a principal investigator (PI). The PI is primarily responsible for the use of the allocations provided by the CHPC and is entirely responsible for the users within his/her Research Programme.

1. Full time permanent staff members of an academic research institution in the following countries may apply on the CHPC user database (UDB) website to be a PI: Republics of South Africa, Namibia, Botswana, Ghana, Kenya, Madagascar, Mauritius, Mozambique and Zambia. An academic researcher from any other country who wishes to use the CHPC may apply to be a user of a Research Programme lead by a qualifying PI and must be a bona fide collaborator of the PI as indicated by a memorandum of agreement (MoA) submitted to the CHPC. It is the responsibility of the PI to add the foreign user to his/her research programme.
2. In the case of a post-graduate student or post-doctoral fellowship holder, his/her supervisor must be a PI at the CHPC, even in the case where the student or postdoc is the sole member of the Research Programme.
3. The PI for a non-academic or commercial Research Programme will be identified in the memorandum of agreement (MoA) or contract between the CHPC (CSIR) and the non-academic or commercial entity, respectively.
4. A qualifying PI must apply for a research programme on the UDB. Once all requested details are provided online the Research Programme proposal will be evaluated and this programme confirmed.
5. The PI will then identify and allocate users to the Research Programme—only candidate users fully registered on the UDB may be selected. Users within a Research Programme will draw upon the Research Programme's allocation for all their job-runs. Once an allocation is exhausted the PI will be required to provide feedback on research outputs achieved from allocation usage. This feedback is used by the CHPC to determine the setting

of the next biannual allocation.

6. Each Research Programme will receive an allocation of computational resources (measured in cpu-hours).
 - (a) An initial allocation of 100 000 cpu-hours is granted to new Research Programmes for the purpose of evaluation; this expires when replaced by a regular allocation or after 6 months.
 - (b) During the evaluation period, the CHPC may increase the initial allocation.
 - (c) At the end of the evaluation period, the CHPC will provide the Research Programme with a regular biannual allocation for the Programme.
 - (d) A regular allocation is granted for six months and is reset, as part of the evaluation process, every six months[†] thereafter. Unused allocations do not carry over into the next period.
 - (e) Allocations that are used up early do not automatically renew until the start of the following 6-month period.
 - (f) Renewed allocations are based on the average of past usage.
 - (g) The maximum 6-monthly allocation for any Research Programme is 6 000 000 cpu-hours.
 - (h) PIs may apply for an increase in their allocation based on the number of users in the Research Programme, the scaling of their application code, and the work to be done. All allocation amounts are granted at the sole discretion of the CHPC.
7. Each Research Programme will receive a storage quota (measured in GB = 10⁹ bytes) on the fast Lustre scratch storage.
 - (a) An initial allocation of 1 TB = 1 000 GB is granted to new Research Programmes for the purpose of evaluation; this expires when replaced by a regular quota.
 - (b) During the evaluation period, the CHPC may increase the initial quota.
 - (c) At the end of the evaluation period, the CHPC will provide the Research Programme, *on application by the PI*, with a regular storage quota valid until the end of the Programme.
 - (d) The storage quota is divided amongst and shared by all the users in a Research Programme. It is therefore very important that an adequate amount is applied for.
 - (e) The standard storage quota for a Research Programme that has access to the *large* queue is 10 TB = 10 000 GB.
 - (f) PIs may apply for an increase in their storage based on the number of users in the Research Programme, the size of input/output data of their application code, and the

[†] Allocations are granted for six months and are reset at midnight on the day of the month when initiated. When a calendar month is short, the reset occurs on the last day of that month. For example, if an allocation was initiated on 30 August, it would reset on 28 February (or 29 February in the case of a leap year).

work to be done. All storage quotas are granted at the sole discretion of the CHPC. The standard storage quota *maximum* is 100 TB. Quotas are reset to the minimum size after 365 days unless an application is made as per (c) above.

- (g) NOTE: the storage quota represents scratch space intended for transitional data that exists for the life time of a processing job. Files older than 90 days will be automatically deleted. It is the responsibility of a PI to ensure that each user in their Research Programme has adequate offsite (outside the CHPC) storage to save the results of their CHPC runs and moves their data to this offsite storage timeously before expiry. No back-up of data on the Lustre scratch storage is made by the CHPC.
8. The PI is responsible for ensuring that each user is sufficiently competent at using the applications codes to be run at the CHPC. Where available, the user should attend CHPC training, otherwise the research group should provide user training.
 9. Each PI will be required to provide the CHPC with regular feedback (every six months) on the progress of the Research Programme and on all outputs (journal articles, conference proceedings or slides, student dissertations, student graduations, technical reports, etc.) produced as a result of the computational work performed on the CHPC's systems.
 10. All users are required to give full acknowledgement to the CHPC in all public presentations of their work (see below) and the PI is required to confirm this for all users within their Research Programme.
 11. In the case that a PI would like to use the CHPC systems for teaching or workshops, the PI must apply for a separate programme for this activity. The CHPC requires PIs to keep separate their research and teaching programmes.

2.2 User Accountability

All users are expected to use their CHPC account, storage, network and computational resources in a responsible, ethical and professional manner.

1. Every user must apply online to the CHPC User Database and provide all requested details.
2. The user must join an existing CHPC Research Programme and apply to the Research Programme's PI. The PI must then add that user as a member to the Research Programme on the CHPC User Database web site. Once that is verified by the CHPC, the user's account will be created.
3. Each user's email address will be subscribed to the `chpc-users` mailing list. Users must not unsubscribe from this list for the duration of their account as this list is used for important system announcements.
4. A user account is active for 365 days. At the end of that period, the account will expire, and must be renewed by a PI adding the user to a research programme. Any account not renewed by the PI within 30 days of expiring will be deactivated. The account and all data associated with it will be permanently deleted 30 days after deactivation.

5. Users may not allow other people to use their account login.
6. Users must apply appropriate security best practices to secure their account, including keeping their password secret. Passwords must be hard to guess and the CHPC imposes minimum limits on length and complexity.
7. Any files, results or data in a user account will be considered to belong to the Research Programme and the PI. It is the responsibility of the PI to contact the CHPC within 30 days of an account's expiry to retrieve the data associated with the expired account.
8. Copyright and/or other applicable intellectual property laws will be followed by the CHPC. In the case of student users, it is required that the PI ensure that all rules and regulations of the student's university with the regards to copyrights, patents, et al., are followed.
9. In particular it is expected that all proprietary and open source code licenses are fully complied with. Users who write their own program code must ensure the appropriate licensing is respected and all licensing conditions followed.
10. Users in an academic Research Programme may not use their CHPC account for any commercial or paying work, nor undertake any work for a third party who is not their PI or member of their Research Programme.
11. Users of a commercial or non-academic Research Programme are required to follow the parameters of the contract or memorandum-of-agreement with the CHPC (CSIR).
12. Users must be competent in the use of the application codes they will run on the CHPC systems, and must be familiar with the Linux command line. Any user who wastes CHPC resources through incompetence may have their account suspended until they demonstrate to the satisfaction of the CHPC that they have (through training or supervised self-study) gained the needed skills and competence.
13. The CHPC follows security best practices on all systems and networks, including active security measure that may lead to temporary locking of accounts.
- 14. Any abuse of the CHPC account or facilities, or violation of these policies (including the associated scheduler, storage and network policies) will result in termination of the user account.**

2.3 Obligations of Academic PIs and Users

PIs and members of academic Research Programmes are not charged for their use of CHPC resources as these are funded by the DST and provided to support academic research and teaching in South Africa. However, there is a *quid pro quo* obligation on academic users of this national facility to provide the CHPC with:

1. Full acknowledgement of the CHPC in all public presentations and publications of their work:

The authors acknowledge the Centre for High Performance Computing (CHPC), South Africa, for providing computational resources to this research project.

2. The CHPC requires regular feedback from a PI to account for their allocation usage. This feedback will be entered into the User DB and linked to the Research Programme. This feedback is required whenever an allocation is 75% consumed, or every six months (whichever occurs first).
3. A short synopsis (in press release format) of the PI's users' research, written at a technical level accessible for an educated lay-person, with suitable illustrations of major results along with permission for the CHPC to use these in our newsletters, website, pamphlets, slides, and any other outlet or publication where we publicise and promote the work done at the CHPC.
4. Student users are required to submit an abstract for a poster on their work to the CHPC's annual National Meeting and Conference in December of each year.
5. Non-student users are required to submit an abstract for a short talk to the above conference each year.

3 Systems Policies

3.1 File Systems and Storage

The primary file systems on the CHPC cluster are:

Mount point*	File System	Size*	Quota	Backup	Access
/home	NFS	40 TB	20 GB	Limited†	Yes
/mnt/lustre/users	Lustre	3 PB	application	None	Yes
/apps	NFS	20 TB	none	None	On request
/mnt/lustre/groups	Lustre	1 PB	application	None	On request

* The mount point and size are subject to change: use `df` to determine the actual mount point and sizes for the above file systems.

† Backups of /home are limited to a whole file system snapshot that can be restored in the event of a system or hardware failure. The CHPC does not provide archiving of files (and we cannot restore individual files).

/home/ contains users' home directories.

/mnt/lustre/users/ is used as work space for running jobs. Files that have not been used for 90 days will be deleted.

Note: the working directory for running jobs must not be on `home/`. All files written during a running job must be on the Lustre file system.

/apps/ is used for *shared* installations of software tools and applications. Code installed under here may be supported to some extent by the CHPC.

/apps/user/ is provided to Research Programmes to store shared libraries and compiled codes. The

CHPC will not provide technical support for these codes. This space is provided to users as an alternative to using /home which is intended for job files and input data. The CHPC may contact users that are storing libraries or application codes in their home directories and request that they be moved to apps/user/ or be purged from home/. Application for use of apps/user/ is to be made by the PI or programme's technical expert (TE) through the CHPC Helpdesk.

/mnt/lustre/groups/ is used for *shared* data storage of files. This is for long term storage and is not subject to automatic deletion. Access is by application only, and is limited by quotas.

File System Quotas

The quota on /home is 15GB for each user. This is a soft quota that may be exceeded temporarily for a grace period of up to 20 days. When a user exceeds the 15GB limit for more than 20 days, write access to that user's home directory is disabled. The hard quota on /home is 20GB: write access is disabled immediately when a user reaches 20GB.

The storage limit for each user directory in /mnt/lustre/users/ is derived from the storage quota for that user's research programme.

Files that are unused after 90 days on /mnt/lustre/users/ will be automatically deleted. Users will be warned 14 days in advance.

Access to /mnt/lustre/groups/ is by application only, and that application must include a detailed motivated request for the group storage quota. Access to lustre/groups/ is granted for 365 days and may be renewed on application at the end of that period.

3.2 Application Codes

Support Levels

0. **Unsupported:** Code installed in /apps/user/<Research_Programme>. This code is installed by the TE designated by the Research Programme PI. The TE is solely responsible for the installation, updating, use and documentation of this code. The TE will be given an account on the CHPC wiki to provide full documentation on the use of this code at the CHPC. The TE is required to follow CHPC standards and practices for installation of software to work with the CHPC scheduler and to assist the CHPC technical staff in monitoring use of the software.
1. **Limited support:** Code installed in central shared directory (/apps/chpc/). Minimal documentation on wiki describing how to run code on CHPC. Code only updated on user request, or when major version released (1.X → 2.0). Only basic user questions answered. CHPC recognises that more than one user requires this code, but it is outside of CHPC staff expertise. At the discretion of the CHPC support may be provided by a 3rd party.
2. **Standard support:** Documentation on wiki describing how to run code on CHPC, along with example scripts. Code updated when new stable version released (X.2 → X.4). User questions on running code answered. Technical queries supported. Introduction to code included in induction course. CHPC recognises this is an important code for the SIG or the

community.

Proprietary and commercial provided codes

Commercial and proprietary codes must be install under `/apps/chpc/` or `/apps/user/projectname/`, and have their own GID to restrict access to users with permission to use that license. For example, Ansys (Fluent) has GID of `ansys` and appears under `/apps/chpc/compmech/`.

3.3 Resource Limits

By default, users on the CHPC systems are limited to:

1. Maximum 240 cores total simultaneous running and queued jobs.
2. Maximum of 10 running and 10 queued jobs at any given time.
3. Standard priority for jobs.
4. Maximum of 20GB storage in `/home/`.
5. Limited working scratch storage in `/mnt/lustre/users/` and maximum file lifetime of 90 days.
6. No access to `/mnt/lustre/groups/` or `/apps/user/`.

Application for Additional Resources

An application for the use of additional resources is required for any modification, departure or exemption from the standard CHPC user policies and limits. This application can be made to request:

1. access to the large queue (i.e., use more than 240 cores, up to 2400 cores, per job)
2. access to the bigmem queue (the large memory nodes)
3. access to the `gpu_n` queues (nodes with GPUs)
4. larger storage quota
5. extended file lifetime in `/mnt/lustre/users/`
6. access to `/mnt/lustre/groups/`
7. access to install codes in `/apps/user/`
8. or any non-standard request.

3.4 Standards

Users are required to follow the conventions of the Linux Standards Base (LSB) and other open source software.

In your home directory you should install software in `local/` or `opt/` subdirectories.

`$HOME/local/`

is used to store software according to the LSB for software that shares the standard common directory tree of:

`bin/` `doc/` `etc/` `include/` `lib/` `man/` `share/`

and so on.

`$HOME/opt/`

is used to store applications that require their own directory tree. Subdirectories under `opt/` will be the name of the software and version, for example

`openfoam-2.3.1/`

Other directories to use include:

`src/` for source code (unpack source code tarballs here)

`build/` for compiling code (where supported)

The LSB conventions are also required for codes installed in `/apps/user/`.

3.5 Bandwidth

Each Research Programme is subject to a network transfer limit of 1 GB per week (the “cap”), which applies to all users. Research programmes that require greater caps must apply with motivation. Bandwidth is limited at the CHPC to SANReN connected sites and specific international sites.

A whitelist of allowed external servers that can be accessed from the login node and compute nodes is maintained. If users of a Research Programme need to connect to an external server from the CHPC systems, then the PI or TE must apply to Helpdesk to add the hostname of the external server to the whitelist.

The CHPC is unable to provide dedicated bandwidth to Research Programmes. The current limit is 10 Gb/s for *all* SANReN traffic and 200 Mb/s for *all* international and national (non-SANReN) traffic. All network traffic is monitored and logged. Bandwidth to non-SANReN sites is restricted and capped. This applies to interactive `ssh` connections to the login nodes and file transfers via the `scp` node.

4 Scheduler Policy

The scheduler policy is applied on the cluster. The implementation is subject to change: we will monitor and evaluate the queueing system and adjust parameters where necessary.

1. There are **allocations**.
 - (a) Jobs can be queued if there is sufficient available allocation at the time of submission;
 - (b) There are no pending allocations;

- (c) When a job is about to execute and the remaining allocation is insufficient, then the job is requeued until the allocation is increased.
- (d) Allocations are reduced for each standard job using the formula:

$$\text{number_of_cores} \times \text{wallclock_time}$$
- (e) Allocations are reduced for each **GPU** job using the formula:

$$40 \times \text{number_of_GPUs} \times \text{wallclock_time}$$
2. Fairshare is applied to each user individually: the Fairshare measurement unit is **cpu·hours** (this may be changed to **node·hours** or **hours** to encourage more efficient use).
 3. Priority of a job is determined by the queue, and the user's individual Fairshare score.
 4. Access to the priority express queue is available for contract (paying) projects upon request.
 5. Reservations are exceptional and rare: reservations are restricted to crucial and time-critical projects. For example, a reservation may be allocated to a workshop where the participants will need access to the cluster only during a limited and specific calendar period. Application for a reservation is to be made via the helpdesk.
 6. Users are required to ensure that their jobs exit with a correct job-exit-status.
 7. Allocation consumption, efficiency, and other usage statistics will be used as part of the evaluation of a research programme and will determine the allocations granted for the new allocation 6-month period.
 8. PIs are required to complete the online evaluation form for their research programme every 6 months or when 75% of their allocation is consumed, whichever occurs first.
 9. Access to the GPU queues is limited and by application only; and for codes that perform exceptionally well on GPUs compared to CPUs.

Queues

Queue name	[Min~]Max. size of job	Max. duration	Max. Running Jobs	Priority	Comment
test	1-24 cores [shared]	3 hours*	1	increased	For testing and development (compile, profile or debug) only.
serial	1-24 cores [shared]	48 hours	20	reduced	For very small core count jobs; shared nodes.
seriallong	1-12 cores [shared]	144 hours	10	reduced	Very long run time; shared nodes.
smp	1 node: 24 cores	96 hours	20	reduced	For single node SMP jobs.
normal	2-10 nodes: 25-240 cores	48 hours	20	standard	Standard and default queue.
accelrys	24-96 cores	96 hours	–	standard	Material Studio users only
large	11-100 nodes: 241-2400 cores	96 hours	10	increased	By application only.
express	2-10 nodes: 25-240 cores	96 hours	2	increased	Contract projects.

Queue name	[Min~]Max. size of job	Max. duration	Max. Running Jobs	Priority	Comment
bigmem	5 fat nodes: 1~280 cores [shared]	48 hours	4	standard	For big memory nodes.
vistest	1 VNC node [shared]	3 hours*	1	increased	For testing visualisation.
vis	1 VNC node [shared]	12 hours*	1	standard	Visualisation node.
gpu_1	1 GPU node: 1 device only [shared]	12 hours†	2	increased	GPU codes.
gpu_2	1 GPU node: 2 devices only [may be shared]	12 hours†	2	increased	GPU codes.
gpu_4	1 GPU node: 4 devices only	12 hours†	2	standard	GPU nodes.

*The `test` and `vis` queues are used for interactive jobs.

†The `gpu_n` nodes provide 1,2 or 4 GPU devices per node.

Regulations

The following restrictions apply:

1. All Research Programmes have allocations and all jobs draw from this. A user belonging to a Research Programme which has exhausted its allocation can not submit jobs.
2. Most jobs are only limited by allocations and scheduled by PBS Pro using the backfill scheduler and Fairshare score according to resources requested and available. Users are to use *job arrays* for multiple small jobs.
3. There is a global limit on *simultaneous queued and running jobs* for job size (cores) and total number of jobs per user. Both maximum total cores (240) and maximum jobs (20) will be enforced for queued jobs. Jobs in the `large` queue have a limit of 2400 cores.
4. Only *contract* (commercial) jobs have access to the express queue.
5. Interactive jobs will be monitored for abuse. A 30 min idle timeout will also be imposed. Their intended purpose is for pre-processing, post-processing, debugging and compiling.
6. Long and large file transfers must use the dedicated `scp` login node for file transfers (`scp/sftp`, `rsync` and `wget/curl`).
7. The `accelrys` queue is for Material Studio (MS) jobs only, and is limited to 2400 cores only. When these cores are all occupied with running jobs, subsequent MS jobs will be queued.
8. The `gpu_n` nodes are all shared: jobs have access to 1, 2 or 4 GPU devices on one node.

Monitoring

In order to accurately monitor usage we require that each user must be a member of a Research Programme in order to run jobs on the CHPC clusters.

All activities on the CHPC systems are monitored and logged.

5 Special Circumstances

At the Director's discretion, in case of emergency, special circumstances, or where deemed

necessary to fulfil the mandate of the CHPC, exceptions may be made to the above conditions.