

HPC accounting

- System Billing Unit (SBU)
- User Usage
- Accounting and account management
 - Computing resource allocation rules
 - Special Projects
 - Daily Accounting Controls
 - Period (4-weekly) account reporting
 - Registering a new account
- Administration of project accounts in the Member States
 - Running the account administration program

System Billing Unit (SBU)

To ensure that computing resources are distributed equitably and to discourage irresponsible use users' jobs are charged for the resources they have used against a project account. The project accounts have computer usage units allocated at the beginning of each accounting year, which, following ISO 8601, starts on a Monday in the period 29 December - 4 January.

The accounting on the Cray HPCF is currently based on the total elapsed time for which cores are used. With that, the HPCF unit of account, the **System Billing Unit (SBU)**, for both serial and parallel jobs is defined as follows.

$$SBU = P * (ELAPSED_TIME)$$

where

$P = 50365140 / (3617 * 36 * 86400)$	is a fixed proportionality factor, adjusted so that the system will deliver 50365140 units per day on a "typical" ECMWF job mix on the HPC system (3617 * 36 is the number of cores of the Cray XC 40 system and 86400 the number of seconds in a day)
ELAPSED_TIME	is the elapsed time (real time in seconds) recorded for a job, multiplied by the number of cores used by the job

ECMWF reserves the right to alter the above formula in the light of experience.

Examples:

1. A serial job using 1 hour (3600 seconds) of elapsed time will be charged about 16.11 SBU.
2. A parallel job requesting 2 nodes, running for 3 hours elapsed time will be charged 3481.14 SBU.

More information on the HPC system, which might be helpful to compare it to other platforms and to estimate computational cost in terms of SBUs, is available at <https://www.ecmwf.int/en/computing/our-facilities/supercomputer>.

User Usage

The following two commands are available on ecgate to help determine your day to day usage.

account [-l uid]

This lists the accounts under which a user-id is registered. **If more than one account is listed, then by default, job usage will be debited to the first account listed, the primary account.** If you wish to debit the usage of a job to an alternative account (out of your list of assigned accounts) you may do so by using the following PBS statement

```
#PBS -l EC_billing_account=<other_account>
```

acct_status

The acct_status command allows users to check the usage of their project accounts. The usage is:

```
acct_status [-p week/period] [-r] [-u] [-h hostname] [-y year] [project(s)]
```

where

- p : gives the usage for the current week or 4-week period, e.g. -p week
- r : gives information for sub-projects
- u : gives information for each user of the project account

-h : gives the usage for the specified hostname (e.g. cca). If omitted the cumulative total for all supercomputers is reported

-y : accounting year YYYY (default: current ISO year)

project(s) : the project(s) or node(s) to be reported on. The default is the project account of the logged in identifier.

The reply shows the usage in terms of SBUs (System Billing Units), it is largely self explanatory. You can only check the usage of those accounts which you are registered to use. To see the list of those accounts, use the command "account".

Note the following:

1. "Total use" is the cumulative usage since the beginning of the year.
2. The "current" week begins on a Monday
3. The reply for the current week or period (using -p) shows:
 - the SBUs used for the current week/period (SBUs:)
 - the total used to date this year (total:)
 - the current allocation (alloc:)
 - the percentage use made to date of the allocation.
4. The node report may list "Responsible User(s) and/or Privileged User(s)" - these are the persons who are permitted to see information at this node level and all levels below it. A "Privileged User" is a person who can change allocations for nodes and projects below this level. The "Privileged User" for a Member State could be its Computing Representative.
5. The usage figures are usually updated daily (before 0700) and include all jobs that finished before 00:00 on the current day. Operational problems may cause this update to be delayed for any given day.

Some examples:

acct_status abc xyz	will report on the two projects or nodes abc, xyz
acct_status -u	will report on the individual usage of the account by the logged in identifier
acct_status -p week -r ms	will report the current week's usage for the node ms and sub-projects under it. Only a "Responsible User" or the "Privileged User" for a node can obtain this report.

ej

The ej (end of job) command gives some basic statistics about your job including the usage of SBUs. It has been included in the system epilogues but can also be used interactively at any time to check your running job. Login to the HPC and type:

```
ej job-ID
```

To find the job ID of your job use the qscan command. The ej command without argument can be placed at various points in an HPC batch job to show job accounting information up to that point within a job. For example, ej will provide a summary report from the beginning of the job up to that ej command. To obtain such a report just add the line

```
ej
```

to your job script. It is also possible to obtain a report on a specific sequence of commands by getting an ej report before executing the sequence, another one when the sequence is finished, and computing the difference of the relevant terms. This information can then be used to identify blocks of commands within a job which take a lot of SBUs to execute.

Please note that

- The information given by ej is updated every 10 seconds at the earliest
- For parallel jobs the reported CPU time is not correct

For more information on ej use `ej -h` on the HPC system.

Accounting and account management

Every user of the ECMWF computer system is given a separate **user identifier**. Because of both security requirements and the way ECFS files are stored and identified, a given identifier must not be shared by two or more users. The first 2 characters of the user identifier are common to a given section or group of people, (e.g. all User Support identifiers begin US). The following characters are then assigned by the ECMWF manager responsible for the user, by the Member State Computing Representative, or by the Principal Investigator of a Special Project, as appropriate.

Groups of users are assigned to a **computer project** (or account). Each computer project has a unique account name. Reports listing usage against each computer project are generated at regular intervals (daily, weekly, 4-weekly, annually), the 4-weekly ones are available, via e-mail, upon request to User Support.

Computing resource allocation rules

Member States are allocated computing resources for their own purposes according to the [Council guidelines on distribution of computer resources](#). Member State Computing Representatives may distribute those resources among projects registered at ECMWF as they feel appropriate. Should a redistribution of resources between the projects of a Member State be required during the course of a year, this may be carried out by the Member State Computing Representative or another designated "privileged user" (see section [Administration of project accounts in the Member States](#) below).

Special Projects

Special Projects are experiments or investigations of a scientific or technical nature undertaken by one or more Member States, likely to be of interest to the general scientific or technical community. In the spring of each year, Directors of national meteorological services are invited to submit requests for Special Project resources to be available for use in the following year. Requests are considered by the Council at its winter session, in conjunction with recommendations from the Scientific Advisory Committee and Technical Advisory Committee. Since Special Project accounts are separate from Member State resources, their account administration is performed by ECMWF. For more details see <http://www.ecmwf.int/en/research/special-projects>.

Daily Accounting Controls

Information is extracted daily from HPCF report files to produce a cumulative total of resources used per user and per account. Use of each account is checked daily and once usage reaches 90% of its resources available, a warning message is sent out, to enable users to arrange for the transfer of additional resources, if necessary. Once usage reaches 100% no further use of that account should be made.

Period (4-weekly) account reporting

Member State Computing Representatives or Special Project Principal Investigators, receive a regular summary of accounting information relating to all their Member State's projects or their Special Project(s) respectively at the end of each four week accounting period. The format of the report is shown below. If a user does not use the system in any period, no entry is made.

Format of HPC account report

```
1. Member State:
2. Combined 4-week report for period ending (MMDD) : 0128
3. Node:      ms          SBUs: 31.9      Total: 51.9      Limit: 20000.0 (0%)
4. Acct:      msacct1     SBUs: 31.9      Total: 31.9      Limit: 10000.0 (0%)
5.           User: ms1    SBUs: 21.1
5.           User: ms3    SBUs: 10.2
5.           User: msd    SBUs: 0.6
4. Acct:      msacct2     SBUs: 0.0      Total: 20.0      Limit: 10000.0 (0%)
```

where

```
Line 1      Name of Member State

Line 2      Period concerned, where the period is identified by its last day
            (in this example January 28).

Line 3      Overall node for the Member State "ms" showing
            SBUs:      total used this period
            Total:     cumulative total for this year
            Limit:     overall allocation, with percentage used to date

Line 4      The summary for a given account (msacct1 and msacct2 in this
            example). The entries have the same meaning as in line 3.

Lines 5     Usage per user (ms1, ms3, msd in this example) during the period
            in question.

            User:      user identifier
            SBUs:      total used this period
```

Registering a new account

New accounts are registered by ECMWF only on the authority either of an ECMWF Department Head (for ECMWF staff accounts) or of a Member State Computing Representative (for Member State accounts) or the Principal Investigator (for Special Projects). For further information contact User Support.

Administration of project accounts in the Member States

Member States may nominate "privileged users" to perform limited administration of Member State accounts using the *acct_admin* program, see below. By default, the Centre will grant this status to all Member State Computing Representatives with a registered user identifier, but Member State Computing Representatives may propose an alternative person and/or nominate a backup if they wish.

A "privileged user" may be authorised to modify all or a reduced set of Member States' accounts. For instance, Member States are free to allocate control to several "privileged" users, each having control of a particular sub-node. It should be noted, however, that the nomination of privileged users should be kept to a minimum.

For security reasons, the establishment of new project accounts and the addition of new users remains under the control of ECMWF.

Running the account administration program

The ECMWF accounting administration program is started (on *ecgate*) by the command:

```
acct_admin
```

The following list of commands is available on entering *acct_admin*:

```
read allows only testing and checking  
edit to modify account allocations (privileged users only)  
exit to terminate having made permanent changes  
quit to terminate without any change
```

The modes "edit" or "read" may be specified at command invocation, e.g.

```
acct_admin edit
```

or after entering the *acct_admin* program.

Within "edit" or "read" mode the prompt is

```
Type command:
```

A list of subcommands may be displayed by the subcommand "help" and is given below.

Subcommands

```
exit  
quit  
help [commands]  
? [commands] (help)  
comment  
add-responsible-user account user  
remove-responsible user account user  
set-sbu-allocation account allocation  
add-sbu-allocation account allocation  
remove-sbu-allocation account allocation  
transfer-sbu-allocation from-account to-account allocation  
list-account account [and-sub-accounts] [and-users]  
list-all-account account [and-users]  
list-users account [all]  
list-tree account
```

Notes regarding these subcommands:

- Account can be a node or a sub-node name.
- Allocation is replaced by an actual value (likewise allocation-increment, allocation-decrement and allocation-change).
- Items within [] are optional qualifiers.
- The value to be inserted for [and-sub-accounts] & [and-users] are yes or no. The value yes means that that option will be turned on (no is default).
- Additional commands besides these may be listed but these are only available to the ECMWF Account administrator.

As the invocation of the `edit` subcommand will lock the system, please **make sure to leave the `acct_admin` command with `exit` or `quit`** as otherwise it will not be available to others.

Listing allocations for Member State projects/nodes

To list all users on all projects for Member State Sweden (country code SE):

```
list-account se yes yes
```

To list all users on Special Project spmsabcd

```
list-users spmsabcd all
```

Changing allocations for Member State projects

To change the allocation of a project, use the 'edit' menu and the add/remove/transfer command

```
acct_admin edit
```

- a. To decrease the allocation of a project/node use:

```
remove-sbu-alloc project amount
```

- If sub-nodes are used to sub-divide the allocation of a Member State, the allocation for the sub-node may also be decreased.
b. To increase the allocation of a project/node use:

```
add-sbu-alloc project amount
```

- It must be ensured that the node above the project/node has enough resources 'not allocated'. If not, some resources will have to be freed (as under A) before the increase can be done.
c. To transfer an allocation from one project/node to another use:

```
transfer-sbu-alloc from-account to-account amount
```

This command will transfer resources from one project/node and its ancestor nodes to the destination project/node and its ancestors.

Example 1

Country 'mo' uses the following nodes and accounts:

```
Node:      mo      Allocation: 10000, Not allocated: 3000
Acct:      moabcd  Allocation: 5000, Total use: 2000
Acct:      moxyz   Allocation: 2000, Total use: 2239
```

The account 'moxyz' has used more than its allocation.

To add an extra 2000 units:

```
acct_admin edit <<EOF
add-sbu-alloc moxyz 2000
exit
EOF
```

Supposing this project runs over its allocation again, the status becomes:

```
Node:      mo      Allocation: 10000, Not allocated: 1000
Acct:      moabcd  Allocation: 5000, Total use: 2000
Acct:      moxyz   Allocation: 4000, Total use: 4239
```

Now, the 2000 units must be transferred from the other account, as only 1000 units are unallocated.

```
acct_admin edit <<EOF
remove-sbu-alloc moabcd 2000
add-sbu-alloc moxyz 2000
EOF
```

The following command would also have accomplished the last change more simply:

```
transfer-sbu-alloc moabcd moxyz 2000
```

Check status using:

```
list-acc mo yes
exit
```

Note that resources must be freed first, or else the error message "not enough unallocated resources to change allocation" is given.

Example 2

Country 'su' has 2 nodes subalt and suokr each with 2 projects:

```
Node:      su      Allocation: 80000, Not allocated: 0
Node:      subalt  Allocation: 10000, Not allocated: 0
Acct:      subaest  Allocation: 5000, total use: 2000
Acct:      subalit  Allocation: 5000, total use: 3000
Node:      suokr   Allocation: 70000, Not allocated: 0
Acct:      suuksib  Allocation: 20000, total use: 3000
Acct:      suukrus  Allocation: 50000, total use: 48000
```

The project 'suukrus' is running close to its limit. To increase the allocation of this project, some resources have to be freed first, as the node 'suokr' has no free allocation.

```
acct_admin edit <<EOF
```

To free allocation from the project 'subaest':

```
remove-sbu-alloc subaest 2000
```

To give the allocation to 'suukrus':

```
add-sbu-alloc suokr 2000
add-sbu-alloc suukrus 2000
```

Check status using:

```
list-acc su yes
exit
EOF
```

To first run it as a test, replace the 'exit' command with 'quit'.

N.B. An easier way of doing the series of removes and adds is to just use:

transfer-sbu-alloc subaest suukrus 2000

The transfers at the node levels, above the projects, are then taken care of automatically.