



Introductory User Guide

CSAR Introductory User Guide

Introduction

This material is designed to provide new users with all the information they need to access and use the SGI systems provided by CSAR. Please note that while we endeavour to keep this information up to date users should refer to the CSAR web site (www.csar.cfs.ac.uk) for the latest news, bulletins and systems information.

Project Investigators are provided with an additional guide to managing their projects (www.csar.cfs.ac.uk/project_management/pi_guide.shtml).

Contents

1. Self Registration	Page 4
<ul style="list-style-type: none">• Getting a userid• General notes on the registration system	
2. Accessing the Machines	Page 5
<ul style="list-style-type: none">• How to log in• SSH Clients• X access	
3. Login Shell	Page 6
<ul style="list-style-type: none">• Available login shells• Changing your shells	
4. Files and Filestore	Page 8
<ul style="list-style-type: none">• Home directories• Local temporary filestores (/tmp)• The global temporary filestore• Hold/long term storage• Data transfer to and from CSAR machines• Other information	
5. Interactive Service and Editing	Page 9
<ul style="list-style-type: none">• Interactive machines• Available editors• Version control• Makefiles	
6. Software and Modules	Page 10
<ul style="list-style-type: none">• Introduction• Modules	
7. Batch Jobs	Page 10
<ul style="list-style-type: none">• An introduction to bsub• Submission script example• Command line example• Further Information	
8. Porting and Parallelising Codes	Page 13
<ul style="list-style-type: none">• Choice of machines• Types of parallelism• Porting codes• Further help and information	
9. Documentation	Page 15
<ul style="list-style-type: none">• UNIX man pages• Cfs systems man pages• CSAR monthly bulletins	
10. Monitoring Your Resources	Page 17
<ul style="list-style-type: none">• An introduction to resources• Checking your user allocations	



- Viewing usage for a specified time period
- Automated warnings and stopping mechanisms

11. Keeping up to date

Page 19

- Status Information
- Scheduled Maintenance
- System Updates
- News and the CSAR Mailing List
- Support and Feedback



1. Self Registration

Getting a Userid

When your project has been approved and set up, your project's Principal Investigator (PI) will be given an access username and password to distribute to users to enable them to self register on the web.

To self-register, go to www.csar.cfs.ac.uk/project_management/register.shtml and follow the instructions there in order to register a username for each machine to which you require access. Your PI will receive an email informing them of your registration and asking them to authorise it. Once they have done this you will receive an email confirming that your registration has been completed.

General notes on the Registration System

The registration database holds all details of users and projects, and is the primary project management tool used by PIs, users and CSAR staff. Project PIs are provided with a comprehensive guide (www.csar.cfs.ac.uk/project_management/pi_guide.shtml) detailing the various tasks required in order to set up and manage their project. Once registered users can log in to the database from the registration page and amend their personal details.

2. Accessing the machines

How to log in

CSAR has two interactive machines - wren and newton. We recommend that you use ssh to log into these machines as follows:

```
ssh wren.cfs.ac.uk
```

```
ssh newton.cfs.ac.uk
```

You will be prompted for a username and password.

Interactive access to fermat and green is not permitted. These are batch only systems, all file editing, job submission etc for these machines should be done from wren.

SSH Clients

If you are using a UNIX workstation you can obtain the source code and README file from <ftp://ftp.ssh.com/pub/ssh/>. Linux distributions generally come with ssh and will either be automatically installed or available via your package management facility.

Windows users can use either putty which is obtained from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html> or SSH Secure Client which is available to academic users free of charge from <http://www.ssh.org/support/downloads/secureshellwks/non-commercial.html>.

For further information on ssh see:

- `man ssh`
- <http://www.ssh.org/>
- <http://www.openssh.org/>

X access

If you are using a UNIX workstation, X terminal or PC running X server software (e.g. Exceed), you can access wren and newton and use various X facilities. From an X session you can login as described above. X access should automatically have been enabled. The parameter, DISPLAY, is set correctly when you login unless you override it in a .profile, .cshrc or .login file. You may need to allow access to your display using the `xhost` command

```
xhost wren.cfs.ac.uk
```

```
xhost newton.cfs.ac.uk
```

in a terminal window on the workstation. This command tells the workstation to allow other systems to open windows on your screen. This command is not normally necessary on X terminals or PCs.



3. Login shell

Available Login Shells

At login to a UNIX machine, your commands are interpreted by a "login shell". The login shells supported by CSAR are the standard Bourne shell (sh), Korn shell (ksh), the C shell (csh), and the extended (or "turbo") C shell (tcsh). In addition the GNU Bourne Again SHell (bash) is supported on newton, and available but unsupported on the Origins. zsh is available on newton but not supported.

When your account is set up you will be allocated one of these shells as your login shell. To change this to another supported login shell contact your PI who can change it using the CSAR Registration System.

When you log in various default configuration files are executed. These set up the default environment. After the default configuration has been set up, your personal environment is configured if the relevant file exists in your home file store:

Shell	Environment File
sh	.profile
csh	.cshrc and then .login
ksh	.profile
tcsh	.cshrc and then .login
bash	.bash_profile or .bash_login or or .bashrc or .profile

When your account was created you will have been given a standard version of the relevant file(s) for your login shell.

Different files may be executed when a shell is started that is not a login shell, and also when a shell exits. More information can be found in the manual page for the shell you are using.

Changing your Shell

To change shells temporarily enter the executable name of the new shell at your prompt. You will then be using this shell until you exit from it to the former shell.

To determine your current login shell enter `echo $SHELL` at the prompt. Alternatively, you can enter `grep {username} /etc/passwd` and the last entry is `{username's} login shell`.

To permanently change your login shell to a supported shell please contact your PI or the helpdesk.

If you wish to use an unsupported shell as your default you will need to edit the relevant login file (see above table) to firstly re-assign the SHELL environment variable and then execute it. Examples are given below:

For sh based shells:

```
OS=`uname -s`
if [ ${OS} = "IRIX64" ]
then
export SHELL=/share/apps/unsupported/bin/bash
exec /share/apps/unsupported/bin/bash $*
fi
```



For csh based shells:

```
set OS=`uname -s`
if ($OS == IRIX64) then
setenv SHELL /share/apps/unsupported/bin/bash
exec /share/apps/unsupported/bin/bash $*
endif
```

Modules (www.csar.cfs.ac.uk/software/modules.shtml) will be setup automatically for your default shell. However if you change shell often modules are not setup (for example bash on the origins), in this case you will need to reinitialise modules.

Initialisation files appear in the directory /opt/Modules/default/init on Newton and /sanopt/modules/modules/init/ on the origins. These must be sourced in order to initialised. This is achieved by using the source command for tcsh and csh, and using the . (dot) operator for bash, ksh and sh.

For example, to initialise modules for bash on the origins you would need to do

```
./sanopt/modules/modules/init/bash
```

Warning: You should always test changes such as those above by, for example, saving your changes and opening another new, different window (i.e. without closing the window where you are editing your login files) to see what happens. If you don't test then be aware that a minor mistake could mean that you get logged out immediately each time you try to login.

4. Files and Filestores

Home Directories

The home filestore is the most important of all file systems. This is where the system places you when you initially log in. It is located on the SAN and is where you will carry out most of your work. This area is backed up. For further information please see www.csar.cfs.ac.uk/user_information/san/san_home.shtml

Local Temporary Filesystems (/tmp)

As is traditional on UNIX machines you can use /tmp for temporary filespace. A secondary directory exists for the duration of jobs located in /tmp which is called \$TMPDIR. There is no backup performed on this area and files are purged regularly. Note that \$TMPDIR is cleared automatically when the job terminates. For more information on the usage of this filesystem please see www.csar.cfs.ac.uk/user_information/san/san_local_tmp.shtml

The Global Temporary Filesystem (/santmp)

Home directories are often insufficient for running batch jobs. The global temporary filesystem known as santmp provides a cheap resource for storing files that may be required over multiple batch jobs. \$SANTMP can be used as a synonym for your santmp directory. No back up of this area is performed. Please see www.csar.cfs.ac.uk/user_information/san/san_tmp.shtml for more information and suggested usage.

Hold/Long Term Storage

The hold directory is an offline tape storage device but data is served by a filesystem on the SAN. Most data will be offline, when requested data is migrated to the online disk storage and is held in a dual state. Data when placed in hold will migrate to the tape system at a later date. To use this area you require a suballocation of resource from your PI and they can increase this as necessary. There is no backup on these directories. The number of files you can store here is typically limited to 300 and the smallest file you should store here is 40MB.

The command to fetch data from the tape system is `dmget`. The status of files can be interrogated with `dm1s -l`. Users are expected to pre-fetch files required for batch jobs or analysis. Detailed advice on usage can be found at www.csar.cfs.ac.uk/user_information/san/san_files_hold.shtml

Data Transfer to and from CSAR Machines

There are a number of ways to transfer data to/from the machines:

- scp (preferred)
- sftp
- ftp (wren only)
- rcp
- tape

Please see www.csar.cfs.ac.uk/user_information/san/san_files_copy.shtml for more help.

Other information

- Other filesystem options www.csar.cfs.ac.uk/user_information/san/san_other.shtml



- I/O and SAN performance
www.csar.cfs.ac.uk/user_information/san/san_files_performance.shtml

5. Interactive Services and Editing

Interactive Machines

Interactive services are provided via wren and newton. You cannot log in to ferret and green which are designated as batch only. If you wish to submit a job to ferret or green you should do so from wren. Job submissions to newton can be done by logging in to newton.

Users should note that the default limit on the number of processes on wren is set to 8. This can be increased, up to a maximum of 24, using the following command on the command line.

```
jlimit processes 16
```

You can put this command in your login file, but you must ensure you set it for wren only otherwise it will limit the processes that your batch jobs can use. For example for sh based shells:

```
if [ ${HOST} == wren ]
then
jlimit processes 24
fi
```

For csh based shells:

```
if [ $HOST == wren ]
then
jlimit processes 24
endif
```

Available editors

The main editors on both wren and newton are vi, which is part of unix, and emacs which is provided by the GNU Free Software Foundation. If you use emacs you should remember to set your DISPLAY environment variable to point to your workstation or PC. On wren you must load the sgi-freeware module to access emacs. Xedit is supported on all the machines.

Although you can edit files directly on wren and newton (consuming negligible CPU cycles), you may find it easier to edit files locally. To do this you will need to scp files from wren or newton to your local machine, perform the edits, then scp the files back again.

Version Control

If you are working on a complex code with other people, you may wish to consider using some form of version control. This allows just one person to edit a file at a time, records all the changes, and asks for a comment before saving the file and giving back access to other users. On both wren and newton the public domain CVS is available. RCS is available on newton. Note that both of these packages are unsupported and on wren you will need the sgi-freeware module loaded to access CVS. Please consult the relevant manual pages for more information.

Makefiles

To save unnecessary recompilation and to aid program development you should consider using a make file. This will contain inter-file dependencies, and running `make` will result in only modified sources files being compiled, before linking to make a final executable. For more information consult a standard unix reference book and `man make`.



6. Software and Modules

Introduction

We endeavour to provide an environment suitable for the development of high quality scientific software. Our machines therefore provide highly tuned C(++) and Fortran compilers, along with equally efficient mathematical (BLAS, etc) and communication libraries (MPI, SHMEM, etc). We also provide a range of debugging and software analysis tools to aid code development and some visualisation packages for post-processing work.

Our machines also have a wide range of application software packages installed, applicable to such fields as chemistry, environmental modelling, engineering and bioinformatics. For users interested in accessing our machines "over the grid", a range of grid middleware is on the systems and for metacomputing, suitable implementations of MPI are available.

More detailed information on the software we provide can be found at www.csar.cfs.ac.uk/user_information/software. Tools, compilers and utilities have their own section at www.csar.cfs.ac.uk/user_information/tools/. Please let us know if there is anything else that you would like to see installed on any of the CSAR systems.

Modules

Software access on the CSAR machines is controlled by modules. Loading a given module sets the paths and environment necessary to use the software associated with it. A number of the more important modules are loaded automatically on login (use the `module list` command to see which). The more useful module commands are listed below. A complete list can be seen via the `module help` command.

Command	Usage
<code>module list</code>	lists the modules you currently have loaded
<code>module avail</code>	lists the modules that are available to be loaded in
<code>module load module_name</code>	loads the specified module so that the software can be used
<code>module unload module_name</code>	unloads the specified module
<code>module switch old_module_name new_module_name</code>	unloads the old module and replaces it with the new module

Further information on modules is available at www.csar.cfs.ac.uk/user_information/software/modules.shtml



7. Batch Jobs

An introduction to bsub

LSF (Load Sharing Facility) is the batch system utilised by the CSAR machines. The command used to submit jobs to LSF is `bsub`. In order to submit to ferret and green you should issue commands on wren. Jobs to be run on newton can be submitted by logging into newton and running commands.

You must give LSF options to inform it about the needs of your job. The basic options are described below.

- **Number of cpus:** requested using the `-n` flag.
- **Wall clock time:** requested with the `-W` flag. This means that your job will automatically finish after that amount of time is used up if it has not already finished. Measured and specified in minutes.
- **Total CPU time:** specified with the `-c` option and is similar to `-W` in that it is a way of restricting the amount of time your job runs for. However `-c` is the total amount of CPU time used. Measured and specified in minutes.
- **Where to run the job:** to specify where your job runs can be achieved with the `-q` and `-m` options. The `-m` option specifies which machine to run on such as green. The `-q` option tells it which queue to run on for that particular machine.
- **Job name:** to give your job a name which can be useful to identify which of your jobs are running when using some of the LSF monitoring commands use the `-J` option.
- **Output:** the stderr and stdout can be controlled with the `-e` and `-o` options. The `-e` options specifies where the stderr should be put and the `-o` option specifies where the stdout should be put. If neither option is specified the output should be emailed to you, if only the `-o` option is specified the stdout and stderr are merged into the specified file. It is also useful to put "%J" at the end of the filename as this will create a unique file for each job output.

Submission script example

The following script can be submitted which will run `a.out` on 8 processors asking for one hour of green and writing the output to a file called `output` and calling the job `my_job`. It should be submitted using `bsub < scriptname`

```
#BSUB -n 8
#BSUB -W 1:00
#BSUB -m green
#BSUB -J my_job
#BSUB -q normal
#BSUB -o output

mpirun -np 8 ./a.out
```

Command line example

```
bsub -n 8 -W 1:00 -m green -J my_job -q normal -o output scriptname
```

Only the `mpirun -np 8 ./a.out` then needs to be in the script.

Further Information

Note that the IRIX machines and the Altix do differ slightly. Please consult the following webpages for more detailed information:



Job Submission (inc machine specific pages):
http://www.csar.cfs.ac.uk/user_information/job_info/

LSF: http://www.csar.cfs.ac.uk/user_information/lsf/



8. Porting and Parallelising Codes

Choice of Machines

The CSAR machines have very similar architectures, all being shared memory machines built around SGI's NUMalink interconnect. However, due to differences between the MIPS processors in the Origins, and the Itanium 2 processors in the Altix, the choice of machine can depend a lot on the type of code that you have and how much work you are prepared to put in (with our help) in modifying it so that it runs efficiently. You may wish to refer to the CSAR helpdesk about the suitability of the different machines for your type of code. When you have decided which machine you wish to run your code on then there are two aspects to consider for code development which are described below.

Types of Parallelism

The CSAR machines are all ccNUMA shared memory machines with an underlying distributed-memory architecture. ccNUMA stands for cache-coherent non-uniform memory access, and this means that the processors should all have the same view of memory at any given moment, but memory access times are not the same across the machine for any processor, with some parts of memory considered "closer" than others.

On a shared memory ccNUMA machine such as the Origin and Altix computers, parallel programs can be written in one of two ways:-

- **Distributed Memory Model:** In the distributed memory model, programs are written so that independent processes communicate with each other by inserting into the code calls to message passing libraries such as MPI or SHMEM. Computer codes using the MPI library are generally highly portable between different parallel machines.
- **Shared Memory Model:** Using the shared memory model, programs are parallelised by inserting compiler directives into the source files which tell the compiler to parallelise loops or sections of code. The most common, and portable, set of these directives is defined in the OpenMP standard.

Porting Codes

Apart from the programming paradigm, there are other issues that you should be aware of when porting a code between machines. The major ones are detailed below.

- **Default Data Sizes:** Different machines can have different default sizes for intrinsic numerical types such as a Fortran `Real` or a C `float`. For example on a Cray T3E a C `float` is 8 bytes long whereas on an SGI Origin it is 4 bytes long. If you do not take this into account, your program might still run without any changes but it might give incorrect results.
- **Compilers:** Compiler behaviour and compiler options can change between different systems and in some cases compiler flags can completely change their meaning. For example on the MIPSpro compiler on an Origin machine `-mp` indicates that there are OpenMP parallelisation directives in the code which should be compiled, whereas to the Intel compiler on the Altix the `-mp` option is an instruction to maintain precision by reducing the optimisation level.
- **Non-standard Fortran Extensions:** Often Fortran code is written to use functions or subroutines which are not part of the Fortran standard, but which are implemented by the compiler vendor. These routines might be timing routines, utilities to extend the input/output capabilities or extra numerical functions. For routines that are not part of the Fortran standard it may be necessary to use different function or subroutine calls on different machines.

- **External Libraries:** External libraries of routines which are used by the code may have different properties on different machines or may not even exist when moving from one architecture to another. Routine names may change, input arguments to functions may be different or completely new routines in different libraries may need to be selected.

Further Help and Information

This information is repeated on www.csar.cfs.ac.uk/user_information/porting/ where links to more detail can be found.

If you are thinking of porting a code to one of the CSAR machines, then you may wish to consider looking into the application support/optimisation services (www.csar.cfs.ac.uk/services/support.shtml) and training courses (www.csar.cfs.ac.uk/services/courses/) which you are entitled to apply for when submitting an application for CSAR resources to the research councils.

9. Documentation

UNIX man Pages

Information on all commands on the systems is available using the standard unix 'man' tool. Additionally information on all the routines that are callable from your Fortran and C programs is available. For example `man mpi` will give a good introduction to the message passing interface (MPI).

CfS Systems man Pages

A list of man pages available on the systems is given below. These provide specific information on the configuration and the way in which the systems can be most effectively used, as well as the current limits on the size of batch jobs etc.

man page	Content
<code>man cfs</code>	General advice
<code>man cfsqueue</code>	General queue information
<code>man cfslimits</code>	Queue limits and other queue information
<code>man cfstimes</code>	Times allocated for job types
<code>man cfsqs</code>	Job status for all jobs
<code>man cfsshecule</code>	Job scheduling
<code>man cfsfiles</code>	File storage areas and file transfer speeds
<code>man dmq</code>	Displays counts of outstanding dmf requests
<code>man lac</code>	List accounting information utility

CSAR Monthly Bulletins

General service and system information or changes are published in the monthly bulletins (www.csar.cfs.ac.uk/about/bulletin/). This bulletin is emailed once a month to all registered users.



10. Monitoring your Resources

An Introduction to Resources

Your project will have an allocation of home directory disk quota on the SAN and some or all of the following resources:

- Green CPU
- SMP PE hours (fermat)
- Wren CPU
- Disk other than that allocated for home directories, e.g. High Volume disk
- HSM/tape capacity in Gbyte years
- Training days
- Optimisation/Application support days

Please see the resources webpage at www.csar.cfs.ac.uk/project_management/resources.shtml for detailed information on the above.

PI's must sub-allocate some of the project's disk quota to all users, and some of the HSM/tape resource to users who require this facility. Sub-allocation of other resources is entirely optional; if the PI chooses not to sub-allocate other resources, users will have unrestricted access to the project's resources. Your PI has primary control over user's allocations; in general we do not have authority to change user level allocations. Please contact your PI with any queries about your allocation levels.

If a PI needs to exercise more control over usage there are two methods: (i) give allocations to individual users, or (ii) use the facilities to sub-allocate resources to groups of users. Your PI will choose the best method depending on the size and nature of your project.

Checking your user allocations

Your own user allocation levels and usage can be seen in your user record in the registration database (see <http://www.csar.cfs.ac.uk/admin/register.shtml>). Note that usage will be visible here only if a suballocation has been made. Also note that allocation/usage figures are cumulative from the start of the project.

Viewing usage for a specified time period

For viewing usage over particular periods of time (regardless of whether a sub-allocation exists) there are two methods:

1. The 'lac' command on wren (see `man lac`).
2. Web-based Usage Reports facility (see <http://www.csar.cfs.ac.uk/admin/accounts>).

The second option is more flexible and is intended to be the primary tool for usage reporting. Users can also see their project's (or subproject if applicable) usage via either of these means.

Automated warnings and stopping Mechanism

An automated warning and stopping mechanism exists to enforce user and project level allocations. Cron jobs run nightly to check usage against allocations. Two warnings will be issued via email prior to any stops/sanctions being enforced. Normally the first one of these will be when usage exceeds 90% of the allocation; the second at 95%. Should allocations not

be increased and the usage subsequently exceeds 100%, then sanctions will be imposed until the situation is rectified.

Warning emails are sent to the user (if applicable) and the project PI. The CSAR Helpdesk will also receive a copy for information. Upon receiving warnings users should contact their PI to discuss their allocation levels.

Note that sanctions may be imposed on all users in a project if the project itself runs out of any resource. Your PI will have been informed if this is the case, and in general it is the PIs responsibility to rectify the overuse by adjusting the resource balance if possible, and to keep users informed.

11. Keeping Up to Date

Staus Page Information

Information on the status of the CSAR systems can be found at www.csar.cfs.ac.uk/mstatus/status.shtml. This information is automatically updated at regular intervals, and occasionally messages are placed here if there are any particular problems. If a machine is unavailable for any reason you can add your email address to a list to be automatically informed when the machine is back up (valid for one occasion only). This page also contains batch queue information.

Scheduled Maintenance

The systems may be unavailable during the scheduled maintenance sessions (generally Tuesday evenings, up to twice a month), and all users will be notified in advance of these sessions via email, on the systems and via the status page.

System Updates

Brief details of compiler updates and other system changes are available at www.csar.cfs.ac.uk/mstatus/system_updates.shtml

News and the CSAR Mailing List

Use is made of the news facility on wren and newton. When you log in you will see a list of unread news items relating to system matters. Read an item by typing `news item_name`. These news items may also be put on the news page on the web, along with more general CSAR information.

All aspects of the service will be described via the CSAR web pages and in a monthly email bulletin. When you register you are automatically added to our email list. We will occasionally provide additional information via email. Please ensure that you inform us if you change your contact details. You can see past monthly bulletins at web address www.csar.cfs.ac.uk/about/bulletin/

Support and feedback

All queries (or comments or suggestions) should be directed to the CSAR Helpdesk (email csar-advice@cfs.ac.uk, telephone 0161 275 6824/5997, fax 0161 275 6040). The Helpdesk is staffed between 8.30am and 6.00pm every working day. All calls are logged using a sophisticated automatic logging system.

The CSAR Service incorporates a 'Service Quality Token' scheme whereby users can give their views on any aspect of the service by allocating gold or black tokens, along with comments. Allocating tokens is a simple procedure via the registration system -- full instructions are given at http://www.csar.cfs.ac.uk/project_management/sqt.shtml. You may change your allocation at any time to reflect your current view.

