

Short Note

The “unwritten” computing rules at SEP

Alexander M. Popovici, Dave Nichols and Dimitri Bevc¹

INTRODUCTION

This short note is intended for the fresh SEP-er. It is a collection of rules and SEP trivia that are obvious to any SEP student who spent more than a year working on our computers. Since many of the “senior” SEP students plan to graduate this year, we thought a formal introduction to our computing environment could be in place. The title is taken from a seminar handout given to the first author by John Etgen a long time ago ...

Home directory

The directory `/homes/sep/username` also called your “home directory” contains mainly source files and text files, generated using an editor. In other words, files that were typed by you or by somebody else. The size limit is 250 Kbytes, because it is assumed you don’t have source code or text files that exceed this size. If you do, you should split them into smaller files as the generic rules of code writing (Kernighan and Plauger, 1978) recommend that code modules contained in a single file should be in the 300-500 lines range. The reasoning for keeping only small files that were generated by typing in your home directory is that `/homes/sep` is backed up every day and all the files that changed since the previous day are stored on disk. It would be a waste of time and tape to back up executables that can be always regenerated or data files. To restore files that you changed one or two days before you can use the command `restoresep [1-2]`, which will interactively prompt you to choose the directory and file you want to restore. Your home directory is not intended for storing executables. Executables are kept on “room devices,” filesystems called `/r0,/r1,...,/r5` in a directory called

```
/r[0-5]/username/bin/machinetype
```

where `username` is your login name and `machinetype` can be one of

```
SUN4,RS6000,HP700,DEC3100, OR CM5 .
```

¹email: not available

The path variable is automatically reconfigured whenever you login on a new machine to contain the appropriate executable directory. In addition the path also contains the directory `/r[0-5]/username/bin/all` intended to keep Unix shells which are machine independent. Recently we have added a new rule to the “hog shells” (shells that run overnight and check for disk space usage) to delete object files in the home directory that were not accessed or modified for 15 days. There are no mechanisms to enforce the policy of storing only small files in your home directory. The general principle at SEP is to keep all the system administration fascist interventions to a minimum. The truth is that system administrators are the first to exemplify the maxim that some grow with responsibility, others just swell. The “hog shells” generate daily messages sent to the owners of the files exceeding the maximum size, reminding them of the location of the aforementioned reprehensible files. Such a message will look like:

```
From: Operator <root>
Message-Id: <199404092320.AA19378@oas.Stanford.EDU>
To: christin@sep.Stanford.EDU
Subject: /home/oas/sep/christin/advance/rs6000/exe/SEPinput.exe
Status: R
You have large files on /sep, please check /var/sephogs and remove them.
```

Scratch space

The `/scr[1-16]` directories are intended for large dataset storage. The scratch disks are designated temporary disk space and files residing on it which are not modified or **accessed** in a given time period are automatically erased. The time period is usually 7 days. The body of SEP-datafiles is automatically placed on a scratch disk when using `seplib` utilities, unless specified differently by a `.datapath` file or your `DATAPATH` variable. The default `.datapath` file residing in your home directory usually defaults to a scratch disk local to the machine that you are running on. A typical `.datapath` file will contain the following scratch disk default locations:

```
datapath=/scr4/mihai/
oas datapath=/scr4/mihai/
pele datapath=/SDA/mihai/
robson datapath=/scr3/mihai/
spur datapath=/scr5/mihai/
```

The “hog shells” also keep track of the amount of space left on the scratch devices and will notify the top four users of the proportional amount of disk space occupied by each of them. The message is also copied to `dave`, `stew`, `mihai`, together with a quote from Jon, intended to shame or just instill solid guilt in the perpetrators...

Such an e-mail message will have the form:

```
Date: Wed, 9 Mar 1994 06:24:36 -0800
From: Operator <root>
Message-Id: <199403091424.AA17142@oas.Stanford.EDU>
To: dave@sep.Stanford.EDU, dimitri@sep.Stanford.EDU,
    jun@sep.Stanford.EDU, martin@sep.Stanford.EDU,
    mihai@sep.Stanford.EDU, stew@sep.Stanford.EDU
Subject: space hogging on /scr
Status: R
You were singled out as a major hog on /scr.
/scr is more than 85 percent full.
/dev/id004c      836106  662652   89844   88%    /home/oas/scr
It is not rude to top the list occasionally,
but do not top the list on successive days. -J.
/dev/id004c      836106  662652   89844   88%    /scr
/scr/dimitri    201664  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/scr/jun        198209  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/scr/martin     147554  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/scr/ftp        53931   XXXXXXXXX
```

Touching policy

Though no “official” policy is in place regarding “touching” the scratch files, historically some students were threatened that if they were caught doing it repeatedly, their accounts would be suspended for one week.

Job priorities

In the old days when a Convex C-1 was our main server, with all the students in SEP running jobs on the same CPU, there were more restrictions on the job size, length and priority. Nowadays the load is spread on tens of workstations and several bigger compute-servers and these rules are somewhat obsolete. We state them mostly for historical reasons, but also to create a basis for the new set of rules regarding running programs on our parallel computer, the CM-5. The priority of the jobs on the Convex was established as a function of the estimated job length. The basic principle was that jobs longer than 10 minutes should be “reniced” (after 10 minutes) to priority 10. Jobs under 30 seconds were considered interactive, jobs between 30 seconds and 10 minutes were considered short, and jobs over 10 minutes were considered long. Interactive jobs could run at priority 0. Short jobs could run at priority 0 unless the load average was over 5 in which case the priority was changed to 4. Long jobs ran at priority 0 or 4 for the first 10 minutes and priority 10 afterwards.

In addition to the priority scheme a number of one-liners defined the rest of the job policy:

Rarely if ever run more than one long job at a time.

If you are not logged in, renice your job.

If a shell runs a succession of long or short jobs, only the first total 10 minutes gets priority 0; after 10 minutes renice the job and the shell.

If you run a long job and an interactive job in parallel, the long job should start reniced.

CM-5

The Convex c-1 was replaced several years ago by a Connection Machine CM-2 parallel computer which in turn was replaced by a CM-5. We do not have a definitive usage policy for the CM-5, though our experience so far is materialized in the following message which appears each time a user logs in:

```
*****
SEP CONNECTION MACHINE CM5 USER POLICY:
    Development time:      Mon - Sat 10am to 10pm
    Production time:      Mon - Fri 10pm to 10am
                        Sat 10pm to Mon 10am
Production: Multi-hour jobs which take up more than 50% of memory
Development: Jobs which run less than 1/2 hour with most of memory and
are not being run too often!
        -- OR --
Jobs which run less than 1 hour with no more
than 50% of memory.
People pushing these limits must be considerate of others:
Don't submit multiple 50% memory jobs one after another
if there are other people who want to run 50% memory jobs.
Don't submit multiple >50% memory jobs in a short period of
time.
*****
User signup for production jobs:
    user          date & time          time estimate
    ----          -
*****
ATTENTION!
    The CM5 will go down every Tuesday 5am-7am for diagnostics.
    To disable the diagnostics for a particular week, type:
    NoCMdiag.
*****
```

We plan timely revision of our usage policy on the CM-5, trying to adapt it to whatever problems will appear in the future. So far though, keep in mind that the CM-5 is mainly a batch machine; it doesn't gracefully handle multiple large jobs. Two large jobs running simultaneously will interfere with each other and sometimes the total time to finish them will be four times longer than the sum of the separate runtimes. The CM-5 doesn't swap on disk, all the jobs are kept in core and therefore some large memory jobs will not fit unless the whole machine is allocated for their execution.

Scalable Disk Array (SDA) policy

The Scalable Disk Array (SDA) is the scratch disk for the CM-5. It is intended to store only datasets that are strictly processed with the CM-5. Actually the disk is mounted read-only by all the other machines, with the obvious exception of the CM-5. The autoerase policy is similar to the one on the `/scr` disks, with a time limit of 15 days. Any files not accessed or modified in a 15 day period are deleted. The only exception to this rule are the files kept in a directory starting with a capital letter. The high-watermark of the filesystem is 85%, and the e-mail system is activated when this limit is exceeded, generating messages of the type:

```
From root Fri Feb 18 17:10:19 1994
Date: Sat, 19 Feb 1994 01:10:19 -0800
From: Operator <root>
Message-Id: <199402190910.AA00190@pele.Stanford.EDU>
To: 3D@sep.stanford.edu, biondo@sep.stanford.edu, dave@sep.stanford.edu,
    dimitri@sep.stanford.edu, mihai@sep.stanford.edu,
    stew@sep.stanford.edu
Subject: space hogging on /SDA
Status: R
/SDA is 86% full. You were singled out as a major HOG.
Delete your unnecessary files from the SDA directory.
It is not rude to top the list occasionally,
but do not top the list on successive days. -J.
/SDA/dimitri    3637456 XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/SDA/3D        3024112 XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/SDA/biondo    2302912 XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
/SDA/mihai     1494768 XXXXXXXXXXXXXXXXXXXXXXX
/SDA/dave      1053392 XXXXXXXXXXXXX
/SDA/david     882640 XXXXXXXXX
/SDA/martin    814848 XXXXXXXX
/SDA/thorbjor  664256 XXXXXX
```

Other miscellaneous information

The passwords used to login have a usable maximum length of 8 characters. Any characters after the first 8 are ignored. Therefore it's not worth to have any password longer than 8 characters, because it is just extra typing. On the other hand we try to crack our user's passwords every month, and the shorter the password the easier to crack. The recommended length of a password is hence 8 characters. The password cracking program we run uses a 70000 word dictionary, and it matches all the lower capital letters to your password in a couple of minutes. For 24 hours it tries to match different combinations. As a result, a better password should contain special characters (!@#%&*^&*()), numbers, and combinations of lower case and upper case. It also helps if the basic word used to compose your password is not found in English dictionaries.