

# Disk Space and Disk Usage

1. [What Disks to Use](#)
2. [How to Copy Files to/from Hydra](#)
3. [Disk Quotas](#)
4. [Disk Configuration](#)
5. [Disk Usage Monitoring](#)
6. [NetApp Snapshots: How to Recover Old or Deleted Files](#)
7. [Public Disks Scrubber](#)
8. [SSD and Local Disk Space](#)

## 1. What Disks to Use

All the useful disk space available on the cluster is mounted off a dedicated device (aka appliance or server), a NetApp filer.

The available disk space is divided in several area (aka partitions):

- a small partition for basic configuration files and small storage, the `/home` partition ,
- a set of medium size partitions, one for SAO users, one for non-SAO users, the `/data` partitions,
- a set of large partitions, one for SAO users, one for non-SAO users, the `/pool` partitions,
- a second set of large partitions for temporary storage, the `/scratch` partitions.

Note that:

- we impose quotas: limits on how much can be stored on each partition by each user, and
- we monitor disk usage;
- `/home` should not be used to keep large files, use `/pool` instead;
- `/pool` is for active temporary storage (i.e., while a job is running).
  - If you need even more disk space, ask to be allowed to use `/scratch`.
  - Both partitions (`/pool` and `/scratch`) are **scrubbed (see below)**: old stuff is deleted to make sure there is space for active users.
- None of the disks on the cluster are for long term storage, please copy your results back to your "home" computer and delete what you don't need any longer.
- While the disk system on Hydra is highly reliable, none of the disks on the cluster are backed up.
- Once you reach your quota you won't be able to write anything on that partition until you delete stuff.
- A few nodes have local SSDs (solid state disks), and for special cases it may be OK to use disk space local to the compute node.

[Contact us](#) if your jobs can benefit from more disk space, SSDs or local disk space.

## 2. How to Copy Files to/from Hydra

⚠ When copying to Hydra, especially large files, be sure to do it to the appropriate disk (and not `/home` or `/tmp`).

### 2a. To/From Another Linux Machine

- You can copy files to/from hydra using `scp`, `sftp` or `rsync`:
  - to Hydra you can only copy from *trusted* hosts (computers on SI or SAO/CfA trusted network, or VPN'ed),
  - from Hydra to any host that allows external `ssh` connections (if you can `ssh` from Hydra to it, you can `scp`, `sftp` and `rsync` to it).
- For large transfers (over 70GB, sustained), we ask users to use `rsync`, and limit the bandwidth to 20 MB/s (70 GB/h), with the "`--bwlimit=`" option:
  - `rsync --bwlimit=20000 ...`  
If this pose a problem, contact us (Sylvain or Paul).
  - Baseline transfer rate from SAO to HDC (Herndon data center) is around 300 Mbps, single thread, or ~36 MB/s or ~126 GB/h (as of Aug. 2016)  
The link saturates near 500 Mbps (50% of Gbps) or 62 MB/s or 220 GB/h
- Remember that `rm`, `mv` and `cp` can also create high I/O load, so consider to
  - limit your concurrent I/Os: do not start a slew of I/Os at the same time, and
  - serialize your I/Os as much as possible: run one *after* the other.

### NOTE for SAO Users:

🔑 Access from the "outside" to SAO/CfA hosts (computers) is limited to the *border control hosts* ([login.cfa.harvard.edu](http://login.cfa.harvard.edu) and [pogoN.cfa.harvard.edu](http://pogoN.cfa.harvard.edu)), instructions for tunneling via these hosts is explained on

- the CF's [SSH Remote Access](#) page, or
- the HEAD Systems Group's [SSH FAQ](#) page.

## 2b. From a Computer Running MacOS

A trusted or VPN'd computer running MacOS can use `scp`, `sftp` or `rsync`:

- Open the Terminal application by going to `/Applications/Utilities` and finding Terminal.
- At the prompt, use `scp`, `sftp` or `rsync`, after `cd`'ing to the right place.
- For large transfers limit the bandwidth and use `"rsync --bwlimit=4000"`.



Alternatively you can use a GUI based `ssh/scp` compatible tool like [FileZilla](#). Note, Cyberduck is not recommended because it uses a lot of CPU cycles on Hydra.

You will still most likely need to run VPN.

## 2c. From a Computer Running Windows

You can use `scp`, `sftp` or `rsync` if you install [Cygwin](#) - Note that Cygwin includes a X11 server.

Alternatively you can use a GUI based `ssh/scp` compatible tool like [FileZilla](#) or [WinSCP](#). Note, Cyberduck is not recommended because it uses a lot of CPU cycles on Hydra.

You will still most likely need to run VPN.

## 2d. Using Globus

(instructions missing)

## 2e. Using Dropbox

Files can be exchanged with Dropbox using the script [Dropbox-Uploader](#), which can be loaded using the `tools/dropbox_uploader` module and running the `dropbox` or `dropbox_uploader.sh` script. Running this script for the first time will give instructions on how to configure your Dropbox account and create a `~/dropbox_uploader` config file with authentication information.

Using this method will not sync your Dropbox, but will allow you to upload/download specific files.

## 3. Disk Quotas

To prevent the disks from filling up and hose the cluster, there is a limit (aka quota) on

- how much disk space and
- how many files (in fact "inodes": the sum of number of files and number of directories) each user can keep.

Each quota type has a soft limit (warning) and a hard limit (error) and is specific to each partition. In other words exceeding the soft limit produces warnings; while exceeding the hard limit is not allowed, and results in errors.

## 4. Disk Configuration

	Maximum	Quotas per user		NetApp	
	disk	disk space	no. of files	snapshots	
Disk name	capacity	soft/hard	soft/hard	enabled?	Purpose
<a href="#">/home</a>	10TB	50/100GB	1.8/2M	yes: 4 weeks	For your basic configuration files, scripts and job files - your limit is low but you can recover old stuff up to 4 weeks.
<a href="#">/data/sao</a> or <a href="#">/data/nasm</a>	40TB*	1.9/2.0TB	4.75/5M	yes: 2 weeks	For important but relatively small files like final results, etc. - your limit is medium, you can recover old stuff, but disk space is not released right away. For SAO or NASM users.

<a href="#">/data/genomics</a>	30TB*	0.45/0.5TB	1.19/1.25M	yes: 2 weeks	For important but relatively small files like final results, etc. - your limit is medium, you can recover old stuff, but disk space is not released right away. For non-SAO/NASM users.
<a href="#">/pool/sao</a> or <a href="#">/pool/nasm</a>	37TB	1.9/2.0TB	4/5M	no	For the bulk of your storage - your limit is high, and disk space is released right away, for SAO or NASM users.
<a href="#">/pool/genomics</a>	50TB	1.9/2.0TB	4.75/5M	no	For the bulk of your storage - your limit is high, and disk space is released right away, for non-SAO users.
<a href="#">/pool/biology</a>	7TB	1.9/2.0TB	4.75/5M	no	For the bulk of your storage - your limit is high, and disk space is released right away, for non-SAO/NASM users.
<a href="#">/scratch</a>	100TB	9.5/10.0TB	23.75/25M	no	For temporary storage, if you need more than what you can keep in <a href="#">/pool</a> - SAO, NASM or non-SAO/NASM users should use <a href="#">/scratch/sao</a> , <a href="#">/scratch/nasm</a> or <a href="#">/scratch/genomics</a> , respectively
<b>Project specific disks</b>					
<a href="#">/pool/kistler1</a>	21TB	20.0/21.0TB	49.9/52.5M	no	NMNH/Logan Kistler
<a href="#">/pool/kozakk</a>	11TB	10.5/11.0TB	26.1/27.5M	no	STRI/Krzysztof Kozak
<a href="#">/pool/nmnh_ggi</a>	21TB	15.0/15.8TB	37.4/39.4M	no	NMHN/GGI
<a href="#">/pool/sao_access</a>	21TB	15.0/15.8TB	37.4/39.4M	no	SAO/ACCESS
<a href="#">/pool/sao_atmos</a>	36TB	8.0/10TB	9/10M	no	SAO/ATMOS
<a href="#">/pool/sao_rtdc</a>	10TB*	2.8/3.0TB	2.5/3.0M	no	SAO/RTDC
<a href="#">/pool/sao_cga</a>	8TB	7.9/8TB	20/19M	no	SAO/CGA
<a href="#">/pool/sylvain</a>	15TB	14/15TB	63/65M	no	SAO/Sylvain Korzennik
<b>Extra</b>					
<a href="#">/pool/admin</a>	10TB*	5.7/6.0TB	14.3/15.0M	no	Sys Admin
<a href="#">/pool/galaxy</a>	15TB,	10.7/11.3TB	26.7/28.1M	no	Galaxy
*: maximum size, disk size will increase up to that value if/when usage grows					

(as of May 1, 2018)

## Notes

- The notation
  - 1.8/2.0TB means that the soft limit is 1.8TB and the hard limit is 2.0TB of disk space, while
  - 4/5M means that the soft limit is 4 million `inodes` and the hard limit is 5 million.
- It is inefficient to store a slew of small files and if you do you may reach your `inodes` quota before your space quota (too many small files).
  - Some of the disk monitoring tools show the `inode` usage.
  - If your `%(inode)>%(space)` your disk usage is inefficient, consider archiving your files into `zip` or `tar-compressed` sets.
- While some of the tool(s) you use may force you to be inefficient while jobs are running, you should remember to
  - remove useless files when jobs have completed,
  - compress files that can benefit from compression (with `gzip`, `bzip2` or `compress`), and
  - archive a slew of files into a `zip` or a `tar-compressed` set, as follows:
 

```
% zip archive.zip dir/
```

 or
 

```
% tar -czf archive.tgz dir/
```

 both examples archive the content of the directory `dir/` into a single `zip` or a `tgz` file. You can then delete the content of `dir/` with
 

```
% rm -rf dir/
```
- You can unpack each type of archive with
 

```
% unzip archive.zip
```

or

```
% tar xf archive.tgz
```

- The sizes of some of the partitions (aka the various disks) on the NetApp will "auto-grow" until they reach the listed maximum capacity, so the size shown by the traditional Un\*x command, like `df` does not necessarily reflect the maximum size.

We have implemented a FIFO (first in first out) model, where old files are deleted to make space, aka scrubbed.

- There is an age limit, meaning that only files older than 180 days (or 90 days) get deleted.
- Older files get deleted before the newer ones (FIFO),
- We run a [scrubber](#) on a regular interval.
- In any case, we ask you to remove from `/pool` and `/scratch` files that you do not need for active jobs.
- For projects that want dedicated disk space, such space can be secured with project's specific funds when we expand the disk farm ([contact us](#)).

## 5. Disk Monitoring

The following tools can be used to monitor your disk usage.

- You can use the following Un\*x commands:

<code>du</code>	show disk use
<code>df</code>	show disk free

or

- you can use Hydra-specific home-grown tools, (these require that you load the `tool/local` module)

<code>dus-report.pl</code>	run <code>du</code> and parse its output in a more user friendly format
<code>disk-usage.pl</code>	run <code>df</code> and parse its output in a more user friendly format

- You can also view the disk status at the cluster status web pages, either
  - [here](#) (at [cfa.harvard.edu](http://cfa.harvard.edu))
  - or
  - [here](#) (at [si.edu](http://si.edu)).

Each site shows the disk usage and a quota report, under the "Disk & Quota" tab, compiled 4x a day respectively, and has links to plots of disk usage vs time.

### Disk usage

The output of `du` can be very long and confusing. It is best used with the option "`-hs`" to show the sum ("`-s`") and to print it in a human readable format ("`-h`").

⚠ If there is a lot of files/directory, `du` can take a while to complete.

📄 For example:

```
% du -sh dir/
136M  dir/
```

The output of `df` can be very long and confusing.

📄 You can use it to query a specific partition and get the output in a human readable format ("`-h`"), for example:

```
% df -h /pool/sao
Filesystem      Size  Used Avail Use% Mounted on
10.61.10.1:/vol_sao  20T   15T   5.1T  75% /pool/sao
```

You can compile the output of `du` into a more useful report with the `dus-report.pl` tool. This tool will run `du` for you (can take a while) and parse its output to produce a more concise/useful report.

For example, to see the directories that hold the most stuff in `/pool/sao/hpc`:

```

% dus-report.pl /pool/sao/hpc
612.372 GB /pool/sao/hpc
capac. 20.000 TB (75% full), avail. 5.088 TB
447.026 GB 73.00 % /pool/sao/hpc/rtdc
308.076 GB 50.31 % /pool/sao/hpc/rtdc/v4.4.0
138.950 GB 22.69 % /pool/sao/hpc/rtdc/vX
137.051 GB 22.38 % /pool/sao/hpc/rtdc/vX/M100-test-oob-2
120.198 GB 19.63 % /pool/sao/hpc/rtdc/v4.4.0/test2
120.198 GB 19.63 % /pool/sao/hpc/rtdc/v4.4.0/test2-2-9
83.229 GB 13.59 % /pool/sao/hpc/c7
83.229 GB 13.59 % /pool/sao/hpc/c7/hpc
65.280 GB 10.66 % /pool/sao/hpc/sw
64.235 GB 10.49 % /pool/sao/hpc/rtdc/v4.4.0/test1
49.594 GB 8.10 % /pool/sao/hpc/sw/intel-cluster-studio
46.851 GB 7.65 % /pool/sao/hpc/rtdc/vX/M100-test-oob-2/X54.ms
46.851 GB 7.65 % /pool/sao/hpc/rtdc/vX/M100-test-oob-2/X54.ms/SUBMSS
43.047 GB 7.03 % /pool/sao/hpc/rtdc/vX/M100-test-oob-2/X220.ms
43.047 GB 7.03 % /pool/sao/hpc/rtdc/vX/M100-test-oob-2/X220.ms/SUBMSS
42.261 GB 6.90 % /pool/sao/hpc/c7/hpc/sw
36.409 GB 5.95 % /pool/sao/hpc/c7/hpc/tests
30.965 GB 5.06 % /pool/sao/hpc/c7/hpc/sw/intel-cluster-studio
23.576 GB 3.85 % /pool/sao/hpc/rtdc/v4.4.0/test2/X54.ms
23.576 GB 3.85 % /pool/sao/hpc/rtdc/v4.4.0/test2-2-9/X54.ms
23.576 GB 3.85 % /pool/sao/hpc/rtdc/v4.4.0/test2/X54.ms/SUBMSS
23.576 GB 3.85 % /pool/sao/hpc/rtdc/v4.4.0/test2-2-9/X54.ms/SUBMSS
22.931 GB 3.74 % /pool/sao/hpc/rtdc/v4.4.0/test2/X220.ms
22.931 GB 3.74 % /pool/sao/hpc/rtdc/v4.4.0/test2-2-9/X220.ms
report in /tmp/dus.pool.sao.hpc.hpc

```

You can rerun `dus-report.pl` with different options on the same intermediate file, like

```
% dus-report.pl -n 999 -pc 1 /tmp/dus.pool.sao.hpc.hpc
```

to get a different report, to see the list down to 1%. Use

```
% dus-report.pl -help
```

to see how else you can use it.

The tool `disk-usage.pl` runs `df` and presents its output in a more friendly format:

```

% disk-usage.pl
Filesystem              Size      Used      Avail Capacity  Mounted on
NetApp.2:/vol_home      4.00T    1.72T    2.28T   43%/14%    /home
NetApp.2:/vol_data_genomics 18.00T   673.63G  17.34T    4%/1%    /data/genomics
NetApp.2:/vol_data/sao   27.00T    5.25T   21.75T   20%/14%    /data/sao
NetApp.2:/vol_data/nasm  27.00T    5.25T   21.75T   20%/14%    /data/nasm
NetApp.2:/vol_data/admin 27.00T    5.25T   21.75T   20%/14%    /data/admin
NetApp.2:/vol_biology    7.00T    8.64G    6.99T    1%/1%     /pool/biology
NetApp.2:/vol_genomics   50.00T   33.82T   16.18T   68%/11%    /pool/genomics
NetApp.2:/vol_sao        45.00T   14.15T   30.85T   32%/5%     /pool/sao
NetApp.2:/vol_sao/nasm   45.00T   14.15T   30.85T   32%/5%     /pool/nasm
Isilon.10:/ifs/nfs/hydra 60.00T   33.12T   26.88T   56%/89%    /pool/isilon
NetApp.2:/vol_scratch/genomics 100.00T  45.66T   54.34T   46%/38%    /scratch/genomics
NetApp.2:/vol_scratch/sao 100.00T  45.66T   54.34T   46%/38%    /scratch/sao
NetApp.2:/vol_scratch/nasm 100.00T  45.66T   54.34T   46%/38%    /scratch/nasm
NetApp.5:/vol/a2v1/genomics01 31.25T   4.62T   26.63T   15%/11%    /scratch/genomics01
NetApp.5:/vol/a2v1/sao01  31.25T   4.62T   26.63T   15%/11%    /scratch/sao01
NetApp.2:/vol_pool_nmnh_ggi 21.00T   4.35T   16.65T   21%/1%     /pool/nmnh_ggi
NetApp.2:/vol_pool_kistlerl 21.00T   1.98T   19.02T   10%/1%     /pool/kistlerl
NetApp.2:/vol_pool_kozakk 11.00T   7.06T   3.94T    65%/1%     /pool/kozakk
NetApp.2:/vol_sao_atmos  36.00T   15.30T   20.70T   43%/3%     /pool/sao_atmos
NetApp.2:/vol_sao_rtdc   2.00T   167.51G  1.84T    9%/1%     /pool/sao_rtdc
NetApp.2:/vol_pool_sao_access 21.00T  654.49G  20.36T    4%/1%     /pool/sao_access
NetApp.2:/vol_sylvain    30.00T   12.67T   17.33T   43%/23%    /pool/sylvain
NetApp.2:/vol_pool_admin  4.00T   912.88G  3.11T   23%/1%     /pool/admin
NetApp.2:/vol_pool_galaxy 10.00T    0.00G   10.00T    1%/1%     /pool/galaxy

```

Use

```
% disk-usage.pl -help
```

to see how else to use it.

You can, for instance, get the disk quotas and the max size with:

```

% disk-usage.pl -quotas
Filesystem      Size      Used      Avail Capacity  soft/hard   soft/hard   size Mounted on
NetApp.2:/vol_home      4.00T    1.72T    2.28T   43%/14%    50G/100G   1.80M/2.00M   10T /home
NetApp.2:/vol_data_genomics
/genomics      18.00T   673.63G   17.34T    4%/1%    486G/512G   1.19M/1.25M   30T /data
NetApp.2:/vol_data/*
nasm:admin      27.00T    5.25T   21.75T   20%/14%    1.9T/2.0T   4.75M/5.00M   40T /data/sao:
NetApp.2:/vol_biology
/biology      7.00T    8.64G    6.99T    1%/1%    1.9T/2.0T   4.75M/5.00M   n/a /pool
NetApp.2:/vol_genomics
/genomics     50.00T   33.79T   16.21T   68%/11%    1.9T/2.0T   4.75M/5.00M   n/a /pool
NetApp.2:/vol_sao
nasm      45.00T   14.33T   30.67T   32%/5%    1.9T/2.0T   4.75M/5.00M   n/a /pool/sao:
Isilon.10:/ifs/nfs/hydra
/isilon      60.00T   33.12T   26.88T   56%/89%    nyi/nyi     nyi/nyi     n/a /pool
NetApp.2:/vol_scratch/*
/genomics:sao:nasm
NetApp.5:/vol/a2v1/*
/genomics01:sao01
NetApp.2:/vol_pool_nmnh_ggi
/nmnh_ggi     21.00T    4.35T   16.65T   21%/1%    15.0T/15.8T 37.41M/39.38M n/a /pool
NetApp.2:/vol_pool_kistlerl
/kistlerl     21.00T    1.98T   19.02T   10%/1%    20.0T/21.0T 49.88M/52.50M n/a /pool
NetApp.2:/vol_pool_kozakk
/kozakk      11.00T    7.06T    3.94T   65%/1%    10.5T/11.0T 26.13M/27.50M n/a /pool
NetApp.2:/vol_sao_atmos
/sao_atmos    36.00T   15.30T   20.70T   43%/3%    25.7T/27.0T 64.13M/67.50M n/a /pool
NetApp.2:/vol_sao_rtdc
/sao_rtdc     2.00T   167.51G    1.84T    9%/1%    2.9T/3.0T    7.13M/7.50M   10T /pool
NetApp.2:/vol_pool_sao_access
/sao_access   21.00T   654.49G   20.36T    4%/1%    15.0T/15.8T 37.41M/39.38M n/a /pool
NetApp.2:/vol_sylvain
/sylvain     30.00T   12.67T   17.33T   43%/23%    28.5T/30.0T 71.25M/75.00M n/a /pool
NetApp.2:/vol_pool_admin
/admin        4.00T   912.88G    3.11T   23%/1%    5.7T/6.0T   14.25M/15.00M 10T /pool
NetApp.2:/vol_pool_galaxy
/galaxy     10.00T    0.00G   10.00T    1%/1%    10.7T/11.3T 26.72M/28.13M 15T /pool

```

## Monitoring Quota Usage

The Linux command `quota` is working with the NetApp filers (old and new), although not the Isilon.

For example:

```

% quota -s

Disk quotas for user hpc (uid 7235):
  Filesystem  blocks  quota  limit  grace  files  quota  limit  grace
10.61.10.1:/vol_home
                2203M  51200M  100G
                46433  1800k  2000k
10.61.10.1:/vol_sao
                1499G  1946G  2048G
                1420k  4000k  5000k
10.61.10.1:/vol_scratch/genomics
                48501M  2048G  4096G
                1263  9000k  10000k
10.61.200.5:/vol/a2v1/genomics01
                108M  14336G  15360G
                613  10000k  12000k
10.61.10.1:/vol_home/hydra-2/dingdj
                2203M  51200M  100G
                46433  1800k  2000k

```

reports your quotas. The `-s` stands for `--human-readable`, hence the 'k' and 'G'. While

```
% quota -q
```

will print only information on filesystems where your usage is over the quota. ([man quota](#))

## Other Tools

We compile a quota report 4x/day and provide tools to parse the quota report.

The daily quota report is written around 3:00, 9:00, 15:00, and 21:00 in a file called `quota_report_YYDDMM_HH`, located in `/share/apps/adm/reports`.

The string `YYDDMM_HH` corresponds to the date & hour of the report: "160120\_09" for Jan 20 2016 9am report.

The format of this file is not very user friendly and users are listed by their user ID.

The Hydra-specific tools, (i.e., requires that you load the [tools/local](#) module):

- [show-quotas.pl](#) - show quota values
- [parse-quota-report.pl](#) - parse quota report

## Examples

[show-quotas.pl](#) - show quota values:

```
% show-quotas.pl -u sylvain
Limited to user=sylvain

----- quota -----
filesystem          type      name      space    #files
/data/sao:nasm:admin user      sylvain   8.0TB   40.000M
/home                user      sylvain   100.0GB  2.000M
/pool/sao:nasm       user      sylvain   2.0TB   5.000M
/scratch/genomics:sao:nasm user      sylvain   10.0TB  25.000M
/pool/sylvain        user      sylvain   30.0TB  75.000M
```

Use

```
% show-quotas.pl -h
```

for the complete usage info.

[parse-quota-report.pl](#), will parse the quota report file and produce a more concise report:

```
% parse-quota-report.pl
Disk quota report: show usage above 75% of quota, (warning when quota > 95%), as of Wed Nov 22 09:00:04 2017.

disks=/data/admin or /data/nasm or /data/sao (volume=vol_data)
-- disk -- -- #files -- default quota: 2.00TB/5M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_data    1.88TB 94.0%   0.01M  0.1%   Hotaka Shiokawa, SAO/RG - hshiokawa

disk=/pool/genomics (volume=vol_genomics)
-- disk -- -- #files -- default quota: 2.00TB/5M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_genomics 1.84TB 92.0%   0.00M  0.1%   H.C. Lim, NMNH/IZ - limhc
vol_genomics 1.58TB 79.0%   0.15M  3.0%   Bastian Benthage, NMNH - benthageb
vol_genomics 707.4GB 34.5%   4.68M  93.5%  Molly M. McDonough, CCEG - mcdonoughm
vol_genomics 1.52TB 76.0%   0.00M  0.1%   Krzysztof Kozak, STRI - kozakk
vol_genomics 1.70TB 85.0%   0.04M  0.8%   Logan Kistler, NMNH/Anthropology - kistlerl
vol_genomics 2.00TB 100.0%  0.00M  0.0%   *** Xu Su, NMNH/Botany - sux

disk=/home (volume=vol_home)
-- disk -- -- #files -- default quota: 100.0GB/2M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_home    80.3GB 80.3%   0.27M  13.5%  Tileman Birnstiel, SAO/RG - tbirnstiel
vol_home    77.4GB 77.4%   0.18M  9.1%   Rebecca Dikow, NMNH/NZP - dikowr
vol_home    88.3GB 88.3%   0.01M  0.7%   Gabriela Procópio Camacho, NMNH - procopiocamachog
vol_home    100.0GB 100.0%  0.02M  1.1%   *** Logan Kistler, NMNH/Anthropology - kistlerl

disks=/pool/nasm or /pool/sao (volume=vol_sao)
-- disk -- -- #files -- default quota: 2.00TB/5M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_sao     3.63TB 181.5%  0.19M  3.8%   *** Guo-Xin Chen, SAO/SSP-AMP - gchen
vol_sao     1.54TB 77.0%   0.55M  11.0%  Anjali Tripathi, SAO/AST - atripathi
vol_sao     1.66TB 83.0%   0.20M  4.1%   Hotaka Shiokawa, SAO/RG - hshiokawa
vol_sao     2.00TB 100.0%  0.00M  0.1%   *** Chengcai Shen, SAO/SSP - chshen
```

reports disk usage where it is at 75% above quota.

Or you can check usage for a specific user (like yourself) with

```
% parse-quota-report.pl -u <username>
```

for example:

```
% parse-quota-report.pl -u hpc
Disk quota report: show usage (warning when quota > 95%),
for user 'hpc', as of Wed Nov 22 09:00:04 2017.

disks=/data/admin or /data/nasm or /data/sao (volume=vol_data)
-- disk -- -- #files -- default quota: 2.00TB/5M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_data    43.2GB  2.1%    0.01M  0.1%    HPC admin - hpc

disk=/home (volume=vol_home)
-- disk -- -- #files -- default quota: 100.0GB/2M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_home    4.9GB   4.9%    0.04M  2.0%    HPC admin - hpc

disk=/pool/admin (volume=vol_pool_admin)
-- disk -- -- #files -- default quota: 6.00TB/15M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_pool_admin 907.8GB 14.8%   0.44M  2.9%    HPC admin - hpc

disks=/pool/nasm or /pool/sao (volume=vol_sao)
-- disk -- -- #files -- default quota: 2.00TB/5M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_sao      0.0MB  0.0%    0.00M  0.1%    HPC admin - hpc

disks=/scratch/genomics or /scratch/nasm or /scratch/sao (volume=vol_scratch)
-- disk -- -- #files -- default quota: 10.00TB/25M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
vol_scratch  47.4GB  0.5%    0.00M  0.0%    HPC admin - hpc

disk= (volume=a2v1)
-- disk -- -- #files -- default quota: 15.00TB/12M
volume      usage  %quota  usage  %quota  name, affiliation - username (indiv. quota)
-----
a2v1        78.1MB  0.0%    0.00M  0.0%    HPC admin - hpc
```

Use

```
% parse-quota-report.pl -h
```

for the complete usage info.

Users whose quotas are above the 75% threshold will receive a warning email one a week (issued on Monday mornings).

This is a warning, as long as you are below 100% you are OK.

Users won't be able to write on disks on which they have exceeded their hard limits.

## 6. NetApp Snapshots: How to Recover Old or Deleted Files.

Some of the disks on the NetApp filer have the so called "snapshot mechanism" enabled:

- This allow users to recover deleted files or access an older version of a file.
- Indeed, the NetApp filer makes a "snapshot" copy of the file system (the content of the disk) every so often and keeps these snapshots up to a given age.
- So if we enable hourly snapshot and set a two weeks retention, you can recover a file as it was hours ago, days ago or weeks ago, but only up to two weeks ago.
- The drawback of the snapshot is that when files are deleted, the disk space is not freed until the deleted files age-out, like 2 or 4 weeks later.

### How to Use the NetApp Snapshots:

To recover an old version or a deleted file, `foo.dat`, that was (for example) in `/data/genomics/frandsen/important/results/`:

- If the file was deleted:

```
% cd /data/genomics/.snapshot/XXXX/frandsen/important/results
% cp -pi foo.dat /data/genomics/frandsen/important/results/foo.dat
```

- If you want to recover an old version:



```
% cd /data/genomics/.snapshot/XXXX/frandsen/important/results
% cp -pi foo.dat /data/genomics/frandsen/important/results/old-foo.dat
```

- The "-p" will preserve the file creation date and the "-i" will prevent overwriting an existing file.
- The "XXXX" is to be replaced by either:
  - `hourly.YYYY-MM-DD_HHMM`
  - `daily.YYYY-MM-DD_0010`
  - `weekly.YYYY-MM-DD_0015`where `YYY-MM-DD` is a date specification (i.e., `2015-11-01`)
- The files under `.snapshot` are read-only:
  - they be recovered using `cp`, `tar` or `rsync`; but
  - they cannot be moved (`mv`) or deleted (`rm`).

## 7. Public Disks Scrubber

In order to maintain free disk space on the public disks, we are about to implement disk scrubbing: removing old files and old empty directories.

### What is Scrubbing?

We remove old files and old empty directories from a set of disks on a weekly basis.

Old empty directories will be deleted, old files will be, at first, moved away in a staging location, then deleted.



#### Please Note

Since the scrubber moves old files away at first, and delete them later,

- there is a grace period between the scrubbing (move) and the *permanent* deletion to allow users to request for some scrubbed files to be restored;
- reasonable requests to restore scrubbed files must be sent no later than the Friday following the scrubbing, by 5pm;
- scrubbed files still "count" against the user quota until they are *permanently* deleted.

Requests to restore scrubbed file should be

- rare,
- reasonable (*i.e.* no blanket request), and,
- can only be granted while the scrubbed files are not yet *permanently* deleted.

Past the grace period, the files are no longer available, hence users who want their scrubbed files restore have to act *promptly*.

The following instructions explain

- What disks will be scrubbed.
- What to do to access the scrubber's tools.
- How to
  - look at the scrubber's report;
  - find out which old empty directories were scrubbed;
  - find out which old files were scrubbed;
  - create a recovery request.

### What disks will be scrubbed?

The disks that will be scrubbed are:

- `/pool/biology` - 180 days
- `/pool/genomics` - 180 days
- `/pool/sao` - 180 days
- `/scratch/genomics` - 90 days
- `/scratch/genomics01` - 90 days
- `/scratch/sao` - 90 days
- `/scratch/sao01` - 90 days

### How to access the scrubber's tools

- load the module:

```
module load tools/scrubber
```

- to get the list of tools, use:

```
module help tools/scrubber
```

- to get the man page, accessible after loading the module, use:

```
man <tool-name>
```

## How to check what will be scrubbed

- To check what files will be scrubbed, use:

```
find-scrub [-in <dir>] [-age <age>]
```

this will look for files older than <age> days in <dir>, by default dir=current working directory, and age=173 or 83 days.

- This search taxes the file system (aka disk server), especially if you have a lot of files, so use as needed only.

## How to look at the scrubber's results

- To look at the report for what was scrubbed on Jul 21 2016 under /pool/genomics/frandsenp:

```
show-scrubber-report /pool/genomics/frandsenp 160721
```

- To find out which old empty directories were scrubbed:

```
list-scrubbed-dirs [-long|-all] /pool/genomics/frandsenp 160721 [<RE>|-n]
```

where the <RE> is an optional regular-expression to limit the printout, w/o an RE you get the complete list, unless you specify -n and you get the number of scrubbed directories.

The -long or -all option allows you to get more info (like age, size and owner)

- To find out which old files were scrubbed:

```
list-scrubbed-files [-long|-all] /pool/genomics/frandsenp 160721 [<RE>|-n]
```

where again the <RE> is an optional regular-expression to limit the printout, w/o an RE you get the complete list, unless you specify -n and you get the number of scrubbed files;

the -long option will produce a list that includes the files' age and size, -all will list age, size and owner.

- 📌 The <RE> (regular expressions) are PERL-style RE:
  - . means any char,
  - .\* means any set of chars,
  - [a-z] means any single character between a and z,
  - ^ means start of match,
  - \$ means end of match, etc (see [gory details here](#)).
- for example:

```
'^/pool/genomics/blah/project/.*\log$'
```

means all the files that end in '.log' under '/pool/genomics/blah/project/'

## How to produce a list of files to restore

- To produce the list of files to restore as some of the files scrubbed under /pool/genomics/frandsenp/big-project, you can:

1. create a list with

```
list-scrubbed-files /pool/genomics/frandsenp 160721 /pool/genomics/frandsenp/big-project > restore.list
```

this will list all the scrubbed files under 'big-project/' and save the list in restore.list

⚠️ Note that /pool/genomics/frandsenp/big-project means /pool/genomics/frandsenp/big-project\*, if you want to restrict to /pool/genomics/frandsenp/big-project, add a '/', i.e.: use /pool/genomics/frandsenp/big-project/

2. edit the file 'restore.list' to trim it, with any text editor (if needed),

3. verify with:

```
verify-restore-list /pool/genomics/frandsenp 160721 restore.list
```

or use

```
verify-restore-list -d /pool/genomics/frandsenp 160721 restore.list
```

if the verification produced an error.

4. Only then, and if the verification produced no error, submit your scrubbed file restoration request as follow:
  - SAO users: email the file(s) or the location of the files to Sylvain at [hpc@cfa.harvard.edu](mailto:hpc@cfa.harvard.edu)
  - non-SAO users: email the file(s) or the location of the files to [SI-HPC@si.edu](mailto:SI-HPC@si.edu)

## 8. SSD and Local Disk Space

We are in the process of making the local SSDs (solid state disks) available on a few nodes available, and for special cases it may be OK to use disk space local to the compute node.

You should [contact us](#) if your jobs can benefit from either SSDs or local disk space.

How to use the SSD is explained [here](#).

---

Last Updated 21 Dec 2018 SGK/PBF.