

Disk Space and Disk Usage

1. [What Disks to Use](#)
2. [How to Copy Files to/from Hydra](#)
 - a. [To/From another Linux machine](#)
 - b. [From a computer running MacOS](#)
 - c. [From a computer running Windows](#)
 - d. [Using Globus](#)
 - e. [Using Dropbox](#)
 - f. [Using Firefox Send](#)
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 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 volumes, filesets or 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

- we impose quotas: limits on how much can be stored on each disk (partition/volume/fileset) 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 more disk space or your job(s) use(s) a lot of I/Os, 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), [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 and 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.

2f. Using Firefox Send

- [Firefox Send](#) is a free online file-sending service (or a file exchange mechanism).
- Using this system along with the command `ffsend` available on Hydra (`module load bioinformatics/ffsend`), you can transfer files to/from Hydra *without needing VPN*.
- Firefox Send is a two-step process,
 1. you first upload a file (or a set of files packed in an archive) to the Firefox Send server which will generate a unique URL for the upload, and
 2. you download the file using that URL and recover the file's original name.

You can upload up to 1GB at a time, and if you sign up for a Firefox account, that limit increases to 2.5GB.

Example 1: Sending to Hydra

a. Uploading files from your local machine (workstation/laptop) using the [Firefox Send website](#):

1. Open the [Firefox Send website](#) (`send.firefox.com`) from any browser.
2. Choose a file to upload, and optionally:
 - a. use `tar`, `zip` etc. to upload an archive of several files.
 - b. modify the expiration of the file (number of downloads or number of days), the default is to allow only one download and it expires within one day.
 - c. add a password that's needed to download the file.
3. Copy the URL generated for your upload.

⚠ You need to save that unique URL to get that file later. Unlike Dropbox or Google Drive, Firefox Send will not show you what you uploaded.

b. Downloading on Hydra from Firefox Send using `ffsend`

on Hydra

```
$ module load bioinformatics/ffsend
$ ffsend download https://send.firefox.com/download/7800f8272ba5ef7b/#cNSwgMaNqmdsdwG6RxM71A
Download complete
```

Example 2: Sending *from* Hydra

a. Uploading from Hydra using ffsend

on Hydra

```
$ module load bioinformatics/ffsend
$ ffsend upload test.tar.gz
Upload
complete

https://send.firefox.com/download/0324d02485dc9a02/#cxER28yNyf2dcwzwlIla6g
```

b. Downloading to your local machine (workstation/laptop) Optional: `ffsend` has options for setting a password and expiration. See `ffsend help` for more information.

Open the URL created on Hydra in a web browser to download the file to your local machine.

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

[Updated 20 Dec 2019]

	Maximum	Quotas per user			
	disk	disk space	no. of files	Snapshots	
Disk name	capacity	soft/hard	soft/hard	enabled?	Purpose
/home	20T*	100/200G	3.6/4M	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.
/data/sao or /data/nasm	45T	1.9/2.0T	4.8/5M	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.
/data/genomics	45T	0.8/1.0T	2.4/2.5M	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 non-SAO/NASM users.
/pool/sao or /pool/nasm	80T	1.9/2.0T	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.

/pool/genomics	80T	1.9/2.0T	4.8/5M	no	For the bulk of your storage - your limit is high, and disk space is released right away, for non-SAO users.
/pool/biology	200G	100/200G	0.45/0.5M	no	For the bulk of your storage - your limit is high, and disk space is released right away, for non-SAO/NASM users.
/scratch/genomics	350T	9/10T	25/26M	no	For temporary storage, if you need more than what you can keep in /pool for non-SAO/NASM users
/scratch/sao /scratch/nasm	350T	9/10T	25/26M	no	For temporary storage, if you need more than what you can keep in /pool for SAO, NASM users
					Project specific disks (/pool)
/pool/kistlerl	21T	20.0/21.0T	49.9/52.5M	no	NMNH/Logan Kistler
/pool/kozakk	11T	10.5/11.0T	26.1/27.5M	no	STRI/Krzysztof Kozak
/pool/nmnh_ggi	21T	15.0/15.8T	37.4/39.4M	no	NMNH/GGI
/pool/sao_access	21T	15.0/15.8T	37.4/39.4M	no	SAO/ACCESS
/pool/sao_rtdc	10T*	2.8/3.0T	2.5/3.0M	no	SAO/RTDC
/pool/sylvain	30T	29/30T	71/75M	no	SAO/Sylvain Korzennik
					Project specific disks (/scratch)
/scratch/bradys	25T	-	-	no	NMNH/Seán Brady/BRADY_LAB
/scratch/usda_sel	25T	24/25T	52M/62M	no	NMNH/Christopher Owen/USDA_SEL
/scratch/nzp_ccg	25T	24/25T	52M/62M	no	NZP/Michael Campana/CCG
/scratch/kistlerl	50T	-	-	no	NMNH/Logan Kistler
/scratch/meyerc	25T	24/25T	52M/62M	no	NMNH/Christopher Meyer
/scratch/nmnh_ggi	25T	24/25T	52M/62M	no	NMNH/GGI
/scratch/nmnh_lab	25T	4/5T	10M/12M	no	NMNH/LAB
/scratch/stri_ap	25T	4/5T	10M/12M	no	STRI/W. Owen McMillan/STRI_AP
/scratch/sao_atmos	186T	98/100T	252M/261M	no	SAO/ATMOS
/scratch/sao_cga	25T	7/8T	18M/20M	no	SAO/CGA
/scratch/sao_tess	50T	36/40T	94M/210M	no	SAO/TESS
/scratch/sylvain	50T	48/50T	115M/128M	no	SAO/Sylvain Korzennik
/scratch/schultz	25T	-	-	no	NMNH/Ted Schultz/SCHULTZ_LAB
/scratch/wrbu	40T	38/40T	99M/100M	no	WRBU
					Extra
/pool/admin	10T*	5.7/6.0T	14.3/15.0M	no	Sys Admin
/pool/galaxy	15T*	10.7/11.3T	26.7/28.1M	no	Galaxy
					Near line (/store)
/store/public	270T	5/5T	n/a	yes: 8 weeks	Public, available upon request
/store/admin	20T	-	n/a	yes: 8 weeks	Sys Admin
/store/bradys	40T	-	n/a	yes: 8 weeks	NMNH/Seán Brady/BRADY_LAB
/store/nmnh_ggi	40T	-	n/a	yes: 8 weeks	NMNH/GGI
/store/sao_atmos	300TB	-	n/a	yes: 8 weeks	SAO/ATMOS
/store/sylvain	100TB	-	n/a	yes: 8 weeks	SAO/Sylvain Korzennik
/store/schultz	40TB	-	n/a	yes: 8 weeks	NMNH/Ted Schultz/SCHULTZ_LAB
/store/wrbu	40TB	-	n/a	yes: 8 weeks	WRBU
	*: maximum size, disk size will increase up to that value if/when usage grows				

(as of Nov 2019)

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 `tools/local` or `tools/local+` modules)

<code>dus-report</code>	run <code>du</code> and parse its output in a more user friendly format
<code>disk-usage</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`)
 - or
 - [here](#) (at `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
```

or try

```
% df -h --output=source,fstype,size,used,avail,pcent,file /scratch/genomics
Filesystem      Type  Size  Used Avail Use% File
gpfs01          gpfs  400T   95T  306T   24% /scratch/genomics
```

You can compile the output of `du` into a more useful report with the `dus-report` 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 /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` with different options on the same intermediate file, like

```
% dus-report -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 -help
```

to see how else you can use it.

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

```

% disk-usage -d all+
Filesystem              Size      Used      Avail Capacity  Mounted on
netapp-n1:/vol_home    6.40T    3.05T    3.35T   48%/38%  /home
netapp-n2:/vol_data_genomics 36.00T   4.83T   31.17T  14%/2%   /data/genomics
netapp-n2:/vol_data/sao 27.00T   8.65T   18.35T  33%/19%  /data/sao
netapp-n2:/vol_data/nasm 27.00T   8.65T   18.35T  33%/19%  /data/nasm
netapp-n2:/vol_data/admin 27.00T   8.65T   18.35T  33%/19%  /data/admin
netapp-n1:/vol_pool_bio 200.00G  30.25G  169.75G  16%/1%   /pool/biology
netapp-n2:/vol_pool_genomics 55.00T  37.98T  17.02T  70%/15%  /pool/genomics
netapp-n1:/vol_pool_sao 37.00T   7.68T  29.32T  21%/1%   /pool/sao
netapp-n1:/vol_pool_sao/nasm 37.00T   7.68T  29.32T  21%/1%   /pool/nasm
emc-isilon:/ifs/nfs/hydra 60.00T  39.82T  20.18T  67%/1%   /pool/isilon
gpfs01:genomics       400.00T  94.60T  305.40T  24%/9%   /scratch/genomics
gpfs01:sao            400.00T   5.04T  394.96T   2%/1%   /scratch/sao
netapp-n1:/vol_pool_kistler1 21.00T  18.50T   2.50T  89%/1%   /pool/kistler1
netapp-n2:/vol_pool_kozakk 11.00T   7.82T   3.18T  72%/1%   /pool/kozakk
netapp-n1:/vol_pool_nmnh_ggi 21.00T  14.79T   6.21T  71%/8%   /pool/nmnh_ggi
netapp-n1:/vol_pool_sao_access 21.00T   2.37T  18.63T  12%/2%   /pool/sao_access
netapp-n2:/vol_pool_sao_rtdc  2.00T   62.13G   1.94T   4%/1%   /pool/sao_rtdc
netapp-n1:/vol_pool_sylvain 30.00T  24.83T   5.17T  83%/36%  /pool/sylvain
gpfs01:nmnh_bradys    25.00T  58.71G  24.94T   1%/1%   /scratch/bradys
gpfs01:usda_sel       25.00T  651.81G  24.36T   3%/4%   /scratch/usda_sel
gpfs01:nzp_ccg        25.00T  924.33G  24.10T   4%/1%   /scratch/nzp_ccg
gpfs01:nmnh_kistler1  50.00T  11.93T  38.07T  24%/1%   /scratch/kistler1
gpfs01:nmnh_meyerc    25.00T   0.00G  25.00T   0%/1%   /scratch/meyerc
gpfs01:nmnh_ggi       25.00T   4.85T  20.15T  20%/1%   /scratch/nmnh_ggi
gpfs01:nmnh_lab       25.00T   0.00G  25.00T   0%/1%   /scratch/nmnh_lab
gpfs01:stri_ap        25.00T   0.00G  25.00T   0%/1%   /scratch/stri_ap
gpfs01:sao_atmos      186.00T  51.15T  134.85T  28%/6%   /scratch/sao_atmos
gpfs01:sao_cga        25.00T   8.14T  16.86T  33%/4%   /scratch/sao_cga
gpfs01:sao_tess       50.00T   3.29T  46.71T   7%/4%   /scratch/sao_tess
gpfs01:sao_sylvain    50.00T   6.63T  43.37T  14%/2%   /scratch/sylvain
gpfs01:nmnh_schultzt  25.00T  376.87G  24.63T   2%/3%   /scratch/schultzt
gpfs01:wrbu           40.00T   3.00T  37.00T   8%/1%   /scratch/wrbu
netapp-n1:/vol_pool_admin  3.92T   2.71T   1.21T  70%/5%   /pool/admin
netapp-n1:/vol_pool_galaxy 400.00G  194.15G  205.85G  49%/1%   /pool/galaxy
gpfs01:admin          20.00T   1.96T  18.04T  10%/21%  /scratch/admin
gpfs01:bioinformatics_dbs 10.00T  868.14G   9.15T   9%/1%   /scratch/dbs
nas:/mnt/pool_01/admin 20.00T   1.67T  18.33T   9%/1%   /store/admin
nas:/mnt/pool_02/nmnh_bradys 40.00T  306.52G  39.70T   1%/1%   /store/bradys
nas:/mnt/pool_02/nmnh_ggi 40.00T  22.09T  17.91T  56%/1%   /store/nmnh_ggi
nas:/mnt/pool_03/public 270.00T  22.55T  247.45T   9%/1%   /store/public
nas:/mnt/pool_01/sao_atmos 299.97T  68.73T  231.24T  23%/1%   /store/sao_atmos
nas:/mnt/pool_01/sao_sylvain 100.00T   8.39T  91.61T   9%/1%   /store/sylvain
nas:/mnt/pool_02/nmnh_schultzt 40.00T   2.49T  37.51T   7%/1%   /store/schultzt
nas:/mnt/pool_02/wrbu  40.00T  618.24G  39.40T   2%/1%   /store/wrbu

```

Use

```
% disk-usage -help
```

to see how else to use it.

You can, for instance, get the disk quotas and the max size, for all the disks, including `/store`, with:

Monitoring Quota Usage

The Linux command `quota` is working with the NetApp (`/home`, `/data` & `/pool`), but not on the GPFS (`/scratch`) or the NAS (`/store`).

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](#))

🔧 The command `quota+` (need to load `tools/local`) return disk quota for all the disks (see the [quota+](#) section in [Additional Tool](#)).

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.txt`, located in `/data/sao/hpc/quota-reports/unified/`.
 - 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):

- `quota+` - show quota values
- `parse-disk-quota-reports` - parse quota reports

Examples

- `quota+` - show quota values:

```
% quota+
Disk quotas for user sylvain (uid 10541):
Mounted on          Used    Quota  Limit  Grace  Files  Quota  Limit  Grace
-----
/home                11.00G  50.00G  100.0G    0  73.13k  2.00M  2.00M    0
/data/sao            1.92T   7.60T   8.00T     0  37.53M  78.00M  80.00M    0
/pool/sylvain        8.79T  12.50T  14.00T     0  57.93M  71.00M  75.00M    0
/scratch/sao        10.00G  11.00T  12.00T     0     2    25.17M  26.21M    0
/scratch/sylvain    6.63T  50.00T  50.00T     0   1.89M  99.61M  104.9M    0
/store/admin         1.00G   none    none
/store/sylvain       8.39T   none    none
```

Use `quota+ -h`, or read the man page ([man quota+](#)), for the complete usage info.

- `parse-disk-quota-reports` will parse the disk quota report file and produce a more concise report:

```
% parse-disk-quota-reports
Disk quota report: show usage above 85% of quota, (warning when quota > 95%), as of Wed Nov 20 21:00:05 2019.
```

```
Volume=NetApp:vol_data_genomics, mounted as /data/genomics
-- disk -- -- #files -- default quota: 512.0GB/1.25M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/data/genomics 512.0GB 100.0% 0.17M 13.4% *** Paul Frandsen, OCIO - frandsenp
```

```
Volume=NetApp:vol_data_sao, mounted as /data/admin or /data/nasm or /data/sao
-- disk -- -- #files -- default quota: 2.00TB/5M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/data/admin:nasm:sao 1.88TB 94.0% 0.01M 0.1% uid=11599
```

```
Volume=NetApp:vol_home, mounted as /home
-- disk -- -- #files -- default quota: 100.0GB/2M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/home 96.5GB 96.5% 0.41M 20.4% *** Roman Kochanov, SAO/AMP - rkochanov
/home 96.3GB 96.3% 0.12M 6.2% *** Sofia Moschou, SAO/HEA - smoschou
/home 95.2GB 95.2% 0.11M 5.6% *** Cheryl Lewis Ames, NMNH/IZ - amesc
/home 95.2GB 95.2% 0.26M 12.8% *** Yanjun (George) Zhou, SAO/SSP - yjzhou
/home 92.2GB 92.2% 0.80M 40.1% Taylor Hains, NMNH/VZ - hainst
```

```
Volume=NetApp:vol_pool_genomics, mounted as /pool/genomics
-- disk -- -- #files -- default quota: 2.00TB/5M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/pool/genomics 1.71TB 85.5% 1.23M 24.6% Vanessa Gonzalez, NMNH/LAB - gonzalezv
/pool/genomics 1.70TB 85.0% 1.89M 37.8% Ying Meng, NMNH - mengy
/pool/genomics 1.45TB 72.5% 4.56M 91.3% Brett Gonzalez, NMNH - gonzalezb
/pool/genomics 133.9GB 6.5% 4.56M 91.2% Sarah Lemer, NMNH - lemers
```

```
Volume=NetApp:vol_pool_kistler1, mounted as /pool/kistler1
-- disk -- -- #files -- default quota: 21.00TB/52M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/pool/kistler1 18.35TB 87.4% 0.88M 1.7% Logan Kistler, NMNH/Anthropology - kistler1
```

```
Volume=NetApp:vol_pool_nmnh_ggi, mounted as /pool/nmnh_ggi
-- disk -- -- #files -- default quota: 15.75TB/39M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/pool/nmnh_ggi 14.78TB 93.8% 8.31M 21.3% Vanessa Gonzalez, NMNH/LAB - gonzalezv
```

```
Volume=NetApp:vol_pool_sao, mounted as /pool/nasm or /pool/sao
-- disk -- -- #files -- default quota: 2.00TB/5M
Disk usage %quota usage %quota name, affiliation - username (indiv. quota)
-----
/pool/nasm:sao 1.78TB 89.0% 0.16M 3.2% Guo-Xin Chen, SAO/SSP-AMP - gchen
```

reports disk usage when it is above 85% of the quota.

Use `parse-disk-quota-reports -h`, or read the man page (`man parse-disk-quota-reports`). for the complete usage info.

Note

- Users whose quotas are above the 85% 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`).

How to Use the NAS/ZFS Snapshots:

- The snapshots on the `/store` disks are:
 - located under `/store/XXX/.zfs/snapshot` (where `XXX` is, for example, `public`) and
 - in sub-directories named `auto-YYMMDD.0230-8w` where `YYMMDD` represent the date of the snapshot.
- Content of NAS/ZFS snapshots can be recovered as described above.

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
 - non-SAO users: email the file(s) or the location of the files to SI-HPC@si.edu

8. SSD Local Disk Space

- Local SSDs (solid state disks) are available on a few nodes available.
- You should [contact us](#) if your jobs can benefit from accessing local SSD.
- How to use the SSD is explained [here](#).