Large Storage at SciNet: Demystifying HPSS

SNUG TechTalk

February 13, 2013

SciNet
compute • calcul
CANADA

Outline

- Storage overview
- HPSS overview from user perspective
- Demonstration

SciNet compute • calcul CANADA

High Performance Storage System

- The High Performance Storage System (HPSS) is a tape-backed hierarchical storage system.
- HPSS provides a significant portion of the allocated storage space at SciNet.
- It is a repository for archiving data that is not being actively used. Data can be returned to the active GPFS filesystem when it is needed.
- Up to 2TB in HPSS is available for all groups upon request.
- Need more? Request it in the yearly resource allocation round.

SciNet
compute • calcul
CANADA

Storage Perspective of SciNet

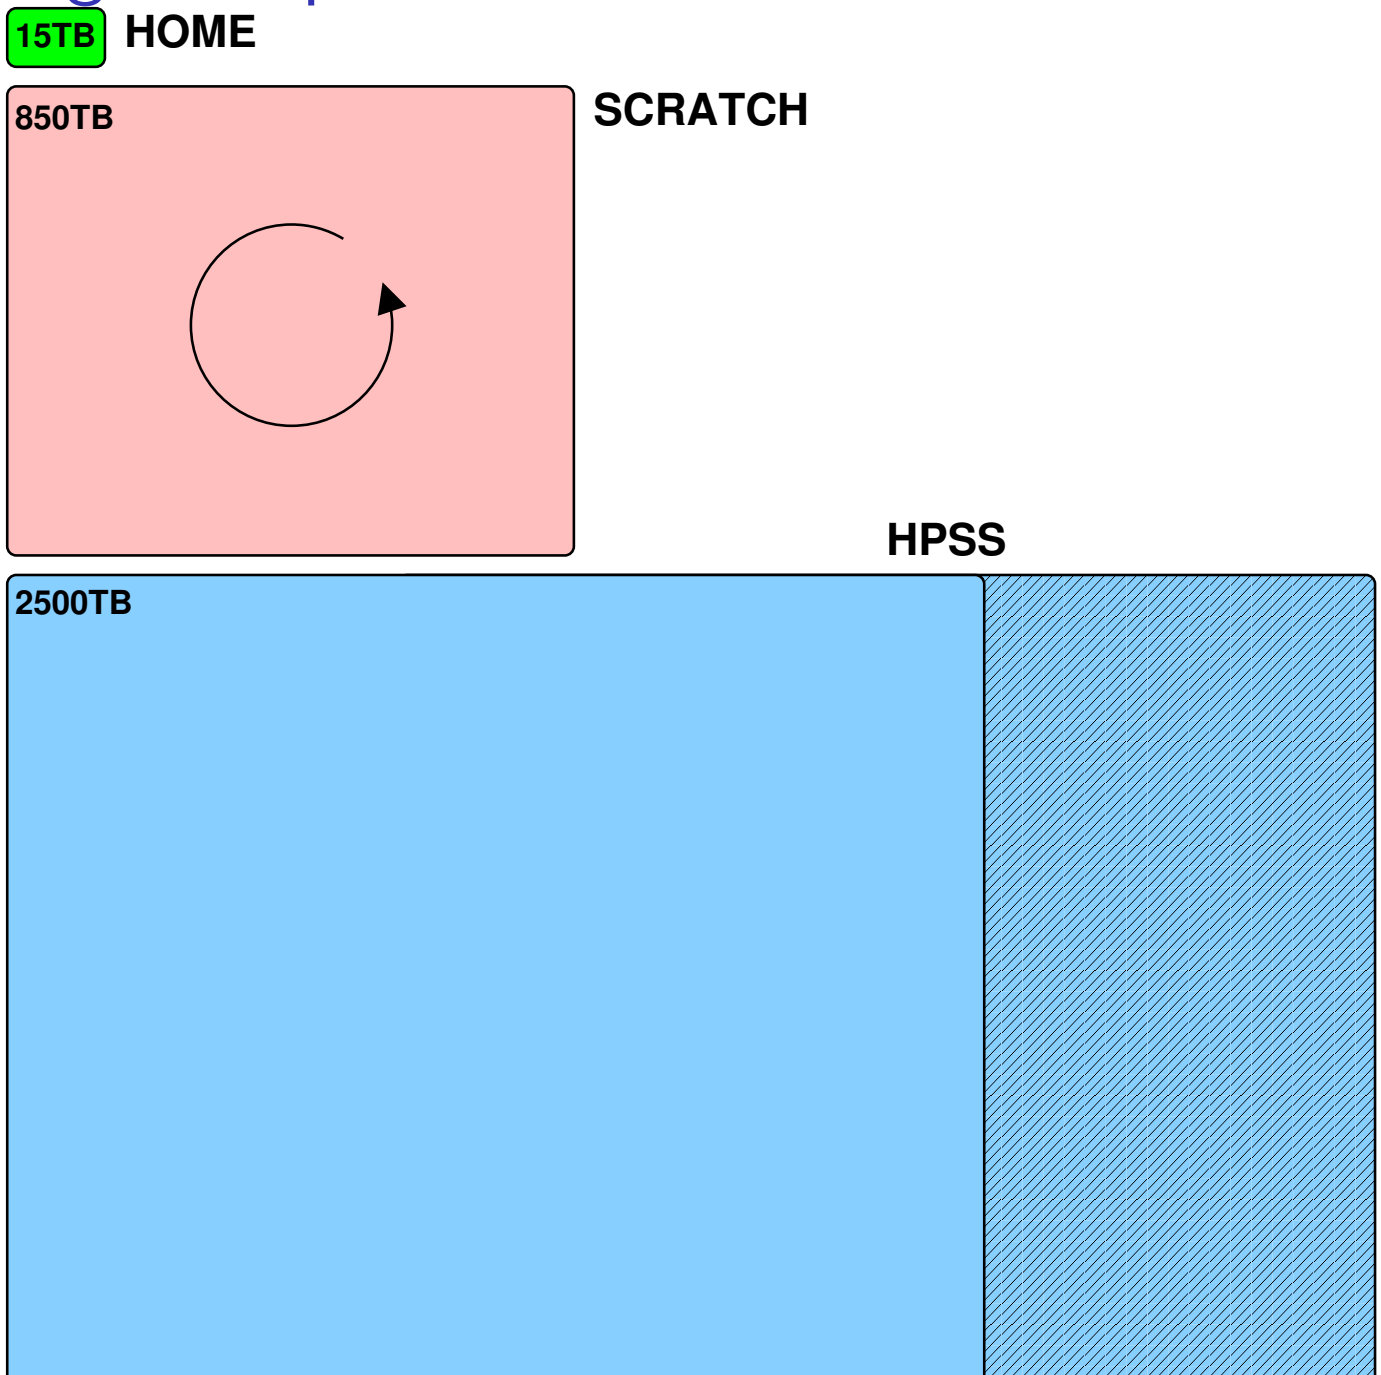

Storage Perspective of SciNet

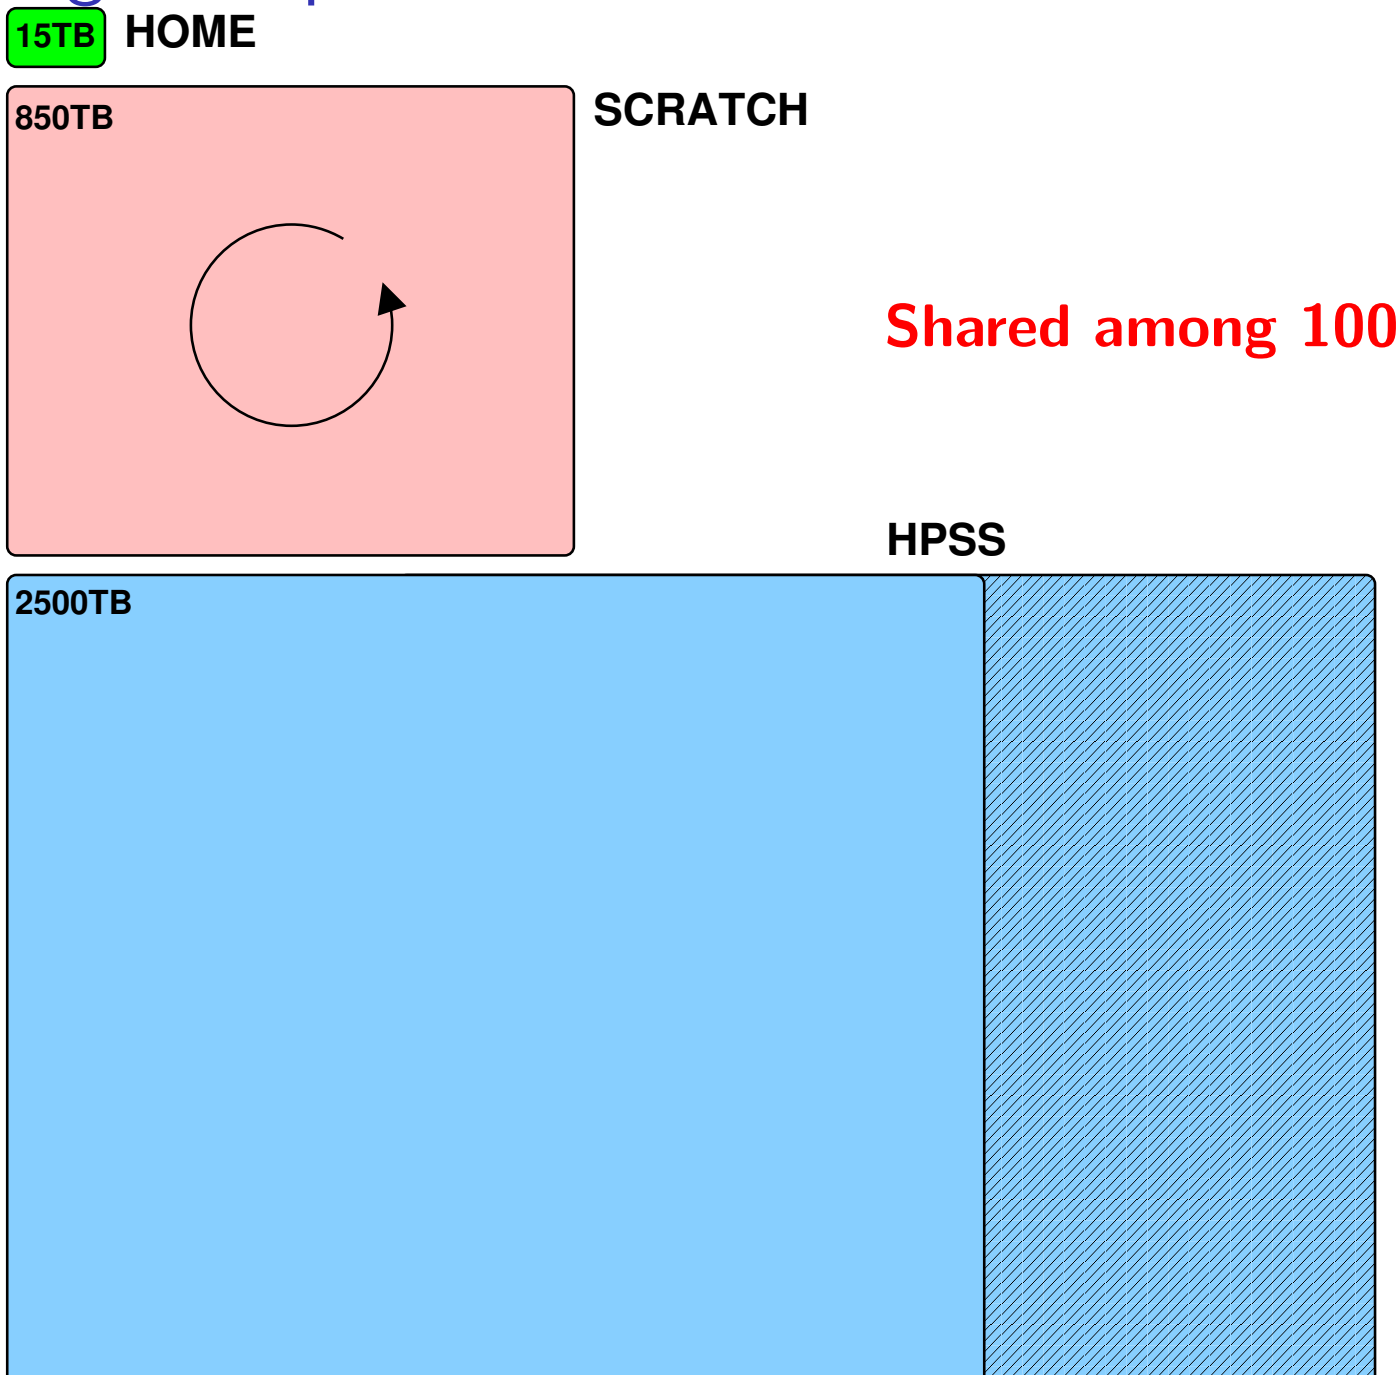

Storage Perspective of SciNet

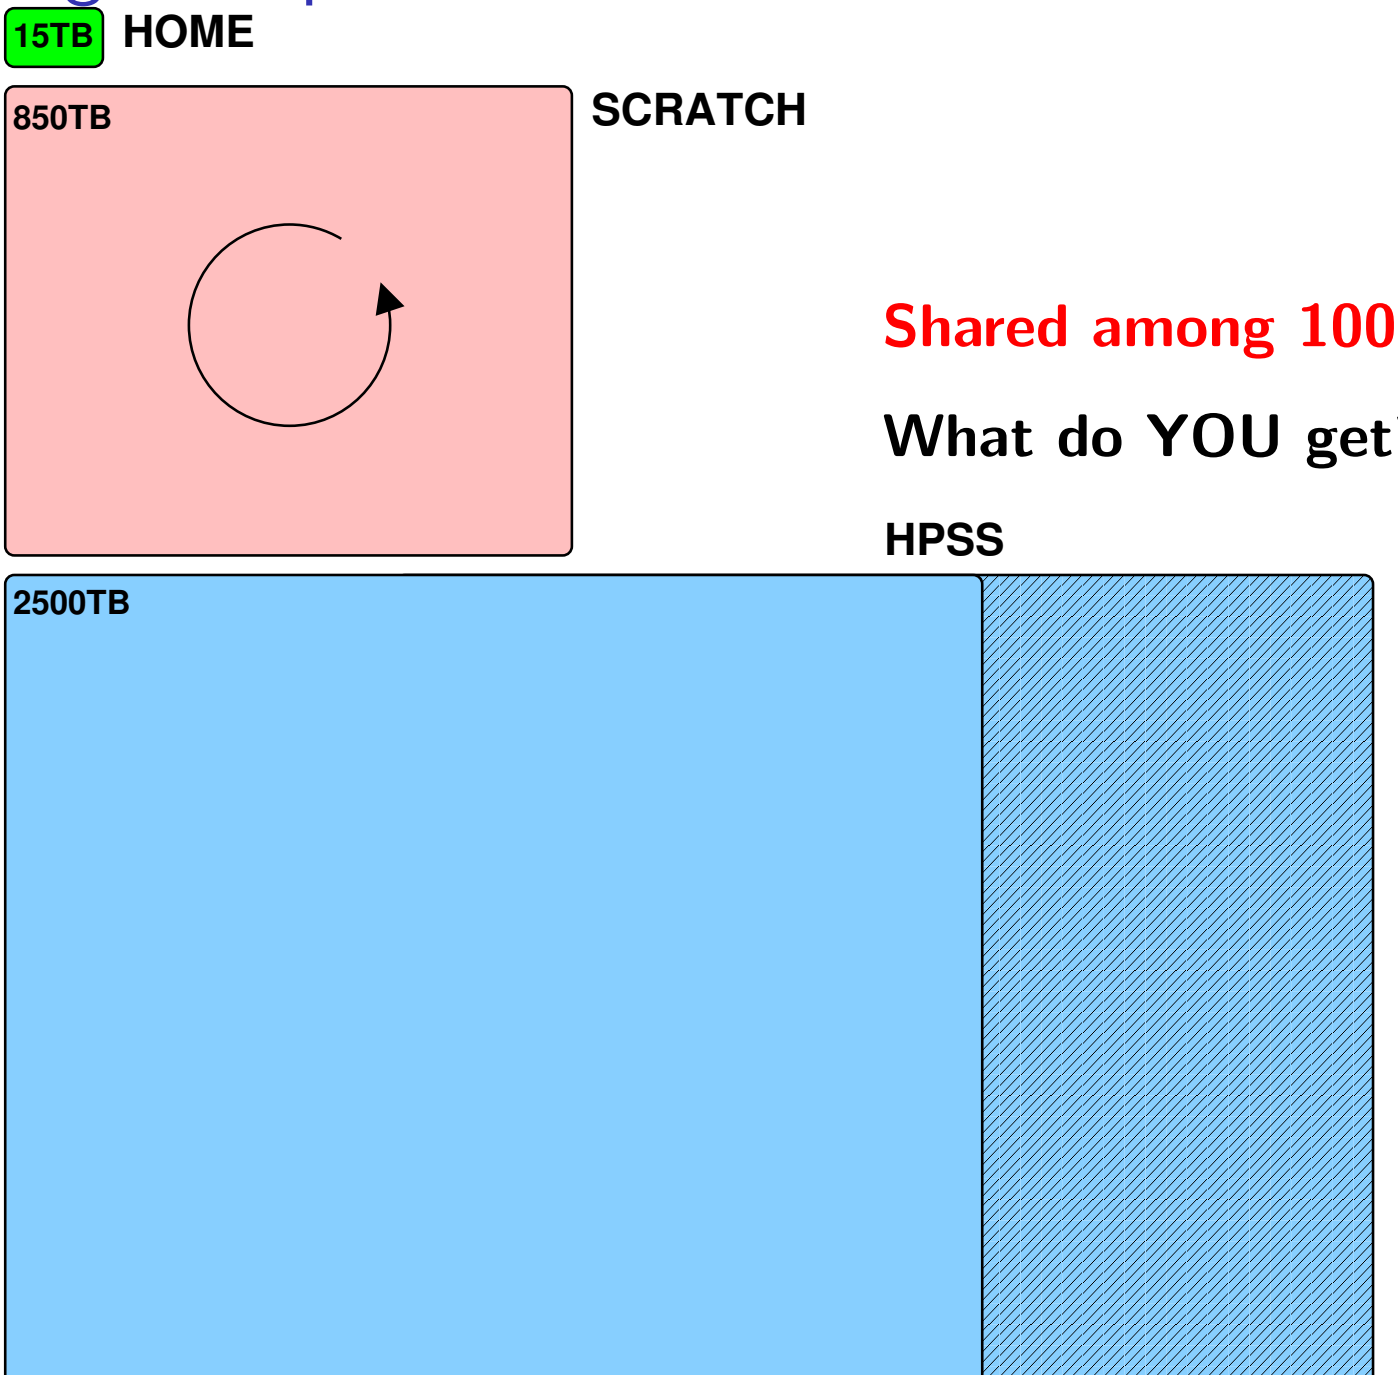

Shared among 1000+ users!

What do YOU get?

Storage Quota on SciNet

SciNet compute • calcul CANADA

Storage Quota on SciNet

Let's look in detail...

SciNet compute • calcul CANADA

Storage Quota on SciNet

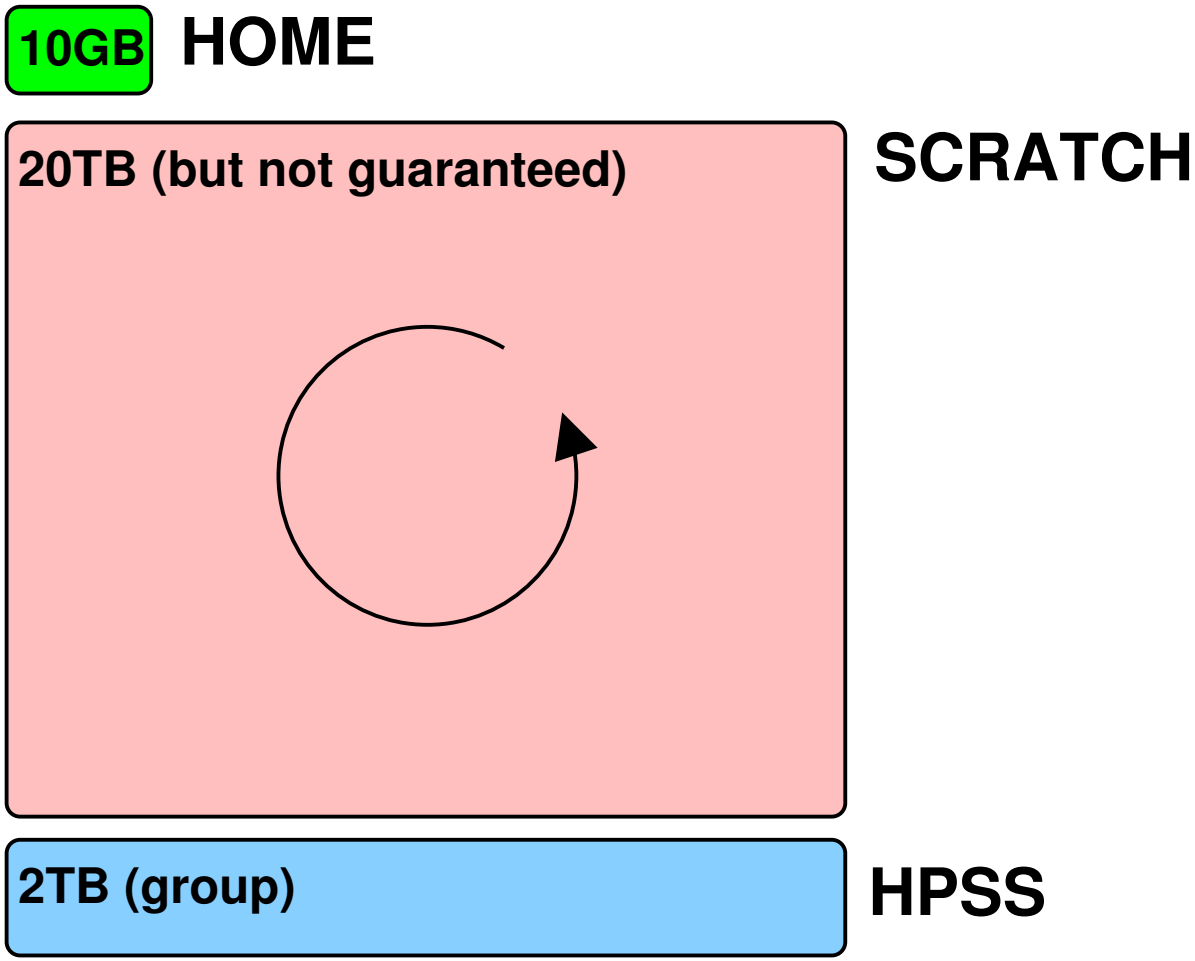

Storage Quota on SciNet

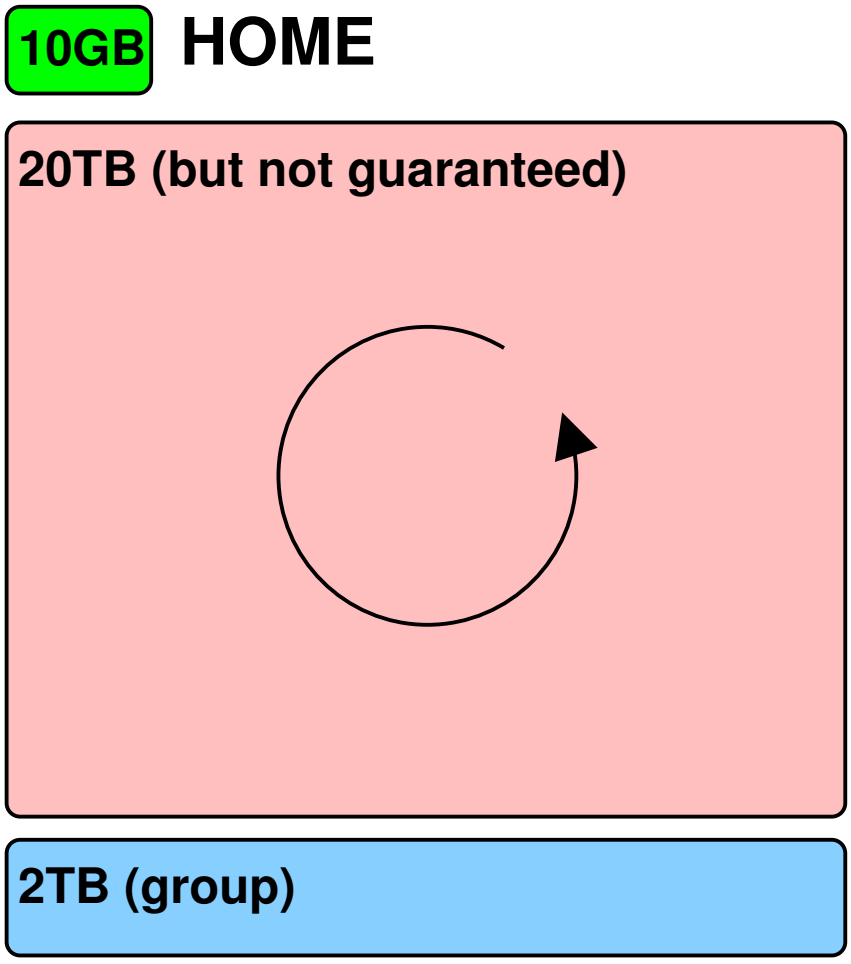

Storage Quota on SciNet

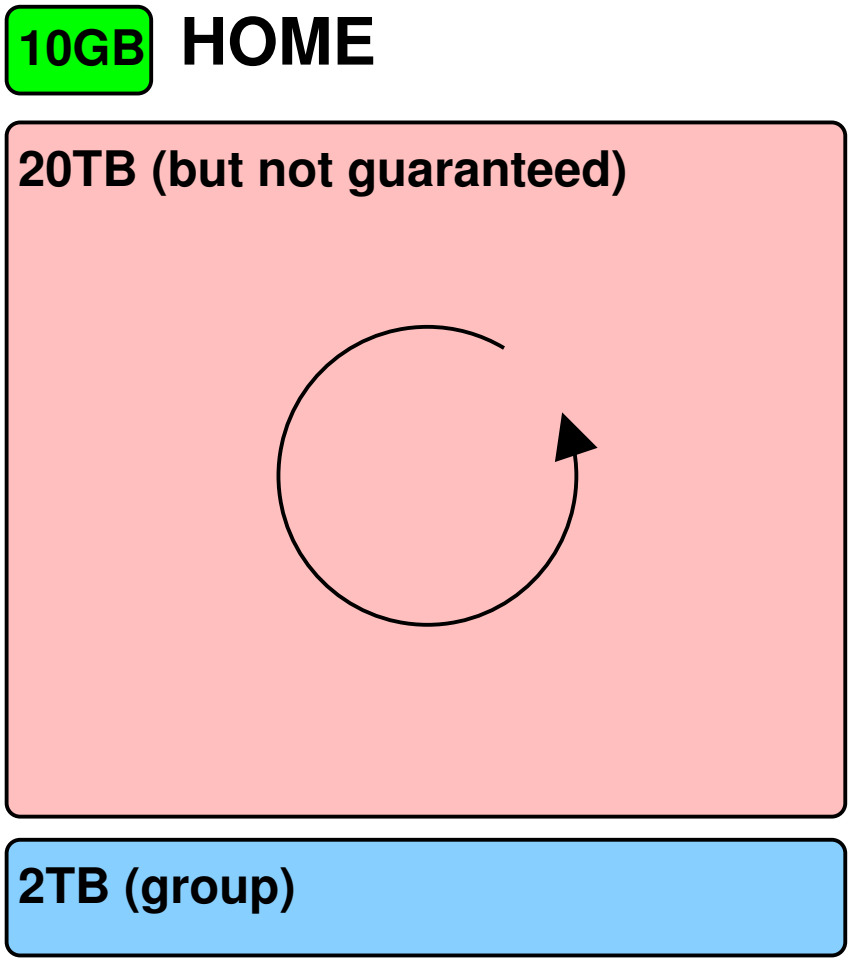

SciNet compute • calcul CANADA

Storage Quota on SciNet

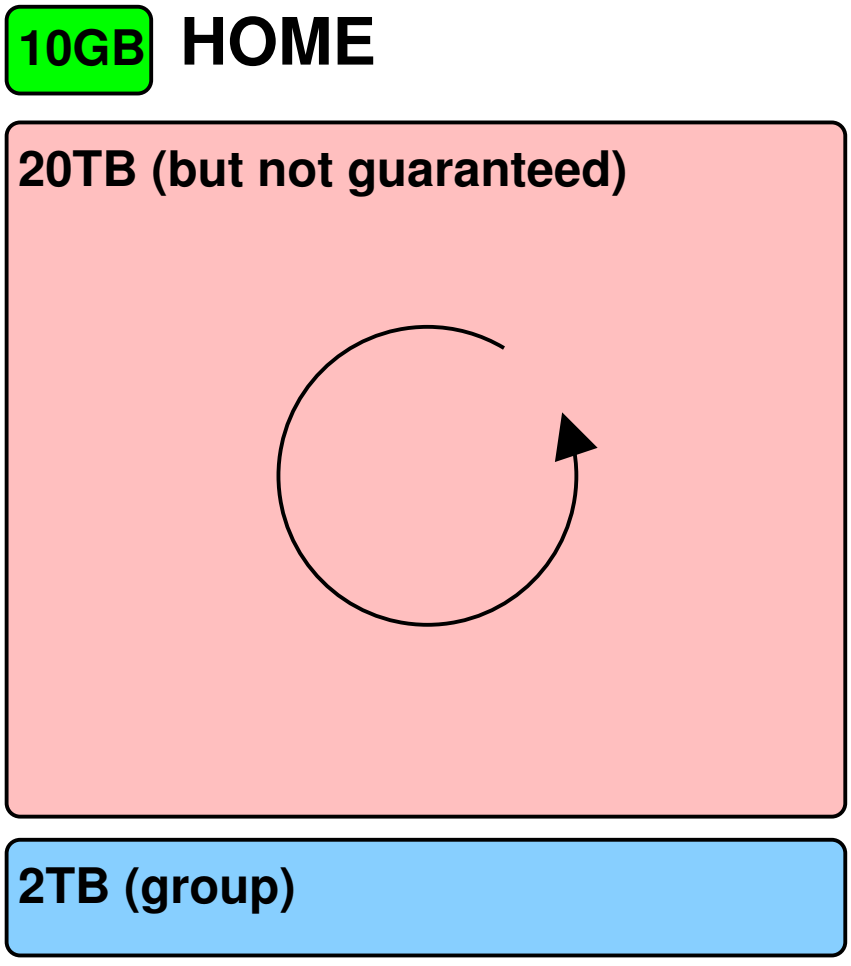

Storage Quota on SciNet

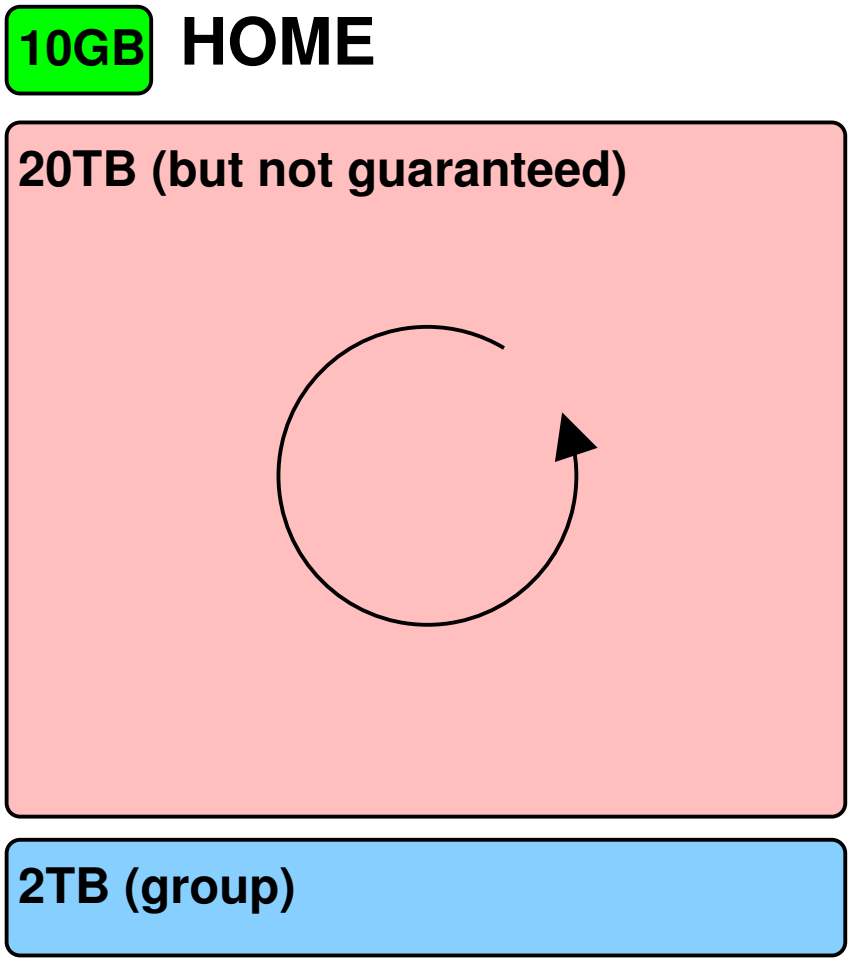

SciNet compute • calcul CANADA

Structure

- HPSS is structured like a file system, with directories, files etc.
- But it is not mounted as a file system!
- 100% POSIX compliant, like in Linux:
  - directory structure
  - namespacing/paths
  - user/group permissions
  - time stamps
  - ...

Utilities

- Access and transfer of data into and out of HPSS using:

  hsi is a client with an ftp-like functionality which can be used to archive and retrieve large files. It is also useful for browsing the contents of HPSS.

  htar is a utility that creates tar formatted archives directly into HPSS. It also creates a separate index file (.idx) that can be accessed and browsed quickly.

  ish is a TUI utility that can perform an inventory of the files and directories in your tarballs.

These utilities are available only on the **hpss-archive01** node.

hsi

- May be the primary client with which some users will interact with HPSS
- Ftp-like interface for archiving and retrieving tarballs or directory trees.
- Provides a number of shell-like commands that are useful for examining and manipulating the contents in HPSS.
- Has built-in checksum algorithm.

htar

- Used for aggregating a set of files and directories from GPFS into HPSS.
- Creates files that conform to the POSIX TAR specification.

  This poses some (soon too be lifted) restrictions:
  - No files larger than 68 GB can be stored.
  - No pathnames with more that 100 characters.
- Does not do gzip compression.
- Has built-in checksum algorithm.

ish

- Index SHell
- Creates offline catalogue of your HPSS content.
- With ish, you can browse this catalogue, using familiar commands such as ls, cd, du and find.
- Ish cannot change your existing files.
- Must update the index by hand from HPSS.
- To see 'inside' tarballs, must index tar files separately.

Two Modes of Operation

1. Batch jobs

2. Interactive sessions of 1 hour

SciNet
compute • calcul
CANADA

Batch jobs

- Archiving and retrieval can take some time
- There is a limited number of tape drives
- Most common mode: submit job script to the archive queue from one of the gpc devel nodes **gpc01-4**:

```
qsub -q archive JOBSCRIPT
```

- The JOBSCRIPT is a file containing script that gets run on the **hpss-archive01** node and has access to hsi, htar and ish.
- Some care should be taken making sure you check for errors.
- Details and examples on the wiki:
  `http://wiki.scinethpc.ca/wiki/index.php/HPSS`

Interactive session

- You may feel more comfortable acquiring an interactive shell, starting an HSI session (especially for deletions).
- Keep in mind, you're restricted to one hour.

```
qsub -I -q archive JOBSCRIPT
```

- puts you on the **hpss-archive01** node.
- Will demonstrate...

SciNet compute • calcul CANADA