C.E.C.I

Consortium des Equipements
de Calcul Intensif
en Fédération Wallonie-Bruxelles

# Introduction to Scientific Data Management

damien.francois@uclouvain.be
November 2020

http://www.ceci-hpc.be/training.html

# Goal of this session:


"Share tools, tips and tricks related to the **storage**, **transfer**, and **sharing** of scientific data"

# 1. Data storage

Filesystems – Object stores – Databases

# Storage Technologies

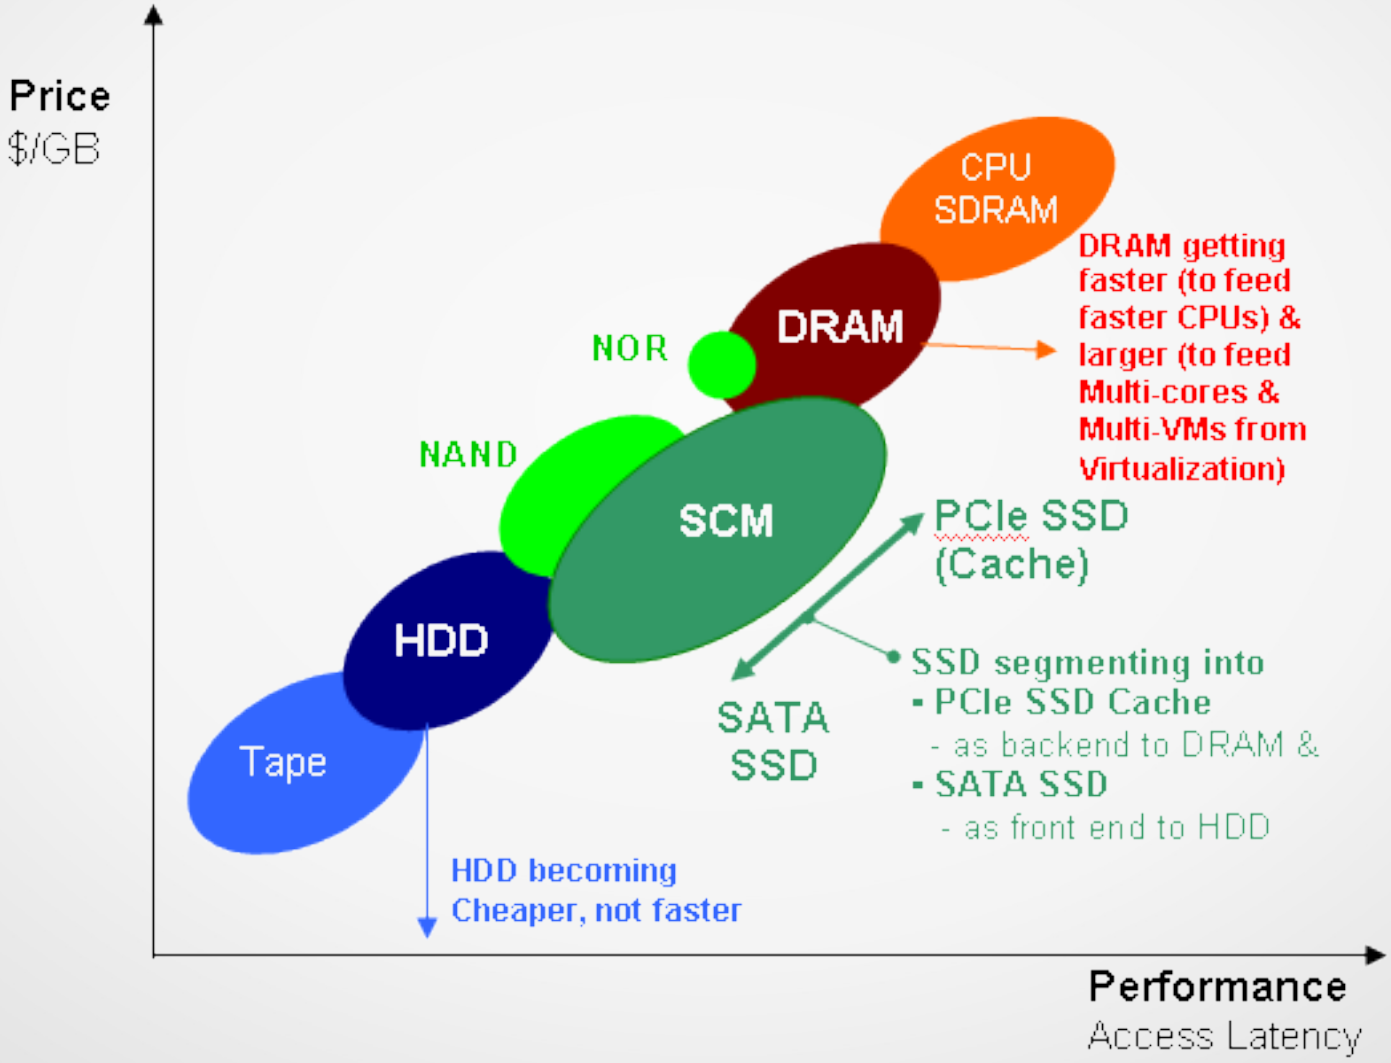




http://www.slideshare.net/IMEXresearch/ss-ds-ready-for-enterprise-cloud

# Storage paradigms

## Filesystem

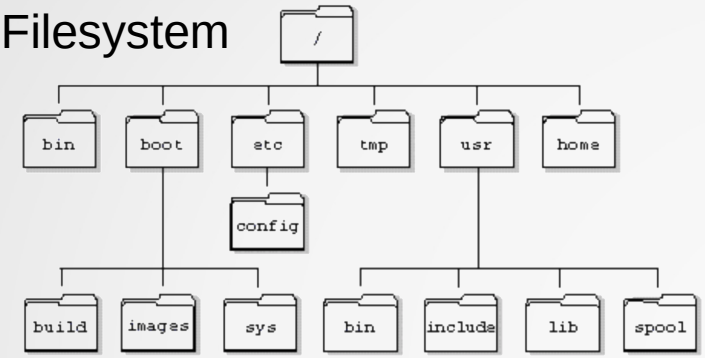


## RDBMS

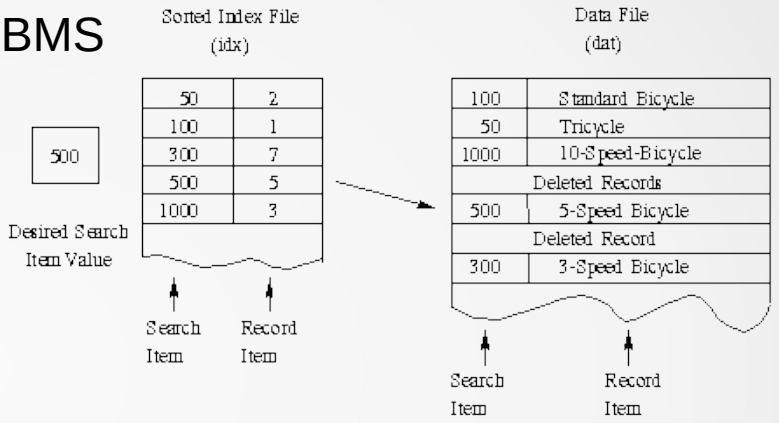


## Objects store

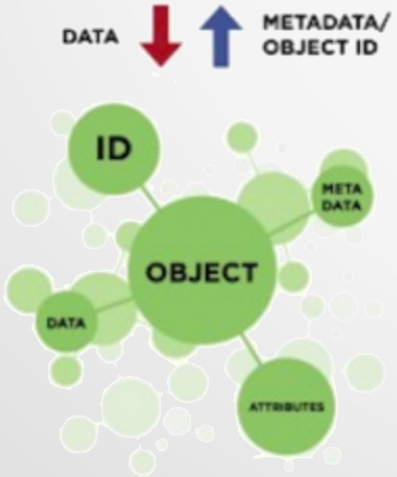


## NoSQL

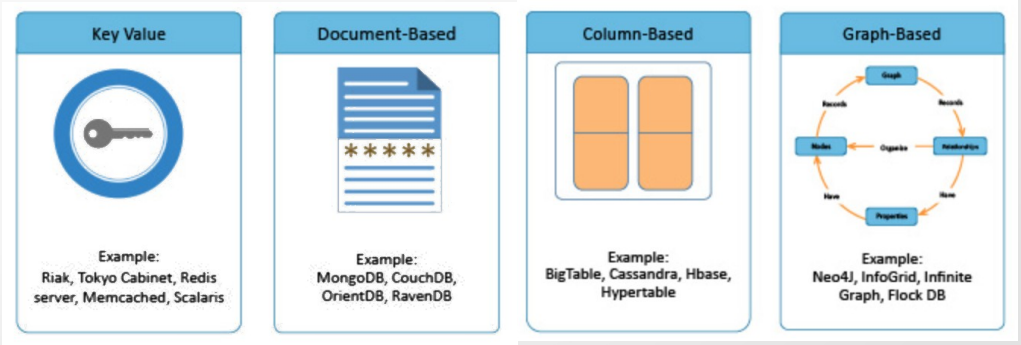

# 1.1 Filesystems

# (local) Filesystems

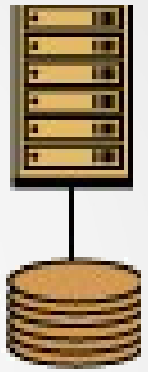

**Generation 0:** No system at all. There was just an arbitrary stream of data. Think punchcards, data on audiocassette, Atari 2600 ROM carts.

**Generation 1:** Early random access. Here, there are multiple named files on one device with no folders or other metadata. Think Apple ][ DOS (but not ProDOS!) as one example.

**Generation 2:** Early organization (aka folders). When devices became capable of holding hundreds of files, better organization became necessary. We're referring to TRS-DOS, Apple //c ProDOS, MS-DOS FAT/FAT32, etc.

**Generation 3:** Metadata—ownership, permissions, etc. As the user count on machines grew higher, the ability to restrict and control access became necessary. This includes AT&T UNIX, Netware, early NTFS, etc.

**Generation 4:** Journaling! This is the killer feature defining all current, modern filesystems—ext4, modern NTFS, UFS2, XFS, you name it. Journaling keeps the filesystem from becoming inconsistent in the event of a crash, making it much less likely that you'll lose data, or even an entire disk, when the power goes off or the kernel crashes.

**Generation 5:** Copy on Write snapshots, Per-block checksumming, Volume management, Far-future scalability, Asynchronous incremental replication, Online compression. Generation 5 filesystems are Btrfs and ZFS.

Typical usage: Operating system & Local scratch space



http://arstechnica.com/information-technology/2014/01/bitrot-and-atomic-cows-inside-next-gen-filesystems/

# Network filesystem

One source many consumers

NAS: ex. NFS  SAN: ex. GFS2

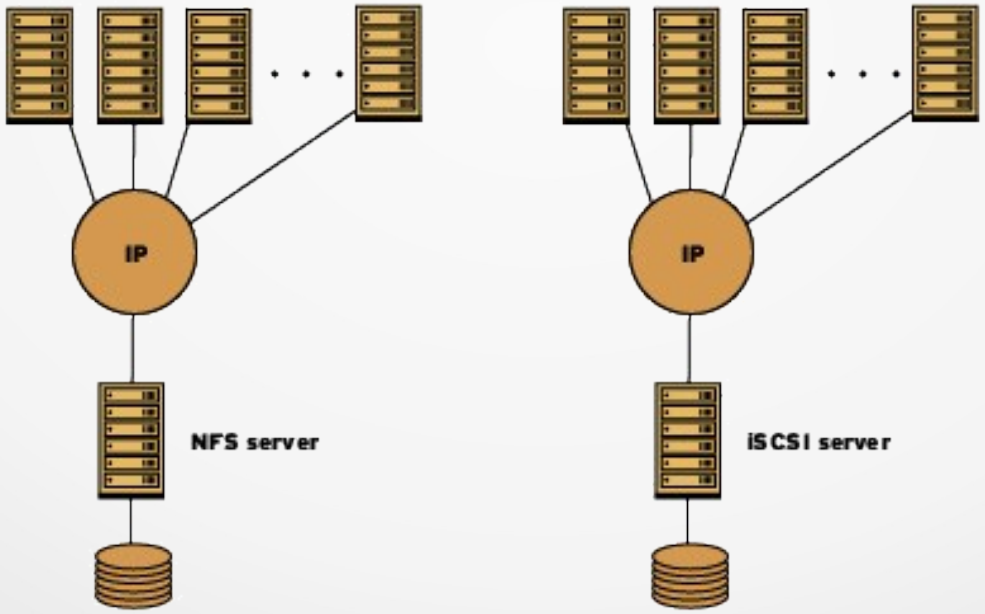


Typical usage: Home directories, Mass storage



Pictures from https://www.redhat.com/magazine/008jun05/features/gfs_nfs/

# Parallel / distributed filesystem

Many sources many consumers
ex: Lustre, GPFS, BeeGFS, GlusterFS

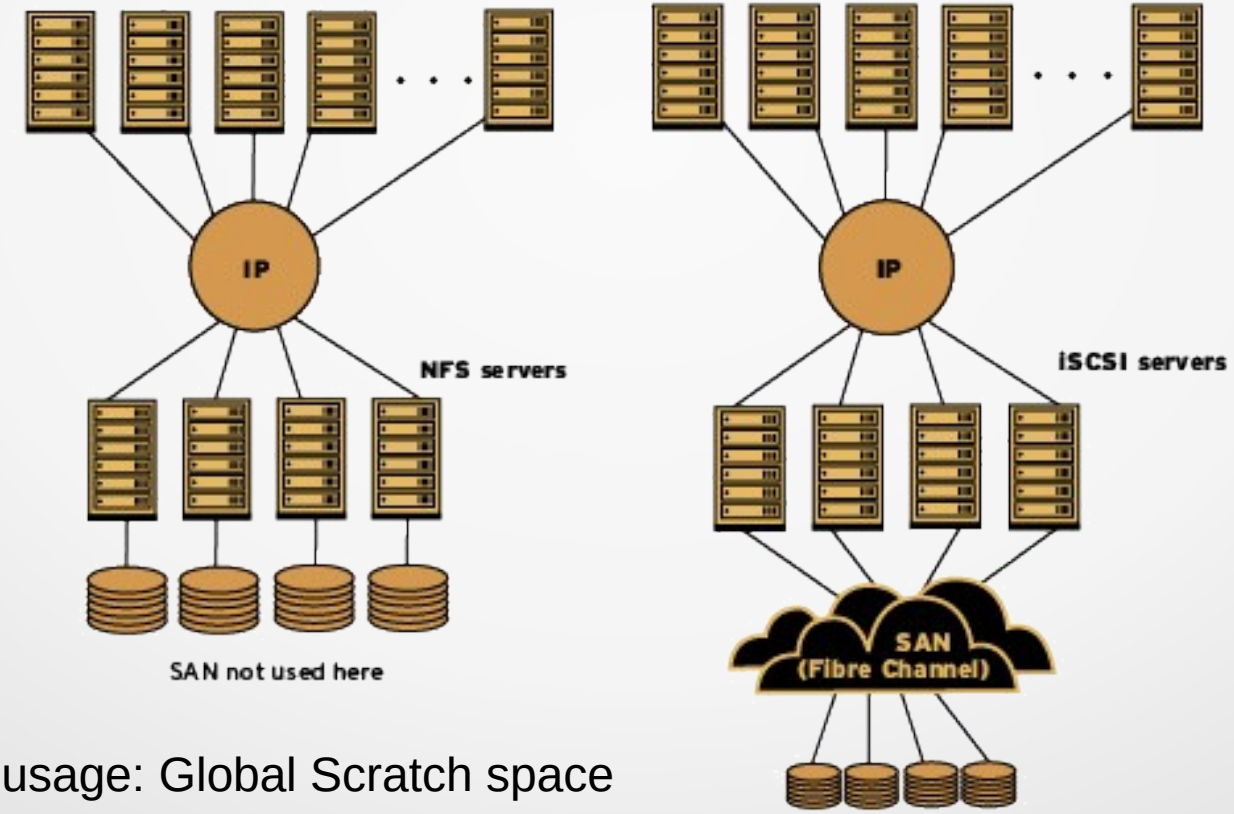


Typical usage: Global Scratch space



Pictures from https://www.redhat.com/magazine/008jun05/features/gfs_nfs/

# Special filesystems – in memory

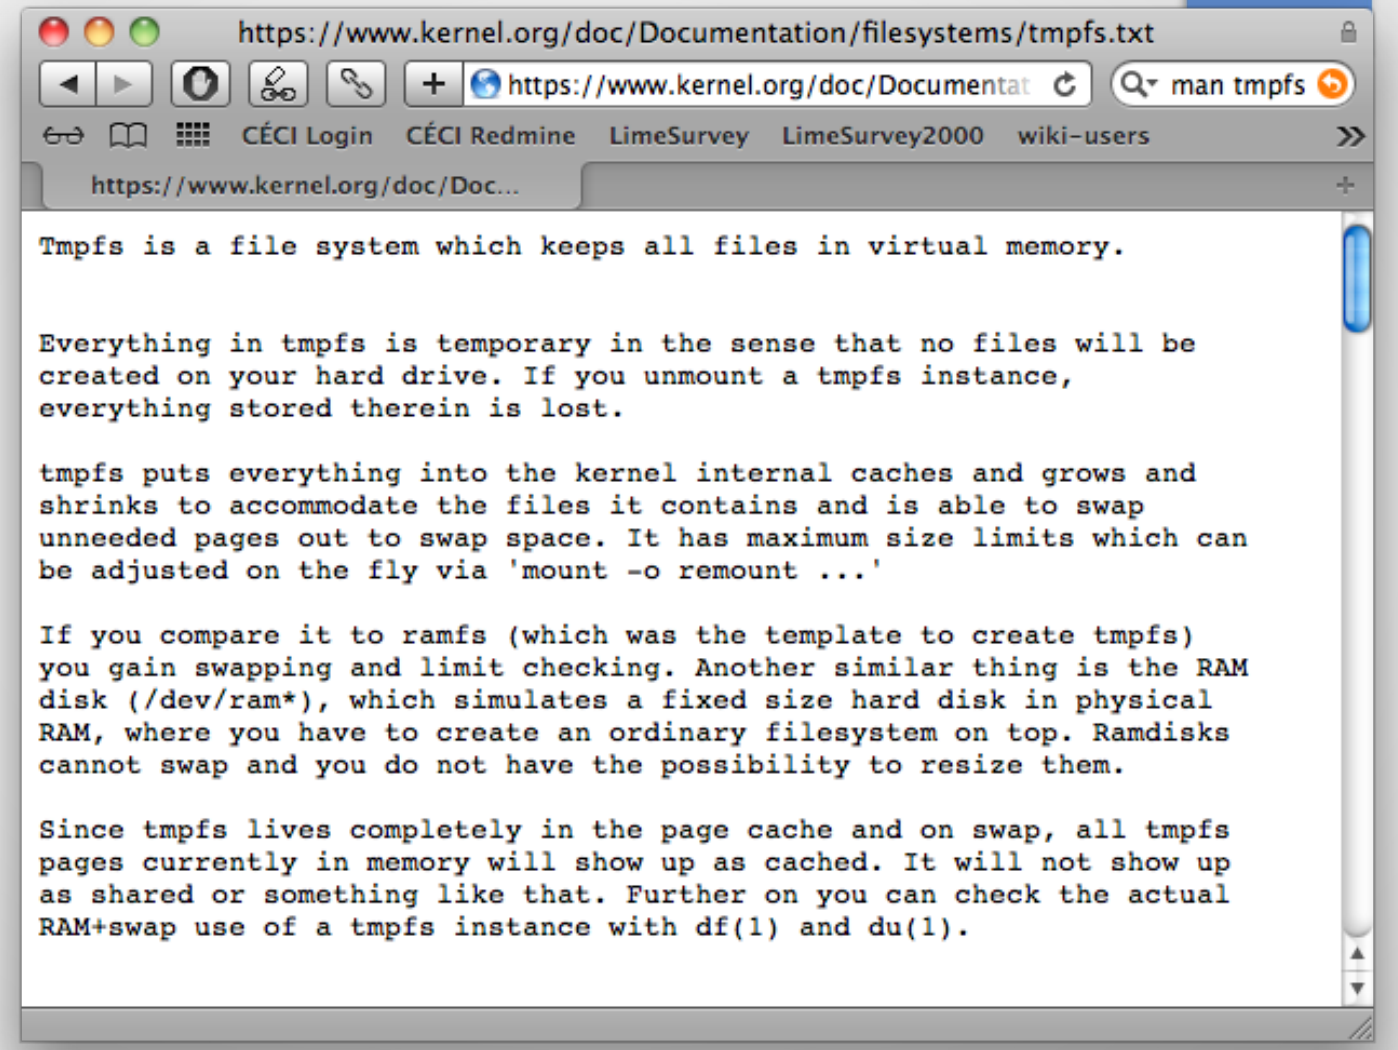


Typical usage: Temporary filesystems

# Filesystems on Lemaitre3

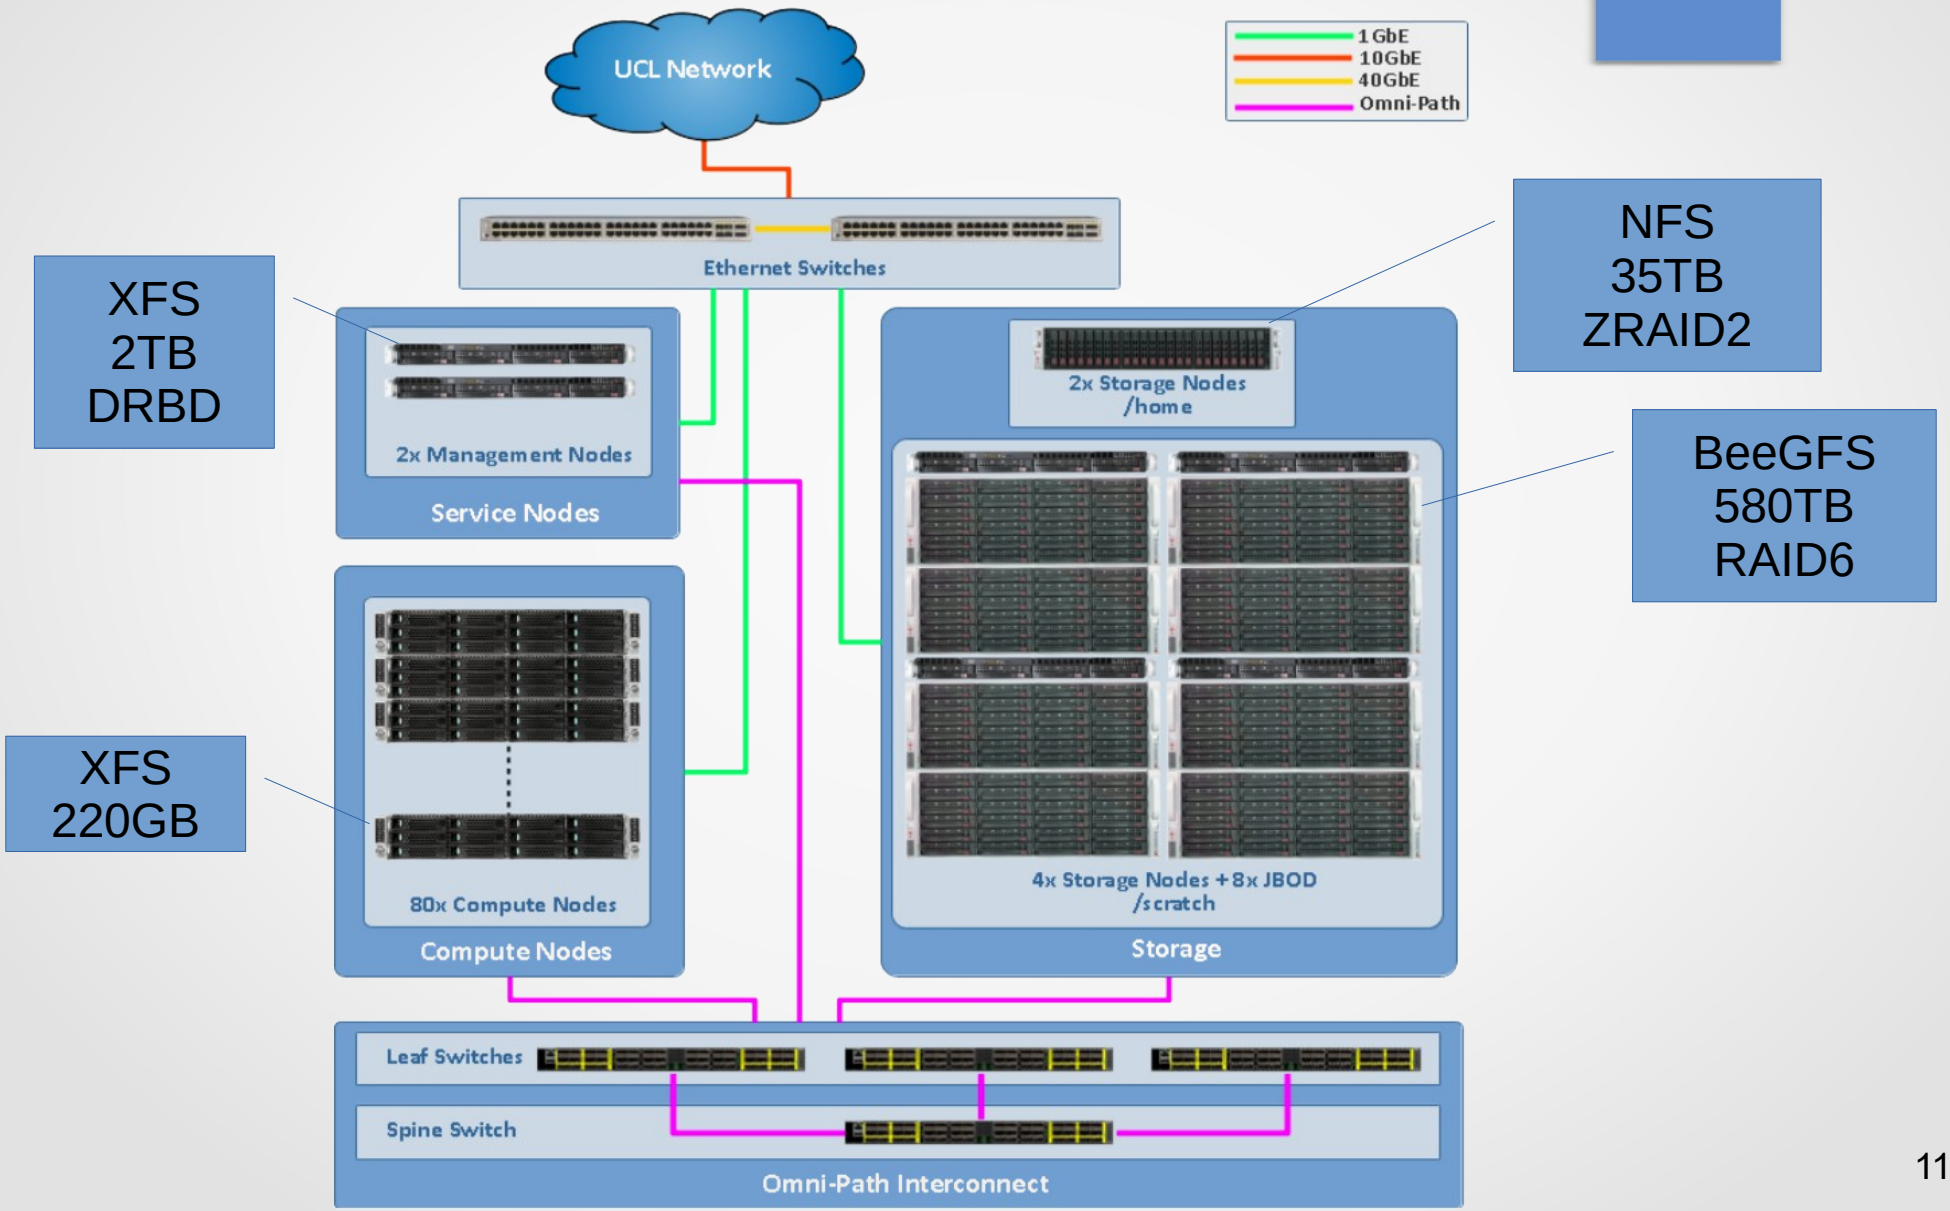

# Filesystems on Lemaitre3

Volume

```
[dfr@lemaitre3 ~]$ df -khT -x tmpfs
Filesystem                                              Type      Size  Used Avail Use% Mounted on
/dev/mapper/centos_controller-root                      ext4      493G  372G   96G  80% /
devtmpfs                                                devtmpfs   23G     0   23G   0% /dev
/dev/sda2                                               ext4      976M  117M  793M  13% /boot
/dev/sda1                                               vfat       50M  9.9M   41M  20% /boot/efi
vsalt:/srv/salt                                         nfs4      9.8G  532M  8.7G   6% /srv/salt
vsalt:/srv/pillar                                       nfs4      9.8G  532M  8.7G   6% /srv/pillar
beegfs_nodev                                            beegfs    583T  203T  380T  35% /scratch
/dev/drbd1                                              xfs       3.2T  162G  3.0T   6% /trinity
gw-ucl:/CECI/gateway/home                               nfs        40T  2.5T   38T   7% /CECI/home
lm3-n:/storage                                          nfs4       35T  9.1T   26T  27% /home
gw-ucl:/CECI/gateway/soft/cecisw/RedHat-7_6-85-4_Omnipath nfs      1.0T  498G  527G  49% /opt/cecisw/arch
gw-ucl:/CECI/gateway/soft                               nfs       1.0T  498G  527G  49% /CECI/soft
```

Source:
- /dev/sd… → local disk
- /dev/mapper… → LVM
- <machine>:<path> → NFS
- other (e.g. beegfs_nodev) →specific filesystem

Local mount point

Inodes

```
[dfr@lemaitre3 ~]$ df -ikhT -x tmpfs
Filesystem                                              Type     Inodes IUsed IFree IUse% Mounted on
/dev/mapper/centos_controller-root                      ext4       32M  816K   31M    3% /
devtmpfs                                                devtmpfs  5.8M   506  5.8M    1% /dev
/dev/sda2                                               ext4       64K    32   64K    1% /boot
/dev/sda1                                               vfat          0     0     0     - /boot/efi
vsalt:/srv/salt                                         nfs4      640K   13K  628K    3% /srv/salt
vsalt:/srv/pillar                                       nfs4      640K   13K  628K    3% /srv/pillar
beegfs_nodev                                            beegfs        0     0     0     - /scratch
/dev/drbd1                                              xfs       320M  681K  319M    1% /trinity
gw-ucl:/CECI/gateway/home                               nfs        98M   14M   85M   14% /CECI/home
lm3-n:/storage                                          nfs4       51G   39M   51G    1% /home
gw-ucl:/CECI/gateway/soft/cecisw/RedHat-7_6-85-4_Omnipath nfs      98M   14M   85M   14% /opt/cecisw/arch
gw-ucl:/CECI/gateway/soft/cecisw/noarch                 nfs4       98M   14M   85M   14% /opt/cecisw/noarc
gw-ucl:/CECI/gateway/soft                               nfs        98M   14M   85M   14% /CECI/soft
```

# A word about inodes (simplified)

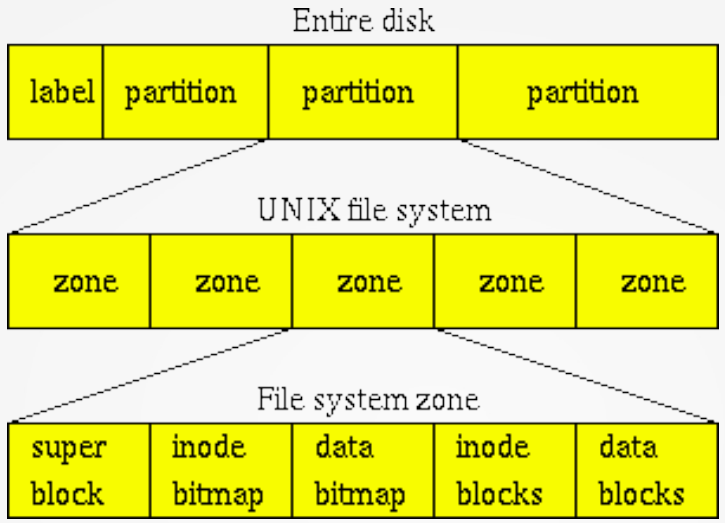


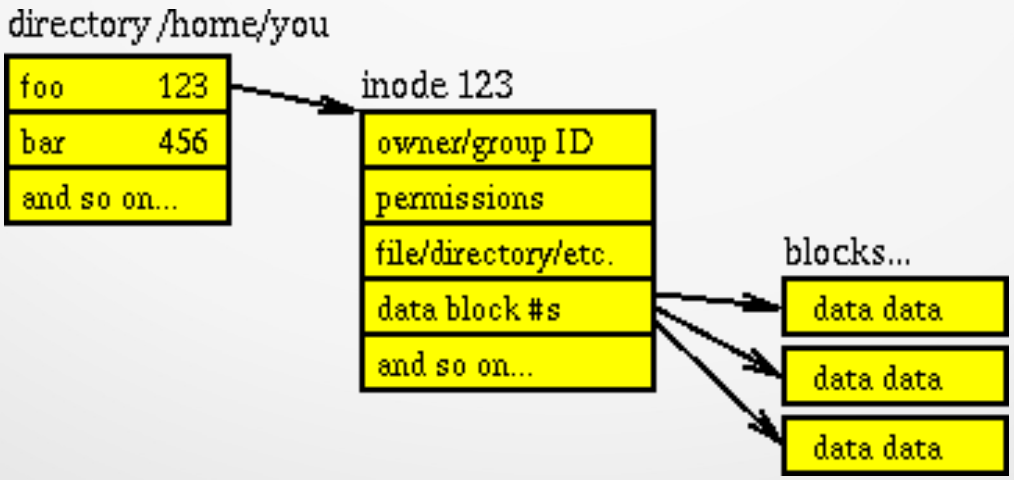




http://www.porcupine.org/forensics/forensic-discovery/chapter3.html

# What filesystem for what usage

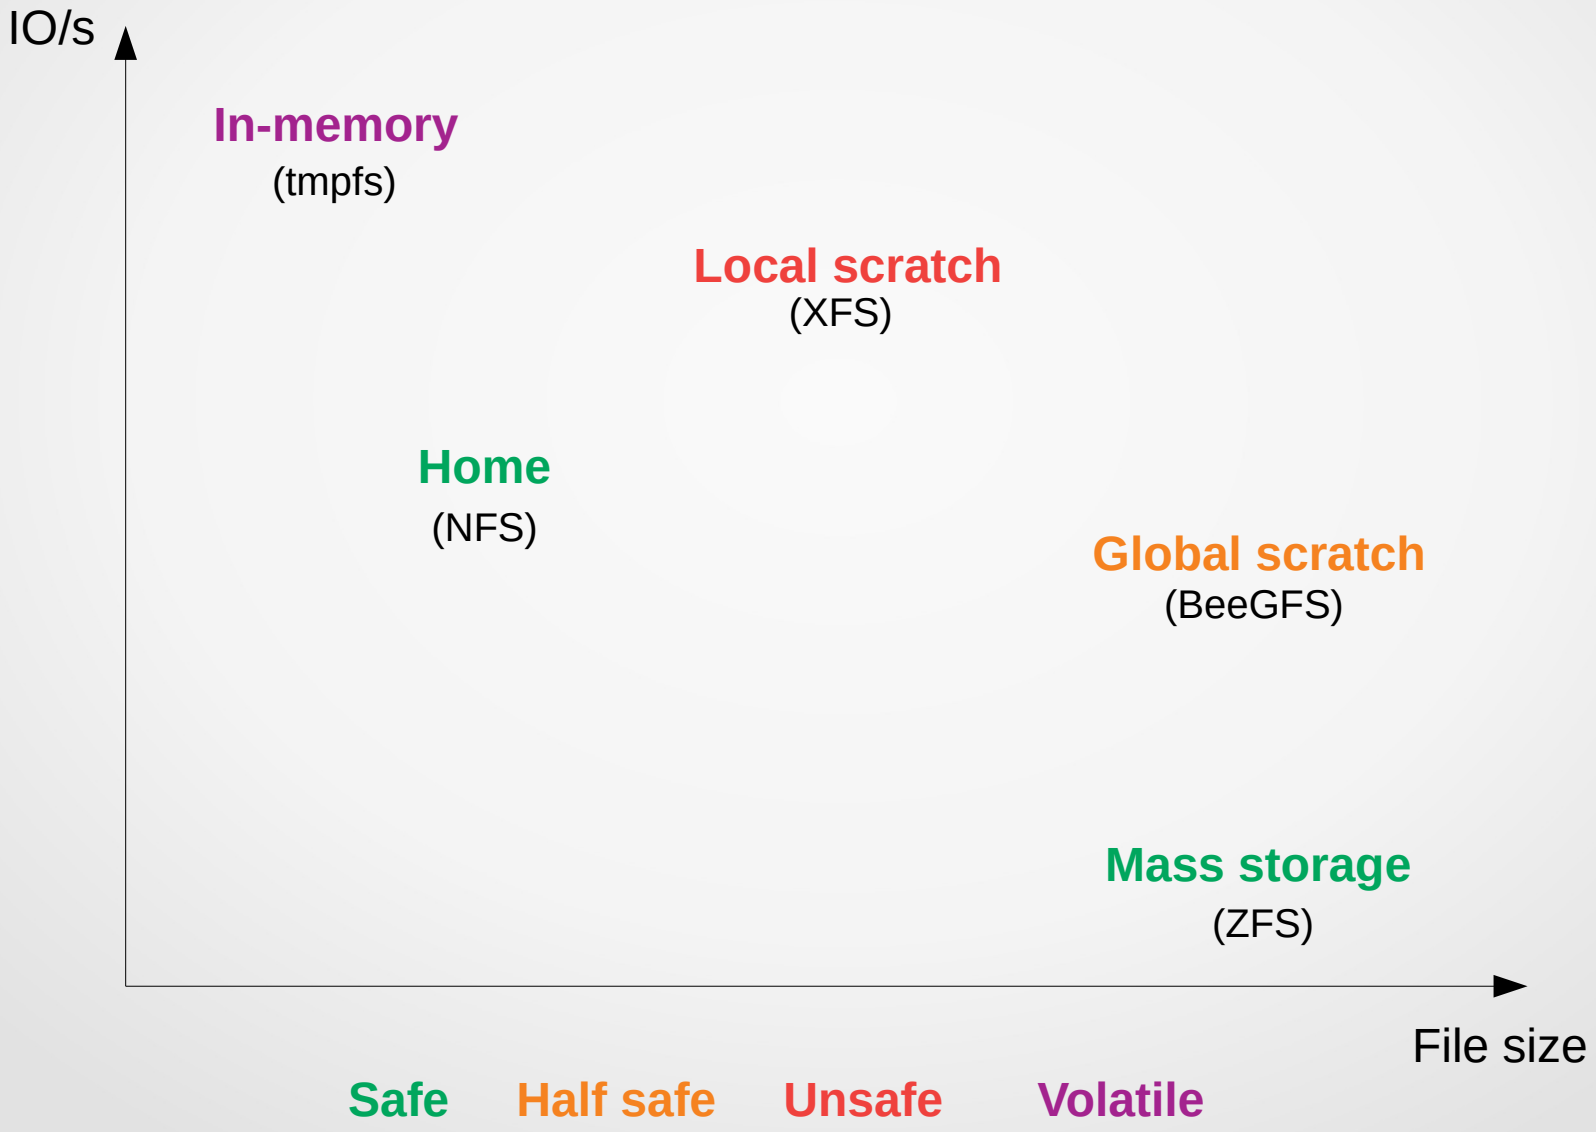

# File formats

# Text File Formats – JSON, YML, XML, INI


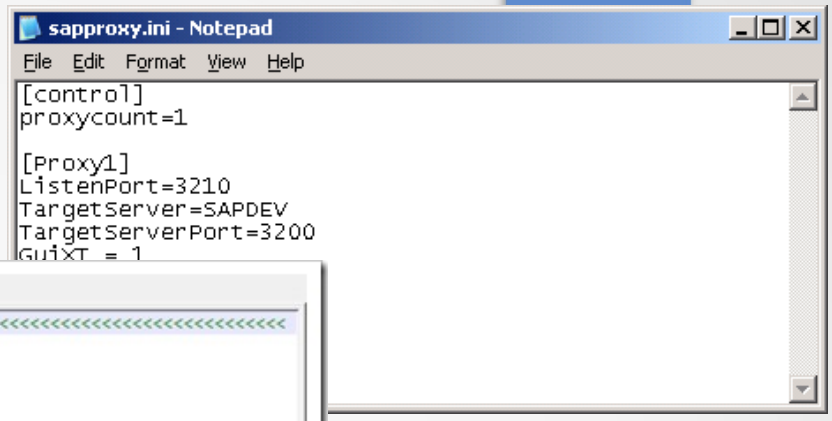


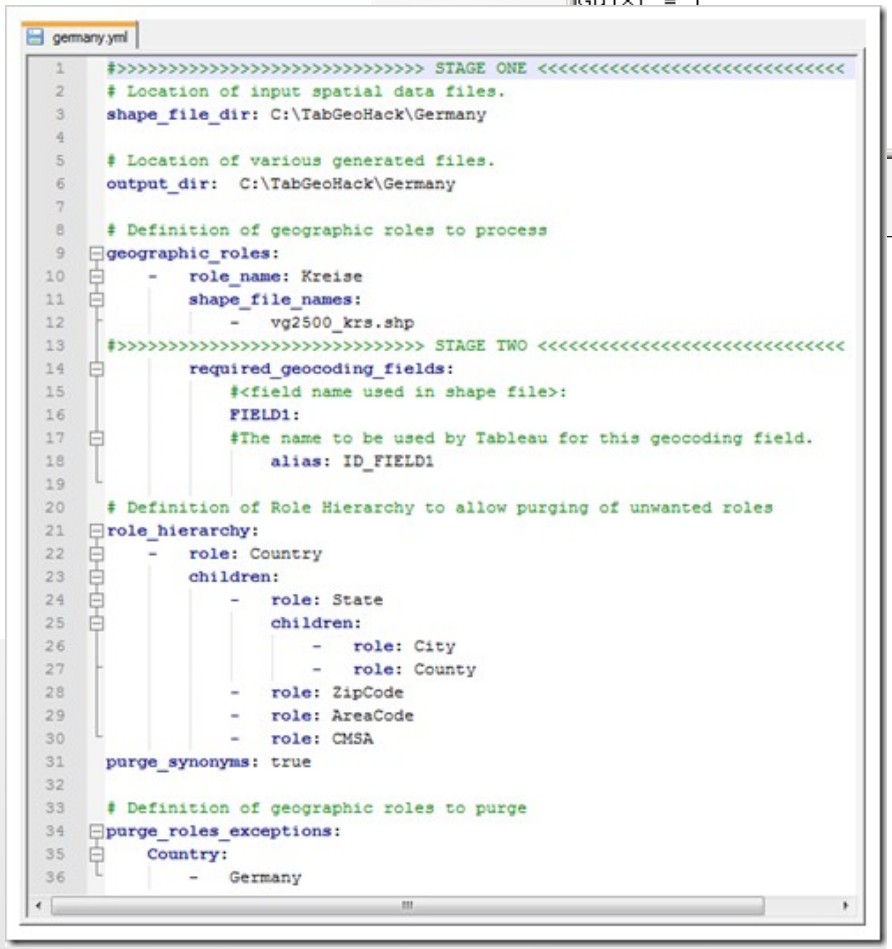


# Text File Formats – CSV

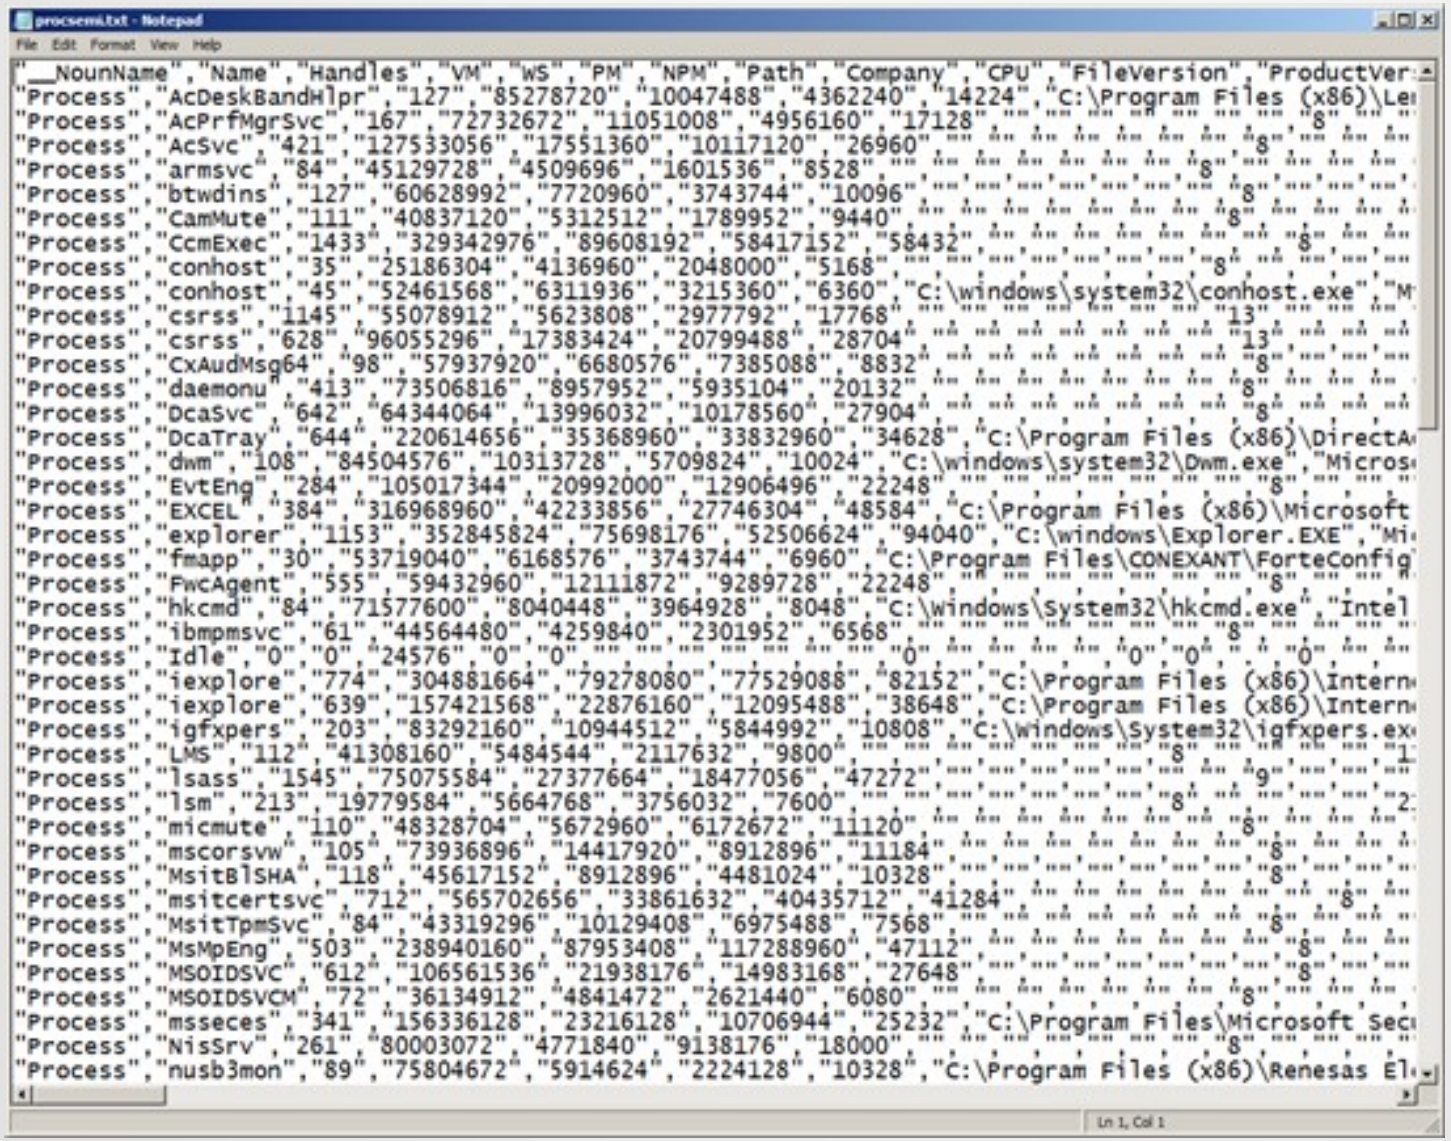

# Binary File Formats – CDF, HDF



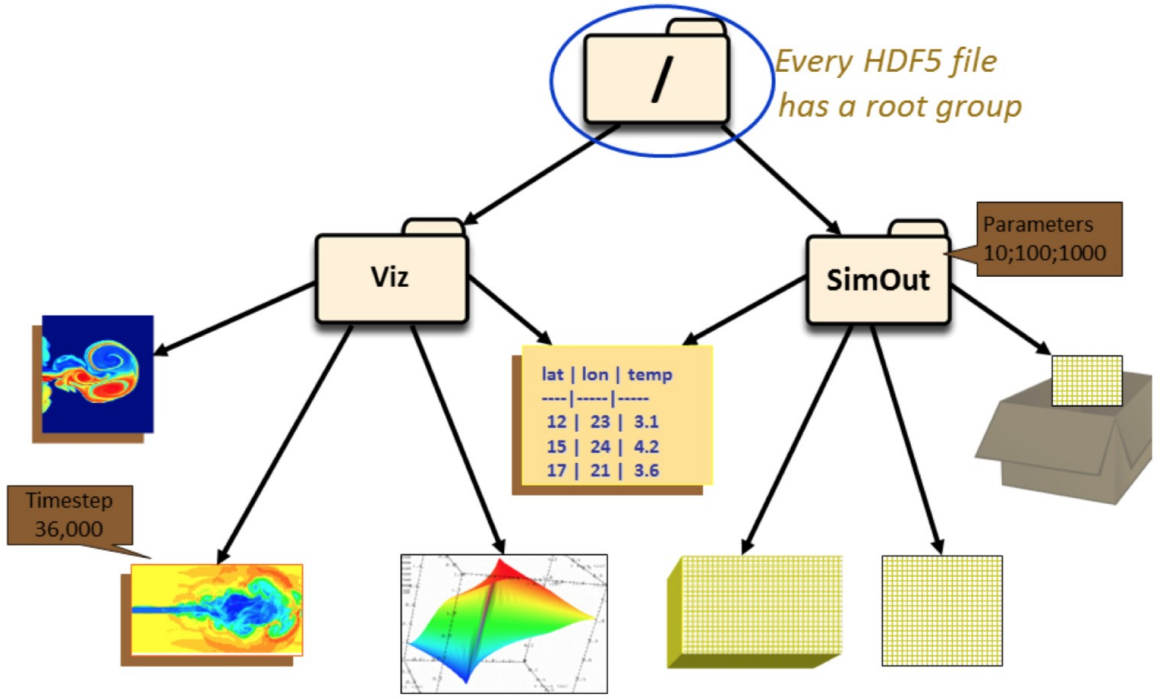


There are two groups in the HDF5 file depicted above: Vis and SimOut. Under the Viz group are a variety of images and a table that is shared with the SimOut group. The SimOut group contains a 3-dimensional array, a 2-dimensional array and a link to a 2-dimensional array in another HDF5 file.



https://support.hdfgroup.org/HDF5/Tutor/HDF5Intro.pdf

# Binary File Formats – CDF, HDF

```
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
 * Copyright by The HDF Group.                                               *
 * Copyright by the Board of Trustees of the University of Illinois.         *
 * All rights reserved.                                                      *
 *                                                                           *
 * This file is part of HDF5.  The full HDF5 copyright notice, including     *
 * terms governing use, modification, and redistribution, is contained in    *
 * the COPYING file, which can be found at the root of the source code       *
 * distribution tree, or in https://support.hdfgroup.org/ftp/HDF5/releases.  *
 * If you do not have access to either file, you may request a copy from     *
 * help@hdfgroup.org.                                                        *
 * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * */

/*
 *  This example illustrates how to write and read data in an existing
 *  dataset.  It is used in the HDF5 Tutorial.
 */

#include "hdf5.h"
#define FILE "dset.h5"

int main() {

   hid_t       file_id, dataset_id;  /* identifiers */
   herr_t      status;
   int         i, j, dset_data[4][6];

   /* Initialize the dataset. */
   for (i = 0; i < 4; i++)
      for (j = 0; j < 6; j++)
         dset_data[i][j] = i * 6 + j + 1;

   /* Open an existing file. */
   file_id = H5Fopen(FILE, H5F_ACC_RDWR, H5P_DEFAULT);

   /* Open an existing dataset. */
   dataset_id = H5Dopen2(file_id, "/dset", H5P_DEFAULT);

   /* Write the dataset. */
   status = H5Dwrite(dataset_id, H5T_NATIVE_INT, H5S_ALL, H5S_ALL, H5P_DEFAULT,
                     dset_data);

   status = H5Dread(dataset_id, H5T_NATIVE_INT, H5S_ALL, H5S_ALL, H5P_DEFAULT,
                    dset_data);

   /* Close the dataset. */
   status = H5Dclose(dataset_id);

   /* Close the file. */
   status = H5Fclose(file_id);
}
```



https://support.hdfgroup.org/ftp/HDF5/current/src/unpacked/examples/h5_rdwt.c

# Binary File Formats – CDF, HDF

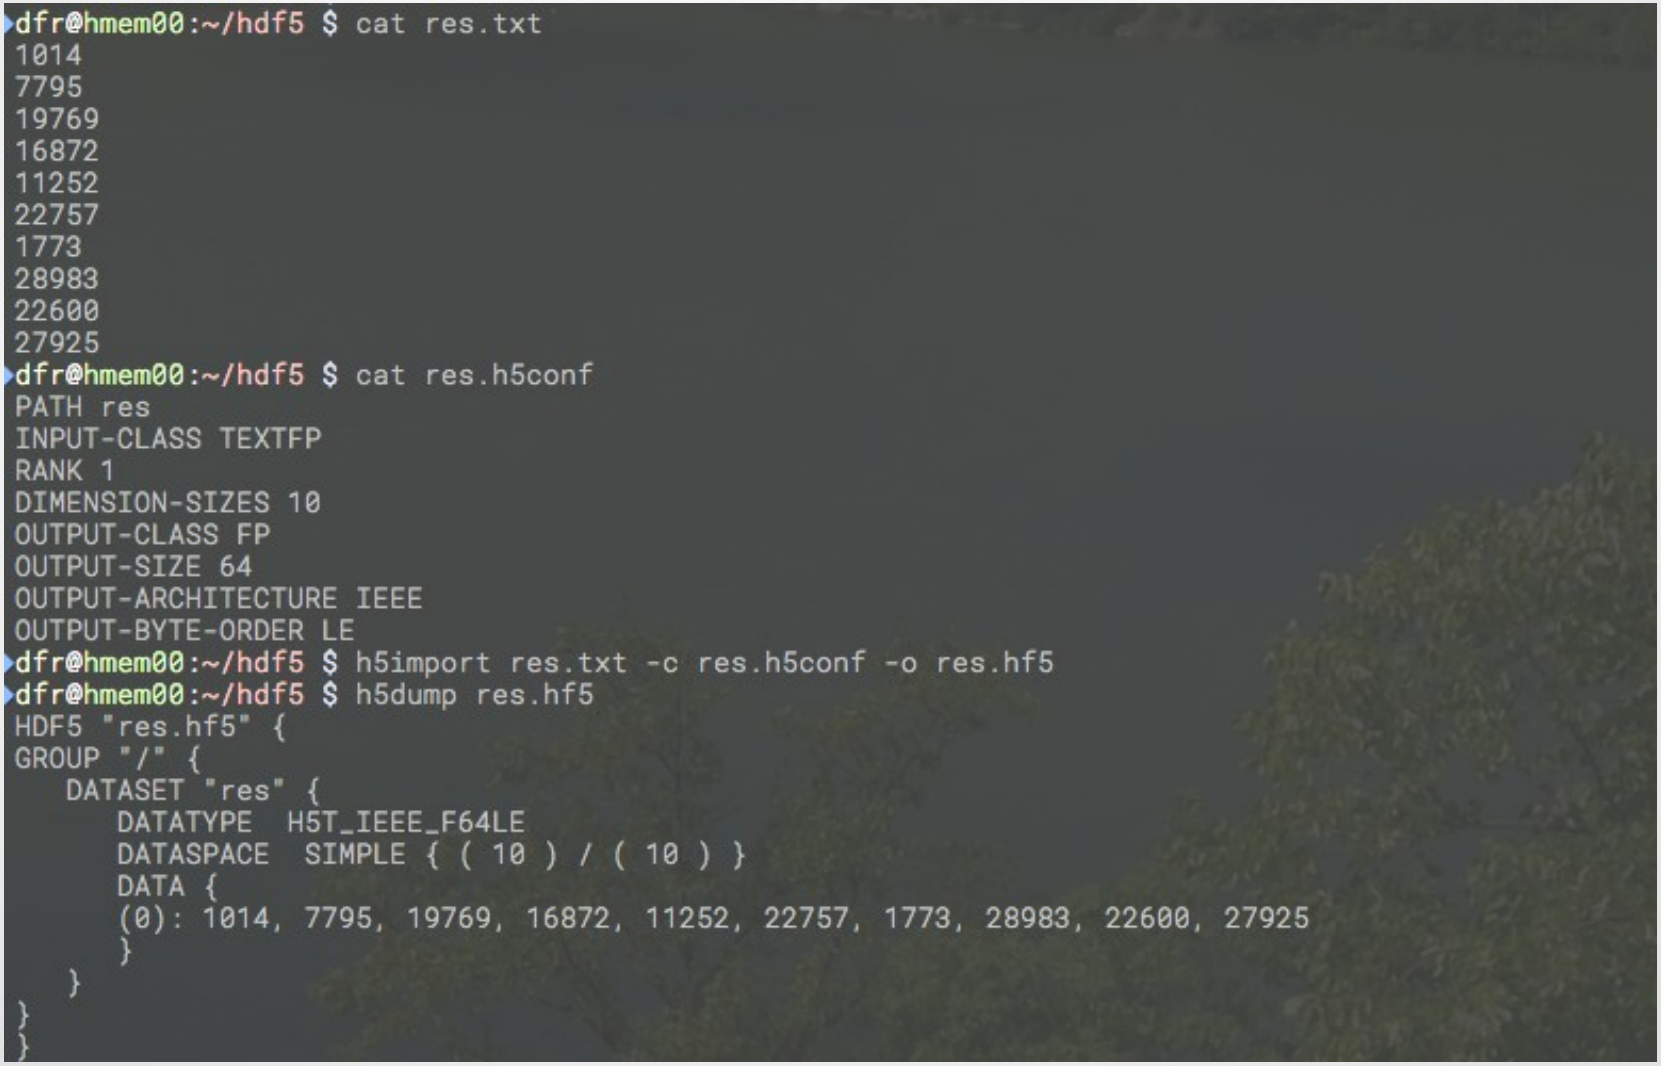

# What file format for what usage

- Meta data
  - Configuration file: INI, YAML
  - Result with context information: JSON
- Data
  - Small data (kBs): CSV
  - Medium data (MBs): compressed CSV
  - Large data (GBs): netCDF, HDF5, DXMF
  - Huge data (TBs): Database, Object store

(“loss of innocence”)

Use dedicated libraries to write and read them

# 1.2 Object storage

# Object storage

- Object: data (e.g. file) + **custom** meta data




- Often built on erasure coding (“software RAID”)
- Scale out easily
- Useful for web applications but coming to scientific world
- Access with REST API (through HTTP)

# S3 Python example


```
import boto.s3.connection

access_key = [redacted]
secret_key = [redacted]
conn = boto.connect_s3(
        aws_access_key_id=access_key,
        aws_secret_access_key=secret_key,
        host='192.168.64.51', port=7480,
        is_secure=False,
        calling_format=boto.s3.connection.OrdinaryCallingFormat(),
       )

bucket = conn.create_bucket('my-new-bucket')
```

```
## Create an object
key = bucket.new_key('hello.txt')
key.set_contents_from_string('Hello file in CEPH S3!')
key.set_metadata('Owner', 'dfr')

key = bucket.new_key('smstoolsS3-2.2.20.tar.gz')
key.set_contents_from_filename('smstools-2.2.20.tar.gz')

hello_key = bucket.get_key('hello.txt')
hello_url = hello_key.generate_url(3600, query_auth=False, force_http=False)
```



https://pypi.org/project/boto/

# S3 Python example

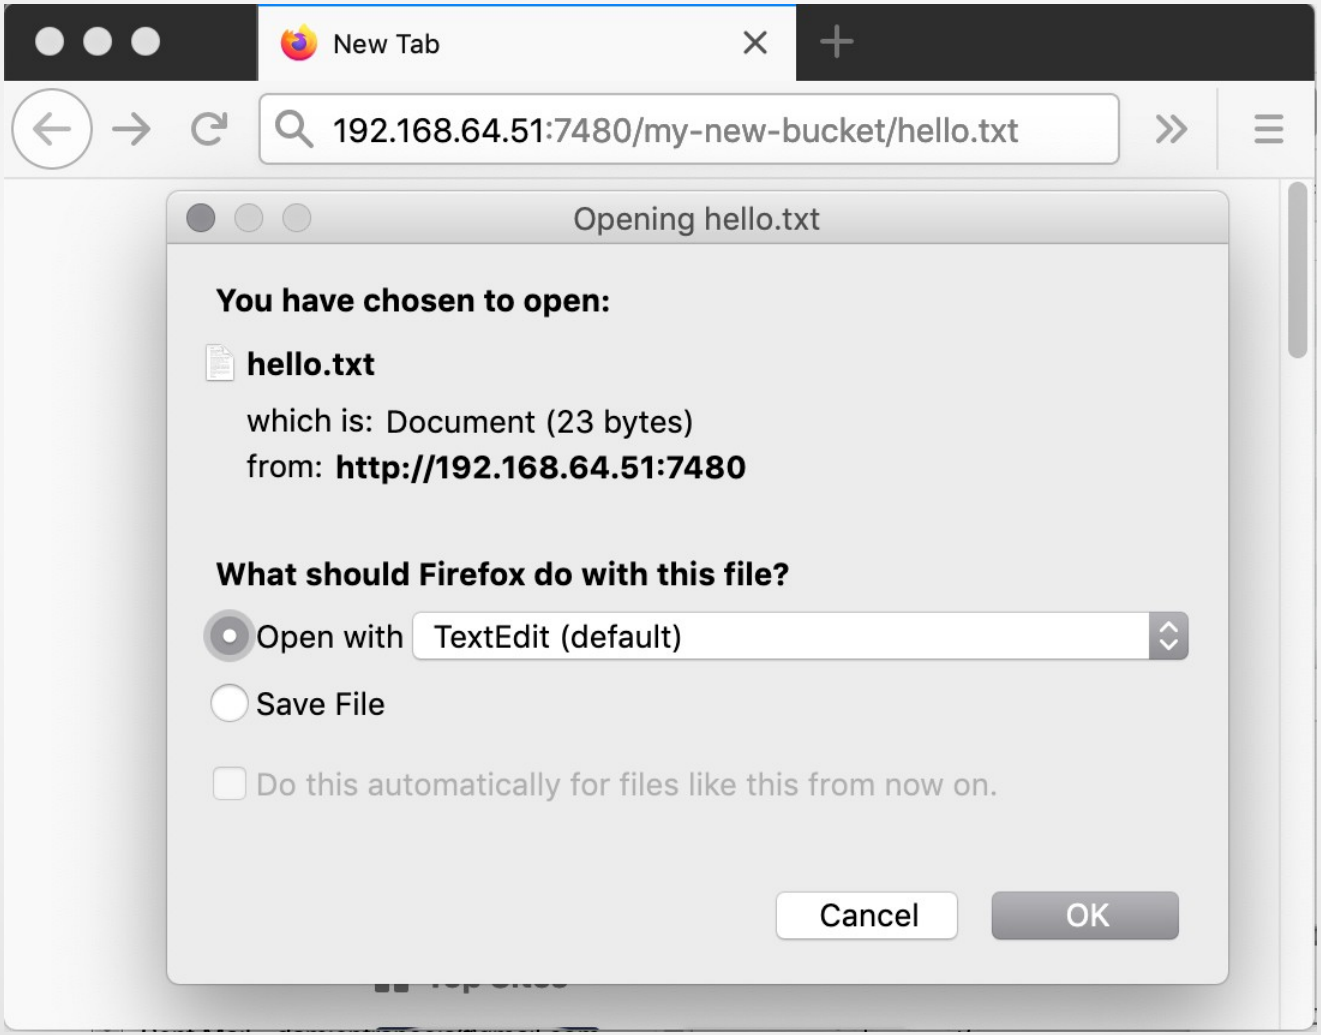




https://pypi.org/project/boto/

# S3 Object tagging (meta data)

## Object Tagging

Use object tagging to categorize storage. Each tag is a key-value pair. Consider the following tagging examples:

- Suppose an object contains protected health information (PHI) data. You might tag the object using the following key-value pair, as shown following:

```
PHI=True
```

  or

```
Classification=PHI
```

- Suppose you store project files in your S3 bucket. You might tag these objects with a key called `Project` and a value, as shown following:

```
Project=Blue
```

- You can add multiple tags to an object, as shown following:

```
Project=x
Classification=confidential
```



https://docs.aws.amazon.com/AmazonS3/latest/dev/object-tagging.html

# 1.3 Databases

# RDBMS

- Mostly needed for categorical data and alphanumerical data (not suited for matrices, but good for end-results)
- Indexes make finding a data element is very fast (and computing sums, maxima, etc.)
- Encodes relations between data (constraints, etc)
- Atomicity, Consistency, Isolation, and Durability








Pictures from http://www.ibm.com/developerworks/library/x-matters8/

# Tables and SQL

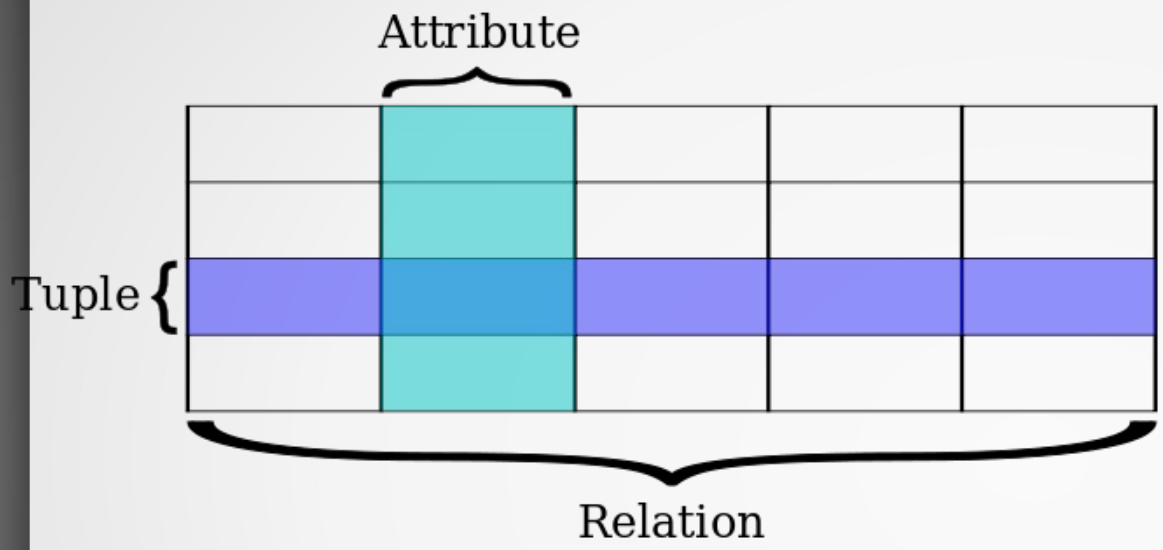


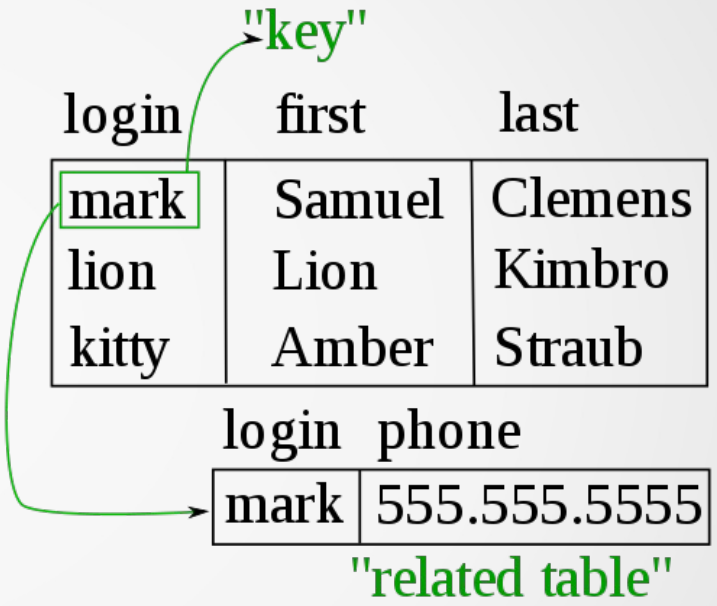


create table Users (login varchar(255), first varchar(255), last varchar(255));

insert into Users values (“mark”, ‘Samuel”, “Clemens”);

select first,last from Users where login=’lion’;
select login, phone from Users join PhoneNb on Users.login=Phone.login;

# File-based RDBMS



- The features of a relational database without the need for a complete setup
- Simply based on files
- Command line interface + API (Python, etc.)
- Still much more efficient than writing a million small files

# NoSQL


- Mostly needed for unstructured, semi-structured, and polymorphic data
- Scaling out very easy
- Basic Availability, Soft-state, Eventual consistency



Pictures from http://www.tomsitpro.com/articles/rdbms-sql-cassandra-dba-developer,2-547-2.html

# File-based NoSQL


- The features of a relational database without the need for a complete setup
- Simply based on files
- Command line interface + API (Python, etc.)
- Still much more efficient than writing a million small files



https://github.com/msiemens/tinydb

# When to use a database?


- when you have a large number of small files
- when you perform a lot of direct writes in a large file
- when you want to keep structure/relations between data
- when software crashes have a non-negligible probability
- when files are updated by several processes

# Example: Danger of NFS

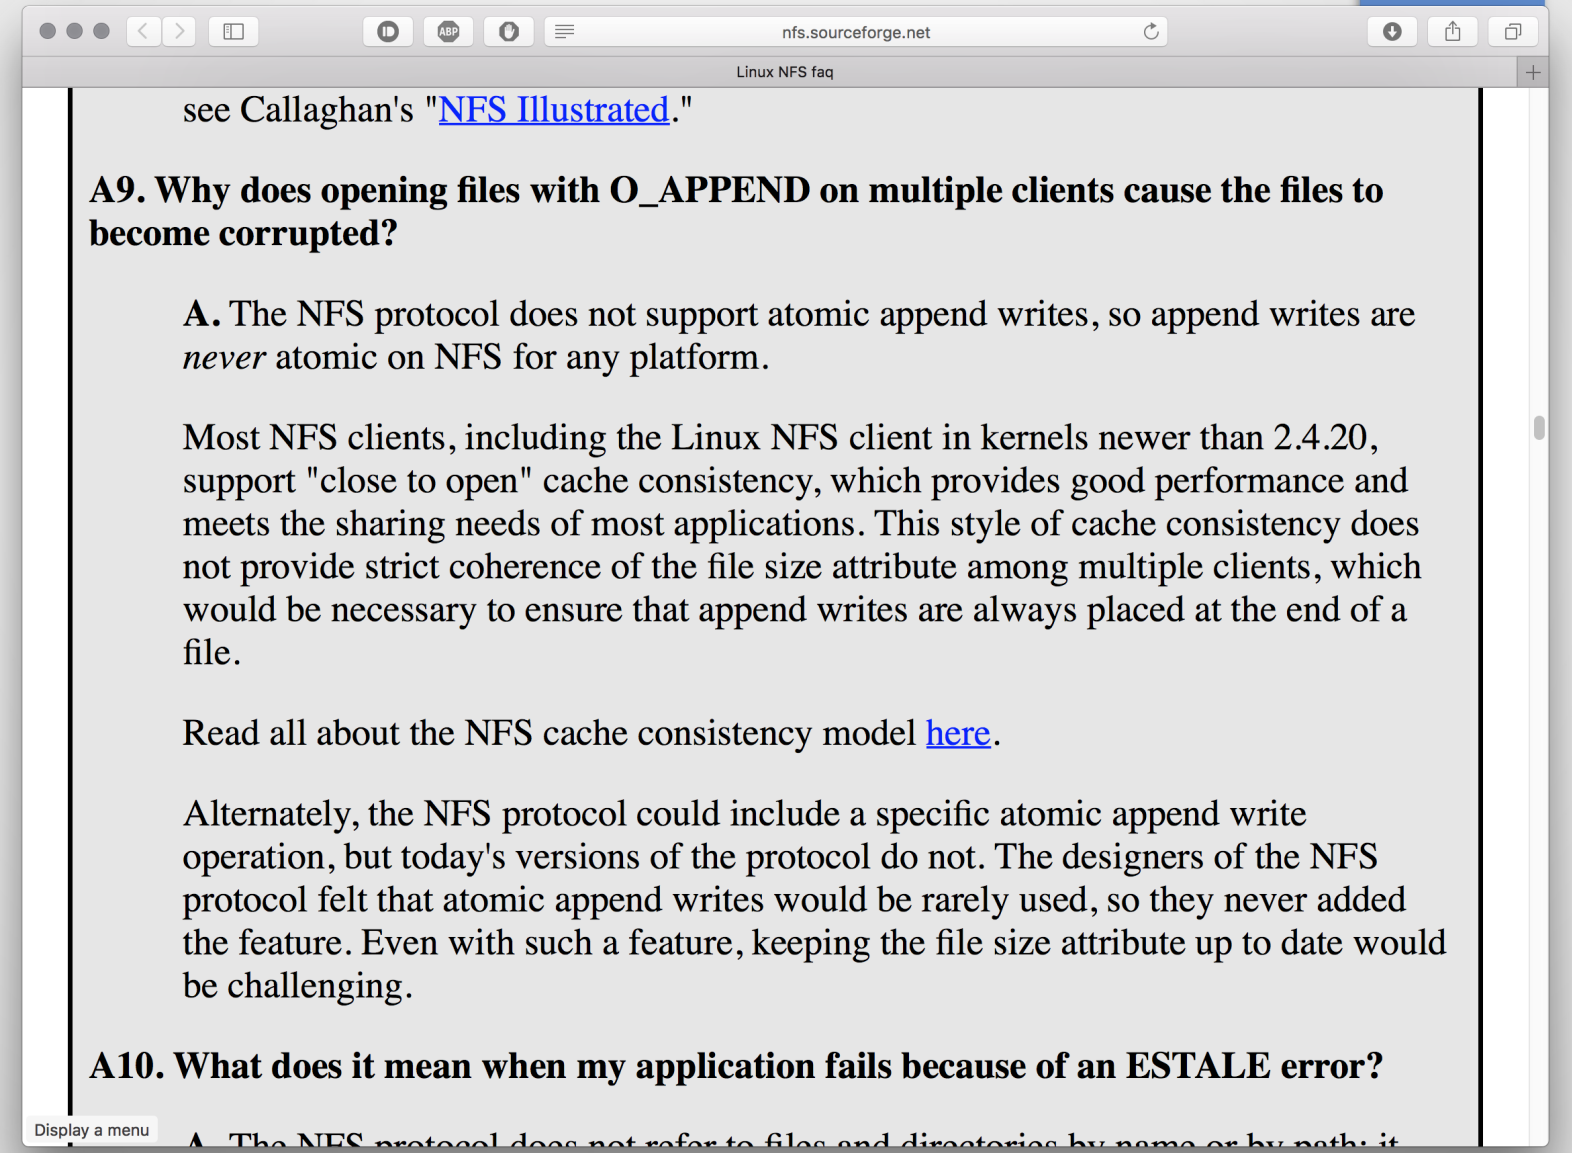

# Example: Danger of NFS

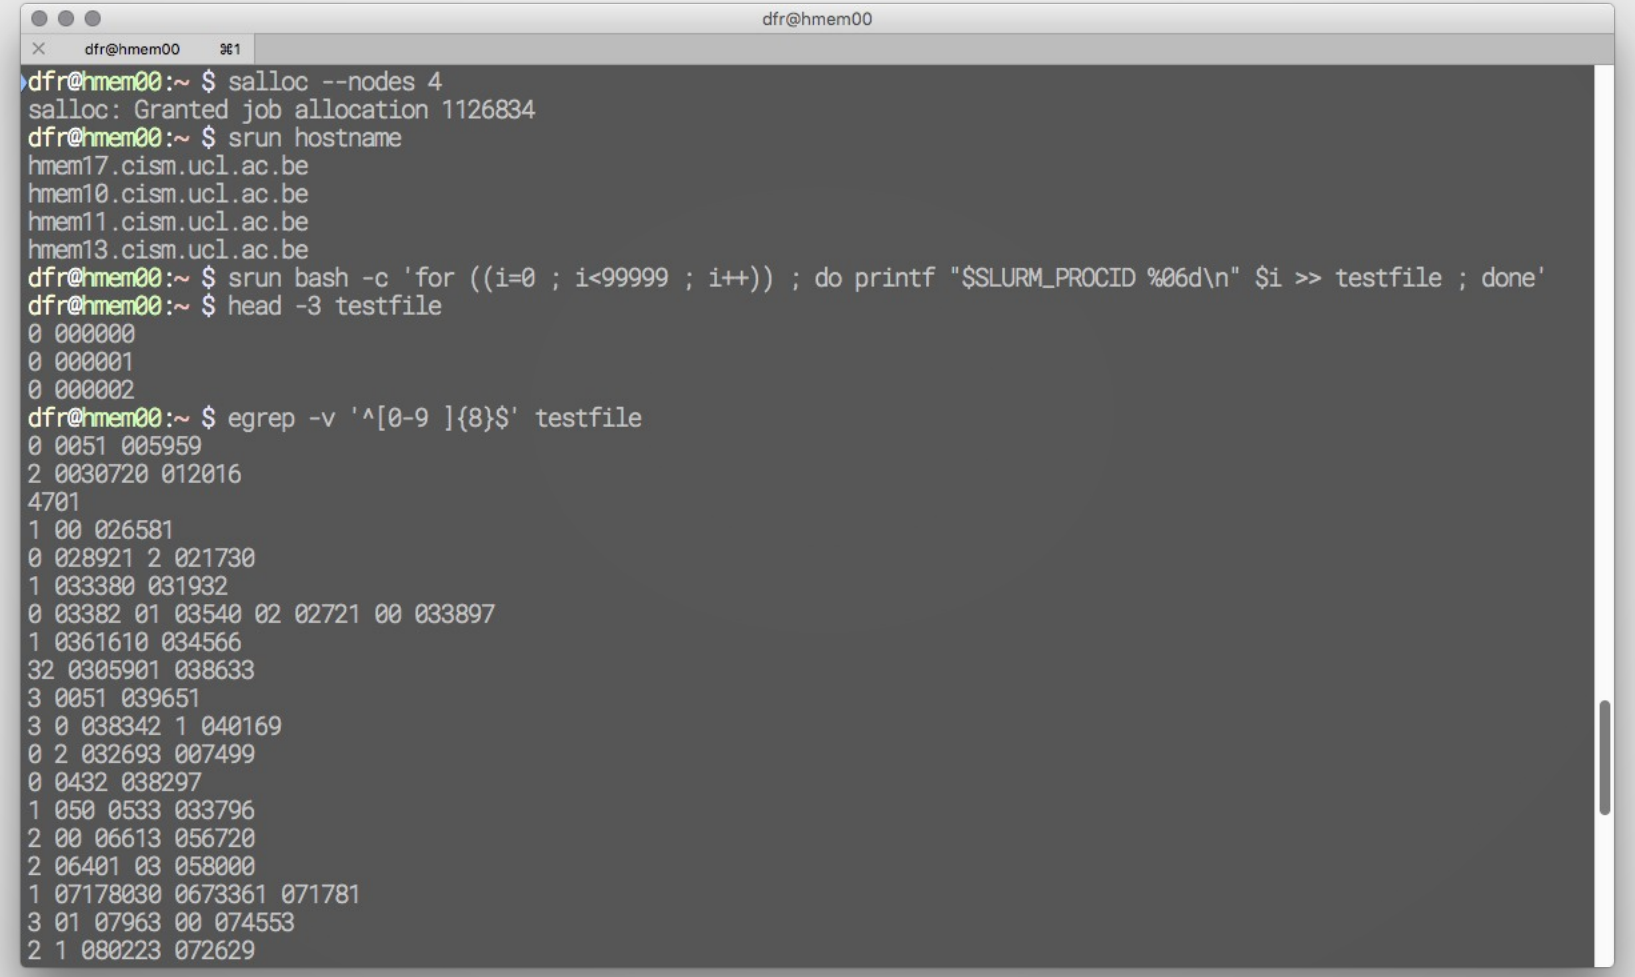

# When NOT to use a database?


- only sequential access
- simple matrices/vectors, etc.
- direct access on fixed-size records and no structure

# 2.

## Data transfer

# Network technologies

Laptop

100Mbps

Building switch

10Gbps

Datacenter switch

1Gbps

Cluster frontend

100Gbps

Compute node

# SCP

```
scp local_file.txt remote_host:/home
```

- Most direct way to copy data from/to UNIX/Linux machines
- Inefficient (sequential and synchronous)
- “outdated, inflexible and not readily fixed” (OpenSSH 8.0 release notes)



https://gravitational.com/blog/scp-familiar-simple-insecure-slow/

# SFTP

```
dfr@ncois:~ $ sftp lm3
Connected to lm3.
sftp> ?
Available commands:
bye                                Quit sftp
cd path                            Change remote directory to 'path'
chgrp grp path                     Change group of file 'path' to 'grp'
chmod mode path                    Change permissions of file 'path' to 'mode'
chown own path                     Change owner of file 'path' to 'own'
df [-hi] [path]                    Display statistics for current directory or
                                   filesystem containing 'path'
exit                               Quit sftp
get [-afPpRr] remote [local]       Download file
reget [-fPpRr] remote [local]      Resume download file
reput [-fPpRr] [local] remote      Resume upload file
help                               Display this help text
lcd path                           Change local directory to 'path'
lls [ls-options [path]]            Display local directory listing
lmkdir path                        Create local directory
ln [-s] oldpath newpath            Link remote file (-s for symlink)
lpwd                               Print local working directory
ls [-1afhlnrSt] [path]             Display remote directory listing
lumask umask                       Set local umask to 'umask'
mkdir path                         Create remote directory
progress                           Toggle display of progress meter
put [-afPpRr] local [remote]       Upload file
pwd                                Display remote working directory
quit                               Quit sftp
rename oldpath newpath             Rename remote file
rm path                            Delete remote file
rmdir path                         Remove remote directory
symlink oldpath newpath            Symlink remote file
version                            Show SFTP version
!command                           Execute 'command' in local shell
!                                  Escape to local shell
?                                  Synonym for help
sftp>
```

# RSYNC


```
rsync(1)                                                        rsync(1)

NAME
       rsync - faster, flexible replacement for rcp

SYNOPSIS
       rsync [OPTION]... SRC [SRC]... DEST

       rsync [OPTION]... SRC [SRC]... [USER@]HOST:DEST

       rsync [OPTION]... SRC [SRC]... [USER@]HOST::DEST

       rsync [OPTION]... SRC [SRC]... rsync://[USER@]HOST[:PORT]/DEST

       rsync [OPTION]... SRC

       rsync [OPTION]... [USER@]HOST:SRC [DEST]

       rsync [OPTION]... [USER@]HOST::SRC [DEST]

       rsync [OPTION]... rsync://[USER@]HOST[:PORT]/SRC [DEST]

DESCRIPTION
       rsync is a program that behaves in much the same way that rcp does, but has many more
       options and uses the rsync remote-update protocol to greatly speed up file  transfers
       when the destination file is being updated.

       The  rsync  remote-update  protocol  allows  rsync  to  transfer just the differences
       between two sets of files across the network connection, using an efficient checksum-
       search algorithm described in the technical report that accompanies this package.

       Some of the additional features of rsync are:

       o      support for copying links, devices, owners, groups, and permissions

       o      exclude and exclude-from options similar to GNU tar

       o      a CVS exclude mode for ignoring the same files that CVS would ignore

       o      can use any transparent remote shell, including ssh or rsh

       o      does not require super-user privileges

       o      pipelining of file transfers to minimize latency costs

       o      support for anonymous or authenticated rsync daemons (ideal for mirroring)

GENERAL
       Rsync  copies  files  either to or from a remote host, or locally on the current host
       (it does not support copying files between two remote hosts).

       There are two different ways for rsync to contact a remote system:  using  a  remote-
       shell  program  as  the  transport (such as ssh or rsh) or contacting an rsync daemon
       directly via TCP.  The remote-shell transport is used whenever the source or destina-
       tion path contains a single colon (:) separator after a host specification.  Contact-
:
```

# Update only what changed: rsync

```
rsync [OPTIONS]... SRC [SRC]... [USER@]HOST:DEST
```

- Always use:  -az
- Other interesting arguments:
  - -v and --progress
  - --include or --exclude
  - --delete and/or --remove-source-file
  - --dry-run
  - --size-only or –checksum
- Works well with GNU parallel

# Resuming transfers

- When nothing changed but the transfer was interrupted
  - append: do not re-check partially transmitted files and resume the transfer where it was abandoned assuming first transfer attempt was with scp or with rsync --inplace

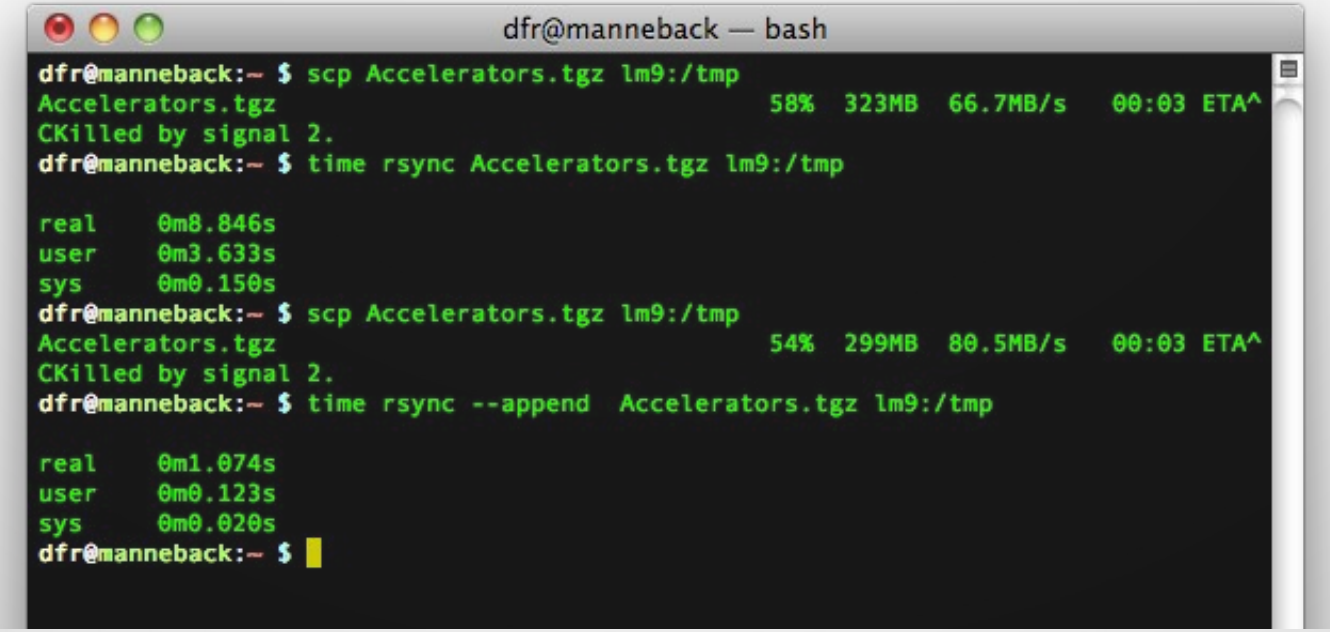

# TAR | SSH

- From local to remote

  $ tar zvzf - /path/to/data | ssh server "cat > /srv/data_server1.tar.gz"

- From remote to local

  $ ssh server tar czf - /path/to/data/ > ./data_server.tar.gz

- Avoid a lot of communication overhead linked to inodes
- Use the pv command to get a progress bar

# Parallel data transfer: bbcp

- Better use of the bandwidth than SCP
- Needs to be installed on both sides (easy to install)
- Needs friendly firewalls

```
dfr@manneback:~ $ time scp Accelerators.tgz lm9:/dev/null
Accelerators.tgz                              100%  551MB  50.1MB/s   00:11

real    0m13.849s
user    0m5.043s
sys     0m2.330s
dfr@manneback:~ $ time bbcp -P 2 Accelerators.tgz lm9:/dev/null/t
bbcp: Creating /dev/null/t
bbcp: 151113 22:06:00  39% done; 112.3 MB/s
bbcp: 151113 22:06:02  79% done; 112.0 MB/s

real    0m9.000s
user    0m0.232s
sys     0m0.926s
```



https://github.com/Stormwind99/bbcp

# Parallel rsync

- Parsyncfp

```
[root@compute-3-11 calof]# parsyncfp --NP=8 --chunk=10G -i ib0 --startdir /mnt/calof lander-calof  /dfs3/staff/hmangala/
INFO: You've specified what looks like an Infiniband interface [ib0]...
INFO: .. and you have 'perfquery installed, so RDMA bytes will be reported as well.
WARN: About to remove all the old cached chunkfiles from [/root/.parsyncfp/fpcache].
 Enter ^C to stop this.
 If you specified '--nowait', cache will be cleared in 3s regardless.
 Otherwise, hit [Enter] and I'll clear them.
Press [ENTER] to continue.

INFO: The fpart chunk files [/root/.parsyncfp/fpcache/f*] are cleared .. continuing.
INFO: Forking fpart.  Check [/root/.parsyncfp/fpcache/fpart.log.20.10.31_2018-12-12] for errors if it hangs.
INFO: Starting the 1st [8] rsyncs ..
       | Elapsed |   1m   |    [    ib0]   MB/s  | Running || Susp'd  |      Chunks       [2018-12-12]
 Time  | time(m) |  Load  |     TCP / RDMA  out  |   PIDs  ||  PIDs   | [UpTo] of [ToDo]
20.10.39    0.07     3.00       0.00 / 132.12          5    <>   0          [8] of [50]
20.10.43    0.12     3.08       0.00 / 426.94          8    <>   0          [13] of [50]
20.10.47    0.18     3.40       0.00 / 818.96          8    <>   0          [15] of [50]
20.10.50    0.23     3.40       0.00 / 848.65          8    <>   0          [15] of [50]
20.10.54    0.30     3.77       0.00 / 835.68          8    <>   0          [15] of [50]
20.10.57    0.35     3.86       0.00 / 781.66          8    <>   0          [15] of [50]
20.11.01    0.42     3.86       0.00 / 734.39          8    <>   0          [15] of [50]
20.11.04    0.47     4.12       0.00 / 618.96          8    <>   0          [15] of [50]
20.11.08    0.53     4.51       0.00 / 572.82          8    <>   0          [15] of [50]
20.11.11    0.58     4.51       0.00 / 534.97          8    <>   0          [15] of [50]
20.11.15    0.65     4.79       0.00 / 537.75          8    <>   0          [15] of [50]
20.11.18    0.70     5.12       0.00 / 525.65          8    <>   0          [15] of [50]
```



https://github.com/hjmangalam/parsyncfp

# Parallel rsync

- GNU Parallel + rsync

## EXAMPLE: Parallelizing rsync

**rsync** is a great tool, but sometimes it will not fill up the available bandwidth. Running multiple **rsync** in parallel can fix this.

```
cd src-dir
find . -type f |
  parallel -j10 -X rsync -zR -Ha ./{} fooserver:/dest-dir/
```

Adjust **-j10** until you find the optimal number.

**rsync -R** will create the needed subdirectories, so all files are not put into a single dir. The **./** is needed so the resulting command looks similar to:

```
rsync -zR ././sub/dir/file fooserver:/dest-dir/
```

The **/./** is what **rsync -R** works on.

If you are unable to push data, but need to pull them and the files are called digits.png (e.g. 000000.png) you might be able to do:

```
seq -w 0 99 | parallel rsync -Havessh fooserver:src/*{}.png destdir/
```

Limit the depth of `find` can help speed things up



https://www.gnu.org/software/parallel/man.html#EXAMPLE:-Parallelizing-rsync

# 3. Data sharing

with other users (Unix permissions, Encryption)

with external users (Owncloud, Dataverse)

# Data sharing

Data sharing with other users

# Sharing with the group

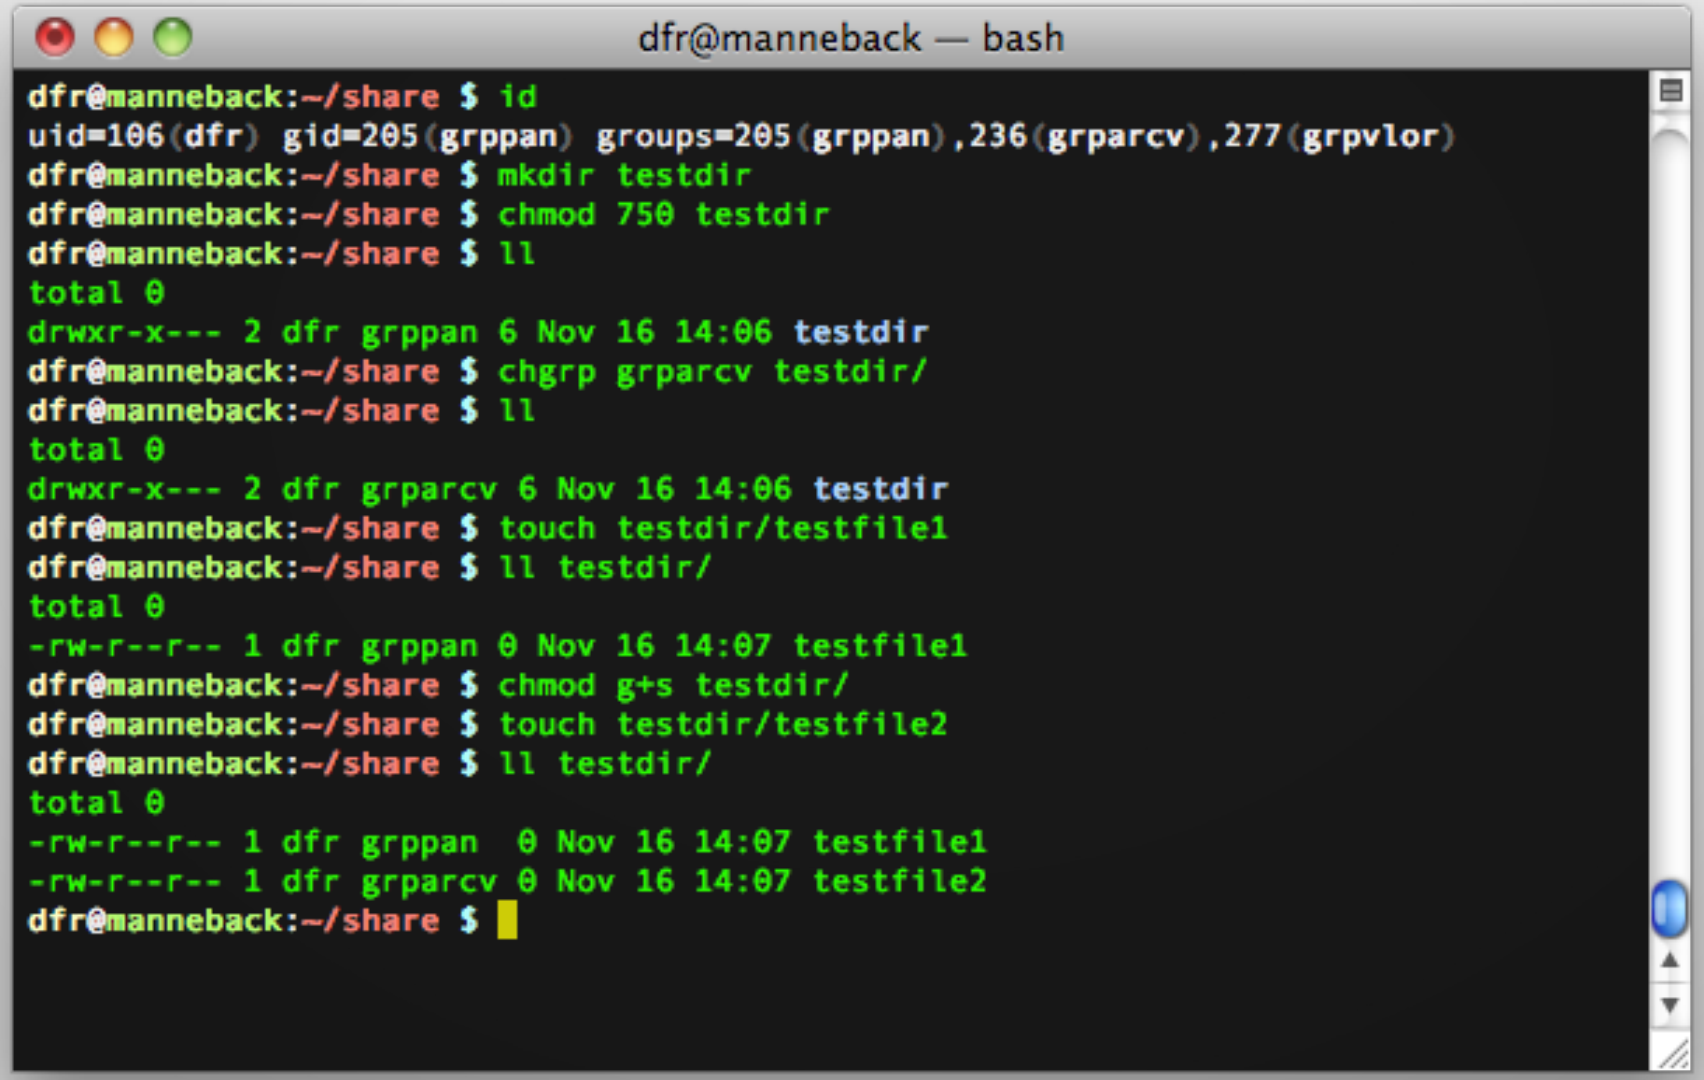

# Sharing and hiding

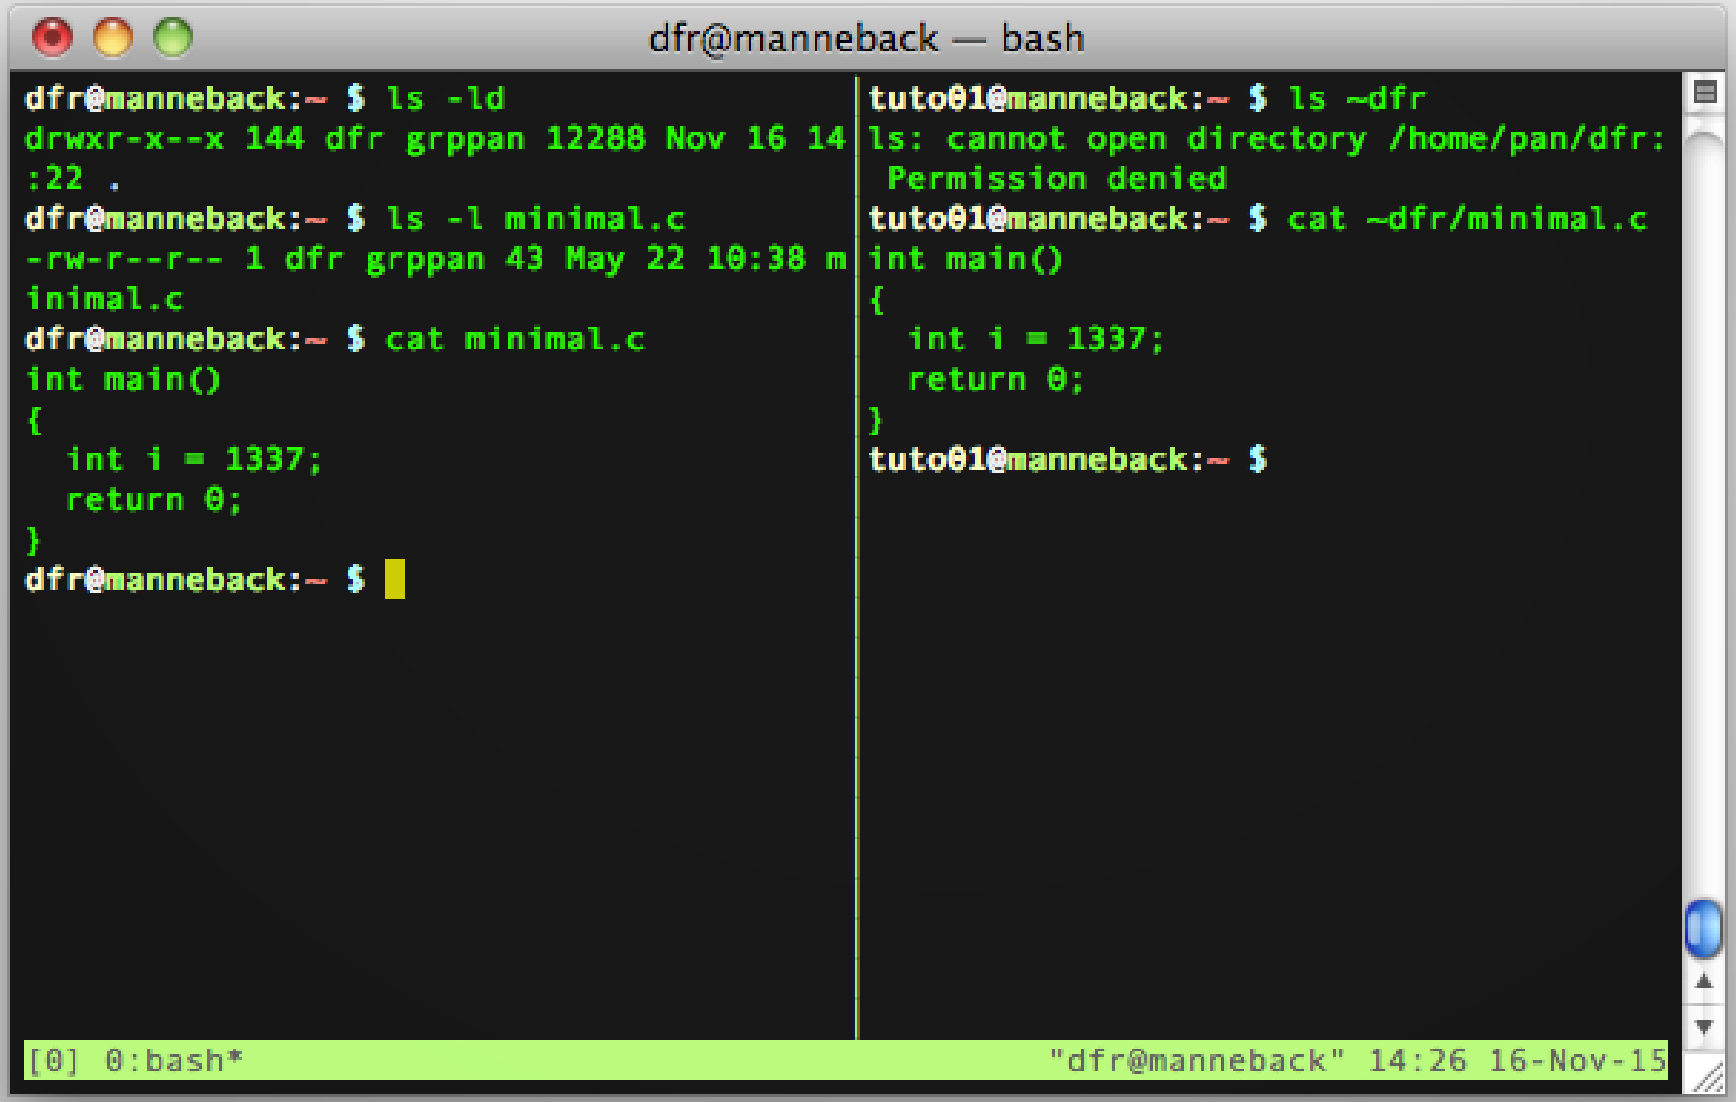

# Sharing and encrypting


dfr@manneback — bash

```
dfr@manneback:~ $ cat minimal.c
int main()
{
  int i = 1337;
  return 0;
}
dfr@manneback:~ $ openssl des3 -salt -in minimal.c -out minimal.c.des3
enter des-ede3-cbc encryption password:
Verifying - enter des-ede3-cbc encryption password:
dfr@manneback:~ $ cat minimal.c.des3
Salted__0?1c_?RjD7
                ?     ??f?c??????X?B?6:=???????!h??.,???ni?dfr@manneback:~ $
dfr@manneback:~ $ openssl des3 -d  -in minimal.c.des3 -out minimal.c.clear
enter des-ede3-cbc decryption password:
dfr@manneback:~ $ cat minimal.c.clear
int main()
{
  int i = 1337;
  return 0;
}
dfr@manneback:~ $
```

# Data sharing

Data sharing with external users

# Data sharing with external users

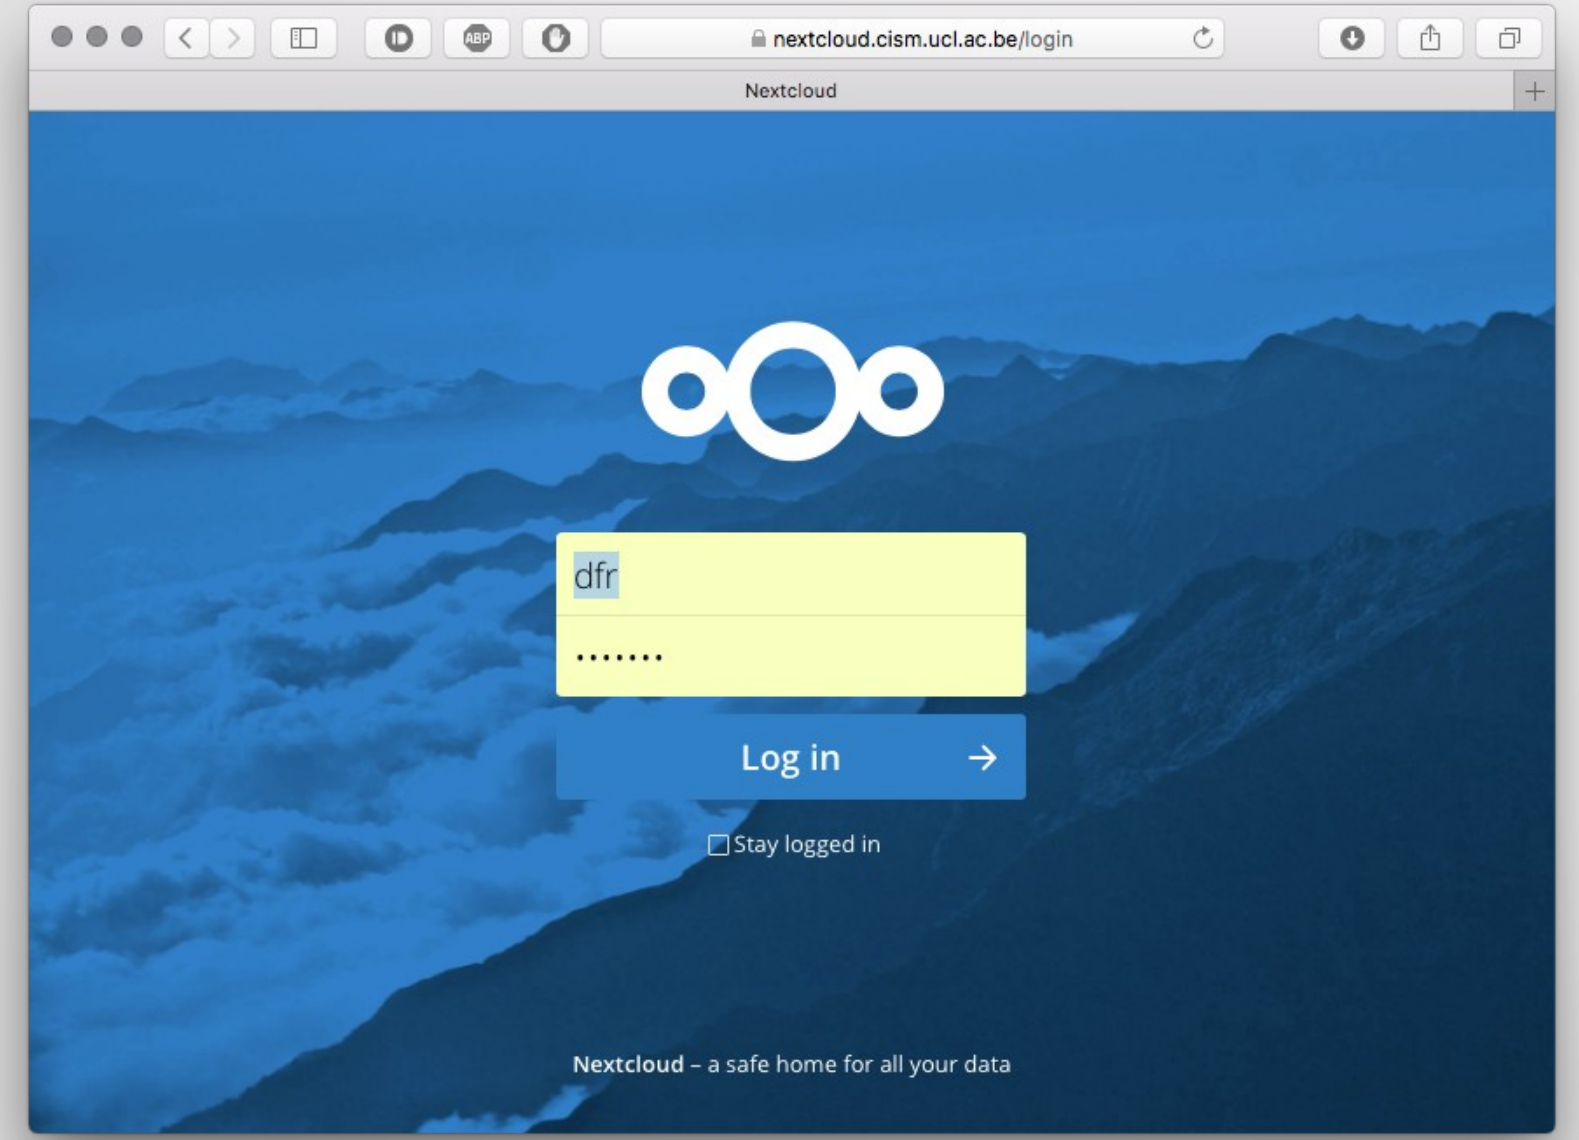


CISM login

# Dropbox-like

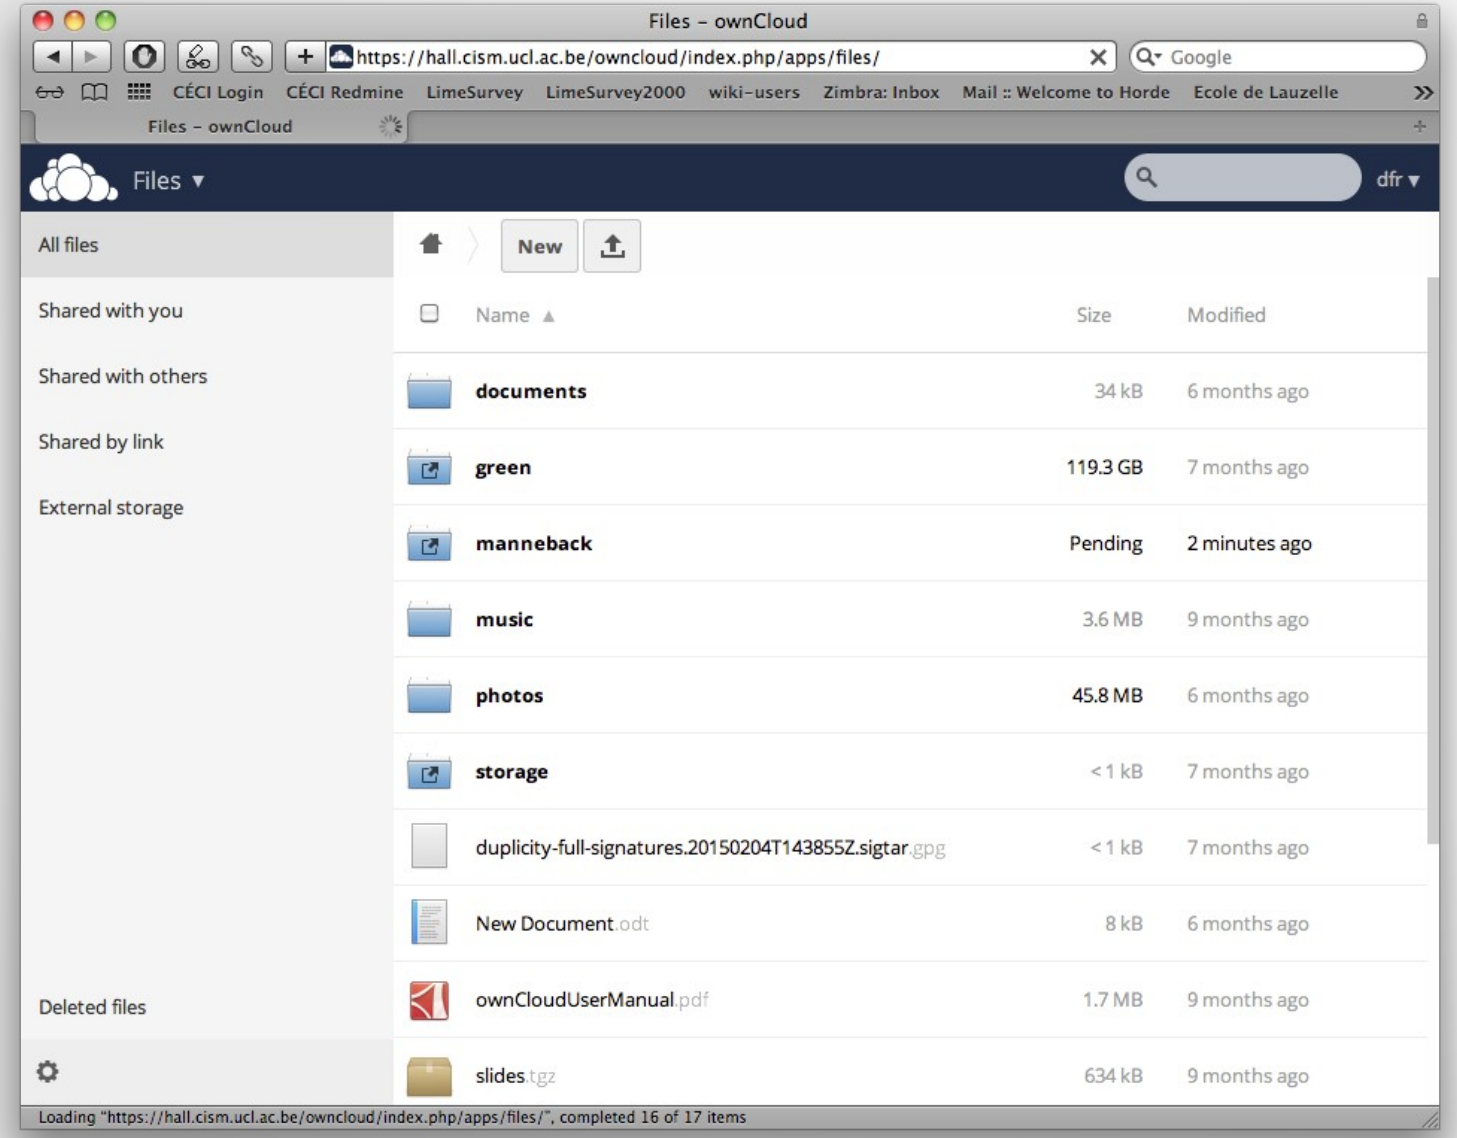

# External SFTP connectors

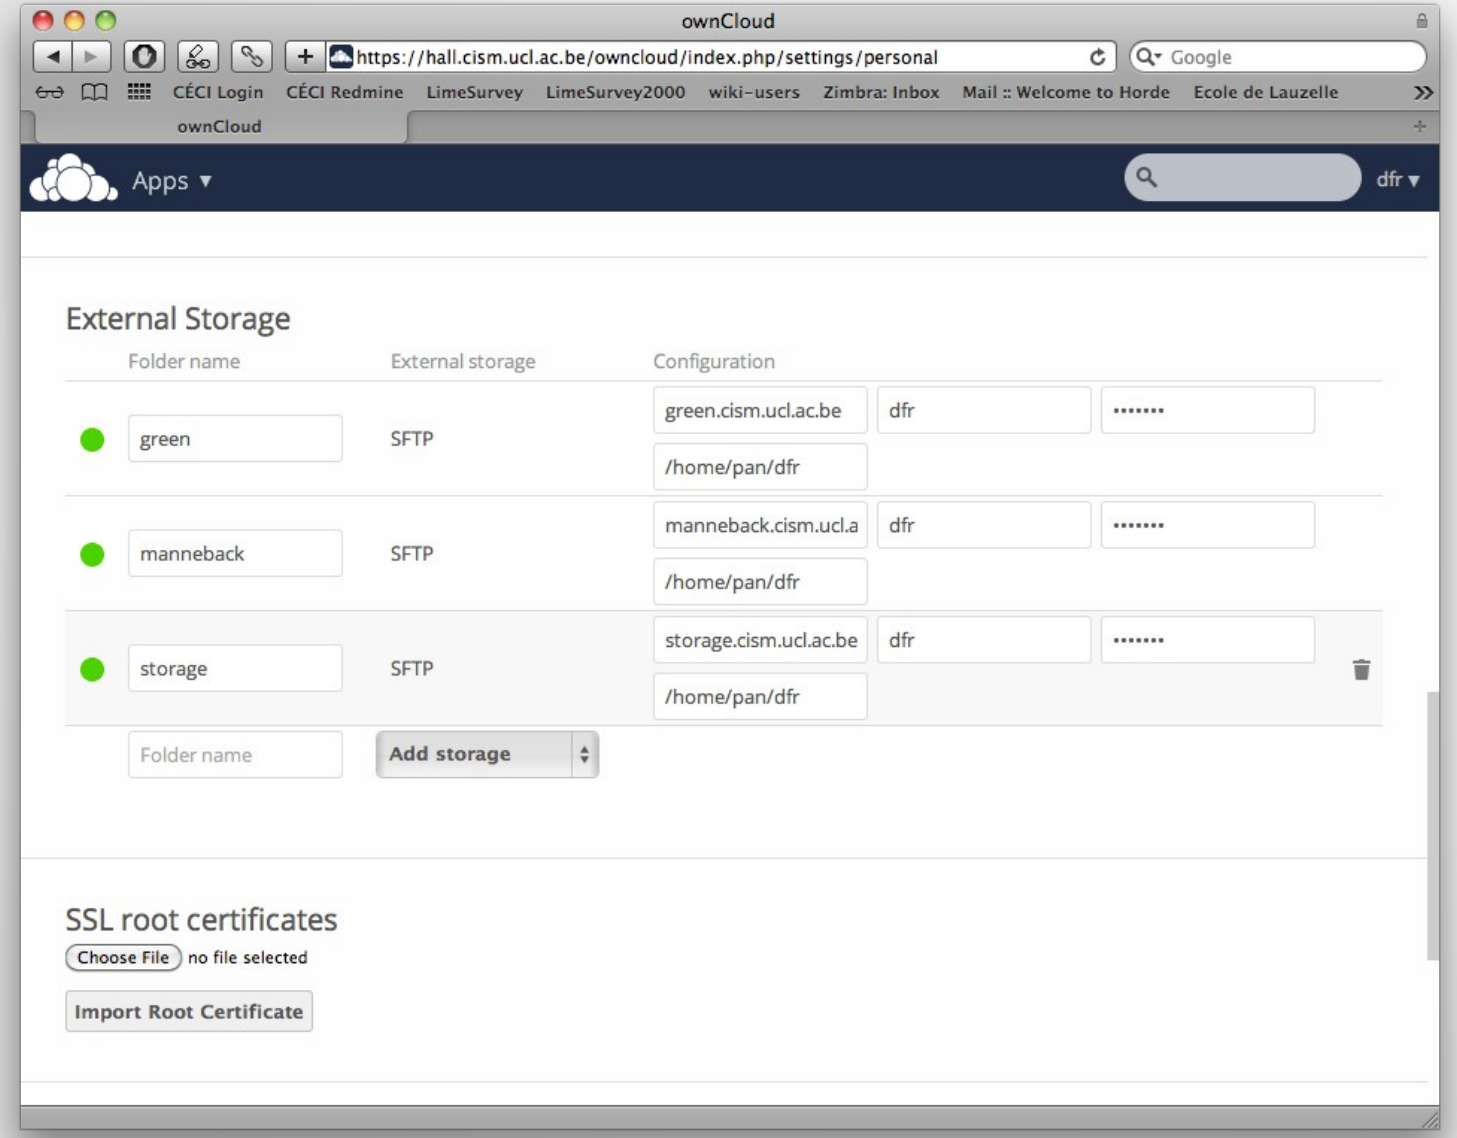

# Dropbox-like

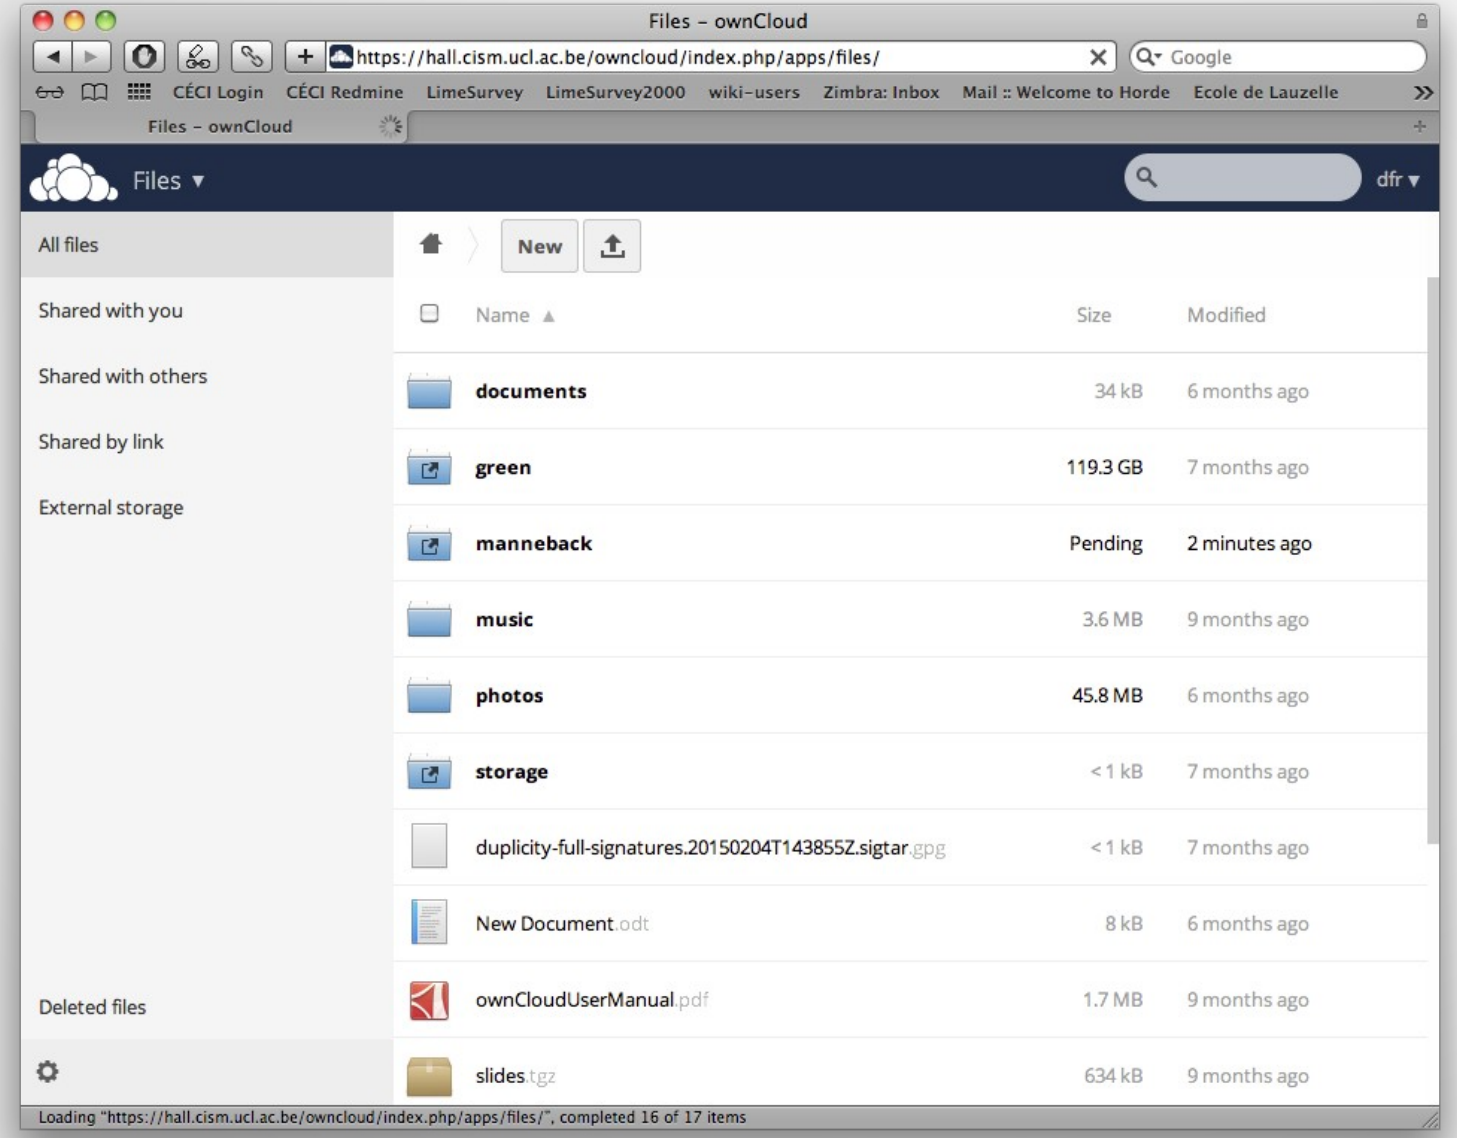

# My home on Manneback

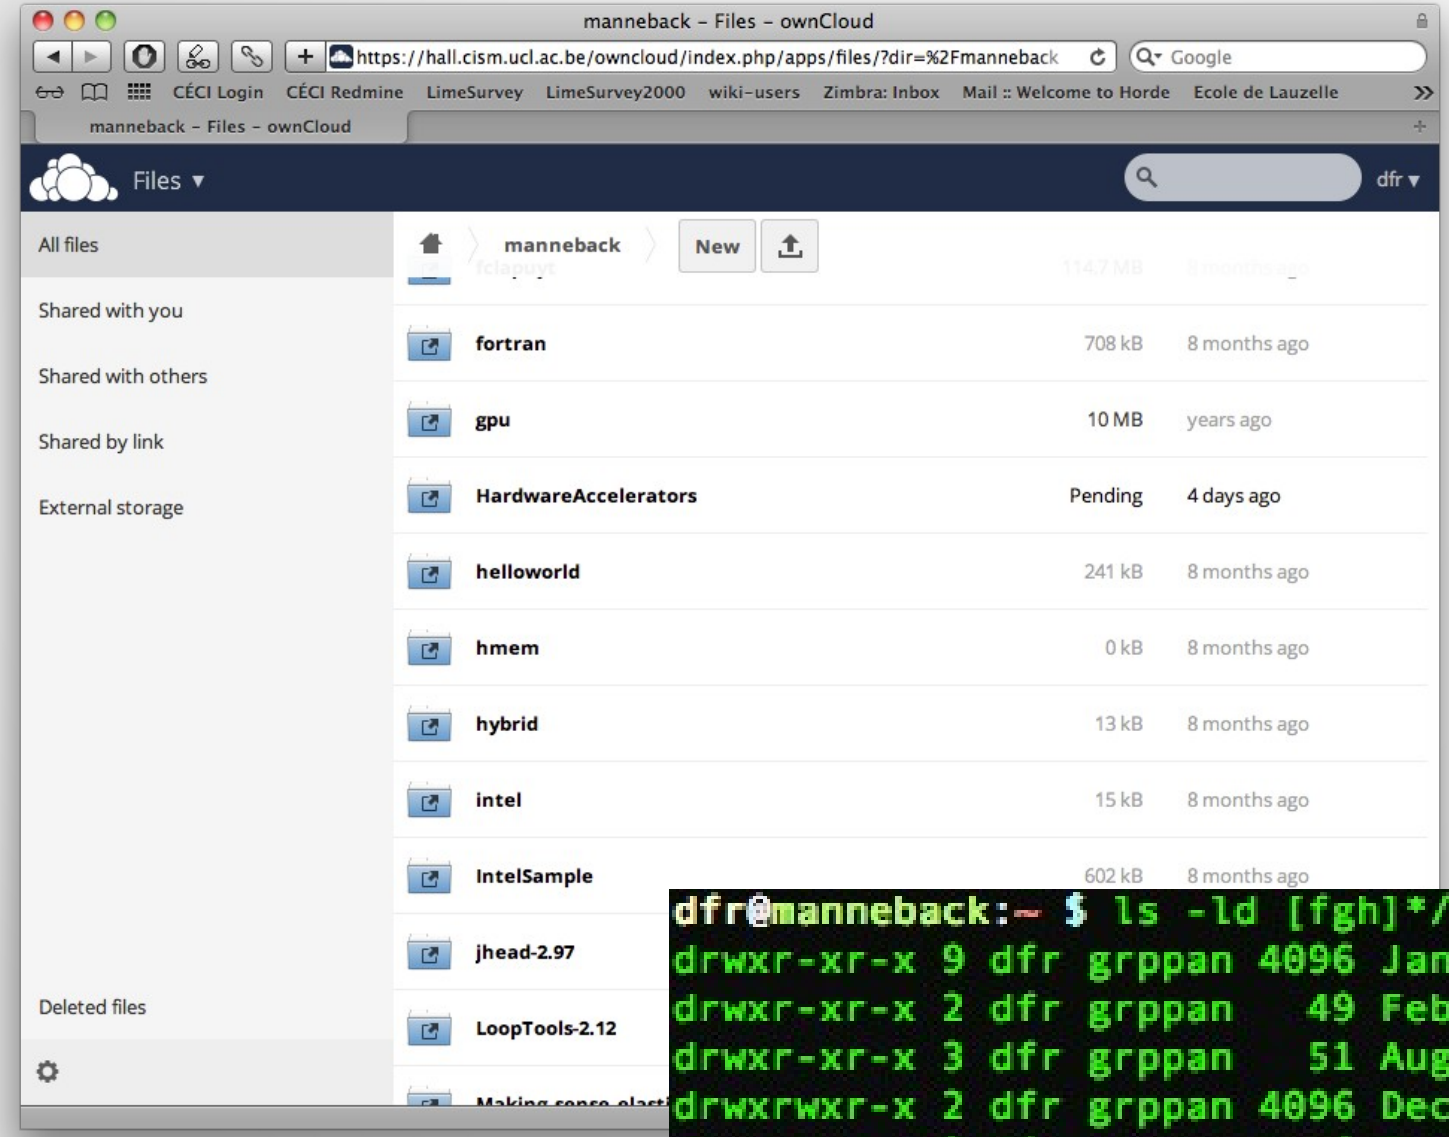


```
dfr@manneback:~ $ ls -ld [fgh]*/
drwxr-xr-x 9 dfr grppan 4096 Jan 28  2015 fffclapuyt/
drwxr-xr-x 2 dfr grppan   49 Feb  1  2013 fortran/
drwxr-xr-x 3 dfr grppan   51 Aug 28  2013 gpu/
drwxrwxr-x 2 dfr grppan 4096 Dec 12  2013 helloworld/
drwxr-xr-x 2 dfr grppan    6 Mar 16  2015 hmem/
drwxr-xr-x 2 dfr grppan   43 Nov  8  2013 hybrid/
```

# Can create a share URL

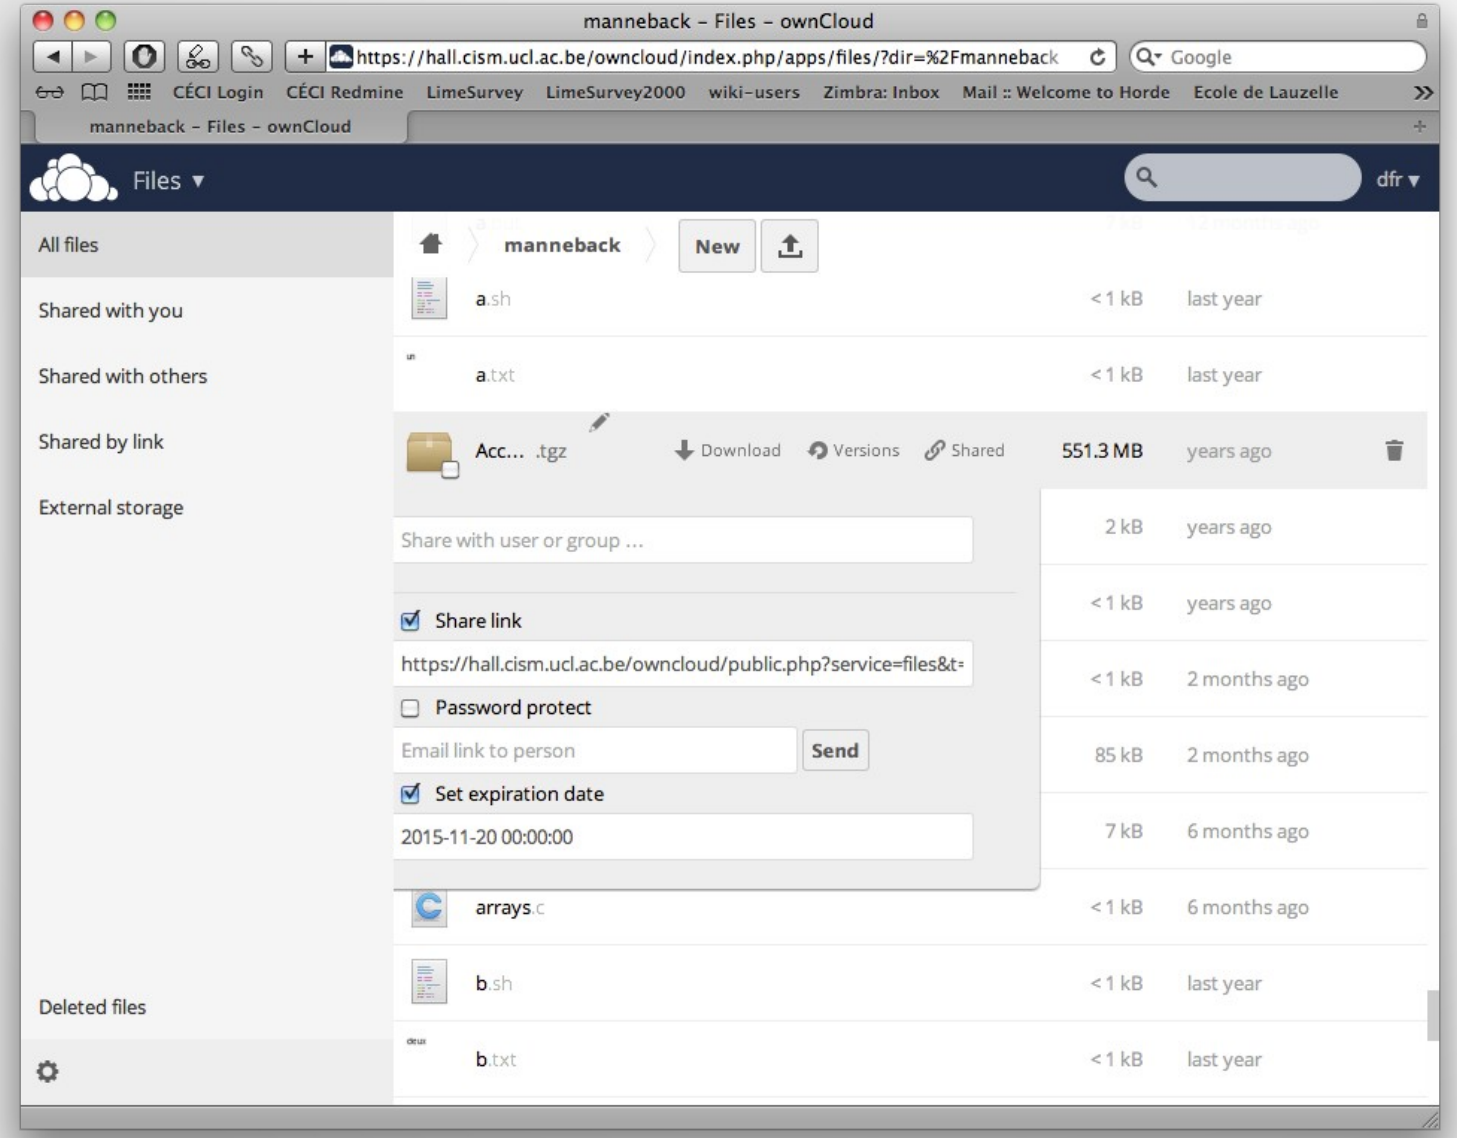

# And distribute it

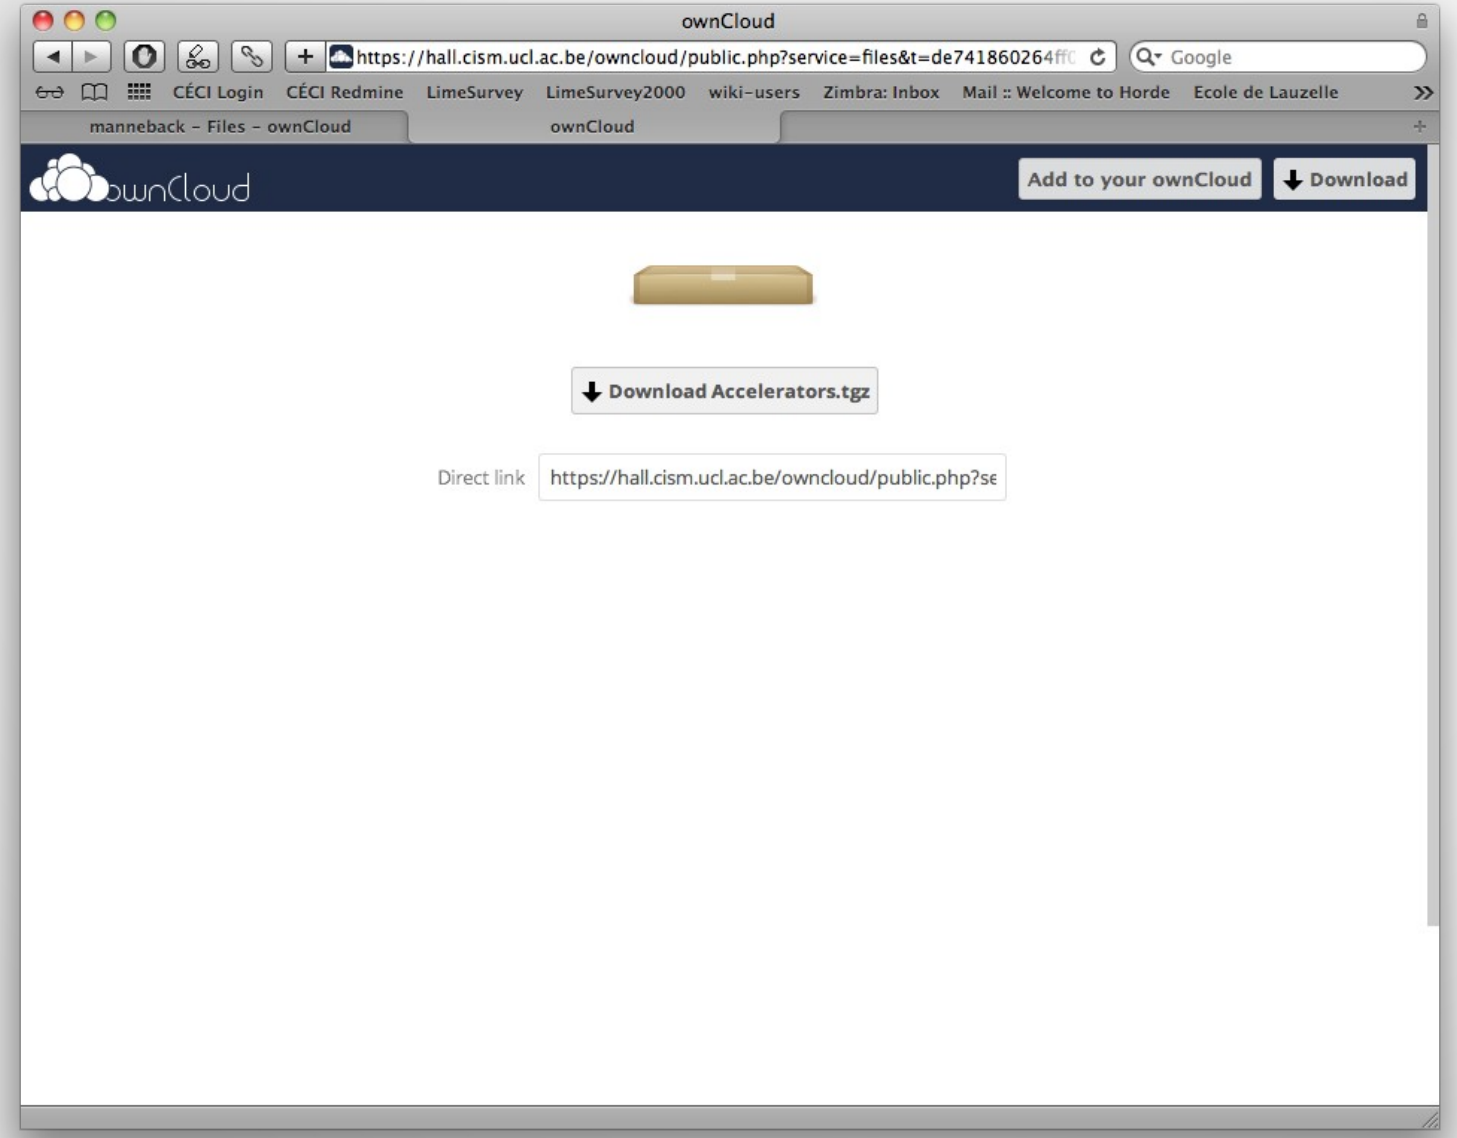

# Open data – FAIR data

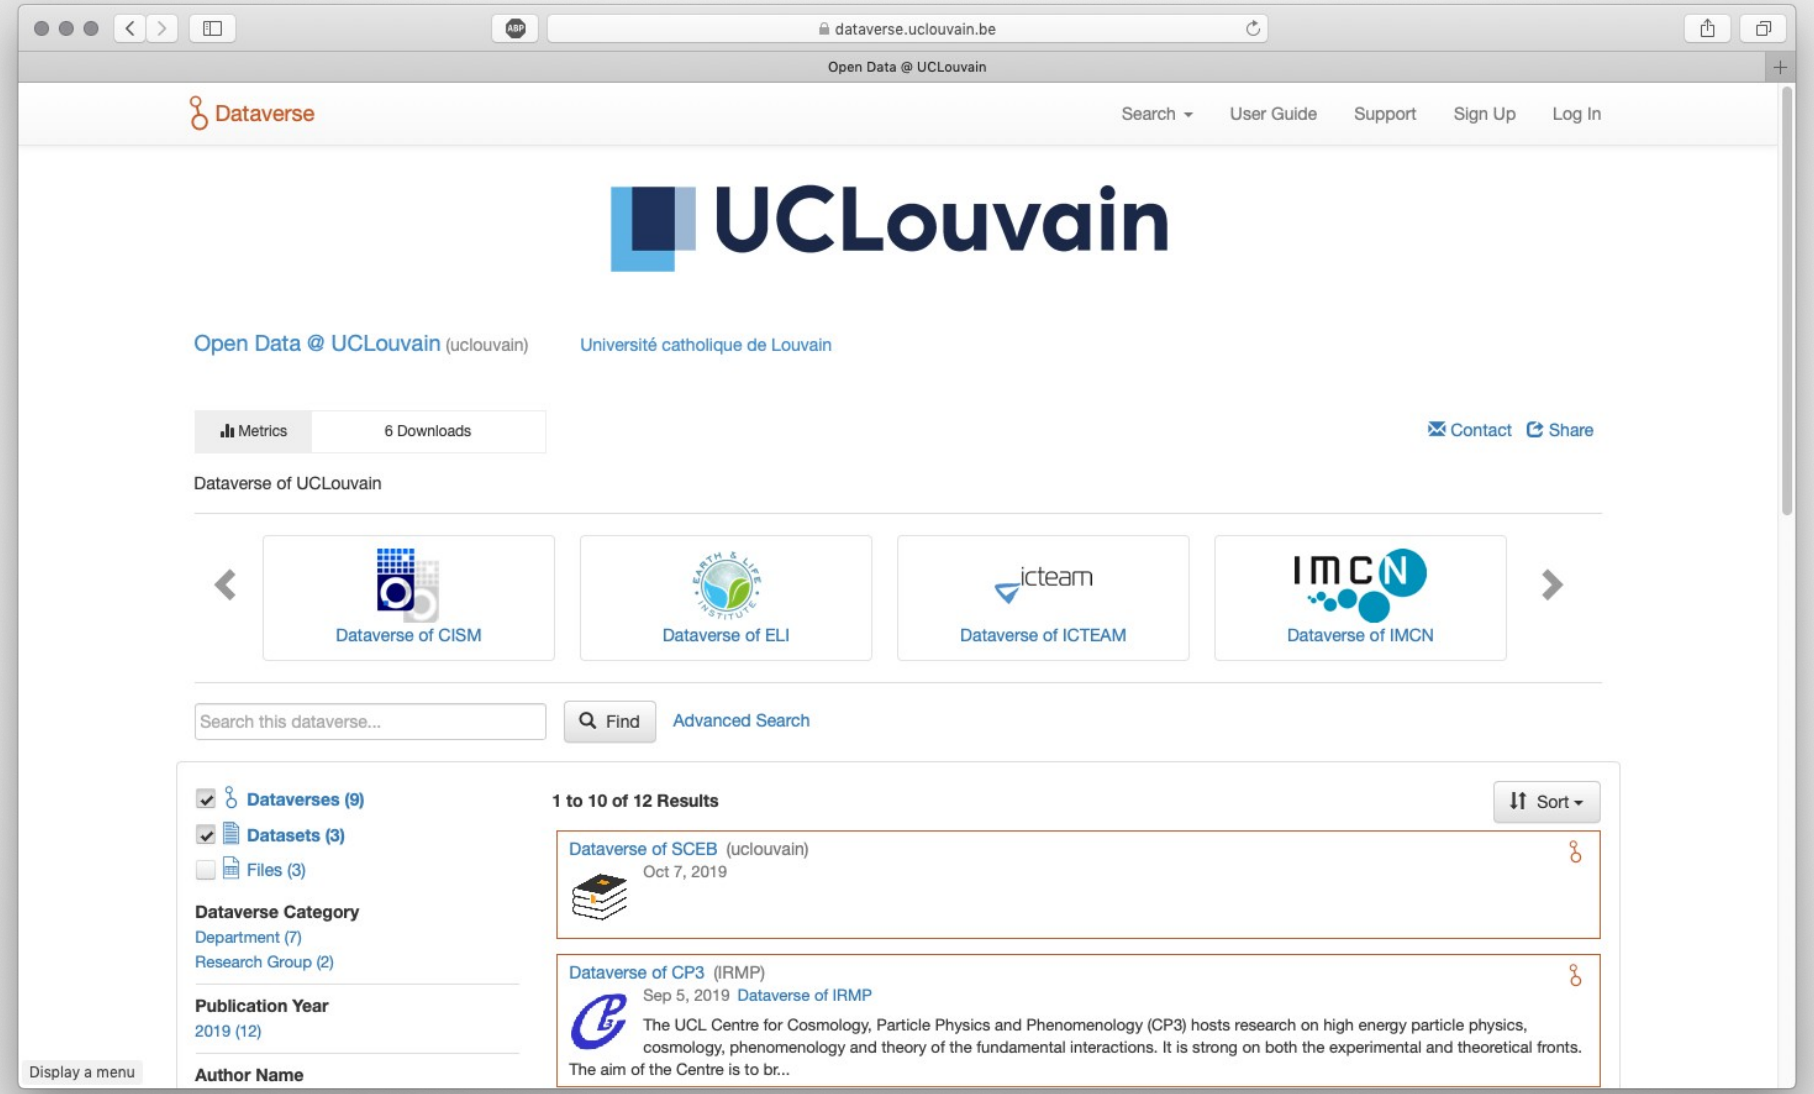

# Open data – FAIR data

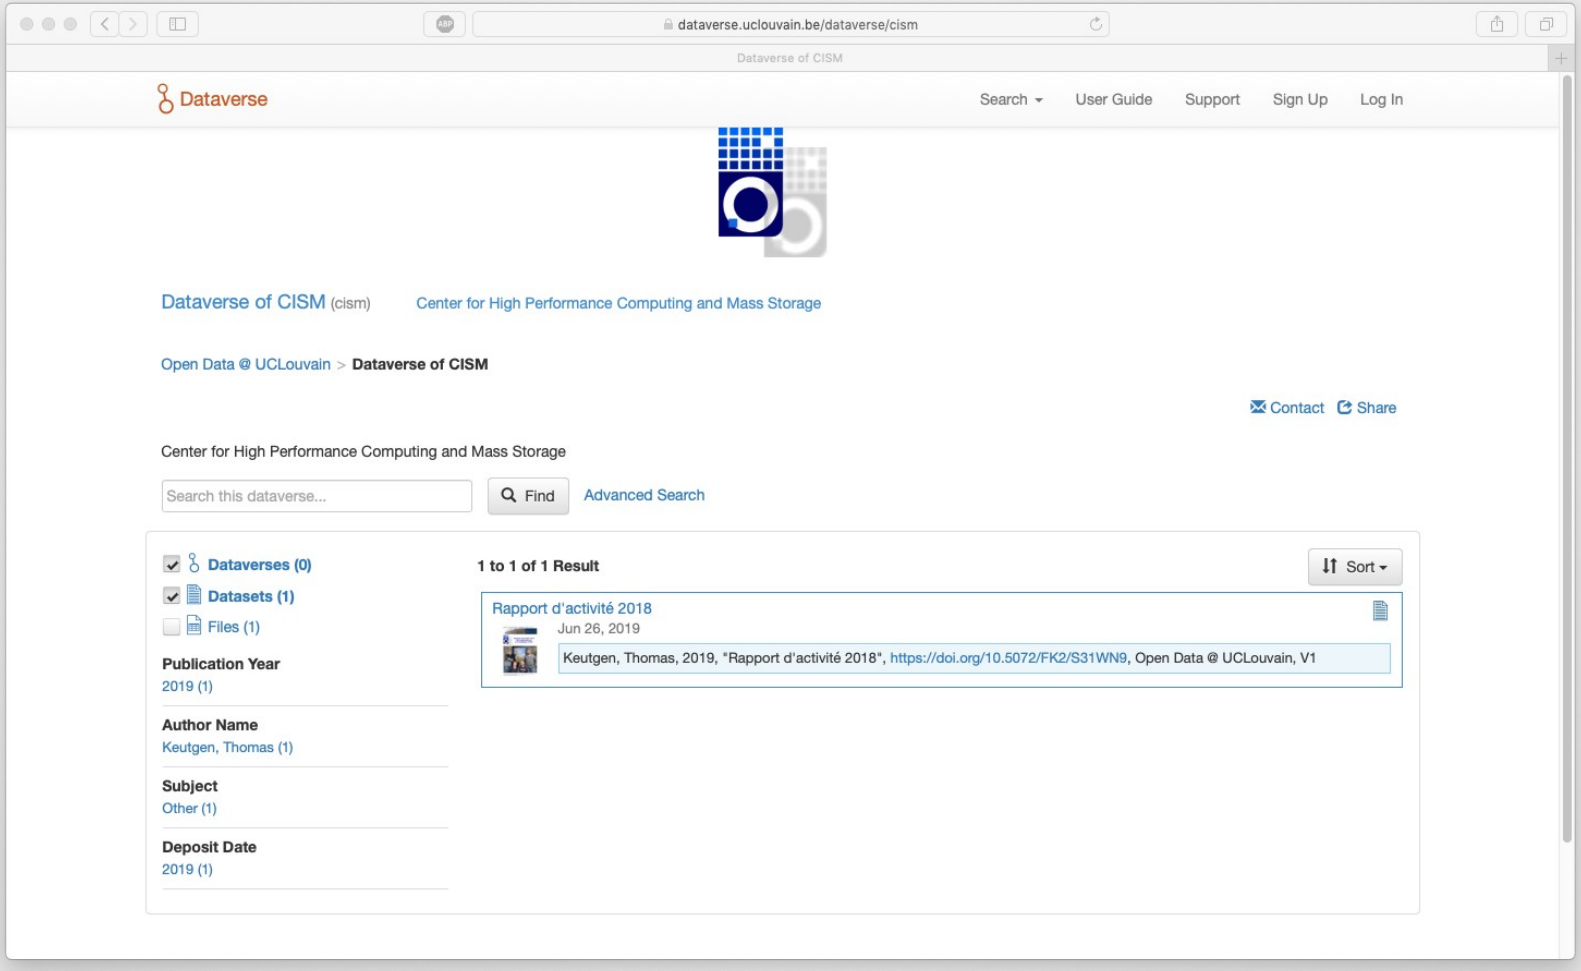

# Summary:

Storage: choose the right filesystem and the right file format

Transfer: use the parallel tools when possible

Sharing: use all the potential of the UNIX permissions and try Nextcloud and Dataverse