

BioHPC User Meeting

February 2019

Storage/Network Structure and Access Groups

Jaroslaw Pillardy

Bioinformatics Facility
Institute of Biotechnology
Cornell University

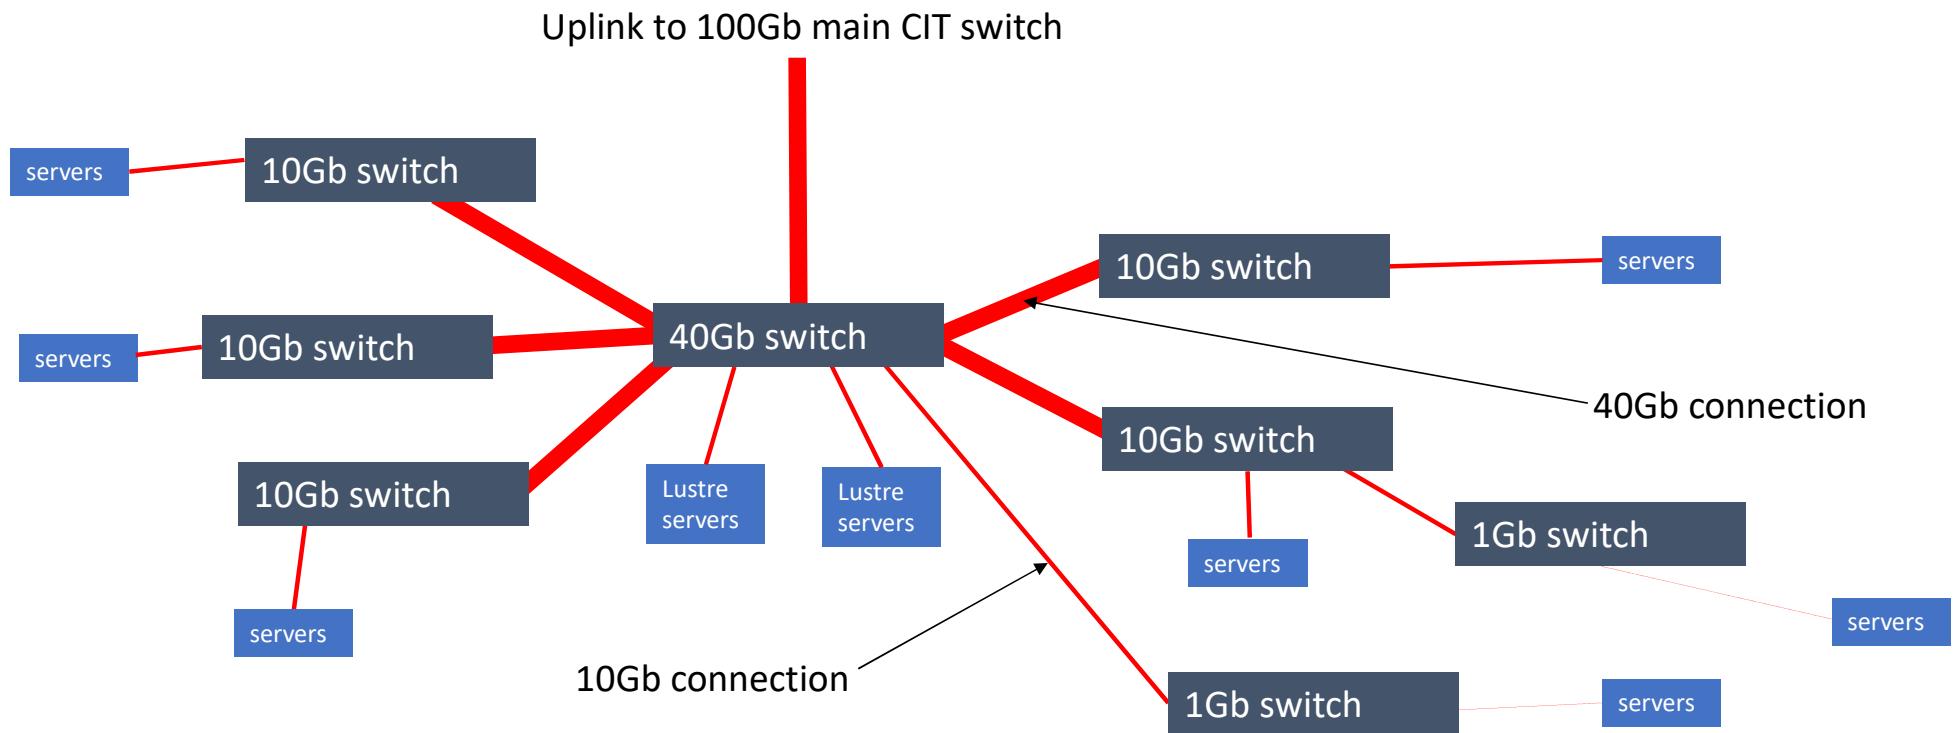
Three storage types available at BioHPC

- **Lustre network storage**
Mounted on all servers by default (/home, /programs, /shared_storage).
Fast, cluster based storage.
- **Gluster network storage**
Mounted on all servers (/glusterfs/home). Slow but reliable cluster based storage. Will be converted to second Lustre cluster for temporary storage.
- **Local storage**
/local on all servers with a symlink to /workdir. Some servers have /SSD.
Exportable via NFS.

BioHPC Cloud Network

Backbone 40Gb switch has 6 free modules, each may connect 8x 10Gb or 4x 40Gb

Typical building network is 1Gb or less. Building to building is 1Gb or 10Gb.



Lustre

Cluster storage system composed of many servers fulfilling different roles.

Resulting storage is visible to users as a single file system.

Data operations are executed in parallel involving various component servers.

Lustre

- **Management servers (MGS: 2)**
Redirect and balance requests, act as point of contact and cache. Two servers cbsumgs1-2.
- **Metadata servers (MDS: 2)**
Handle all file/dir information - a big storage metadata database. Main server (cbsumds1) is block-level mirrored (with blocking) to cbsumds2
- **Object Storage Servers (OSS: 10)**
Store actual files (directories are in MDS). One file per server, no stripping, for file safety. Parallel access to files via multiple servers. ZFS file system, two raidz2 devices per server. Servers cbsuoss01-10.

Lustre

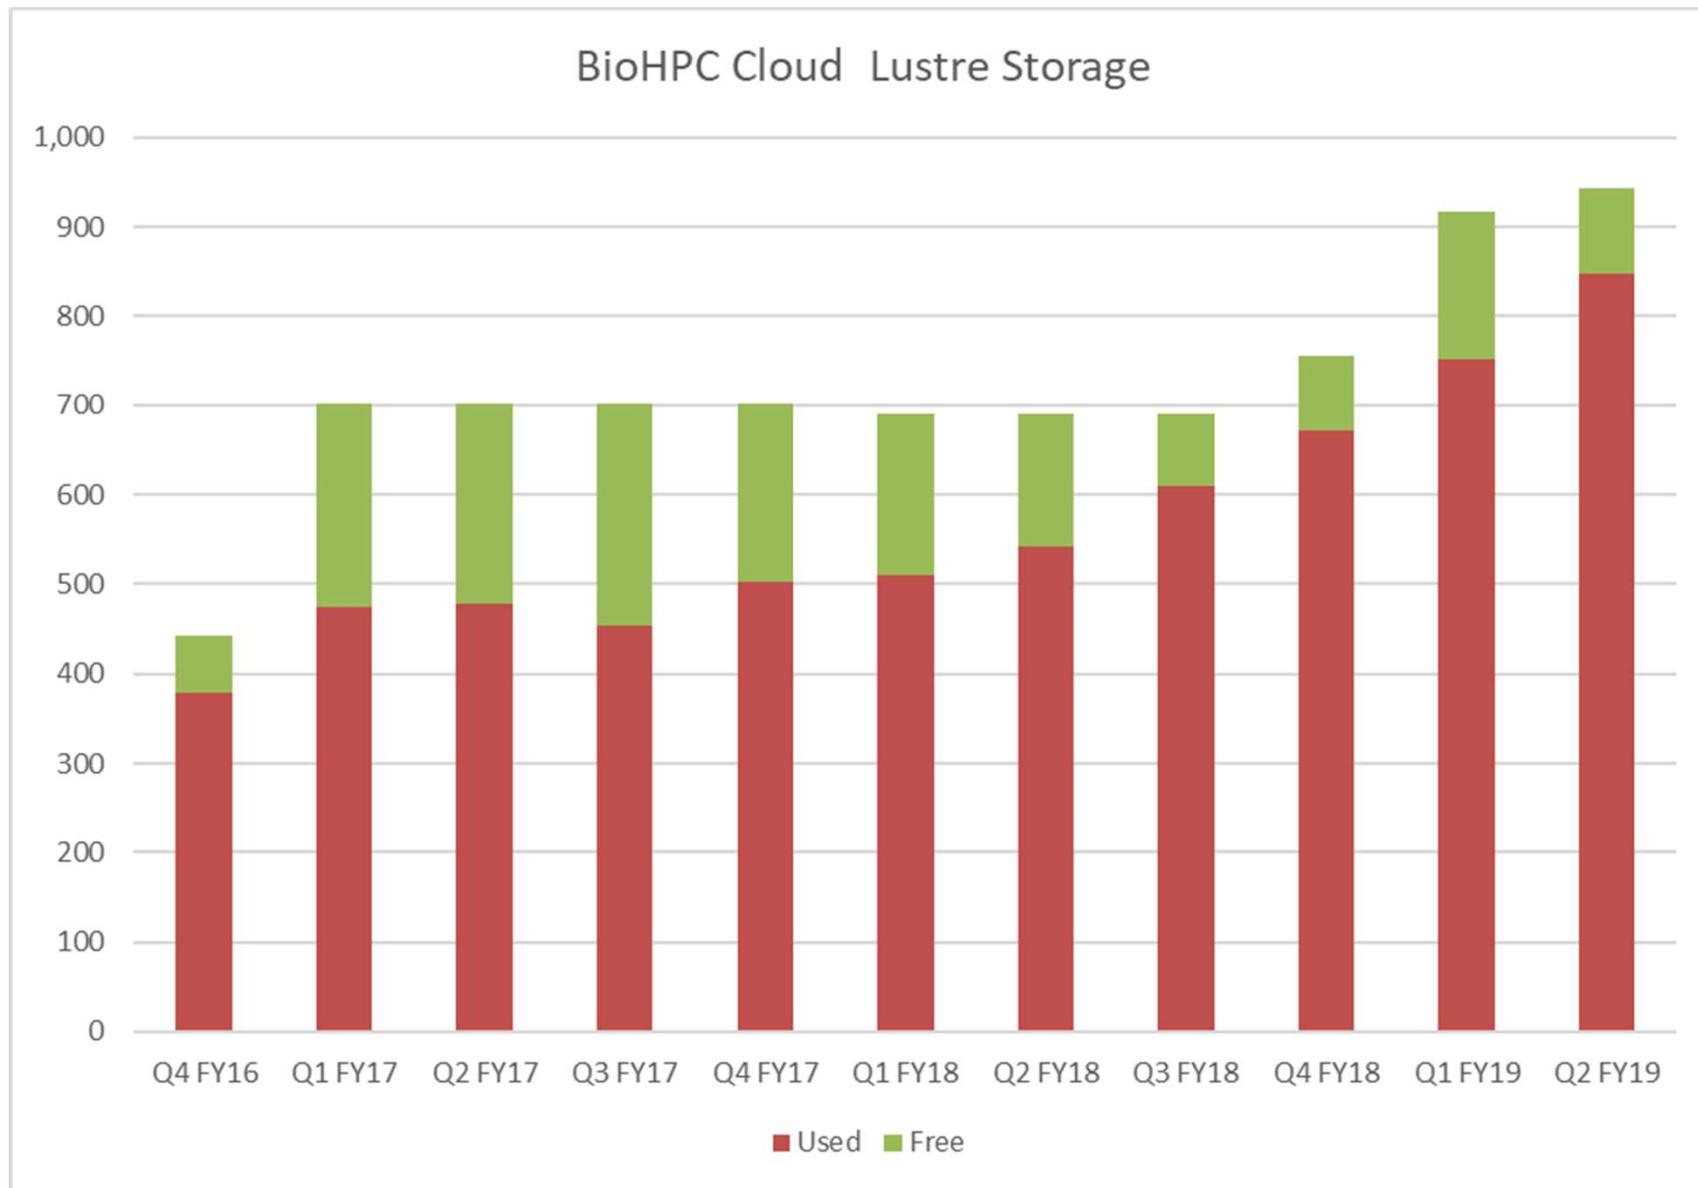
```
[root@cbsumgs2 ~]# lfs df -h /lustre1
  UUID           bytes   Used  Available Use% Mounted on
lustre1-MDT0000_UUID    950.1G   58.6G    827.4G   7% /lustre1[MDT:0]
lustre1-OST0000_UUID    34.6T   32.4T     2.2T  94% /lustre1[OST:0]
lustre1-OST0001_UUID    34.6T   32.1T     2.5T  93% /lustre1[OST:1]
lustre1-OST0002_UUID    34.6T   32.7T     1.9T  95% /lustre1[OST:2]
lustre1-OST0003_UUID    34.6T   32.6T     2.0T  94% /lustre1[OST:3]
lustre1-OST0004_UUID    34.6T   32.6T     2.0T  94% /lustre1[OST:4]
lustre1-OST0005_UUID    34.6T   32.5T     2.1T  94% /lustre1[OST:5]
lustre1-OST0006_UUID    34.6T   33.4T     1.2T  96% /lustre1[OST:6]
lustre1-OST0007_UUID    34.6T   33.4T     1.2T  96% /lustre1[OST:7]
lustre1-OST0008_UUID    47.6T   46.0T     1.6T  97% /lustre1[OST:8]
lustre1-OST0009_UUID    47.6T   46.2T     1.4T  97% /lustre1[OST:9]
lustre1-OST000a_UUID    31.8T   29.7T     2.2T  93% /lustre1[OST:10]
lustre1-OST000b_UUID    31.8T   29.8T     2.1T  93% /lustre1[OST:11]
lustre1-OST000c_UUID    63.7T   58.3T     5.4T  92% /lustre1[OST:12]
lustre1-OST000d_UUID    63.7T   58.0T     5.7T  91% /lustre1[OST:13]
lustre1-OST000e_UUID    63.7T   58.7T     5.0T  92% /lustre1[OST:14]
lustre1-OST000f_UUID    63.7T   58.0T     5.7T  91% /lustre1[OST:15]
lustre1-OST0010_UUID    31.8T   2.2T    29.6T   7% /lustre1[OST:16]
lustre1-OST0011_UUID    31.8T   5.8T    26.0T  18% /lustre1[OST:17]
lustre1-OST0012_UUID    95.2T   93.3T    1.9T  98% /lustre1[OST:18]
lustre1-OST0013_UUID    95.2T   83.6T    11.6T  88% /lustre1[OST:19]

filesystem_summary:      944.7T   831.3T   113.4T  88% /lustre1
```

Lustre

```
[root@cbsumgs2 ~]# lfs df -i /lustre1
      UUID          Inodes      IUsed      IFree  IUse% Mounted on
lustre1-MDT0000_UUID    672301056   266856244   405444812   40% /lustre1[MDT:0]
lustre1-OST0000_UUID    27796494    9430377   18366117   34% /lustre1[OST:0]
lustre1-OST0001_UUID    31020027   9659484    21360543   31% /lustre1[OST:1]
lustre1-OST0002_UUID    25153042   9330381   15822661   37% /lustre1[OST:2]
lustre1-OST0003_UUID    26010982   9354355   16656627   36% /lustre1[OST:3]
lustre1-OST0004_UUID    27132994  10141234   16991760   37% /lustre1[OST:4]
lustre1-OST0005_UUID    27101972   9741489   17360483   36% /lustre1[OST:5]
lustre1-OST0006_UUID    17117485   6888172   10229313   40% /lustre1[OST:6]
lustre1-OST0007_UUID    17733920   7353155   10380765   41% /lustre1[OST:7]
lustre1-OST0008_UUID    22285884   8623826   13662058   39% /lustre1[OST:8]
lustre1-OST0009_UUID    20779752   8959511   11820241   43% /lustre1[OST:9]
lustre1-OST000a_UUID    27060944   8789946   18270998   32% /lustre1[OST:10]
lustre1-OST000b_UUID    25866822   8406767   17460055   33% /lustre1[OST:11]
lustre1-OST000c_UUID    66056673   20673611   45383062   31% /lustre1[OST:12]
lustre1-OST000d_UUID    68925896   21466991   47458905   31% /lustre1[OST:13]
lustre1-OST000e_UUID    63273313   21412209   41861104   34% /lustre1[OST:14]
lustre1-OST000f_UUID    68330848   20531501   47799347   30% /lustre1[OST:15]
lustre1-OST0010_UUID    251557686   3012169   248545517   1% /lustre1[OST:16]
lustre1-OST0011_UUID    223654683   5033866   218620817   2% /lustre1[OST:17]
lustre1-OST0012_UUID    46183470   30625956   15557514   66% /lustre1[OST:18]
lustre1-OST0013_UUID    125925664   28848028   97077636   23% /lustre1[OST:19]

filesystem_summary:   672301056   266856244   405444812   40% /lustre1
```



Lustre

- Currently 944.7TB capacity, will be upgraded soon.
- Small files are bad for all known file systems ... including Lustre. For best access keep number of small files to minimum – preferably archived with tar.
- Fast, but not fast enough to directly compute.

Local storage

- On rental machines only /workdir (/local/workdir), not exportable
- On hosted machines /local/storage is exported to all BioHPC (including login nodes). Access is for hosted group only.

```
/programs/bin/labutils/mount_server server /storage
```

It will be mounted as `/fs/server/storage`

Groups

- We use groups for access control to all BioHPC resources: files/directories, storage, servers, credit accounts
- BioHPC group is stored in central database, it is mapped to Linux (LDAP) group for files/directories access
- BioHPC groups can be combined into “supergroup”, an automatic union of several groups. Each is then mapped to Linux (LDAP) group.

Groups

- Each group has an owner, who can add/remove members and is notified of any membership changes
- Group management via “My Groups” menu
- Group can also be a paid storage group (in addition to other functions), if it has the corresponding directory `/home/groupname` where its storage can be accessed.

institute of biotechnology >> brc >> bioinformatics >> internal >> biohpc cloud: groups management

BIOHPC CLOUD: : GROUPS MANAGEMENT

Filter by:

Name: *

Description: *

Show inactive groups

Order by

Index	Name	LDAP	Description	Status	Acct class	Special account	Owner	ACL	Linked group	Action
1	CBSU Staff		CBSU Staff	active	none	CBSU	jarekp	0	0	Group Users
4	Restricted Collab		Access to standard restricted workstations for collaborators	active	general	none	jarekp	0	0	Group Users
8	Special		This is a group for temporary access to selected machines.	active	none	CBSU	jarekp	0	0	Group Users
13	IGD special		For IGD members to access cbsufsrv4:/data1 and do some calculations.	active	none	IGD data	jarekp	0	0	Group Users
17	cbsulm07		For unlimited access to cbsulm07	active	none	cbsulm07	jarekp	0	0	Group Users
19	bscb00	bscb00	BSCB Department group	active	none	none	jarekp	0	0	Group Users
27	qisun_c4	qisun_c4		active	none	none	jarekp	0	0	Group Users
28	proteomics_share	proteomics		active	none	none	jarekp	0	0	Group Users
64	cbsufsrv10		Group to allow access to cbsufsrv10 for dbGaP	active	none	none	jarekp	64347	0	Group Users
74	genomics_share	genomics_share		active	none	none	jarekp		0	Group Users
123	metabolomics_share	metabolomics		active	none	none	jarekp		0	Group Users
185	dbGaP_jrflab	dbGaP_jrflab		active	none	none	jarekp		0	Group Users

Manage Credit Accounts

My Storage

Profile

Reservations

My Reservations

My Groups

Temporary Accounts

Change Password

Logout



BIOHPC CLOUD: : GROUP USERS MANAGEMENT



Group 'qisun_c4'

Filter by:

Lab ID *

Last name *

Cornell ID *

Created >= MM/DD/YYYY and < MM/DD/YYYY

Email *

Status active

Group --none--

Comment *

Order by Lab ID Ascending

Index	First Name	Last Name	Lab ID	Num Lab ID	Cornell ID	E-mail	Institution	Department	Created	Status	Web indx	Disk Usage (GB)	Disk Usage Update	Home	Groups	Comment	Action
13	Chris	Myers	crm17	4967	crm17	c.myers@cornell.edu	Cornell University	CBSU	12/21/2010 12:58:06 PM	active		17.1	1/26/2019 8:01:01 PM	/home/crm17			Delete Move home to group storage
1	Jaroslaw	Pillardy	jarekp	4965	jp86	jp86@cornell.edu	Cornell University	Biotech	11/18/2010 5:11:35 PM	active	3	496.0	1/26/2019 8:01:01 PM	/home/jarekp			Delete Move home to group storage
66	Qi	Sun	qisun	5034	qs24	qs24@cornell.edu	Cornell University	CBSU	5/20/2011 3:06:00 PM	active	4	4,761.4	1/26/2019 8:01:01 PM	/home/qisun			Delete Move home to group storage
3657	Zehong	Ding	zehong	1251		dingzehong@itbb.org.cn	Cornell University	CBSU	7/3/2013 10:29:16 AM	active	847	49.3	1/26/2019 8:01:01 PM	/home/zehong			Delete Move home to group storage

4 Records found. Show 1-4

Add user with labid to the group

Website credentials: user: jarekp 'jp86@cornell.edu' [BioHPC Cloud]

[View A...](#)

Groups for directory access

- Two regular groups: *group1* and *group2*
- One supergroup *groupall* being a union of *group1* and *group2*.
- Please note the groups have to be created by admins, but can be managed by owners.

Disclaimer: next slide uses colors for better readability. Since I am partially color blind color choice may be disturbing to some viewers. Viewer discretion is advised.

/home
drwxr-xr-x root root
all users read access

/groupall
drwxr-x--- root groupall
all groups read access
other users no access

/directory1
drwxrwx--- root groupall
all groups full access

/directory2
drwxrwx--- root group1
all group1 users full access
other groups no access

/directory3
drwxrwxr-x root group2
all group2 users full access
other groups read access

Thank you!