

Uni.lu HPC School 2019

Keynote/PS9: User environment and storage data management



Uni.lu High Performance Computing (HPC) Team

S. Peter

University of Luxembourg (UL), Luxembourg

<http://hpc.uni.lu>



Latest versions available on **GitHub**:



UL HPC tutorials:

<https://github.com/ULHPC/tutorials>

UL HPC School:

<http://hpc.uni.lu/hpc-school/>

Keynote/PS9 tutorial sources:

ulhpc-tutorials.rtf.d.io/en/latest/





Summary

1 Overview of the data management within UL HPC

[Big] Data components in HPC

Shared Storage on UL HPC

User environment

2 Daily Data Management

Quotas

Backup

Version control with Git

GDPR

Learn more

3 Migration from Gaia & Chaos to Iris

4 Q & A session



[Big]Data Management: FS Summary

- **File System (FS):** Logical manner to *store, organize & access* data
 - ↪ (local) **Disk FS** : FAT32, NTFS, HFS+, ext4, {x,z,btr}fs...
 - ↪ **Networked FS**: NFS, CIFS/SMB, AFP
 - ↪ **Parallel/Distributed FS**: SpectrumScale/GPFS, Lustre
 - ✓ typical FS for HPC / HTC (High Throughput Computing)

[Big]Data Management: FS Summary

- **File System (FS):** Logical manner to *store, organize & access* data
 - ↪ (local) **Disk FS** : FAT32, NTFS, HFS+, ext4, {x,z,btr}fs...
 - ↪ **Networked FS**: NFS, CIFS/SMB, AFP
 - ↪ **Parallel/Distributed FS**: SpectrumScale/GPFS, Lustre
 - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

[Big]Data Management: FS Summary

- **File System (FS):** Logical manner to *store, organize & access* data
 - ↪ (local) **Disk FS** : FAT32, NTFS, HFS+, ext4, {x,z,btr}fs...
 - ↪ **Networked FS**: NFS, CIFS/SMB, AFP
 - ↪ **Parallel/Distributed FS**: SpectrumScale/GPFS, Lustre
 - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

Name	Type	Read* [GB/s]	Write* [GB/s]
ext4	Disk FS	0.426	0.212
nfs	Networked FS	0.381	0.090
gpfs (iris)	Parallel/Distributed FS	11.25	9.46
lustre (iris)	Parallel/Distributed FS	12.88	10.07
gpfs (gaia)	Parallel/Distributed FS	7.74	6.524
lustre (gaia)	Parallel/Distributed FS	4.5	2.956

* maximum **random** read/write, per IOZone or IOR measures, using concurrent nodes for networked FS.

UL HPC Storage capacity



9852.4 TB (incl. 1020TB for Backup)

2425 disks

- 4 distributed/parallel FS
 - ↪ GPFS : 3244 TB
 - ↪ Lustre: 1940 TB
 - ↪ OneFS: 3188 TB...

Understanding Your Storage Options

Where can I store and manipulate my data?

- **Shared storage**
 - ↳ NFS - **not scalable** $\sim \simeq 1.5$ GB/s (R) $\mathcal{O}(100$ TB)
 - ↳ GPFS - **scalable** $\sim \simeq 10$ GB/s (R) $\mathcal{O}(1$ PB)
 - ↳ Lustre - **scalable** $\sim \simeq 5$ GB/s (R) $\mathcal{O}(0.5$ PB)
- **Local storage**
 - ↳ local file system (/tmp) $\mathcal{O}(200$ GB)
 - ✓ over HDD $\simeq 100$ MB/s, over SSD $\simeq 400$ MB/s
 - ↳ RAM (/dev/shm) $\simeq 30$ GB/s (R) $\mathcal{O}(20$ GB)
- **Distributed storage**
 - ↳ HDFS, Ceph, GlusterFS - **scalable** $\sim \simeq 1$ GB/s

⇒ **In all cases:** small I/Os really **kill** storage performances



UL HPC Shared Storage Capacities

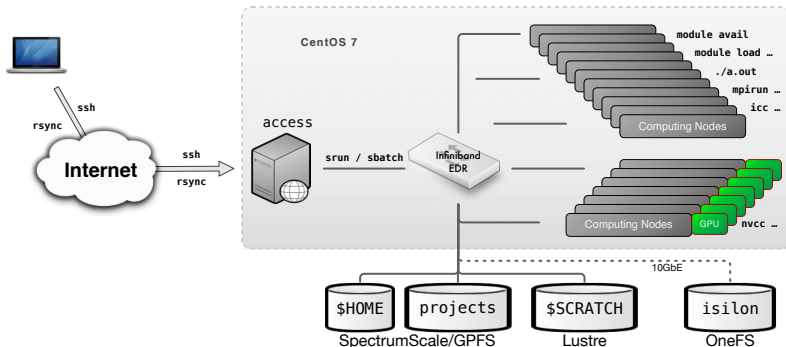
Cluster	GPFS	Lustre	Other	Backup
iris	2284	1280	6/3188 ²	600
gaia ¹	960	660	0/3188 ²	240
chaos ¹	0	0	180	180
g5k	0	0	32.4	0
nyx ¹ (experimental)	0	0	242	0
TOTAL:	3244 TB	1940 TB	3648.4 TB	1020 TB

¹: *Deprecated end-2019!!*

²: *Common Isilon/OneFS shared storage mounted on gaia and iris*

Uni.lu HPC Total Storage Capacity:
9852.4 TB

Compute Nodes Environment





Where is what

Directory	Env variable	Filesystem
/home/users	\$HOME	SpectrumScale
/work/projects	-	SpectrumScale
/scratch/users	\$SCRATCH	Lustre
/mnt/isilon/projects	-	OneFS



How to use

Directory	Usage
/home/users	personal space, software & packages
/work/projects	shared project storage
/scratch/users	intermediate fast storage, work here
/mnt/isilon/projects	archival storage, do not use for processing



Summary

- 1 Overview of the data management within UL HPC
 - [Big] Data components in HPC
 - Shared Storage on UL HPC
 - User environment
- 2 **Daily Data Management**
 - Quotas
 - Backup
 - Version control with Git
 - GDPR
 - Learn more
- 3 Migration from Gaia & Chaos to Iris
- 4 Q & A session



Quotas

Check file size quota with

```
df-ulhpc
```

Check inode quota with

```
df-ulhpc -i
```

Check free space on all file systems with

```
df -h
```

Check free space on current file system with

```
df -h .
```



Warning

88 % usage on Gaia GPFS

- No new projects or quota increase anymore.
- You need to move to Iris!



Default quotas

Directory	size quota	inode quota
\$HOME	500 GB	1,000,000
\$SCRATCH	10 TB	1,000,000
/work/projects/...	16 MB	-
/isilon/projects/...	990 TB globally	-



Backup

- **NO** backup in \$SCRATCH (/scratch or /tmp) directories
- **Cleanup:** files in \$SCRATCH older than 60 days are removed every month
- **Cleanup:** files in /tmp on compute nodes are removed at the end of the job



Backup: Iris

- \$HOME
 - ↪ daily backup to another server in the same data center
 - ↪ rotation: last 7 daily backups, one per month for the last 6 months
- /work/projects
 - ↪ daily backup to another server in the same data center
 - ↪ rotation: last 7 daily backups, one per month for the last 6 months



Backup: Isilon (HPC share)

- `/mnt/isilon/projects`
- weekly snapshot
- rotation: only one snapshot kept
- no true backup, because it's on the same system



Version control

- “backup” for your code
- benefits (from [Atlassian](#)):
 - ↪ complete long-term change history of every file
 - ↪ branching and merging
 - ↪ traceability
- relevant for GDPR compliance



Gitlab.uni.lu

- local **GitLab** instance hosted by HPC
- data stays within UL
- as many private repositories as you want
- access for external collaborators with Github account

Git in practice

Basic workflow

- **Pull** latest changes `git pull`
- **Edit** files `vim / emacs / subl ...`
- **Stage** the changes `git add`
- **Review** your changes `git status`
- **Commit** the changes `git commit`

Git in practice

Basic workflow

- **Pull** latest changes `git pull`
- **Edit** files `vim / emacs / subl ...`
- **Stage** the changes `git add`
- **Review** your changes `git status`
- **Commit** the changes `git commit`

For cheaters: An even more basic workflow

- **Pull** latest changes `git pull`
- **Edit** files `vim / emacs / subl ...`
- **Stage & commit** all the changes `git commit -a`

Git Summary

- **Advices: Commit early, commit often!**

- ↪ commits = save points
 - ✓ use descriptive commit messages
- ↪ Do not get out of sync with your collaborators
- ↪ Commit the sources, not the derived files

- **Not covered here (by lack of time)**




- ↪ does not mean you should not dig into it!
- ↪ *Resources:*
 - ✓ <https://git-scm.com/>
 - ✓ tutorial: IT/Dev[op]s Army Knives Tools for the Researcher
 - ✓ tutorial: Reproducible Research at the Cloud Era

Git Summary

<https://github.com/louim/in-case-of-fire>

In case of fire



-  1. `git commit`
-  2. `git push`
-  3. `leave building`

GDPR and UL HPC



www.eugdpr.org



- EU General Data Protection Regulation (**GDPR**)
 - ↳ replaces the Data Protection Directive 95/46/EC
 - ↳ legislation came into effect May 25th 2018

- **The UL HPC facility handles both:**

- ↳ **data about people** (facility users identification details)
 - ✓ ULHPC Identity Management (IdM) system
 - ✓ Account request form results
- ↳ **large scale data** that may contain Personally Identifiable Info
 - ✓ stored by facility users in networked, parallel & distributed filesystems used across the HPC infrastructure
 - ✓ can be considered as falling under GDPR regulations.



GDPR and UL HPC

- **Personal data is/may be visible, accessible or handled:**

- ↳ directly on the HPC clusters

- ↳ through *Resource and Job Management System* (RJMS) tools

- ✓ glue for a parallel computer to execute parallel jobs

- ✓ **Goal:** satisfy users demands for computation

- ✓ comes with web interfaces

Monika, Ganttchart

- ↳ through service portals

hpc-tracker, XCS, Galaxy

- ↳ on code management portals

GitLab, GitHub

- ↳ on secondary storage systems

DropIT, OwnCloud

Towards a ULHPC QoS Master Plan

Objectives

- **Formalizing the way we tackle security hardening**
 - ↪ Work in progress with **continuous improvement**
 - ↪ Completes other initiatives at SIU, LCSB, SnT etc.
 - ↪ Ongoing adaptation to match GDPR compliance
 - ↪ In line with UL guidelines

Best practices for you

General

- data (pseudo-) anonymisation
- data minimisation
- data partitioning
- secure laptop
 - ↪ enable FileVault / disk encryption
 - ↪ lock your screen when you leave your place
 - ↪ apply security updates
 - ↪ anti-virus / anti-malware software
 - ↪ (encrypted) backup of your laptop
- secure access credentials
 - ↪ consider using a password manager
 - ↪ use 2FA when possible (authenticator better than SMS)

Best practices for you

On ULHPC

- double-check permissions on your \$HOME and \$SCRATCH folders
- secure your SSH key with a passphrase
- empty /tmp at the end of the job
- reserve a full node
- store your data on iris (SED)
- mind backups
- encrypt your files with gocryptfs
- (soon) enable two factor authentication (e.g. with TACC OpenMFA)

Workshop from Elixir Luxembourg

Research data management and stewardship

June 25 - 26, Luxembourg Learning Centre

Lectures and hands-on sessions on:

- Understanding the FAIR Principles for data
- What is data stewardship and how it is done in practice
- Data management planning
- Scientific and computational reproducibility of research
- Working with Human Data and Data Protection obligations



Register today on elixir-luxembourg.org!



Summary

- 1 Overview of the data management within UL HPC
 - [Big] Data components in HPC
 - Shared Storage on UL HPC
 - User environment
- 2 Daily Data Management
 - Quotas
 - Backup
 - Version control with Git
 - GDPR
 - Learn more
- 3 Migration from Gaia & Chaos to Iris
- 4 Q & A session



Decommissioning timeline

- Users:
 - ↪ **July:** job submission will be limited, data fully accessible
 - ↪ **September:** no new jobs, data accessible in read only mode
 - ↪ **December:** your migration to the Iris cluster must be fully completed
- In the background:
 - ↪ improve connectivity between Gaia and Iris
 - ↪ prepare Iris storage for incoming data
 - ↪ transfer of project data from Gaia to Iris
 - ↪ meetings with PIs



Changes

- Scheduler: SLURM instead of OAR
 - ↪ Different command-line options
 - ↪ #SBATCH instead of #OAR in launcher scripts
 - ↪ Updated launcher scripts available at github.com/ULHPC/launcher-scripts
- Operating system: CentOS instead of Debian
 - ↪ You might need to recompile/reinstall your software

Data

- clean-up and pack data
- use rsync to transfer small amounts of data between clusters

```
rsync --bwlimit=10m --rsh='ssh -p 8022' --exclude="/.local" \  
--exclude="/.cache" -avzP . access-iris.uni.lu:~/gaia_home/
```

- make sure you have SSH keys set up
- run inside screen
- if you have many small files, consider packing them into one archive file
- for transfers > 10 TB contact HPC sysadmins



Summary

- 1 Overview of the data management within UL HPC
[Big] Data components in HPC
Shared Storage on UL HPC
User environment
- 2 Daily Data Management
Quotas
Backup
Version control with Git
GDPR
Learn more
- 3 Migration from Gaia & Chaos to Iris
- 4 **Q & A session**



Thank you for your attention...

Questions?

<http://hpc.uni.lu>

High Performance Computing @ uni.lu

Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clement Parisot
Dr. Frédéric Pinel
Dr. Emmanuel Kieffer

University of Luxembourg, Belval Campus
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu



1 Overview of the data management within UL HPC

[Big] Data components in HPC
Shared Storage on UL HPC
User environment

2 Daily Data Management

Quotas
Backup
Version control with Git
GDPR
Learn more

3 Migration from Gaia & Chaos to Iris

4 Q & A session