


Computation resources architecture

Below is the architecture model of our computation and storage resources.



Every LBT computing resources (HPC, storage and IS clusters) are internally designed and built, and based on open-source software and solutions.

Currently in the cluster room we have 1* HPC clusters (a.k.a. Supercomputer), called Baal (Lucifer for the standby login server), which is composed as follows:

- 2x Intel 24 cores for post-processing jobs
- 3x Intel 16 cores and bi Nvidia GTX-1080Ti GPU nodes (nodes 61 to 63)
- 11x Intel 40 cores and bi Nvidia GTX-1080Ti GPU nodes (nodes 64 to 73)
- 1x Intel 40 cores and bi Nvidia GTX-4070 GPU nodes (nodes 74 to 75)
- 1x Intel 40 cores and bi Nvidia RTX-2080Ti GPU nodes (node 76)
- 3x Intel 40 cores and bi Nvidia RTX-3080 GPU nodes (node 77 to 79)
- 1x Intel 40 cores and bi Nvidia RTX-3080Ti GPU node (node 80)
- 1x Intel 48 cores and bi Nvidia RTX-A6000 GPU nodes (node 81)
- 4x Intel 40 cores and bi Nvidia RTX-A5000 GPU nodes (node 82 to 85)

* the 2 first (Lucifer and Hades) have been completely dismantled.

Because it will not do any good to compute without storage capacity, following storage volumes are also available:

- **/workdir**: ~146TB (useful capacity)
- **/archive**: ~127TB* (useful capacity, without data compression or deduplication) replicated every weeks into another server with the same storage capacity
- **/archive/ibpc_team**: ~43.5TB* (useful capacity, without data compression or deduplication) dedicated for non-LBT members
- **/scratch**: except for the first 3 GPU nodes and post-processing nodes that have around 12TB each and all other GPU nodes with 1.6 to 2 SSD TB each, every nodes have around 200GB each

to let you use temporary computing files (mainly used for mono-node jobs)

- **/scratch-dfs**: ~142TB (useful capacity) for temporary computing files in parallelized job contexts

* maybe more considering ZFS compression and deduplication.



Some other disk spaces are not listed above because they are kept in reserve.

Except for **/archive/ibpc_team** and **/scratch** volumes, all above-mentioned volumes are **replicated and/or distributed** on a storage cluster currently composed by 11 servers.

/scratch volume on each node serves only to store temporary computing files. So, because it's cleared every night (deleting old job's directories and all not well-formed directories), you should not try to use it for chaining jobs in a single -but dedicated- scratch directory.



Archive volume (/archive) is not available on computing nodes.

In order to get the best performance as possible, we choose a high throughput and low latency network technology: **Infiniband QDR** (40Gbs).

Below, the evolution year-per-year of the computing performance and storage:



From:

<http://www-lbt.ibpc.fr>, baal.lbt.ibpc.fr/wiki/ - **LBT's Computation Resources wiki**

Permanent link:

<http://www-lbt.ibpc.fr>, baal.lbt.ibpc.fr/wiki/doku.php?id=cluster-lbt:architecture

Last update: **2023/11/30 20:52**

