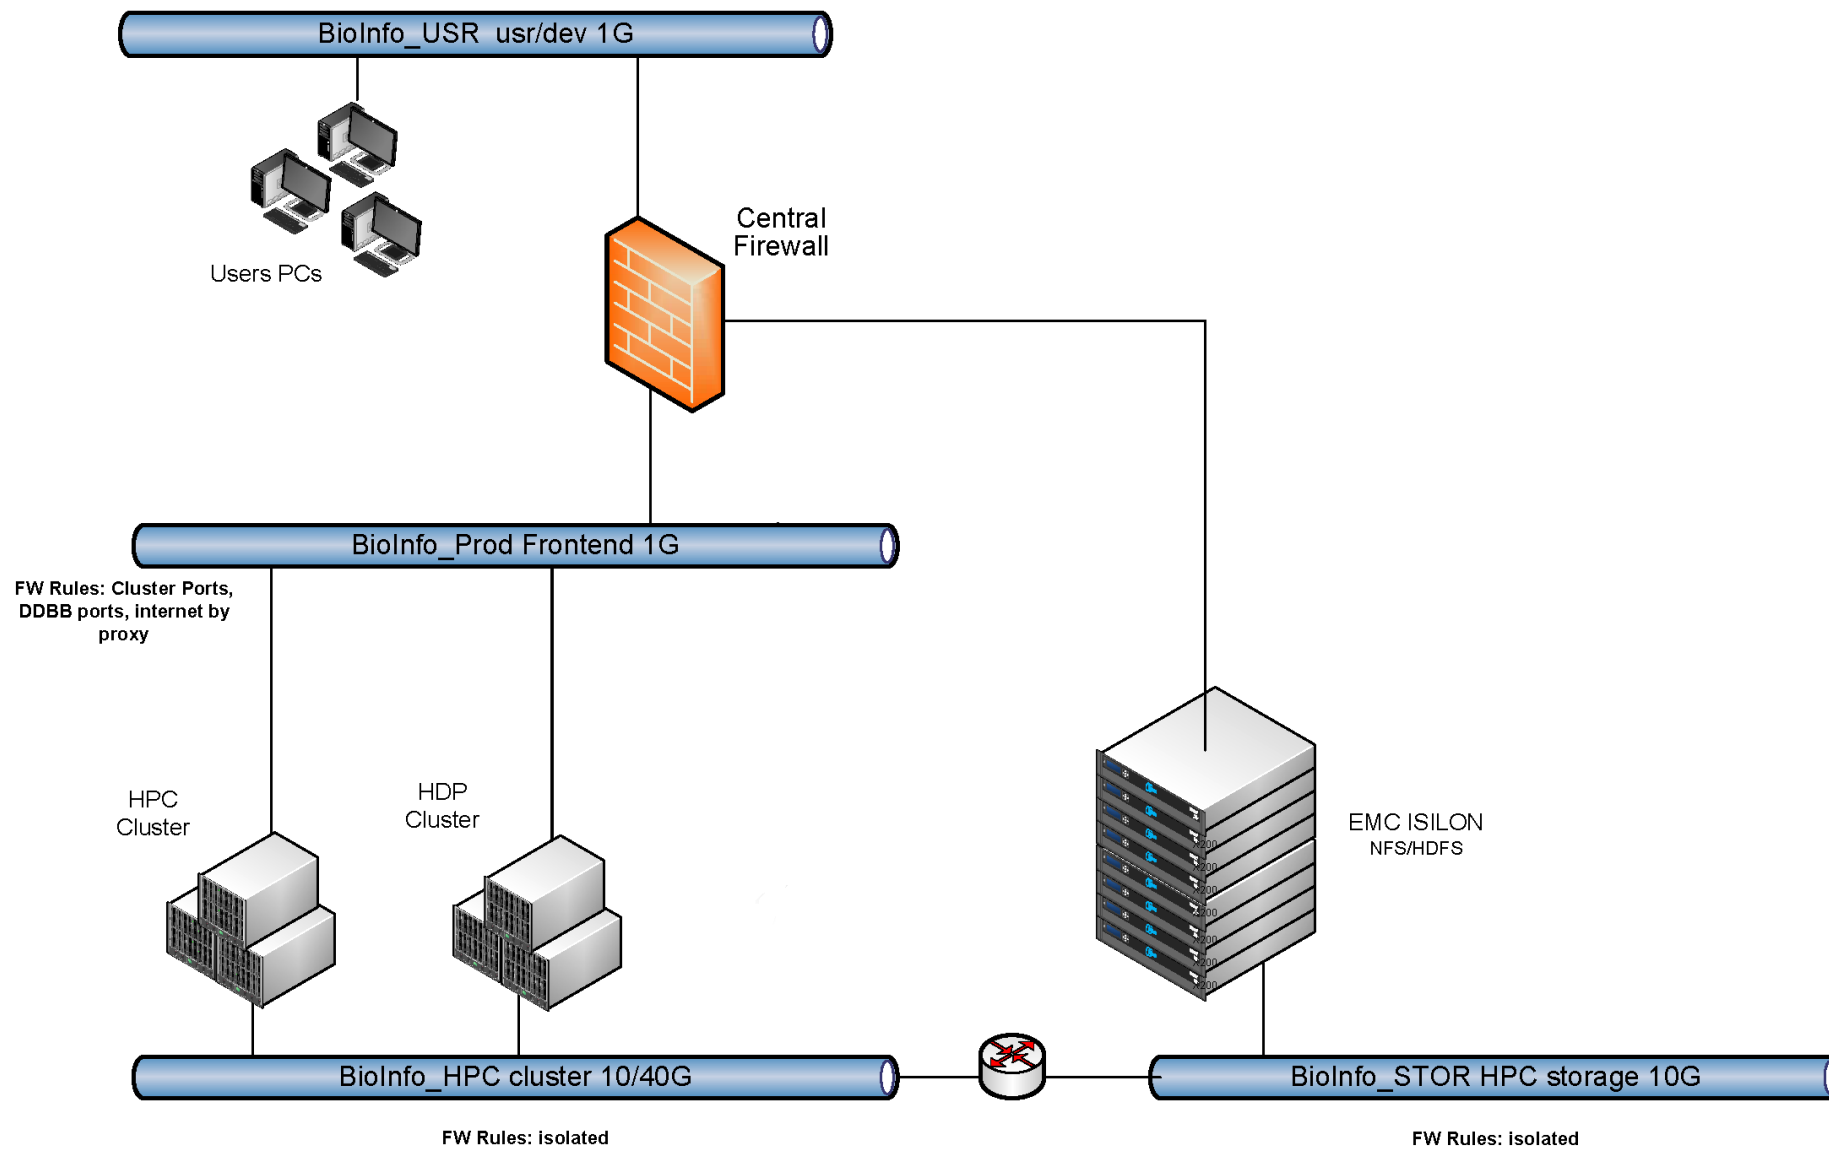


BioInfo HPC



# HPC Cluster

- ❑ A HPC cluster consisting of 37 servers (2 master nodes in HA, 3 submit nodes and 33 computing nodes) that together provide 1590 vCPUs y 4447 GiB of memory.
  
- ❑ All the servers are interconnected by 3 networks:
  - A secure network of 1 Gbits that, through the central trunk firewall connects to the network of users.
  - An isolated 10/40 Gbits network for the ultra-low latency interconnection of servers.
  - An isolated 10 Gbits network for connecting high-performance storage to servers.

# OS, Auth y Storage

- All the servers with CentOS 7.9 as OS
- BioInfo Auth system: NIS (Isolated, it does not sync with the Directorio Activo CNIC (cnic accounts))
- Automount of /home/\$USER in all servers (rw) (quote 120 GB per user)
- Mount on boot in all servers of /programs (ro)
- Mount on boot in all servers of /data\_\$LAB (rw) (default quote 2TB per Lab)

# Distributed Resource Management & Batch Scheduler

- SoGE v 8.1.9\_2 (aka SGE) (<https://arc.liv.ac.uk/SGE/>)
- 2 master nodes in HA (queue manager and batch scheduler)
- 3 submit nodes to login and submit jobs
- 33 exec nodes
- Access ONLY by ssh to submit nodes
- Without external access by users to master or exec nodes, all the exec nodes are internally submit nodes also. (Jobs running in a node can submit a new job to the cluster)

# Queues & Jobs @ BioInfo

- Queues/Projects per user group with access to a fixed number of exec nodes. All the exec nodes are shared between all queues/projects.
- No time limit for jobs
- Jobs can Rerun or reschedule.
- Default parameters: -cwd -r
- Mandatory parameters: -P \$Project -A \$account -N \$job\_name
- Default resources by job: 1 vCPU/2 GB
- You can request more resources with the -l flag (ex: -l h\_vmem=8G)

# Parallel environments

- pthreads (internal at node)
  - omp (internal at node)
  - MPI/MPICH (internal at node, external only mem between nodes)
- 
- You have to request it: `qsub -pe $parallel_env #vCPU`
- 
- NOTE: in parallel env, the requested memory (`-l h_vmem`) is per thread, take in account for the total memory for job requested.