


ORGANIZATION OF THE SNS HPC CENTER

ACCESSING AND INTEGRATING COMPUTATION RESOURCES



Comitato centro HPC

Version 4.1 22/02/2022

INTRODUCTION

The HPC Center of Scuola Normale Superiore (Centro HPC or CHPC hereafter) was created on December 4, 2019 (D.D. n.643, 644) with the purpose of supporting SNS researchers in: (i) hosting computing nodes and storage for scientific simulations; and (ii) hosting web pages and virtual machines to disseminate and collect research results, digital archives and related content (mostly relevant for virtual humanities).

The policies of the CHPC and its development plan are set up by the CPHC Committee (Committee hereafter), installed on April 7 2020. The Committee is formed by three SNS faculty members and two SNS technical staff members. The Committee will take decisions by consensus, or by majority when consensus is not reached. The present document briefly describes the procedures that will take the CHPC to an operational status and the subsequent operational policies. Further information and dedicated forms for submitting requests to the Committee will be prepared for all cases illustrated below. This document constitutes a first draft of the CHPC operating rules and will be expanded upon and modified in the future.

RESOURCES OF THE CHPC

Current state

The CHPC is run by a staff of seven people including a coordination officer (HPC Staff or Staff hereafter). In coordination with SNS researchers, the staff manages the server room and its contents, located in the San Silvestro complex of SNS, as well as additional individual servers in Palazzo della Canonica and Palazzo della Carovana.

The San Silvestro server room (simply server room or SR hereafter) is equipped with two redundant electrical lines (max. power 120 kW each), four “cooling towers “ (max 40 kW each) and hosts 16 server racks. The SR currently hosts computation equipment from Classe di Scienze and Classe di Lettere for an approximate total of 250 servers owned by different groups and administrative structures in SNS. In addition, the SR hosts a clustered Network Attached Storage (NAS) system and the five computing nodes that have been purchased by Classe di Scienze¹; these resources form the core of the current shared resources for intense computation.

Future development

Once the CHPC is operational, the Committee will prepare an expansion plan of the shared resources, subject to budgetary limitations. At the moment of writing the SNS administration

¹ Dipartimento di Eccellenza

has agreed to finance the expansion of the SR to the maximum capacity compatible with the building: electrical lines of 150 kW each and six cooling towers. Future expansions will include resources dedicated to Classe di Lettere.

Expansion of institutional resources for scientific computing is envisioned in the form of a homogeneous, medium-scale distributed clusters for Classe Scienze and Lettere. Both would facilitate easier integration and sharing of resources use of green, energy-efficient technologies; the former would also provide scalability to large scientific problems and allow code development for use on large, national and international HPC centers. Future acquisitions by SNS researchers wishing to apply for HPC hosting should aim to follow this philosophy. Consulting with CHPC staff is strongly recommended before any purchase; in this way the availability of physical space, compatibility of the hardware, availability of support (e. g. cooling and power), and potential integration of resources can be discussed in due time.

Decommissioning

Institutional resources will always have priority in the SR, in case available space is limited. Owners of old, low-usage machines are encouraged to either add them to the pool of shared servers (thereby obtaining guaranteed computational time on institutional resources), or decommission them.

In order to comply with SNS's "green policy" and to maximize the available space in the SR, the CHPC will regularly evaluate machines for decommissioning, as part of the hosting agreement procedures (see below). Old, computationally-intensive (see below, **Non Institutional Resources**) machines that do not meet a minimum usage threshold (nominally set at 10%), will be slated for decommissioning and removal from the SR. The timing and details of this will be discussed with their owners. All machines for which no support/replacement/maintenance is possible will also be slated for either removal or integration with the shared resources.

POLICIES

Institutional (shared) resources

This paragraph specifies the usage policies of current, shared institutional resources. It will be modified in the future if and when new resources are available. The CHPC will run a shared computational cluster formed by:

1. A clustered NAS (total *raw* space available 3.2 PB, 0.2 PB reserved for Classe di Lettere). It will provide the permanent and temporary network storage areas that will be accessible through the cluster.
2. two *fat nodes* equipped with 3.0 TB of memory and 4 CPU sockets and three *thin nodes* equipped with 2 CPU sockets, 1.5 TB of memory and 4 GP-GPUs.
3. login and maintenance nodes (*to be acquired*).
4. networking hardware for running equipment in points 1-3 (partially already present and partially to be acquired).
5. future acquisitions owned by Classe di Scienze.²

The cluster will run on Linux OS and use a workload manager to submit and run computational jobs. The login will be possible only using the SNS VPN or authenticating through SNS AAI (Authentication Authorization Infrastructure). Access to computing nodes will be possible using the batch system or interactively, under specific circumstances, only from the front end nodes.

The Staff will be responsible for:

- hardware maintenance, security and cooling of this equipment and of other nodes integrated in the cluster
- Unix system administration
- centralized user account system (LDAP or IPA server)
- managing software updates
- managing the batch system
- updating the documentation
- installing scientific software of general use (see below for a more complete discussion about software)
- managing backups

Access to the cluster will be provided by a personal account upon filling of the appropriate module provided by the Staff (see **Application Forms** section); non-independent researchers (undergraduate students, Ph. D. students and postdocs) will have to indicate a SNS professor or researcher who is responsible for their research. Using the account a user will be able to:³

- access files and data in his/her home directory (under the specified quota) using the cluster login nodes
- access nodes or data owned by the project team he/she is working with (via the login nodes or directly)
- run light and short pre/post processing jobs on the login node
- submit a job on common or dedicated (see below) resources
- access internal Git, Taiga and other development tools

² For future shared resources not integrated in the cluster (owned by Classe di Lettere?) ad hoc policies will be adopted.

³ actual resource limits and more detailed technical information will be provided in CHPC webpages

Computational time on shared resources will be allocated by the Committee according to SNS general policies, users' needs. Users and research teams will fill a project request form and submit it to the committee for evaluation. Upon approval, users will be able to access the project's allotted computing hours via the batch system; management of hours within the project is responsibility of the project lead. Complete rules concerning projects will be available on the CPHC web site.

All users must adhere to set resource limits, as well as a code of conduct. Breach of the rules may cause the suspension or cancellation of the account, upon decision of the Staff. In such cases the user will be allowed to download his/her data within a given deadline.

User accounts will have a 3 year duration and will be renewed upon request. Upon expiration of the account users will have a 3 month grace period to access the account and download their data and another 12 month period during which it will be possible to contact the Staff for the same reason; *after this term user data will be deleted without further notice and the Staff will not be responsible for any damage under any circumstances.*

Non-institutional resources

Owners of computational resources should apply to the HPC Committee in order to begin the hosting arrangement (see **Application Form** section). The duration of the hosting service is annual, tacitly renewable for a maximum of five years, after which a new application must be submitted. The annual renewals will also allow for changes in the sharing state of the nodes, as well as an evaluations of the usage rates, maintenance status and energy consumption. At the end of the hosting period, if there is no renewal agreement, the owner has 60 days to arrange for the removal of the equipment; failure to do so will result in staff removing the equipment. SNS reserves the right to terminate the hosting service at any time; a minimum 60 days' notice will be provided in such an eventuality. The duration of the hosting service will also take into account the usage rate of the machine, maintenance status, power consumption and cooling requirements to maximise energy efficient use of the resources. To select a hosting policy we distinguish two different categories for group resources, "Computationally intensive" and "Not computationally intensive" that are further discussed below. Non institutional resources will be monitored in terms of the amount of jobs and usage (CPU, memory, disk), in order to provide statistics for renewal and decommissioning discussions.

The login will be possible only using the SNS VPN or authenticating through SNS AAI (Authentication Authorization Infrastructure).

Computationally intensive

This category includes (large) groups of nodes, network or storage resources that are used for scientific data analyses and simulations. The classification of computationally intensive

and non-computationally intensive is part of the hosting service proposal. The CHPC can host this equipment in three modalities:

1. resources are fully integrated in the institutional cluster. The owner group negotiates with the Committee on a proportional amount of computing hours to be spent on the whole cluster as an “in kind” compensation for the service, including possible reservations (nodes normally in general queues are temporarily dedicated to specific users) and priority accesses. This will be the preferred way to integrate equipment as it facilitates the creation of more efficient and powerful clusters, especially once a homogeneous cluster is in place. It also allows machines not to remain idle, when not used by priority access, so as to benefit the SNS community.
2. the resources cannot be shared seamlessly because of administrative reasons (e. g. they have been purchased under an ERC or other grant that requires exclusive use during the amortization period). In this case the computing nodes may be accessed by dedicated queues and/or directly from the login nodes.
3. resources can be *partially* shared. Owners of multiple nodes can designate some of them for the shared arrays, while restricting access to others. In this case, the “in kind” compensation (discussed above in category 1.) will reflect the relative specifications of the shared nodes.

For shared resources the Staff takes full responsibility for their management.

For non shared resources the Staff takes care of:

- hardware maintenance, security and cooling
- network access
- integration with the centralized account system if needed (mandatory to access shared resources, e. g. the storage areas)
- installation of the OS (among the ones selected by the CHPC)
- integration with the batch system if needed

Moreover, for non shared resources the owner must provide (if deemed necessary by the Staff and HPC Committee):

- PDUs
- power cables
- rack mounting equipment
- network switches and adapter cables for the BMC network
- network switches and adapters for the service network

Not computationally intensive

This category includes resources which will not be integrated in any way because their primary purpose is not scientific computation. Examples are hosting web pages, photographic archives, etc. They are intended to be accessed directly by the

owners/developers and be accessible from the Internet. In this case the Staff will provide the same type of support as outlined for computationally intensive equipment regarding the hardware layer; management of the software layer will require additional arrangements, due to increased security issues and law requirements.

Transitional measures

Currently the San Silvestro server room hosts IT equipment and clusters owned by various SNS structures and researchers. People currently responsible for the equipment will be asked to provide information for the hosting service as specified in the previous section and an evaluation (see **Group Resources** and **Future Development** sections for guidelines) will be made on a case by case basis. The goal is to start in 2021 with a situation where all resources installed in the San Silvestro server room have been accounted for and authorized by the Committee.

Equipment purchased by SNS structures that is not currently hosted in the SR will also undergo a case by case evaluation, once the CHPC is operational. This primarily consists of equipment for web pages and virtual machines. These machines will be subject to the policies outlined under the **not computationally intensive** section above. Specific actions will be planned and carried out in agreement with the major stakeholders in Classe di Lettere to update as much as possible the hardware and software layers of research products in order to guarantee security and compliance to Italian law.

Software development and licensing

Users of the institutional resources and owners of dedicated resources who develop software are fully responsible for any misuse of licensed software or Intellectual Property (IP) breach. Neither the Staff nor the Committee will be responsible for any code or executable present in any server hosted in the SR in case of unauthorized use or IP breach. The Staff may provide assistance for the installation of software on non shared resources. For non open source software the request must be documented with all of the relevant licensing information.

Hosting of web pages and virtual machines

If any group or research personnel plans to host web pages or virtual machines accessible from outside the SNS network⁴, in addition to the information provided for hardware hosting the owner/administrator is held responsible for the cybersecurity and protection of personal information and must comply with

- the policies of SNS as described by the SNS ICT

⁴ This section refers to interactive web servers and not to official SNS web-pages, personal or group research pages.

(<https://ict.sns.it/it/policies#TOC-Web-Server>, <https://ict.sns.it/it/policies#TOC-Sicurezza>)

- with the rules set by the SNS Data Protection Officer (DPO)
- with any additional internal procedure set by the CHPC

He/she must provide all the necessary software updates and fixes to prevent any breach and promptly communicate to the SNS DPO and to the Staff any problem. The Committee and Staff will not be responsible for any damage originating from such breaches. In addition the Staff may shutdown without further notice any web site, virtual or physical server if it presents documented security issues or in the event of a breach.

The Staff may provide basic assistance (subject to availability) in setting-up and maintaining platforms for the public dissemination of research results obtained with CHPC resources. Guidelines and a list of acceptable protocols will be provided in the future.

Data Protection and IT Security

With respect to any non shared resources, the owner is fully responsible for the data privacy (according to the GDPR *General Data Protection Regulation*), data protection and security (according to the current legislation, including: Misura Minime di Sicurezza ICT per Pubblica Amministrazione, Regolamento Amministratore di Sistema) of all the IT resources.

Staff can provide basic security, subject to the guidelines stated above. In particular, the staff will be responsible for the hardware layer, while any virtual machine and/or website needs to have a designated administrator, approved annually by the SNS Director, according to national law.

Participation of CHPC staff in research projects

SNS research personnel may also apply to the Committee for the participation of members of the Staff in research projects for an amount not exceeding 50% of total FTE, taking into account existing projects. The Committee is responsible for evaluating these requests and allocating FTEs. For example, such support could consist of collaborating in code optimization, parallelization, or analysis pipeline development.

Application Forms

Users who want to access the SNS shared cluster must send a brief application to the Committee with the following details:

- name and contact email
- (for students or postdocs) the name of an SNS professor or researcher who is responsible for their research

People or groups who want to host equipment in the CHPC not managed by the Staff (see **non computationally intensive** section) must provide contact information to the Staff and the Committee. The Staff will use this information to communicate any issue with the owner's equipment or if any problem with this equipment arises. This information includes:

- name of a system manager / contact person, his role and contact details
- names of the system administrators and webmasters who can operate on the applicant's virtual machines / websites and their contact details
- for each system administrator / webmaster indicated, a one page curriculum vitae suitable to allow the verification of the skills necessary to perform these functions
- the commitment to promptly communicate any changes in the information provided
- the commitment to adhere to SNS regulations for all the users of the equipment
- the safety and warranty documentation according to Italian and EU laws

Technicalities

For additional details on policies or technical issues, please refer to the official HPC web page:
<https://hpccenter.sns.it>