National Facility Computing Infrastructure Data Policy

October 23, 2020

Version	Authors	Date
v 1.0	MC Toribio, R Hammargren, J Conway, M Lindqvist	October 23, 2020

Contents

1 Purpose of this document

The purpose of this document is to define policies for access and for disk space and CPU usage in the National Facility Computing Infrastructure (NaFCI).

> **NaFCI resources are provided primarily to support radio astronomy data reduction from Swedish and Nordic users**.

1.1 Radio/mm instruments supported by the NaFCI

The Onsala National Infrastructure for Radio Astronomy (Onsala Space Observatory, OSO) has a mission to support Swedish and ALMA Nordic astronomers in the use of Radio/mm astronomy observing facilities. The NaFCI computers are primarily made available to help provide this support to users. In principle support is provided to any radio/mm instrument although prioritized due to limited personnel and computer resources. The highest priority is given to support ALMA users (based on a MoU commitment to ESO), followed by the instruments in which Onsala participates in observations, including EVN/Global VLBI and LOFAR, and finally by other other interferometric instruments (JVLA, eMERLIN, PdBI, etc). Single dish data can also be processed on NaFCI computers. Similarly, projects that develop techniques and the user base for future radio astronomy instruments of high priority for the OSO infrastructure (such as i.e. SKA) can also be prioritized. OSO has also a role in supporting Geoscience observations including geodetic VLBI and therefore NaFCI resources can after discussion with the infrastructure Director and NAFCI management also be used for supporting such geoscience use. Operationally the prioritisation of the use of NaFCI resources is set by the OSO infrastructure management and is ultimately the responsibility of its Director.

1.2 Nature of NaFCI Support Activities

The detailed nature of what is classed as a valid *support* task for the national infrastructure in radio/mm astronomy continually develops. Broadly national infrastructure support consists of providing computer and human resources to aid tasks which are generic to multiple projects rather than being specific to science extraction on a single project. At present the dominant support use case is the customized generation of ALMA image cubes from visibility data. Types of generic support tasks that are expected to expand in the future include the processing of spectra, 2D-image or 3D image cubes into Advanced Data Products (ADPs) such as RM cubes, source catalogues etc. This development is consistent with plans for OSO staff to be help run the Swedish node of the SKA Regional Centre (SRC) network which will provide such advanced data product services to users. It is globally efficient that national infrastructure staff become expert in generic support tasks and that a hardware and software

environment is maintained at NaFCI to support this work. In some cases in order to develop expertise in a technique which is expected to become generic in the future National Infrastructure resources can be used to process a single project as a pilot.

In general NaFCI resources are meant to be used in a time-limited way to support a specific stage in the processing of user data. In contrast, the permanent storage and long-term science exploitation of data products should be provided by other platforms that users have access to. A potential exception (in line with plans for OSO/NaFCI to support a SKA Regional Centre node) would be the permanent storage and user accessibility of processed data products that have legacy value for multiple users.

2 Overview of NaFCI computer facilities

National Facility Computing Infrastructure (NaFCI) are available for astronomers for reduction of radio data. NaFCI consists of two computer clusters, one local at Onsala (NaFCI-1) and one at Chalmers campus in Johanneberg (NaFCI-2), consisting of a private part (for OSO use only) of the *Vera* cluster, at Chalmers Centre for Computational Science and Engineering (C3SE, a SNIC centre).

2.1 NaFCI-1 computers at OSO

A brief description of the NaFCI-1 cluster computers at OSO is given in Table 1. Users of this system have a common home directory accessed when they login. They can access computing nodes and storage areas mounted with standard naming `/nmt/xxxx_Y/`. The filesystems are shared so one can access all mounted filesystems from each computing node. Support of NaFCI-1 computers is carried out directly by the IT group at OSO.

2.2 NaFCI-2 Computers at C3SE

NaFCI-2 resides in the Vera cluster at the C3SE and consists of 5 nodes $\times$ 32 cores. Each of the nodes consists of 2x Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz (2x 16 cores) with 384 GB RAM. The total disk space on this system is 583TB in the Cephyr storage server. Based on high performance computing requirements, users can be granted access to this computing facility.

Support of NaFCI-2 computers is carried out by C3SE staff and extra guidance from the IT group at OSO.

3 Supported Software

Currently the supported software packages on NaFCI-1 and NaFCI-2 are shown in Table 2.

Name	Processors	Cores	RAM	mount	disk size
Bele	2x Xeon X5650	12	112 GB		
				/mnt/bele_1/	7.3 TB
				/mnt/bele_2/	11 TB
				/mnt/bele_3/	11 TB
Bure	2x Xeon X5660	12	120 GB		
				/mnt/bure_1/	17 TB
Duv	2x Xeon E5-2620	12	160 GB		
				/mnt/duv_1/	22 TB
Dore	2x Xeon E5-2620	12	160 GB		
				/mnt/dore_1/	37 TB
				/mnt/dore_2/	37 TB
Grim	2x Xeon Gold 6248R	48	1,5 TB	/mnt/grim_1/	150 TB
Parking area					80 TB
TOTAL					372.3 TB

Table 1: NaFCI-1 computers specifications. In addition there is a computer, *Menja*, (with similar technical specifications to *Grim*) federated in the same computer system but which is reserved for the use of OSO staff scientists.

Software packages	NaFCI-1	NaFCI-2
CASA	✓	✓
LOFAR software, including 'FACTOR'	✓	✓
GILDAS	✓	
AIPS (legacy software, best effort)	✓	

Table 2: NaFCI software

4 CPU usage

Processing in NaFCI-1 is prioritized using the `nice` command. All processes start with a `nice`-ness level that defines its priority for CPU resources. NaFCI management can modify the start level to prioritize processing if needed. In case of the need of intense computing, submit a request beforehand.

Users of NaFCI-2 are referred to follow instructions from their contact staff scientist to run their processes on *Vera*.

5 Data Backups

NaFCI-1 and NaFCI-2 disk storage use Raid 6 systems providing a partial security against data loss – which means all files in these systems should be secure against one faulty disk.

Both of the NaFCI-1 and NaFCI-2 users' `home` folders are backed up automatically once per day [1]. There are tools available to the user for setting up regular backups of scripts located in other folders than the `home` folder. The rest of the data, i.e data not on the home disk or selected scripts, is not backed up.

Since only home directories and selected scripts have full redundancy against disk failures it is the responsibility of users to carry out backups of important data and scripts. Conversely data that can be easily reproduced should not be placed in backed up areas.

6 Parking area and Legacy Archive

Use of space on processing nodes should be kept to the minimum in order to maximise the efficiency of processing for all users. For the purpose of temporarily hosting data that is not actively being processed, a *Parking area* has been designated on **NaFCI-1** (see Table 1). NaFCI-1 users should move data to the parking areas as soon as they stop actively processing in order to free resources for other users. In addition, data residing on the processing nodes that has not been processed for a period of **180 days** ($\sim$6 months) will be automatically **moved out** from the processing nodes to the parking area [2]. Data stored in the parking area will remain accessible for a further 180 days ($\sim$6 months). *Files older than 180 days on the parking area will be* **deleted** *from the file system automatically.* A similar procedure is under development for data hosted in **NaFCI-2**. For the time being, NaFCI-2 users are requested to **manually delete** data that has not been processed for a period of **1 year**.

Note that designated parking areas are meant for medium term storage only. Indefinite storage of data is in general not provided on NaFCI; at the completion of the processing phase, final data products of user accounts should be moved

---

[1] For NaFCI-2, further details can be found at the C3SE Documentation site

[2] Currently the procedure is not automatic and involves manual action yet. It requires further development and testing before adopting a safe automatic flow

out of the NaFCI system and stored elsewhere. A potential exception exists of projects that produce **legacy data** (defined as of interest to multiple users over a significant time). For such cases a long term archive solution on NaFCI can be explored upon request.

7   NaFCI Project Accounts

Users are not granted permanent accounts on NaFCI-1 which are then used for multiple projects - instead accounts are set up in connection with specific projects and are removed when that project finishes. Consistent with this principle time-limited user accounts are created with usernames of the type e.g.,`arcguest_X` or `rguest_X` for ALMA and non-ALMA projects respectively.

Projects with intensive computing requirements can be granted access to NaFCI-2. User access for projects run on NaFCI-2 are linked to individuals access credentials on the SNIC/C3SE system (i.e. logon names which are uniquely be identified with an individual). For active NaFCI-2 projects users are given access to a dedicated path on the NaFCI partition in the Vera cluster. After the project is finished, this access is removed. The expiration date is defined at the start according to the timeline of the project. Requests for extensions can be submitted.

7.1   Access Request

In order to request access to computing resources at NaFCI an online access request form must be submitted with the following information:

1. An agreement that the general regulations for using IT resources will be followed

2. Type of project:

   - ARC project –ARC node staff will follow up on their needs and estimated timeline.
   - Other project – Scientific justification and the need for resources (processing, storage) and an estimated timeline.

   In addition to the above, non-Chalmers guests that are physically visiting Onsala need to fill in the OSO visitors form (password-protected) before they visit.

In the form above, non-ARC project users should provide information on the type of project, need for resources (processing, storage) and an estimated timeline in their request. For ARC project accounts, ARC staff will follow up their needs.

Non-Chalmers guests travelling to Onsala for Face-to-Face support **on site** have their personal equipment granted access to the network during their stay

at the Observatory. The visitor's host ensures that the information on the equipment is provided prior to arrival to the Observatory premises.