Cluster Computing Resource Basic User Guidelines
MRI Center, NYSPI
Juan Sanchez and David Semanek
5/2022

The purpose of this document is to outline user guidelines for the MRI Center computing cluster resources, including MRIdoctor and Docker Swarm.

The MRI Center computing resources are provided for the purpose of processing ("pre-processing") approved MRI and related data through pipelines managed by the MRI Center cluster computing staff and PsyIT. Regular users of the systems should never execute their own scripts or create or call Docker containers outside of the managed pipelines following the instructions of the cluster computing staff and documentation. For experienced shell users, running user generated scripts created for the purpose of gathering information about or re-organizing data within a specific project is allowed, provided the script does not change the names or attributes of folders or files that would break the pipelines enabled for your projects; such scripts are run at your own risk regarding the integrity of your data.

Each user is provided a typical Linux Home folder on MRIdoctor that can be used as a workspace. Please do not store imaging data or outputs in this location as space on the Home partition is limited and shared between all users; if you need an additional location where size is a concern please reach out to the cluster computing staff. Please be mindful of space and resource utilization during regular usage of the cluster and attached storage. It's very helpful if all users are mindful of being efficient in using space such not making unnecessary copies, removing stale or malformed data, and limiting DICOM file usage and handling.

Do not under any circumstances run commands on MRIdoctor that query or issue instructions directly to the Docker daemon or Docker Swarm cli. Only cluster users with prior administrative approval may run commands directly through the Docker cli, including starting and stopping containers or services, or building or altering images. All other users access the Docker Swarm resources seamlessly through pipeline commands with provided documentation.

The MRI computing cluster is a work in progress, and all users have a responsibility to maintain the stability and integrity of the system by being conservative about the commands run on the system and reporting any issues or bugs immediately to the cluster computing staff. When in doubt, please check with us!

Juan.Sanchez@nyspi.columbia.edu
David.Semanek@nyspi.columbia.edu
Waseem.Ijaz@nyspi.columbia.edu