

Questions IT should be asking faculty about research computing (in regards to their research):

- What is your research in? Are you interested in using the research computing facilities available? If you aren't, may we ask why not? (If the hesitancy is due to a lack of need, ask if you can continue communicating with them, offering training opportunities, etc because "we don't want you to prematurely limit your research opportunities due to a lack of computational resources. Don't ask questions based on the size of your resources. Ask all questions and we will help you find the resources you need to answer it.")
- If the research computing facility is interested in this: How can we use your research to help diagnose and rehabilitate local infrastructure? That is, what can we take from your research to upgrade pre-existing infrastructure? Is there anything from your research that we can use for our institution as a whole?
- Would you be willing to present your research at this amazing symposium I'm hosting?
- I happen to know another faculty in a different department that has research that would pair well with yours, would you like their contact information?
- Would you be willing to acknowledge the research computing system in your publications and grant submissions?

Questions IT should be asking faculty about research computing (in regards to using university resources):

- What resources will you need (machines, computer accessories, etc.)?
- What level of computing knowledge do the active research participants have? Are you familiar with basic Linux [or whatever OS the system uses] commands? Have you ever used a non-GUI interface?
- What kind of software will be needed for your research? What OS does your software run on? If the software licensing is at a cost, who is your contact to determine if the licensing allows for a research computing environment and if it does, how the licensing works (ex: per system or per core)?

Questions for faculty if they would like to use the local supercomputing resources:

- What's your expected run time?
- How much data are you processing? If it is large, does all of the data need to be read into memory or can the file be randomly accessed by the job as needed? Is your data dependent on itself?
- Are you aware of our policy regarding backups? Our definitions for backup and snapshot are...
- Are you familiar with SLURM (or the scheduler resource uses) and how to submit a job to the system?
 - Are you familiar with our queues/partitions?
 - Do you know how to determine which partitions/queues are you going to use and how to call them within the script?
- Have you been on-boarded to the supercomputer? (If the group does on-boarding)

- Do you need assistance writing code?
- Have you considered parallelizing this code?
- Have you checked out all the resources we have to offer? [Then list any you think may be of importance given the conversation or give them the link to where they can find a list of all resources.]
- What is your output directory once your code has finished?
- Are you comfortable determining how many nodes will you use? How many cores will you need?
- What is the urgency of this data processing?
- Are you aware of the policies we have in place for the supercomputer?

Questions faculty should be asking about research computing

General questions regarding the research computing resources:

- **First, what research computing facilities are available?**
- **Does it cost anything to use the system?**
- **How do I gain time on the system?** [Some systems use allocations where a certain amount of computational hours/core hours/service units are allowed per project or per researcher. Others are open and do not limit you in how much you utilize the resources.]
 - **If the system does not work on allocations, is differential priority given to certain accounts?** [Sometimes those who use the system less have a slightly higher priority, but other ways to manage priority exist.]
- **How much storage do I have access to? Is it possible to gain more? If so, does it require funding?**
- **How is the storage set up and what, if any, of the data is backed up? How are you defining “backed-up”?** [Often there is a /home and /scratch directory. The /home directory is sometimes backed up, whereas the /scratch directory usually isn't. However, the definition of “backed up” often changes between research and IT staff. Be clear in what you mean. IT's definition is often “copied and stored off site” whereas a researcher's definition is often “if I hit ‘rm *’ will I really have lost all of my data?”. A snapshot is often what a researcher is referring to. However, be clear about what this means! If snapshot-ing is what is offered, if the system has a large enough failure, your valuable piece of data may be gone forever. If you have any doubt about what is happening on your system, ask more questions and rephrase what you hear to ensure you truly understand what is happening with the storage on your system.]
- **What scheduler is being used?** [A scheduler is the software that is used to organize the jobs submitted to the system. It pairs the resources available with the resources requested for all of the jobs submitted.]
- **Can I install my own software? Is there software preinstalled on the system that I can use? How do I request software be installed/assistance with installing software?** [Many systems have a module system that allows you to have access to a large pool of software. On some you are prohibited from installing your own software. Exploring this allows you to know more about what you can do on the system.]
- **If I have questions, who do I contact? Will they be able to provide scripting or job submission help?**
- **Is it possible to have a shared space for [lab] students to collaborate together on files?**
- **If I have a collaborator that is off-campus, can they gain access to the system because they are working with me? If so, how?**
- **If I have grant funding, do I get priority on the system? How do I gain priority?** [Many systems have protocols in place to allow researchers to purchase time, resulting in a higher priority, or purchase equipment which then creates a priority queue on that equipment for the researcher.]

- **How can I include the campus resources in a grant?** [In some cases this may include how to purchase equipment or storage, if the resource uses a condo-model for funding and expansion. However, this can also mean a statement for a Facilities&Equipment document.]
- **Can I use the resource for teaching a course? If so, what is the best way to go about collaborating with the research computing staff to ensure a smooth workflow?**
- **How should I denote the use of the system in my publications? Does it help the center to know what grants I am including the system on?**

Nitty-gritty questions to ask about how to use a computational resource:

- **What's the maximum runtime I can have for my job?** [Runtime: The time allocated to a specific script on the computing cluster.]
- **What is currently in an accessible directory for my use?** [Softwares like MatLAB and SPSS]
- **I would like to cancel my job, how do I do so?**
- **How many jobs can I submit at once?**
- **Are/Is there workshops/training for my research assistants on the supercomputer?**
- **When is the supercomputer's downtime? How long is the downtime?**
- **I have a software addition suggestion, who do I contact to discuss this with?**
- **I'm used to computing on CUDA[Gpu Programming vs conventional CPU], does your infrastructure support that?**
-