

Scientific Computing Allocations Process  
Mount Sinai School of Medicine  
April 24, 2013

Introduction

Mount Sinai recognizes the value to researchers of providing core advanced computational facilities. Since these scientific computing resources are both popular and limited, Mount Sinai is establishing an allocations process to help guide the use of these resources so that they are aligned with institutional priorities.

Principal Investigators (PI) will be able to request compute and storage resources for their research projects, and a peer-review committee will allocate these resources based on scientific merit, financial impact, and effective use of resources. PI is defined as the lead scientist or researcher for a project who takes direct responsibility for a project including, but not limited to, compliance issues, finances, overseeing the research and reporting to funding agencies or department chairs on the outcome.

The process will start with allocation requests due on June 1, 2013 and allocations beginning on July 1, 2013.

The Available Scientific Computing Resources

The Scientific Computing Allocations Committee will allocate time and space on the following resources *per year*:

- 64M compute core hours per year on Minerva
- 1 PB of scratch (not backed up) parallel file system space on Minerva
- 1 PB of archival storage consisting of two copies on two geographically distributed tapes

The Allocations Process

The allocations process will be as follows:

1. The PI for each project will visit <http://hpc.mssm.edu>, click on the link for allocations and fill out the web form with the following information:
  - a. Name of PI
  - b. Name of project
  - c. Names of collaborative researchers on the same project (if any)
  - d. Department
  - e. Description of proposed scientific activity
  - f. Description of potential scientific impact
  - g. Description of proposed computational approach including the viability of the approach
  - h. Number of core hours requested
  - i. Amount of scratch storage requested
  - j. Amount of archival storage requested
  - k. Justification of requested amounts
  - l. Funding sources and amounts
  - m. Pending funding sources and amounts
  - n. Publications related to past usage
  - o. Any other special requests

Scientific Computing Allocations Process  
 Mount Sinai School of Medicine  
 April 24, 2013

- p. Agreement to acknowledge Mount Sinai as supporting the research performed on the resources as specified on this webpage:  
<http://hpc.mssm.edu/acknowledge>
  - q. Agreement to notify Scientific Computing by emailing  
[hpcgrants@mssm.edu](mailto:hpcgrants@mssm.edu) about grants awarded as a result of using these resources
  - r. Agreement to notify Scientific Computing by emailing  
[hpcpublications@mssm.edu](mailto:hpcpublications@mssm.edu) about documents published as a result of using these resources
  - s. Indicate whether this is regular or supplemental request.
2. Requests will be due on the first day of the month before the beginning of the next quarter as shown in Table 1 with approximately 70% of the available compute and scratch resources allocated in June and 10% of the resources at the other quarters. The committee reserves the right to oversubscribe or undersubscribe resources based on the quality of the submissions, paid cycles/storage and institutional priorities.

<b>Proposal due</b>	Dec 1	March 1	June 1	September 1
<b>Notifications</b>	Dec 20	March 20	June 20	September 20
<b>Allocation begins</b>	Jan 1	April 1	July 1	October 1
<b>Allocation expires</b>	Dec 31	March 31	June 30	September 30
<b>Minerva cycles allocated</b>	6M	6M	46M	6M
<b>Minerva scratch allocated</b>	50 TB	50 TB	600 TB	50 TB
<b>Archival storage allocated</b>	250 TB	250 TB	250 TB	250 TB

**Table 1 Allocation dates and resources to be allocated per year**

3. Each of the three scientific reviewers on the allocations committee will evaluate each proposal based on scientific merit and financial impact in accordance with the allocation criterion listed in Table 2 below. The representative from the Scientific Computing group will evaluate each proposal based on the “effective use of resources” criteria. Each proposal will receive a final score computed as the average of the reviewers’ scores. The allocation will be based on the score and deliberation by the committee. The requesting PI will receive notice of their allocation.

<b>Criteria</b>	<b>Description</b>	<b>Weight</b>
Scientific merit	Potential impact for new, significant scientific discovery	10
Financial impact	Project has or will have significant financial support	10
Effective use of resources	Researchers are experienced and have used and/or will use the resources in a scalable and responsible manner	5

**Table 2 Allocation Criterion**

Scientific Computing Allocations Process  
Mount Sinai School of Medicine  
April 24, 2013

4. The allocations will start at the beginning of the quarter and will last for one calendar year. The compute allocations will receive high priority in the queue and the storage allocations will be governed by hard quotas.
5. The list of PIs and their allocations will be published.
6. After receiving the results of the proposal from the committee, the PI can appeal to the committee with additional information. The committee will review these appeals and provide a response within a month.
7. PIs will be sent a report on their compute and storage usage monthly and will be encouraged to burn through their time evenly so there is not a “run on the bank” at the end of the allocation by all PIs. The PI will also be reminded about when his/her allocation will end.
8. The committee will review all usage quarterly.
9. Allocations will expire one year after they begin. Extensions can be requested but are not guaranteed to be granted.
10. The PI must notify Scientific Computing about any funding that was granted as a result of using these resources.

Time-sensitive Jobs and Supplemental Requests

If a PI has an urgent project deadline, that PI can send a request to [hpchelp@mssm.edu](mailto:hpchelp@mssm.edu) explaining the following:

- what science needs to run,
- why it is urgent,
- when the job needs to run by and
- how many nodes/cores/disk space is needed and for how long.

At least two members of the committee will review the request and the PI will be notified if their request has been approved or denied within 48 hours of the request. Every effort will be made by the Scientific Computing staff to accommodate the request as the soonest possible time.

If a PI uses up their allocation or needs time during an off-cycle, they can make a request for a supplemental compute or storage allocation using the regular process above and checking the box for supplemental allocation. Only 200,000 hours can be requested as a supplemental allocation. The committee will return an answer within two weeks.

The Allocations Committee

The allocations committee will consist of a rotating group of three scientific members and one scientific computing member with a term of one year that can be renewed. The membership of the committee will be approved by Dr. Charney. The initial committee members are: Eric Schadt, Pamela Sklar, John Morrison and Patricia Kovatch.

Options for Scavenging and Purchasing Resources

All users will be able to:

Scientific Computing Allocations Process  
 Mount Sinai School of Medicine  
 April 24, 2013

- Run jobs at low priority with a maximum job length of 24 hours, receive 500 GB of storage on scratch and 1 TB of archival storage per year. Allocated users can select to use their allocated high priority time, or run as a scavenger.
- Make use of scheduler/resource manager (Torque/Moab) checkpointing facilities to continue computations over multiple 24 hour periods.
- Purchase high priority queue compute time and storage at the rates shown in Table 3.

<b>Cost @ Sinai</b>	<b>What you get</b>	<b>Cost @ Amazon</b>	<b>What you get</b>
\$3,000	100,000 core-hours (\$0.03/core-hour) (must be used in 1 year)	\$0.08 \$0.64	Virtual core hour Tightly-coupled core hour (InfiniBand)
\$100	1 TB scratch space/year with unlimited usage	\$120	1 TB/year plus additional costs for transferring data in/out and usage
\$78	1 TB/year of two copies of tapes (one offsite) with unlimited usage		
\$10,000	New dedicated Minerva node (64 core AMD or 8 core Intel, 256 GB memory node)		

**Table 3 Costs @Sinai and @Amazon for compute and storage**

A note about this process

This process was based on allocation processes for both national resources at the Department of Energy, the National Science Foundation, NASA and campus resources at Yale, University of Chicago, and UCLA.