Final Project Astro 250

1. Describe your project and how it relates to the DOE's Office of Science mission. See: http://science.energy.gov/sc-3/mission-and-functions, http://science.energy.gov/hep, http://science.energy.gov/hep/about, and http://science.energy.gov/np/about. (Keep it to 3 paragraphs or under). If it doesn't, stretch your and my imagination….

2. Project Summary: provide a brief summary easily understood by people outside of your field.

3. Describe the major components of your code algorithmically - is it based on sparse matrix operations, FFT's, AMR, etc. (Keep it to 3 paragraphs). Are the guts of these libraries at NERSC or are they home-brewed?

4. If your code is parallel, does it use MPI, OpenMP, both? How does it scale? For the latter question put some runs through the debug queue trying a variety of combinations and graph or table up the results. Explain possible ways in which you see the code could do more (more physics, larger problem sizes, etc.) and how this would or would not effect the scaling.

5. If your code is not parallel, are there limitations on the number of jobs you could send through the serial queue (from I/O, or just handling the number of results)? Send a dozen in at once and then 2 dozen, etc. and see if the timing changes. Are the steps you could take to make the code parallel and do you think it is worth it? For the latter question, it is often the case of: Do I need to get my science done asap since the experiment will end before I could write a parallel code to do this. If there is a long-term aspect to this analysis, what would you tackle first to make it parallel.

6. Is Edison the best machine for you to run on or in other words what would be your ideal machine? Will Cori be better (with BurstBuffer)?

7. How much time will you need, what disk space will you need and why?

8. Do you need access to any special web portals, databases, etc.