1

Good afternoon. I'm Rudi Eigenmann. I'm the NSF Program Director managing the program called **Petascale Computing Resource Allocations,** or PRAC.

Thanks for taking time to join this webcast today.

2

Before I discuss the solicitation I would like to inform you that the Division of Advanced Cyberinfrastructure in the Directorate for Computer and Information Science and Engineering at NSF is conducting a nationwide search for qualified candidates to serve as Program Director for high performance computing. For more information, visit the website given on the slide:

https://www.usajobs.gov/GetJob/ViewDetails/358941800?org=ACI

I strongly encourage you to consider this position. Or, if you know somebody who might be qualified and interested, please let them or me know. Thank you.

3

Now to the solicitation. I will first give a brief overview of the PRAC program, including some of the most important things you need to know about submitting a proposal. I will then present answers to some frequently asked questions, and end by opening the floor for further questions from the audience.

4

Here is a big-picture slide. For many years, computation has now been a key enabler of science and engineering. To conduct such computational science, you need to have access to computational resources that can satisfy the needs of your applications. How do you obtain such access?

You have probably seen the pyramid representation, where the highest performing computers are at the top. The National Science Foundation sponsors one such system installation. Note that NSF's concern is not about providing a machine with the highest *peak* performance, but one that sustains high performance across a broad range of applications. It is called Blue Waters and stands at the University of Illinois. To obtain computer time on this resource, you write a PRAC proposal.

The middle tier of the pyramid supports an even broader class of problems that also utilize high compute power. NSF has a number of such machines, tied together via the XSEDE virtual organization. To obtain machine access, you write a so-called XRAC proposal to the XSEDE project – visit https://www.xsede.org/ for more information.

The next layer in the pyramid contains many campus-level machines – you could extend the pyramid further to desktops, and even cell phones, if you wish.

Our focus in this webcast is on the top tier.

5

That leads us to solicitation 14-518. By submitting proposals to NSF in response to this solicitation, you request resources on Blue Waters. Key to writing a successful proposal are two things:

Firstly, you must demonstrate that you have a very important science or engineering research problem that requires petascale computing capabilities. Secondly, you must show that your code will effectively exploit the capabilities offered by Blue Waters. Note that the capabilities include more than just compute power. Storage, I/O, and networking are equally relevant.

If you are a junior researcher, you are not at a disadvantage. We welcome such proposals – with you in the lead or as a co-PI.

The next PRAC deadline is coming up in a little over a month, on March 10.

6

While most NSF awards give you money, PRAC gives you primarily computational resources.

Be aware that these resources represent a big investment on NSF's side – Blue Waters is very expensive. Therefore what you get is precious and the proposals will get evaluated as carefully as any other NSF proposal. The review includes both the science and the suitability for execution on the Blue Waters machine.

In addition to computer time, you will also get free consulting by the Blue Waters team. So, you are not entirely on you own in developing and optimizing your computational application.

You can also request up to $40k of travel support – either for you or for the Blue Waters team to collaborate with you.

We expect that there will be 10 to 15 awards.

7

That completes the overview. Let's now look at Blue Waters in a bit more detail.

The goal of the system is to advance the frontiers of science and engineering by providing a computational capability that makes it possible for investigators to tackle much larger and more complex research challenges across a wide spectrum of domains.

In street language this means *you will get access to one of the most terrific research tools "to boldly take science where no-one has taken it before!"*

There is also a more mundane goal: We want to build a community that is capable of utilizing petascale computing.

8

Petascale computing is of interest in many science domains already. Examples of such domains include:

materials science, nano-engineering, fluid dynamics, climate and earth system dynamics, cosmology and astrophysics, chemistry and biochemistry, sustainability, health information technologies, cybersecurity, economics and social science, neuroinformatics and bioinformatics, as well as many different topics within physics, engineering, and increasingly in cross-disciplinary sciences.

If you work in any of these areas, PRAC may be of especially high interest to you.

9

Some of the specific features of the machine are that it aims to *sustain* petaflop performance.

The system is installed at the University of Illinois in Champaign. On the right are pictures of the petascale facility building and the machine room.

Blue Waters allows you to tackle much larger and more complex research challenges than in the past, across a wide spectrum of domains. The program expects to achieve significant impact by creating new knowledge about the natural world, by increasing industrial competitiveness, and by improving national security. Be sure you articulate how the proposed work will lead to such benefits.

Blue Waters entered full-scale operation in 2013.

10

Here are some architectural features of Blue Waters:

The architecture is designed for sustained petascale performance on several 100,000 cores for a broad class of problems.

It has 360,000 cores, structured into 22,000 nodes with two 8-core processors each. The processors are from AMD.

In addition there are 4000 nodes with one processor and one GPU each. The GPUs are NVIDIA Kepler.

This provides a total peak performance of just over 13 petaflops.

There are 1.5 petabytes of memory and the interconnect is a 3D Torus of type Cray Gemini.

The I/O system is massive as well. There is a bandwidth of more than 1 TB/s and the total storage capacity is 26 PB.

Another important capability of Blue Waters is that the data produced by the computation can be visualized on the system directly.

For more information, visit https://bluewaters.ncsa.illinois.edu/hardware-summary.

11

You need to convince the reviewers of your proposal that you have the necessary expertise for optimizing the performance of your computational code on the Blue Waters architecture. This expertise may include:

- Discovering and exploiting parallelism within codes,
- Overlapping different types of operations,
- Exploiting multi-level caches, local and remote main memory,
- Orchestrating intra-node and inter-node communication,
- Performing parallel I/O,
- Exploiting heterogeneous processors, and
- Dealing with petabyte-size memory-resident data and correspondingly large input-output datasets.

12

Clearly, such performance optimization is not trivial. Several years of preparation are usually necessary to get your code in shape so that it begins to scale to the degree expected for Blue Waters.

If you have just started to develop a computational code or you have begun to port a serial code to a parallel machine, it would not be a good idea to write a PRAC proposal. You would use a campus machine, instead, or request an XSEDE allocation.

For computational applications that have already demonstrated high scalability, the Blue Waters team will then help you improve the scalability further. But, you will need to bring in substantial expertise and staff for this job as well.

You also need to present a solid plan for how your code will perform well on Blue Waters. This plan may entail the following:

- You describe the starting point, which is high scalability on machines you have used so far.
- Then in the first year, you describe an initial set of performance optimizations that will scale your code to a part of the full machine.
- Next, in the second year, you may propose to apply further optimizations that will achieve scalability to the full machine.

Note that multi-year proposals are possible. In practice, proposals for more than three years are unlikely to be awarded.

The Blue Waters team will assist you in executing this plan. You are encouraged to contact them during the proposal preparation as well - https://bluewaters.ncsa.illinois.edu/contact-us. In the development plan you should indicate topics and tasks for which you need their help.

13

Now a few more specific tasks for your proposal writing:

In your project description, you will explain the target research problem and describe the intellectual merit and broader impact, i.e. the scientific progress enabled by the proposed work and how it affects a broader community. These two areas are important in any NSF proposal. PRAC is no exception.

You will then describe the computational codes, followed by the development plan (i.e., how you will bring your code to petascale performance on Blue Waters) and the source of funding that lets you do that. This is very important: PRAC does NOT provide funding for developing your code. You need to show that you have separate funding for that.

Then you describe the compute resources you need. They include: the number and type of nodes, the memory usage, the number and duration of runs in each phase of your development and the total node hours you need for all this. A node hour is one node that runs for one hour. You may need up to several million node hours.

We also need you to disclose resource allocations on other machines that you may have obtained (e.g., XRAC or DOE INCITE awards). This is taken very seriously.

14

Finally, here are the review criteria:

There are the standard NSF review criteria: does the proposal have intellectual merit and will it achieve broad impact? Again, these two criteria apply in ANY NSF proposal.

Then there are two solicitation-specific criteria:

1. Is the proposed research leading to breakthrough science and engineering?
2. Does the proposed research really require the capability of a resource such as Blue Waters?

15

Now I'll answer some frequently asked questions about the PRAC program.

The first one is: How can I tell if my proposed work is a good fit for the PRAC program?

We have just discussed this. A successful PRAC proposal shows that:

1. You have a very important science or engineering research problem that requires petascale computing capabilities, and

2. You have a solid plan to effectively exploit the capabilities offered by the Blue Waters machine.

16

The next question is: Can I be a PI on one proposal and a co-PI on another?

The answer is No, but you may be listed as senior personnel on another PRAC proposal.

17

Next question: Should I discuss my proposal with NSF Program Directors?

PIs are welcome to discuss planned proposals with the Program Director. When you call, good knowledge of the solicitation will be appreciated. Note, however, that, once submitted, the substance of proposals will not be discussed by NSF Program Directors. This would constitute unfair competition, or the perception thereof.

18

Another question: If PRAC does not fund the development of the computational application, which program does?

The answer is that you will need to write a proposal to an NSF program or to another agency that sponsors research in your science domain. Once you get it funded, consider writing a PRAC proposal.

19

Next question: When is it better to request resources from an XSEDE-allocated machine?

You should request XSEDE instead of Blue Waters resources if you are only beginning to parallelize your code. In this case request a startup allocation. This is fairly easy. If you don't have funding for developing your application into an efficient parallel code, you should also consider XSEDE resources or, better even, campus machines. A full XSEDE/XRAC proposal is appropriate if your code scales well, but does not need the massive capabilities of Blue Waters.

20

The final question before I open the floor to questions from the audience is: Will there be future PRAC deadlines?

The answer is yes. After the upcoming deadline on March 10, there will be another deadline in November of this year and then every year in November.

21

Thanks for listening to my presentation.

Here is contact information. Again, my name is Rudi Eigenmann, my email is reigenma@nsf.gov and my phone number is 703 292 2598

The presentation, script, and audio recording from today's webinar will be available on NSF's Event website www.nsf.gov/events

And now, let's see if there are additional questions from the audience.