Event ID: 2101381
Event Started: 2/26/2013 12:50:29 PM ET

Please stand by for realtime captions

Thank you for standing by. All participants will have a listen only line. During the Q and a session if you would like to ask a question, press*one on your phone. This call is being recorded. If you have any objections, you may disconnect. I will now turn it over to Barry Schneider.

Good afternoon, everybody. We have a small group gathered from the division of advanced cyberinfrastructure. We are glad you could make it. We will talk a bit about the recent solicitation that was announced -- HPC -- building a more inclusive environment. This is regarding NSF 13-528. I put a hyperlink to the frequently asked questions. In doing this, my colleagues, Eileen and Bob Chadduck talked about the possibility of making this a little bit bigger than just this presentation. I am sure that a lot of people out there know that there have been a lot of things going on with spent time talking about this and also what is going on -- I am getting used to saying ACI, but it's going slow.

Regarding the office of cyber infrastructure, we are now a division in the CISE directorate. We are engaged in a transition process. The website will be transitioning. Things are working so through.

I also wanted to take this opportunity to mention that there has recently penned a memo about open access. This is from John Holdren the science advisor. Open access to publications and data and how to deal with that. We think that people should spend some time..

We wanted to remind people of some preliminaries people there have been some changes in the GPG Starting January 14 , the changes are into areas. Perhaps three areas. Now you have to explicitly call out a 2 merit criteria in the project summary. There are some issues in the bio sketches. This allows you to cite things other than publications. Artifacts -- [indiscernible]. Another reminder is that cost sharing is prohibited. What is in the pool at the bottom -- state narrative in nature and stay away from quantifying things.

We decided since we are all here that we would list from the ACI website other solicitations that people might be interested in knowing about and thinking about. Some are close. If you are interested, take a look.

A general vision statement about ACI -- the overall vision without reading the slide is that we are providing digital resources -- hardware, software, training, education, and outreach to enable the research at all disciplines. The emphasis is that we want to enable science at all levels.

Here is a slide of the big picture. The long range goals. The first one - I already said. We are trying to support a instruction that is flexible and can manage a large collection of heterogeneous digital resources. We want to provide flexibility to accommodate data systems as well. Instruments -- it could be telescopes, etc.

At the same time, we have a responsibility to maintain the HPC environment for the users . Importantly, we want to bring these in new communities. We will shortly have these resources to help.

Also, provide capabilities for the national cyber infrastructure to collaborate with other cyber infrastructures here and [indiscernible] This is a list which I pulled -- is a webpage. This is the NSF Cyberinfrastructure investment. These seem to be the ones that were the most prominent on that page. Exceeds, reports, future grant -- these are things people know about. If you don't, I encourage you to go to this and take a look.

Some words on recent NSF investments -- the tier 1 resource is very high end and, leading-edge resource. This targets most challenging confrontational problems around the country. It is a machine which has a well-balanced mix of CPU and [indiscernible] processes. It has large storage. It is designed to handle small numbers projects -- high demand confrontational projects.

I also want to mention Yellowstone. This is a machine that has been put in by NCAR. It targets the GEO community.

Then, the extreme digital program -- you see on the next slide with the resources are. The main point is that this program targets a broad large class of computational data resources. This is managed by the [indiscernible] project. It will be dedicated March 27.

The new stampede and [indiscernible] resource. It is multicore. By 2015, we will get close to [indiscernible] Here is a list of the computational resources. This is on the [indiscernible]. The point I want to make is that this is a very large group of machines that do different things. We have a view that we want to provide a broad array of resources to the community of all kinds.- High-end things like stampede and previous good partners at [indiscernible]. The [indiscernible] machine. We have Gordon which is a machine for data intensive. Lone Star is for cluster trestles. It has higher throughput and longer running [indiscernible] Blacklight -- [indiscernible] [indiscernible] is a new resource which has a large number of TPU's with the current process processes.

Open science grid is a part of the tier 1 resource. We have high throughput machines and a nice test bed for HPC applications . The security people to go to the website.

-- I encourage people to go to the website.

We also have visualization resources. University of Tennessee Knoxville and Longhorn [indiscernible] which is specifically in place to deal with visualization.

We have a number of storage resources. Soon these resources will also be a part of the location. The project will be allocating resources.

If you measure success in terms of adoption, the next slide shows what has been happening. As you can see, we built the CI and the researchers came. Ever since the project has been turned on,

a number of service units are being given to the community -- close to 3 billion now. The number of users has increased substantially. I believe there are close to 9000 users on the tract to systems.

-- Track 2 systems.

. Then ACI strategic plan. This will set the tone for what we're going to do. You can read the slide. The point I want to concentrate on is the diversification part.

The idea here was not to put in place at this time another stampede. We are trying to diversify the portfolio and complement the capabilities that currently exist. We want to bring more of the long tail of science services into what the NSF is providing in terms of complications. Computations.

That is the main goal.

This is called lamenting the current investments that we have. I listed the solicitation on the slide. There are a number of things. In developing proposals, we do not intend to get for this to be an exhaustive list. We are not suggesting it is cast in concrete. We encourage imaginative thinking. This is just a few things.

Some details about the solicitation -- total funds available -- [indiscernible]. The project description is limited to 30 pages. You can list a proposal but be a sub awardee on others. The outcome really will depend on the proposals received.

I don't think we want to discourage novel ideas. We want to encourage him. We hope that as we see these proposals come in we will have enough flexibility to deal with this.

We did decide to The awards. There are two types of these resources. I will let Bob talk about the data. We did These. cap These. If you submit one that goes over, we will return it. We are happy to look at reports within the scope. The project duration is 14 years. The acquisition and deployment of the final [indiscernible] will be FY 15.

If you have a current resource you want to expand, that's fine. We will provide operating and management and maintenance costs at 20%. The proposals need to address something about the innovative capability. We certainly want to see evidence presented about this.

Thank you, Barry. It is my privilege to join this meeting. Thank you for leading us off this afternoon.

As you can see from the solicitation, this is intended to inspire inspiration and specifically in the areas of the opportunities for data focused infrastructure. This is a response to all of the aspirations that go with the goal and opportunity for an innovative high-performance data processing capability. What I will draw your attention to in the solicitation -- while all aspects are fully essential, one of the key pieces that you will see in the program description -- one of the key aspects of really a data processing capability intended to significantly advance the state of the art in the architecture and components and information processing approaches and as you can

read the solicitation, I won't read it to you. One of the key pieces that goes with this also is in terms of the aspects of the ability of the resource to complement the other advance cyber infrastructure investments and as you will see in the slide, the elements of support for accessing, analyzing, and disseminating data across cyber infrastructure. One of the key pieces is intended to be inspirational in terms of its broad focus.

One of the things that you see -- again, without reading the slide -- consideration of the review processes as they apply to the proposals.

One of the things that goes with this is what I described in terms of responding to these aspirations for data focused cyber infrastructure. The emphasis intended is for novel data intent of high-performance computing capabilities across multiple domains and disciplines. This opportunity has been presented to the community into science. One of the things, also, is that the elements are able to access information appropriately to support science and engineering and research and education. At the appropriate timescale and at the appropriate scale for advanced science.

What you will see if the element of a specific review criteria with which is applied to the consideration of proposed activities in the data resource area.

Thank you, very.

-- Barry.

I understand there are people unable to cannot. I apologize. I don't know what's going on. We will try to find out a way to get the information to you. This will be posted, but you may want to ask questions.

Here are some considerations. Basically, lamenting the capabilities and resisted a few other things here. Again, these are in the solicitation to some degree, but we are not trying to be exhaustive. We want to hear interesting and new ideas.

Capabilities, possibilities. We talked about this. Expanding boundaries. A plane enabling interactive research workflows and processing. You can look at these and think about this. Solicitation specific review criteria is listed here.

This is the last slide. Any questions? In doing the questions, we will have Bob, Irene and I. Thank you all.

If you would like to ask questions, please -- press *1 on your phone.

Good job. Nobody has a question.

We have one coming in. This question is from Gary [last name indiscernible] -- your line is open.

This is Andrew speaking with Gary. On one of your earlier slide, under budget you had an appropriate use of these funds not the whole dividend after full deployment. What does that mean?

We talked about that.

Could we bring up the slide?

Can I ask you follow-up to that?

Sure.

What I am wondering is whether that meant that part of our plan could include using some of the funds to buy commercial cloud cycles on-demand to deal with this. If that is the kind of thing that you imagine under that sentence?

There it is.

Yes, that is one example that would be appropriate. So, yes. That is that the only example that one could imagine. Certainly, that is an example of what that is intended to address.

Figure.

Is that okay?

Yes. I thought that's what it meant, but I didn't want to over interpret it.

There were other things as well.

Next question comes from Paul [last name indiscernible] from University of Notre Dame.

Can you guys discuss the performance metrics relative to what is going to be different architectures from virtualized systems to data bound systems? Your thoughts in general. Will this be the traditional -- what is the HPL benchmarks. Can you talk about that?

Yes.

For those of you familiar with previous solicitations, there has been a set of benchmarks required. Frankly, as we talked about this solicitation, there were 2 aspects that cause us to rethink that. One is that we had a deer -- clear data component. The world of benchmarking in respect to what we intended was not broad enough or specific enough to allow this type of novelty that we would like to have.

Similarly, as Barry mentioned, we recently named acquisition center deployments in Stampede and And car and [indiscernible] that were geared for the types of benchmarks that we know and love. We were trying to go -- allow for much broader interpretation of usage. So, the last bullet

on the slide -- convincing evidence presented -- that leaves things open. If you have benchmarks that you want to submit, this illustrates the capability you are proposing. They could be submitted. Again, we left this in saying that there needs to be some convincing evidence presented that a resource will perform. But, it doesn't have to be a traditional benchmark.

This is not to say that -- if you have something and you want to make a convincing case, you should present the evidence.

Does that answer your question?

Yes. Thank you.

Next question.

Next question comes from show group with -- Sharup from Lockheed Martin.

Yes, there is a slide entitled invigoration for competitive proposals. My question is -- could you please elaborate on what you term as dynamic and innovative workflows? In particular, what spectrum of workflows do you see today? I am referring to the third bullet in the slide.

What I would say is that -- this is my own feeling. Most of the resources that we have in place today have not [indiscernible] in an innovative way with the most complex workflows. Other areas of science. If you have ideas how to put in place those kinds of things, we would welcome this.

At the risk of being too specific, which we are not trying to be, I think if you look at the bigger picture of Cyberinfrastructure at large, all of the investments that we are making in this, one could at least imagine that some workflows could use a combination of the Cyberinfrastructure, maybe even in concert with what is proposed here. So, that area could be ripe for exploration. If that makes sense.

It does. Thank you.

We could imagine someone doing a large scale competition on Stampede and then wanting to take that data over to wanted the visualization sources. Process it. I think it could be done today,

---

There are also many instruments that NSF supports. Imagine workflows that would be more effective for the community. If the workflows between the instrument and analysis and constipation -- commutation -- commutation -- were more [indiscernible] This has been inspiring. In responding to the aspirations of science, I was focused on the fact that it is intended to be innovative, reliable, and dynamic. The aspects as we described -- this is another aspect of this particular solicitation intending to be inspirational in nature.

At the same time, it is a delicate walk of -- these are really for the research community. Production quality resources. And, there are ways to be both innovative and exploratory and still provide reliable and robust information.

Next question is from Eric [last name indiscernible]. Your line is open.

Could you clarify the timeline of the proposal and the four-year maintenance and operation budget to be negotiated?

Hello, Eric. I think I understand your question. The acquisition and deployment -- the submission is April 15. By January By January 2015 you have to provide a high degree of availability and usability of resources. It has to be available by the end of 2015. However, again, given these considerations, one would imagine some aspects being enhanced or further developed or evolved. Over the remaining life, even though the operation is a 2015 time frame. So, the operations and maintenance award, as Barry mentioned, is a separate action. It is up to 20% of the acquisition and it will begin when the resources became a locatable under the [indiscernible] process. The one currently in place. Does that answer your question?

Yes. Thank you.

Okay.

If you would like to ask a question, please press *1.

We have a question coming in. One moment.

This question is from Gary [last name indiscernible]. Your line is open now.

The title of the solicitation, basically, says to expand the communities. The communities that can use these resources. At the same time it is mostly hardware deployment. My question is -- in the operational phase is it your anticipation that money will be spent on things like education and out reach and training in getting people to use these resources and Jeanette to new communities as well? Or primarily focused on keeping the infrastructure going.

I think that is really up to the type of proposal and what you feel is appropriate to engage the communities you are targeting.

That is one interpretation.

As we discussed, [indiscernible]. The actual cost . This could be devoted to certain activities that they are justified. If you want to make sure that the user base can use it, I think that is what some of the money is for as well.

Thank you.

Next question is from Ralph [last name indiscernible]. Your line is open, sir.

I would like to follow up on that question. I read said that the operational part of the funding would start when the resource was allocated under the X XR a T. If part of the proposal is to

develop a new tape ability, would some part of the operational funding be available to support the development activity? This could result in hardware acquisition, but the question is would development costs before the acquisition be an allowable expense under the operations budget?

First, I am glad you got in, Ralph.

Yes, I was one of the people but got the message that the meeting was canceled.

There was an earlier link.

I know.

To get to your question, I think a lot depends on what you require to get to January 2015. Remember, the operations fund only kicks in once it is truly operational and deploy. If there is funding required to make it deployable, that should be in the acquisition and deployment and the operation should really start with the actual operation of the system. This is bound tightly in a proposal. It has to be substantially available by January 2015. Fully available by the end of 2015.

I understand. Thank you.

The next questions from -- comes from the line of Rob Pennington.

We have a question about the maintenance and operating funds. Is that a plan to pay the direct electrical and cooling costs? Or do you have other guidance to share about this?

[indiscernible - multiple speakers] We know who you are. You came in as disguised as Robert. We know who you are.

The question I heard was -- is the electrical power a part of the operations and maintenance -- is that the question?

Is it part of that? Do you plan to pay the actual cost of power and cooling -- whatever the Duke components are that are deployed.

This is your proposal. You tell us.

Are you asking if it is legitimate?

Yes, then that would be an acceptable cost.

Thank you.

If you would like to ask a question, please rest *1 please press *1.

One moment.

Next question is from John Strand he. -- John [last name indiscernible] This is Richard Moore masquerading as Sean. In the solicitation you talked about the 20% operating costs and then say an additional 5% can be requested with strong justification. Could you give an example of the kinds of -- what might constitute from your perspective strong justification for going up?

I don't know. A heart attack? I'm not sure.

What we are trying to say there is that we want to stay open to the idea that if the world is good and we have money, we might be able to do it. It has to be a strong case.

Okay.

The next question is from the line of tenet [last name indiscernible]. Your line is now open.

-- Kenneth [last name indiscernible] Barry mentioned the longtail of science. My experience is that the farther you get on the tail, the more people need interfaces rather than hardware. Yet, this proposal seems to be more attuned to the high-performance side of things. That is the reading that I get. Could you clarify that?

Maybe very will. Barry will. I don't know what parts of the solicitation need to be seen in that way.

I think I understand the question. Why you're asking him. To some extent I think it is true. This was at the higher end. To a certain degree, that is what we are supposed to be doing. ACI has clusters. We have software and a large amount of the funding goes into software. These are the things I think you are talking about.

There are solicitation juke and look at that address those issues. I think this one is less so.

Thanks.

Maybe I misunderstood. So, you consider clouds to be hardware. I guess they really are.

To be understand your question correctly?

I think you answered it. If I interact with something like a social scientist, the issues they have in accessing this equipment is much more user interface issues than specific hardware issues.

I understand. Like API.

Or the innovative and reliable workflows for example.

It is clear, thanks.

Okay.

Next question is from the line of Rob Pennington.

This is [indiscernible]. That doesn't sound like Robin. [laughter] My question is whether budget [indiscernible] is allowable given that this will be experimental.

Can you be clear about this?

I think the answer is no.

We can't do it.

Again, if you would like to ask a question, please press *1.

The next question comes from the line of Paul Brenner at the University of Notre Dame.

After one of the questions I had to go back and check the solicitation. With the NSF cost sharing requirements being prohibited for voluntary, how would one offer to pay the power bill without violating the cost share?

I will have very go back Barry go back to that slide. The first slide where it talks about -- there it is.

Voluntary cost sharing is recommended for this solicitation as it is for all of NSF. That was effective December 2011. There were 2 reasons for that change. One had to do with trying to level the playing field. The second one was to reduce audit expense on poach -- both NSF and the institution. Therefore, anything that the institution was like to contribute in-kind to the project cannot be in the proposal itself. The facility equipment and other resources -- as Berry has said, the description has to be narrative and not quantifiable in any way. And it is not considered in the review and the evaluation. In other words, it is not in the proposal itself. It is really in these other documents. There are ways to do many things, but, does that provide the clarity that you need?

Yes and no.

I think the facility makes sense, but as we all know what the big systems, the power bill alone can't be white substantial. And if I don't mention the power bill in the proposal, and I just make the system bigger, it has a competitive edge, I guess, but it is still a little gray to me that a half $1 million power bill would not get mentioned.

First, it would not be an issue in the acquisition and deployment. It is really in the O and M. Although you would have to clarify this.

[indiscernible] To your point about buying more hardware, I don't see how that would work. They are actually -- the award is separate from the Onan. You can take money from this and put it into the [indiscernible].

Yes, I guess it would have been better post as people versus the power bill for support, etc., in the OHM.

Like I said, it is really up to you. I continue with the guidelines say, but I can't really --

---

I think the specific questionnaire asking is not how to pay for it, but can you pay for it out of your own pocket without cost-sharing?

A proposal where you have an extra 500,000 to get the people rather than pay the power bill seems more attractive to me. Whether the institution or the people submitting proposals to do that, my question is more if it is legal by the NSF guidelines.

I think the issue here is facilities and resources and other equipment. I wouldn't isolate this. There are all kinds of facilities and resources that could be available to some projects but not others depending on circumstance and the place and the focus of what the intended resource is.

I would like to answer it in that context. I don't think there is anything specific about the issue.

These resources -- if you look at the storage and how it may or may not be a consideration.

I will into the question they are.

Again, if you would like to ask a question, please press *1.

One moment.

Our next question is from Thomas Hauser. Your line is open.

I want to follow-up on this. On the facility and cost share question, part of the proposal of the new quite -- criteria is the [indiscernible]. If I have a proposal where everything is ready so that the data center is there, but somebody may have networking, do I need to put that in the section or can I mention that on the link to the [Indiscernible - heavy accent] -- things like that?

I think the intent of this -- I think you are talking about the difference between facility equipment and other resources. Versus the requirement to demonstrate that you actually have facilities to put in whatever your putting. Those, prickly, are two different things. Facilities, equipment and other resources as defined in the GPT and it is not necessarily targeted toward this award. The intent of the review criteria is to ensure that there is some place that will be available where cyber security has the physical infrastructure necessary.

So, I wouldn't want there to be confusion between these two issues.

Okay.

And, as Barry suggested, the FRO will be very familiar -- in the proposal itself. That is different from what we are asking you to describe about the facilities and cyber security.

In some cases, I think you would need to demonstrate that the assets were protected.

Okay.

One moment.

No more questions are coming in at this time.

We were reminded that the presentation has been posted. Bob and Mary Barry and I [indiscernible] Yes, and if some people could get in and they went to ask questions, they could contact us.

Thank you all.

Thank you.

[Event concluded]