This webinar covers solicitation NSF 13-602, The NSF Cloud infrastructure, and its re-issuance.

This solicitation seeks to enable new research in service provisioning via new resource virtualization mechanisms. It also seeks to broaden the application areas that can use cloud computing. We expect that the infrastructure will further support research both into the design, provisioning, and management of a reliable, persistent, real-time cloud architecture, as well as into the design and deployment of large-scale distributed services and applications that use such an environment. Ultimately, the research infrastructure will enable researchers to go beyond the use of existing commercial cloud offerings, allowing them to influence such offerings in the future.

This broadening may include the combination of dynamically configured clouds with traditional high-performance computing, especially for education, simulation and cyber-physical systems.

Because the solicitation is aimed at broadening the use of cloud computing into new application areas, proposers are reminded to think of the relevant communities more broadly than just the traditional CISE/CNS researchers.

We do not believe that research on novel cloud architectures is necessarily distinct from research on novel applications that advance the use of cloud computing. The two aspects complement and drive each other.

We hope that proposals will find areas of common interest which extend, expand and enhance both communities and which lay bases for long-term infrastructure sustainability for both communities. Nevertheless, this solicitation is primarily intended to support cloud computing, and not, for example, high performance computing.

CISE anticipates two phases of activity. Phase I, to be enabled by this solicitation, will support required infrastructure design and ramp-up activities, as well as demonstration of readiness for full-fledged execution. This phase focuses on the concrete design and implementation activities necessary for eventual scaling to a larger, community infrastructure should additional funds become available.

Approximately one year prior to the end of each Phase I award, each funded infrastructure activity will be evaluated by an external committee based on the goals of the NSFCloud program as well as specific metrics outlined in the original proposal. Depending upon the outcome of this review, NSF may choose to ramp down a funded project over the remaining duration of the award. A successful project will continue to receive the planned support through the remainder of the Phase I grant period for that project.

Subject to the availability of funds, CISE expects to issue a solicitation for Phase II of NSFCloud, seeking competitive proposals from among the Phase I awardees describing how the infrastructure activity would transition to full operations. The exact timing may be dependent on NSF's evaluation of Phase I progress and the performance periods specified in the Phase I grants. In Phase II, the infrastructure is expected to become fully staffed and operational, fulfilling the proposed mission of serving as a testbed

that is used extensively by the research community. Subject to the availability of funds, Phase II will have a similar duration and funding level.

Reiterating a key point: the solicitation requests proposals for infrastructure. While there may be research questions associated with the architecture or technology of the infrastructure, individual research is not appropriate for, or conformant with the requirements of, the solicitation. The research that uses the infrastructure will come from existing and/or future solicitations.

We're holding this webinar to discuss the solicitation and to answer questions. If you've not had a chance to read it yet, please be sure to do so.

It is good practice to make lists of the proposal's "shoulds and shalls " and use these lists as part of the process of planning a proposal.

While the total funding is $20M and the number of expected awards is 2, this does not necessarily mean that awards must be exactly $10M. All proposals must present budgets that realistically reflect the work to be done. All proposals may be subject to scope change during award negotiations. NSF will consider the proposal merit and likely community impact in making any final decisions.

While the maximum project length is 4 years, proposals for projects of less than 4 years are welcome. In particular, proposals that can describe a realistic path to community impact at earlier dates are appropriate and may influence panel evaluations.

Proposers are reminded that there may be a tradeoffs between staffing costs and equipment purchases.

NSF will issue the awards as cooperative agreements. This means that, unlike traditional grants, NSF will communicate regularly with the projects and provide both help with community issues and advice on project management. Proposers should carefully consider the impact of the cooperative agreements on project structure and management. In particular, it may influence the choice of whether the proposal is a collaborative proposal between multiple institutions or whether sub-awards are used to bring in other participants.

Proposals can be submitted by any entity eligible under Chapter 1, Section E of the GPG. This list is not restricted to universities and is a change from the original issuance of the solicitation. Please read the GPG for a full list of eligible entities.

NSF will obtain reviews from a range of non-conflicted experts. In addition to cloud computing, the reviewers may have expertise in project management, community engagement, high-performance computing, networking, and other areas as identified by NSF. At NSF's discretion proposers may be requested to come to NSF for a Reverse Site Visit.

We will ask our reviewers to assess the NSF Intellectual Merit and Broader Impact criteria. In addition they will be asked to comment on a set of specific review criteria.

Because this is an infrastructure project these four additional criteria focus on team experience in delivering infrastructure and exposition of a management plan and project process that we believe are appropriate to the project scope and goals.

The first criteria is intended match the facility to what is known of community needs. The second is intended to encourage reuse of existing assets.

Note that the solicitation provides the evaluation criteria; however, this does not mean that capabilities of proposed architectures or positive attributes of submissions beyond those required by the solicitation will be ignored in the evaluation process. NSF will be listening carefully to panel feedback on positive attributes of the proposals.

We expect that NSFCloud projects will need to fund some or all of the Phase I infrastructure. Projects will also likely need to develop software as determined by the proposal architecture. The exact ratio and levels of equipment expenses versus development expenses will be determined by the architectures and by the potential reuse of existing resources and software. We expect the proposers to figure out the appropriate balance.

In all cases budgets must be justified and appropriate for the level of effort.

When buying equipment, it is good practice to attach equipment quotes justifying the budget to the proposal as supplementary documents. We understand that these quotes are subject to revision as prices change over time.

Travel and coordination costs should match and be consistent with the structure of the project management plan.

Remember that the solicitation states that proposal budgets should include funds for domestic travel for two two-day meetings each year.

The past provides a basis for the future.

CISE/CNS has a significant investment in mid-scale infrastructure. The GENI program is deploying “GENI Racks” (actually small computing clusters) at approximately 50 universities across the US. The MRI and CRI programs have funded a significant number of computing clusters dedicated to systems research. Large-scale clusters such as PRObE are potentially available for integration into nation-wide experiments.

NSF realizes that the broader scope of future cloud computing means that there are similarities and potential overlaps with CISE/ACI programs in HPC. These provide a potentially exciting base for new application structures. For example, the use of HPC facilities in training and simulation with cloud-based front ends. Control of cyber-physical systems is another area with exciting potential for the combination of cloud and HPC technologies. As has already been demonstrated in the CC-NIE program, we expect that there will be significant other cyber infrastructure assets that can be shared across the CNS and ACI communities.

Thus while the center of this solicitation is in CISE/CNS we welcome well-conceived plans for leveraging a larger set of existing infrastructure resources. Leveraging existing resources may enable greater impact for less cost. It may also reduce the time and risk to deploy NSFCloud. Thoughtful proposals may also provide a path for existing user populations to rapidly become users of NSFCloud.

As Keith stated, the first phase of the NSFCloud program is about designing the infrastructure, facilitating early experiments to demonstrate the value of the infrastructure and having a solid plan for a wider build out in phase II.

Thus in phase I, that is to say "this solicitation," the proposers must show that they have a solid architectural concept and a plan for implementing and deploying sufficient facilities such that the relevant communities can be engaged and can run "interesting" experiments.

Proposers should describe future experimentation as examples that illustrate how the architectural choices, deployed facilities, and supporting tools, etc., satisfy the solicitations requirement that the infrastructure "…support both experiments involving novel cloud computing architectures and experiments involving novel applications of cloud computing."

While we expect proposals to address the topics described, we do not dictate the order in which they are addressed. Proposers are free to choose the order that makes the most sense in a coherent narrative.

This section is limited to 20 pages.

In the next 9 slides we will discuss the material that must be present in the 20 pages of the proposal project description.

The project vision should provide a clear roadmap of the initial deployments and future growth opportunities. It should also provide the roadmap for the engagement with the community. Details of these topics will likely be present in later sections. In this section we seek the vision against which the details can be assessed.

Innovation in and with cloud computing means that both the classical systems and networking communities and new domains (such as CPS) will use (and possibly fund parts of) the infrastructure

This section needs to describe the components, how they will fit together to build a system and the proposer's assessment of the level of effort required to build new components or modify existing ones. New components, that is software or hardware building blocks that are not simple reuse or minor modifications of existing ones should be described in significantly more detail than existing ones. The description of new components should consider the risk associated with the component development.

The material in this section should allow reviewers to understand what is being developed and how much it will cost. Cost estimates based on actual experience with comparable software modules are encouraged

The purpose of this section is not to define a research agenda for the design of the infrastructure, but rather to show how the infrastructure will support "compelling new research and **education** opportunities."

Keep in mind the solicitation goals with respect to "enabled research"

"... infrastructure will enable exploration of resource sharing in clustered computing, virtualization with software-defined networking technologies, quality of service guarantees, and the interplay among applications, cloud computing architectures, and the physical environment."

It is important to consider how the infrastructure will

-- support multiple experiments of "sufficient size"

-- enable evaluation of different cloud configurations at scale, long-running, and instrumented

-- provide isolation between experiments.

Remember that this solicitation seeks:

"infrastructure that will specifically enable the academic research community to:

a) develop and experiment with novel cloud architectures and

b) pursue architecturally-enabled novel applications of cloud computing"

"We also seek to facilitate interactions between the broader academic community as well as industry, including cloud computing researchers, the high-performance computing community, and cyber-physical systems researchers"

Consider how the "community" should be involved in defining the space of "novel cloud architectures" "architecturally-enabled novel applications"

This is a BIG infrastructure project -- not a research project. To succeed, substantial project management skills will be needed.

Consider the organization of your team carefully (examples of needed skills)

-- systems engineering

-- large scale project management

-- operational management

-- community engagement

Remember that "proposals should describe [...] how this infrastructure will leverage existing NSF investments or commercial software".

Where new development is needed, be clear how the skills of your team support this.

This is a large infrastructure project that will require well organized management and coordination. We expect proposers to show that they understand how to organize the various activities, coordinate the results so that there is a high probability of success, and coordinate the architecture, development and early deployment activities with the external community.

Remember that the awards will be issued as cooperative agreements and that NSF will actively engage the awardees. Proposers are strongly encouraged to describe how their management structures will support the cooperative agreement.

If parts of a proposed NSFCloud project depend on existing infrastructure the required management and coordination functions and responsibilities should be documented.

A project of this scale necessarily has risks. These risks can be associated with research questions associated with the architecture, equipment related, organizational, community. Thus, since the end goal to deliver operational infrastructure to the community, your proposal must show that you understand how to manage risks.

We are not expecting you to run a risk-free program. However, we are expecting that the risks associated with organization, new architectural ideas, software,

hardware or tools be carefully and professionally managed.

We do **not** require conformance to risk management standards.

We **do** seek thoughtful assessment that highlights a critical part of the proposer's project management skills.

Metrics are a fundamental part of evaluating and communicating the ongoing success of the program, for informing the risk management process, and for informing the project changes that are likely to be needed as the project progresses.

Your proposal must identify an initial set of metrics, including short- and long-term progress.

The metrics that are appropriate for early stages of the project may not be appropriate for later stages.

Remember: You need to be able to communicate the value of the infrastructure to the larger CISE community.

The proposers should describe their long term vision for the NSFCloud infrastructure and show how the elements of the proposal support that vision.

This must include a longer-term vision for transition to community governance and sustainability. Proposers should describe the management aspects associated with scaling to a larger infrastructure should the infrastructure advance to Phase II with additional funding.

Remember: Over the long term there will be users/participants and infrastructure providers other than NSF.

Please read the solicitation carefully to:

Understand the goals of this solicitation

Understand the proposal requirements

Understand the review criteria

Questions ?