



**FAILURE TO RETURN THIS ADDENDUM IN ACCORDANCE WITH THE
INSTRUCTIONS MAY SUBJECT YOUR BID TO REJECTION.**

To: Prospective Offerors

Subj: UNC-CH Bid Number: 3000012020
Bid Opening Date/Time: 01-16-2025
Addendum Date: 12-04-2024
Addendum Number: 01

This addendum is being issued to the subject IFB to answer questions raised during the time for questions allowed for in the IFB. Please note that the bid opening date has been changed to read 01-16-2025 vice 01-10-2025. All responses are to be submitted accordingly. All other information contained in the IFB remains the same.

Question 01: Some of the minimum requirements for each faculty member participating in this pilot have been referenced but we don't know what the expected usage will be for each or are there more specifics you could share for each Faculty member's needs?

Answer: The examples of faculty member projects minimum requirements are illustrative and do not necessarily reflect actual projects. We would like the vendor to describe the hardware configuration that they can provision that will meet the minimum requirements of the example projects. The actual usage can then be appropriately scaled based on the eventual mix of projects that will utilize the provided resources to the limit of the appropriate overall budget.

Question 02: Is there an expectation for these workloads to scale as needed by the faculty members or will it be a fixed solution for budget purposes? Is usage going to vary by user, or is the expectation to build out one cluster that meets the needs of each user and everyone would have access to? Or do we need multiple clusters, potentially one for each faculty member?

Answer: The examples of faculty member projects workloads are illustrative and do not necessarily reflect actual projects. Usage will vary by user, however the expectation is the provisioning of a single shared system that can accommodate all users in a shared manner, potentially using appropriate queueing systems for fair sharing. It would be appropriate for submissions to include the options of on-demand versus queued access or a hybrid approach using a combination of queueing and on-demand and also include different pricing structures for different access scenarios.

Question 03: Learning models and machine learning were mentioned. Is there information regarding the size and type of data that will be used to feed the models?

Answer: The data will be of multiple modalities such as text, electronic health records, various types of imaging data, engineering and scientific data, health data. Depending on the particular research problem, data can be as large as hundreds of terabytes.

Question 04: How are they connecting to these machines to feed it the data?

Answer: It is expected that the data will be moved to the machines using a secure process such as sftp. For exceptionally large datasets, other processes such as shipping data on secure hard drives may be needed.

Question 05: Do we need one or more direct connections to sustain the level of data transfer required? If so, what size pipes will be used, and what is the expected amount of data transfer per month?

Answer: It is expected that the data center will provide a minimum of a 40G network connection.

Question 06: What level of coding is currently being performed by researchers in the area of HPC computation?

Answer: The researchers are experts in their domain of study and will have the ability to create code that can solve the problems of interest. However, they may not be experts in HPC and scaling of the code to large model sizes. It is expected that the vendor will provide training and assistance in helping users scale their code to accommodate the HPC system provided.

Question 07: What level of IaC development coding does the technical team have?

Answer: Please see answer above.

Question 08: Queued resources (i.e. jobs), is there any scheduler or scheduling of current workloads being utilized now?

Answer: Since this is a new procurement, there are no existing schedulers of choice. The vendor is encouraged to provide one of the common scheduling systems used in HPC systems such as slurm.

Question 09: What coding languages are currently used by research teams? Technical teams?

Answer: We believe that the research and technical teams use a variety of coding languages common in scientific computing including versions of FORTRAN, C, C++, Python, MATLAB, SAS, GO etc. The vendor should provide a list of languages that they can support on the proposed system.

Question 10: Is there a current IAM structure that is in place that will be able to link to the HPC environment? Active Directory? Office365? ONYEN?

Answer: ONYEN

Question 11: Please define the level of live technical and customer support. Will users be active 24/7/365? Will the technical team require SLA service time leads for support?

Answer: Users will be active 24/7/365. However, the vendor should specify the level of live as well as asynchronous customer and technical support that they can provide within the proposed budget.

Question 12: Are there any current workloads or structures utilizing Microsoft Azure? AWS? GCP?

Answer: The projects targeted for this system will be originating from a variety of sources ranging from desktops, on-premise HPC systems, and all-commercial cloud vendors. However our expectation is that none of the workloads will have dependence on a particular cloud vendor.

Question 13: Is it expected that the HPC provider host a preconfigured set of common software that is available for researchers to use in the HPC environment (e.g., ANSYS, MATLAB, or other common software tools)? If so, is there a list?

Answer: Yes, it is expected that the HPC provider will host a preconfigured list of common software. Such a list will be provided on request once the projects are selected (which likely will be after vendor selection takes place). The vendor may want to address what software they have hosted and/or ported for current/past clients.

Question 14: Is it expected that the detailed code support and training provided to researchers extends beyond “does it run” to “does it run well/does it get the right answer”? Is the expectation that the support be provided by engineers with modeling and simulation backgrounds?

Answer: The vendor can choose to describe the level (or range) of support they can/will provide. At a minimum, the vendor is expected to provide support to port the software to the system and ensure that it is running correctly and provide guidance in scaling the software if needed.

Question 15: If \$2M in funding is provided Y1, is there an expectation for, as appropriate, a “funding rollover” provision in Y2?

Answer: If the question is asking what will happen if less than \$2M is spent during the first calendar and/or fiscal year, will it remain available the following year, then the answer is, “Yes.” This is recurring funding from the NC General Assembly with \$2M provided every fiscal year (July 1 through June 30), and unspent funds do not revert at the end of the fiscal year but carry forward. Because this is a pilot program with numerous unknowns about number of projects that can be supported with the HPC resources being requested, and the full \$2M was not spent during a fiscal year, the Collaboratory likely would add additional projects to utilize the available funding.

Question 16: Is there a requirement or guidance around whether technical support / training must be provided by US persons or based in NC?

Answer: There is no requirement stated on technical support / training. However, the RFP does state that preference may be given to vendors (and related support assets) within the State of North Carolina.

Question 17: Is there a specific number of advanced support hours per project, or specific number of projects expected per year?

Answer: The vendor should specify how many hours of advanced support they can provide annually. This may impact the scope of this pilot project regarding decisions on what project(s) is/are chosen and how many are chosen to participate in utilized the vendor's HPC resources.

Question 18: Is it expected that the Collaboratory or each individual project is being obligated to pre-buy storage allocations or be charged based on consumption?

Answer: The Collaboratory would prefer charges based on consumption rather than pre-buying in a situation where project landscale currently is unknown (i.e., projects/researcher have not yet been chosen). The vendor is expected to specify the amount of storage they can provide with the system (and costs associated with storage). The expectation is that the projects selected by the Collaboratory to utilize these HPC resources will fit within the provided storage. Under exceptional circumstances, should a project require storage beyond this provisional capacity, the Collaboratory can discuss additional resources to accommodate such cases.

Notwithstanding the answers to the questions provided above, we encourage the vendor to consider alternative creative solutions that advance academic research that benefits from digital engineering as defined in the RFP. The Collaboratory will consider and evaluate such approaches.

1. Check ONE of the following options:

- Bid has not been mailed. Any changes resulting from this addendum are included in our Bid.
- Bid has already been mailed. No changes resulted from this addendum.
- Bid has already been mailed. Changes resulting from this addendum are as follows.

2. This addendum must be signed below and submitted with your:

- a. Bid if your Bid has not been mailed.
- b. Separate correspondence if you have already mailed your bid.

-
Execute Addendum

Company Name: _____

Authorized Signature: _____

Name and Title (Typed): _____

Date: _____

Dtd 09-01-2010