

## Communications regarding UCAR RFP000074 (NWSC-3)

Revision:

Version 1.1, 13 April 2020

### Table of Contents

|     |  |   |
|-----|--|---|
| 1   | Overview.....  | 1 |
| 2   | Conventions.....   | 1 |
| 2.1 | Example brief description of question .....  | 1 |
| 3   | RFP Questions and Answers, issued 13 April 2020 .....  | 2 |
| 3.1 | Attachment 1, Technical Specifications, Section 1, Software.....                                 | 2 |
| 3.2 | Attachment 1, Technical Specifications, Section 3.3.4, Production PFS.....                       | 2 |
| 3.3 | Attachment 2, Benchmark Rules, Sections 5.1.3 and 5.1.4, and Benchmark Website Instructions..... | 2 |
| 3.4 | Attachment 2, Benchmark Rules, Section 5.1.1, and Benchmark Website Instructions .....           | 3 |
| 3.5 | Attachment 2, Benchmark Rules, Section 5.3.....  | 4 |

## 1 Overview

This document contains prospective Offeror questions related to UCAR RFP000074 (NWSC-3) and UCAR’s responses to those questions.

## 2 Conventions

Each question and its corresponding response is formatted as shown below, providing a unique question identifier and a brief title for the question, the question itself, and UCAR’s response to the question.

Example:

### 2.1 Example brief description of question

Question            The text of the Respondent’s question will appear here. It may be stated verbatim or modified slightly to remove any irrelevant attributes of the question or any indication of the Offeror’s identity.

UCAR’s response to the question immediately follows.

### 3 RFP Questions and Answers, issued 13 April 2020

The following questions were received by UCAR between the release of the RFP (on 02 April 2020) and 13 April 2020.

#### 3.1 Attachment 1, Technical Specifications, Section 1, Software

**Question** Prior to submitting our "Registration of Interest," we are seeking confirmation on the response requirement. Will NCAR accept a proposal for a software portion only, or does the response need to include all components, i.e. software, hardware, and storage, to be accepted?

An Offeror proposal in response to UCAR RFP000074 must include a complete NWSC-3 solution, comprising all hardware, software, infrastructure, networking, delivery, installation, and five (5) years of software licenses and hardware/software maintenance, support, and other services. An exception, as described in §2 of Attachment 1 of the RFP, is provided for an Offeror who chooses to propose only an HPC or PFS solution. If an Offeror wishes to submit a quotation for a specific hardware or software component of NWSC-3, the Offeror may do so, but it will not be considered a response to UCAR RFP000074.

#### 3.2 Attachment 1, Technical Specifications, Section 3.3.4, Production PFS

**Question** As stated in Section 3.3.4, "*The PFS solution shall have an initial usable file system capacity of 60 PB (petabytes) and a rack infrastructure that allows the usable capacity to be doubled by the simple addition of data storage devices.*" Does this mean it is required that all of the needed additional infrastructure, such as drive enclosures, controllers, cables, racks, and power be in place at the initial installation, so that doubling the capacity is done by merely adding HDDs (and SSDs as specified in 3.3.5)?

UCAR's requirement stipulates that the proposed solution has the ability to increase capacity simply by adding additional HDD/SSD drives. The Offeror's proposed solution should include all of the needed storage infrastructure components, such as drive enclosures, controllers, cables, and rack power in place at the initial installation. If the architecture allows for additional drive enclosures and cabling to easily be added within the rack/controller infrastructure, that is an acceptable alternative, as long as it can be done in a manner that is non-disruptive to the services provided by the initially installed storage.

#### 3.3 Attachment 2, Benchmark Rules, Sections 5.1.3 and 5.1.4, and Benchmark Website Instructions

**Question** For the CESM2\_MG2 kernel benchmark, the last sentence of the first paragraph of page 2 of the PDF containing instructions on the benchmarks website requests: "*Please provide output files for a number of MPI ranks that both fully-subscribed and over-subscribed hardware cores,*" but it is stated on page 10 of the UCAR\_RFP000074\_Attachment\_2\_NWSC-3\_Benchmark\_Rules\_v1.docx in Section 5.1.3 MG2 that "*MG2 should be run*

*on a single node, using all available cores, and using one MPI rank for each of the available cores."*

Analogous to CESM2\_MG2, for the WACCM\_imp\_sol\_vector kernel benchmark, the last sentence of the second paragraph on page 2 requests: *"Please provide output files for a number of MPI ranks that both fully-subscribed and over-subscribed hardware cores,"* but it is stated on page 10 of the UCAR\_RFP000074\_Attachment\_2\_NWSC-3\_Benchmark\_Rules\_v1.docx in Section 5.1.4 WACCM that *"WACCM should be run on a single node, using all available cores, and using one MPI rank for each of the available cores."*

Do the benchmark rules override the PDF so that oversubscribed runs are no longer required? Conversely, if oversubscribed runs are still required or desired, then which achieved figure of merit (FOM) must be entered into the UCAR\_RFP000074\_Attachment\_2A\_Benchmark\_Results\_Spreadsheet\_v1.xlsx; i.e., the best FOM or always the FOM from the fully subscribed (but not over-subscribed) run, even if the oversubscribed run yielded a higher FOM?

UCAR would like the benchmark results to be returned for both the fully subscribed and oversubscribed cases, as requested in the instructions provided on the NCAR HPC Benchmarks website<sup>1</sup>. The result for the fully subscribed case (i.e., one MPI rank for each available core) should be used as the figure of merit (FOM) to enter in the Benchmark Results spreadsheet<sup>2</sup>.

### **3.4 Attachment 2, Benchmark Rules, Section 5.1.1, and Benchmark Website Instructions**

Question                      Based on the following language found in Section 5.1.1 of Attachment 2: *"5.1.1 CLUBB: 'CLUBB should be run on a single node, using all available cores, and using one MPI rank for each of the available cores,'"* the results for this benchmark will be for runs on a node of the proposed system which is fully subscribed with MPI tasks but NOT oversubscribed (that is, with hyper-threads) as requested in previous documentation, correct? The CLUBB benchmark data only provide reference files for pcols=16 and pcols=192. The README and PDF state that results for any value between 16 and 192 would be accepted. Without the reference files, there is no way to verify the results of a different value of pcols between 16 and 192. Is it correct then to assume we can only test with pcols=16 and pcols=192 for CLUBB?

For CLUBB, the fully subscribed result (one MPI rank per core) is required to be returned and should be entered into the Benchmark Results spreadsheet<sup>2</sup> as the figure of merit (FOM). An oversubscribed result may optionally be returned, in addition to the fully subscribed result, if it showcases interesting performance.

The CLUBB benchmark is used outside of the NWSC-3 benchmark suite with other values for pcols, hence the language in the README and instructions. However, for the NWSC-3 procurement, you are correct: we are only requesting results for pcols=16 and/or pcols=192. For CLUBB, the fully subscribed result (one MPI rank per core) is required to be returned and should be entered into the Benchmark Results

spreadsheet<sup>2</sup> as the FOM. An oversubscribed result may optionally be returned, in addition to the fully subscribed result, if it showcases interesting performance.

### **3.5 Attachment 2, Benchmark Rules, Section 5.3**

Question            The benchmark rules document mentions two Microbenchmarks, STREAM and OSU MPI, that vendors need to complete as part of the RFP requirements. However, the results spreadsheet supplied doesn't have provision to include results from these two micro benchmarks. Please clarify.

The primary purpose of the Benchmark Results spreadsheet<sup>2</sup> is to calculate the aggregate Cheyenne Sustained Equivalent Performance (CSEP) value. Since CSEP is intended to be a comparative measure of a system's capacity based upon the relative performance of NCAR applications, the synthetic STREAM and MPI benchmark results are not expected to be entered into the spreadsheet. Nevertheless, the STREAM and MPI benchmark results are important to UCAR's assessment; thus, they should be returned as files capturing STDERR and STDOUT. The STREAM and MPI benchmarks are required to be run, and their results are required to be returned with the Offeror's proposal.

<sup>1</sup> NCAR HPC Benchmarks Website: [https://www2.cisl.ucar.edu/hpc\\_benchmarking](https://www2.cisl.ucar.edu/hpc_benchmarking)

<sup>2</sup> UCAR\_RFP000074\_Attachment\_2A\_Benchmark\_Results\_Spreadsheet\_v1.xlsx