University of Central Florida

Procurement Services

ADDENDUM

IMPORTANT DOCUMENT – INVITATION TO BID ADDENDUM

ITB NUMBER: 2020-04GCS OPENING DATE & TIME: **CHANGED to September 8, 2020 @ 3PM EST.**

ITB TITLE: High Performance Computers

ADDENDUM NUMBER: 1 ADDENDUM DATE: August 21, 2020

The purpose of this addendum is to answer questions asked during the open q/a period and the specifications from page 4 of the bid have been revised. See next page.

PLEASE ACKNOWLEDGE RECEIPT OF THIS ADDENDUM AND RETURN IT WITH YOUR BID. FAILURE TO SIGN AND RETURN WITH YOUR BID COULD RESULT IN REJECTION OF YOUR BID.

_____ _____

PROPOSERS SIGNATURE PRINT OR TYPE PROPOSER'S NAME

_____ _____

COMPANY NAME EMAIL ADDRESS

12479 Research Parkway• Orlando, FL 32826-3248 • (407) 823-2661 • FAX (407) 823-5551
Orlando Tech Center

BEFORE	*REVISED*
• GPUs: 2x Quadro RTX 6000 + NVLink ○ Processor: Intel Core i9-9920X ○ Operating System Drive: 2TB SSD (NVMe) ○ Data Drive: 3.84 TB SSD (SATA) ○ Memory: 128 GB ○ Operating System: Ubuntu 18.04 + Lambda Stack ○ Network: 1 Gigabit Ethernet • GPUs: 4x Quadro RTX 6000 + NVLink ○ Processor: Intel Core i9-9940X ○ Operating System Drive: 2 TB SSD (NVMe) ○ Data Drive: 3.84 TB SSD (SATA) ○ Memory: 128 GB ○ Operating System: Ubuntu 18.04 + Lambda Stack • GPUs: 8x Quadro RTX 6000 + NVLink ○ Processor: 2x Intel Xeon Gold 5218 ○ Operating System Drive: 1.92 TB SSD (NVMe) ○ Extra Storage: 3.84 TB SSD (SATA) ○ Memory: 512 GB ○ Operating System: Ubuntu 18.04 + Lambda Stack	• GPUs: 4x Quadro RTX 2080Ti + NVLink (blower version, pair-wise connections) ○ Workstation towers ○ Processor: Intel Core i9-9920X **or** Xeon server processor ○ Operating System Drive: 2TB SSD (NVMe) ○ Data Drive: 3.84 TB or less SSD (SATA) ○ Memory: 128 GB ○ Operating System: Ubuntu 18.04 + CUDA, cuDNN, Keras, PyTorch, TensorFlow ○ Network: 1 Gigabit Ethernet • GPUs: 4x Quadro RTX 6000 + NVLink (pair-wise connections) ○ Server rack mounted ○ Processor: Intel Core i9-9940X **or** Xeon server processor ○ Operating System Drive: 2 TB SSD (NVMe) ○ Data Drive: 3.84 TB or less SSD (SATA) ○ Memory: 128 GB ○ Operating System: Ubuntu 18.04 + CUDA, cuDNN, Keras, PyTorch, TensorFlow • GPUs: 8x Quadro RTX 6000 + NVLink (pair-wise connections) ○ Server rack mounted ○ Processor: 2x Intel Xeon Gold 5218 ○ Operating System Drive: 1.92 TB SSD (NVMe) ○ Extra Storage: 3.84 or less TB SSD (SATA) ○ Memory: 512 GB ○ Operating System: Ubuntu 18.04 + CUDA, cuDNN, Keras, PyTorch, TensorFlow

1. Vendor Question: Will you consider multiple vendors. I am able to source the machines with the first two specs but have been unable to source a machine with the 3rd set of specs.

   UCF Answer: Yes.

2. Vendor Question: Can multiple drives be sourced to meet the 3.84TB SSD Data drive capacity?

   UCF Answer: Yes, can also be less than 3.84TB SSD

3. Vendor Question: Does the 1 Gigabit Ethernet need to be an independent network card, or can this be met with an onboard ethernet controller for system 001, 002 and 003?

   UCF Answer: Either way is fine. The onboard ethernet controller should support 1 Gb Ethernet.

4. Vendor Question: Is a newer generation processor acceptable if it meets performance criteria for system 001, 002 and 003?

   UCF Answer: Yes.

5. Vendor Question: The Quadro RTX 6000 PCIe cards only support NVlink Bridge with a maximum of two GPUs. UCF is requesting a 2X ,4X, and 8X GPU configurations.

   UCF Answer: We would prefer to keep the NVLink connection, we want pair-wise connections for the 4x and 8x machines.

6. Vendor Question: 4X or 8X RTX 6000 GPUs in one system cannot all be connected together with NVLink. Please advise?

   UCF Answer: We would prefer to keep the NVLink connection, we want pair-wise connections for the 4x and 8x machines.

7. Vendor Question: 4X to 8X GPUs with NVLink would require a NVSwitch system and a Tesla V100 SXM GPUs are recommended in lieu.

   UCF Answer: We would prefer to keep the NVLink connection, we want pair-wise connections for the 4x and 8x machines.

8. Vendor Question: Core i CPUs (desktop processors for PCs) would not be recommended with a GPU computing configuration, a proper server with Xeon processor and ECC memory would be appropriate and more stable.

   UCF Answer: Yes. Please upgrade to Xeon processors and ECC memory.

9. Vendor Question: Are all machines expected to be rack-mount based chassis or desktop / Tower chassis?

   UCF Answer: The 2x should be towers and the 4x and 8x rack-mount

10. Vendor Question: Is the end user open to considering other manufacturers of the listed platforms besides Lambda labs? I noticed it called for "Lambda Stack" in the OS description and only Lambda Labs would have this proprietary software for this purpose.

    UCF Answer: Yes

11. Vendor Question: Are the 4 GPU high performance computers workstations as well?

    UCF Answer: No, to be rack mounted.

12. Vendor Question: Does the university intend to purchase all 3 server types or just one server type?

    UCF Answer: All 3 server type.

13. Vendor Question: Normally do not see Intel i9-9920X processor as this is a workstation class type, any objections to replacing the process with a server class processor?

    UCF Answer: Using a sever class processor like Xeon is fine.

14. Vendor Question: How many computers of each listed configuration are you needing?

    UCF Answer: 2x 6qty, 4x 2qty, 8x 1qty

15. Vendor Question: “Lambda” is a company that provides computers and “stack” refers to a software package. What is the software you are expecting be installed? For example, at AMAX we can install a deep learning stack. Please be specific.

    UCF Answer: Needs to come with basics like CUDA, cuDNN, Keras, PyTorch, TensorFlow.

16. Vendor Question: What warranty are you needing?

    UCF Answer: The default warranty that comes with your machines

17. Vendor Question: What type of form factor(s) should the 3 types of GPU systems be? Ex, 2U rackmount, 4U rackmount, workstations, etc.?

    UCF Answer: 2GPU are workstations, 4GPU 2U rackmount and 8GPU can be 4U or 6U rackmount (as budget permits)

18. Vendor Question: If any should be in a workstation, will they be in an office location or a data center?

    UCF Answer: The 2GPU machines will be workstations in a collaborative office environment. The 4GPU and 8GPU will be rack mounted in a data center.

19. Vendor Question: Are there any network requirements on the 4x and 8x GPU systems? Ex 1GbE, 10GbE RJ45 or SFP+, Infiniband?

    UCF Answer: Needs to support 1 gigabit ethernet network connection

20. Vendor Question: NVLink can NOT be used with RTX 6000 GPUs when combining above 2X GPUs in one machine, Hence, under the 4X and 8X GPU configuration, it is advised to use a NVSwitch + V100 GPUs instead for the configuration to work. Need to know if this substitution is acceptable for the RFP?

    UCF Answer: See answer for question number 5.

21. Vendor Question: The RFP is asking for Core I9 Desktop processors with NON-buffered RAM. It is recommended to use a XEON With ECC Registered memory for best PCIE Lane configuration and Speed. Need to know if this substitution is acceptable for the RFP?

    UCF Answer: See answer for question number 8.

22. Vendor Question: Does the RFP have a preference as to having Rack mounted or Desktop Based Chassis? Rack Mounted machines generally have better Air Flow and Redundant Power Supplies.

    UCF Answer: See answer for question number 9.