



Procurement and Contracting Services

Request for Proposals L302303

ADDENDUM #3

**Please mark all proposal submission
Envelopes with the following information**

**Sealed RFP # L302303
Due on July 25th, 2023 no later than 1:00 PM, MST**

The following questions were received prior to the close of the Technical Question period on July 6th, 2023 at 1:00 PM MST:

1. For the AI/ML and Scientific Compute scalable units, are there requirements on the form factor of the H100 and A100 GPUs? e.g., SXM4/5 or PCIe Gen4/5?

Answer: Our power and space constraints are not flexible; vendors may choose whichever form factor they believe provides the most advantageous solution to our request within those constraints.

2. Local Disk requirements - 5.10.2.5 states All compute nodes must include a minimum of 2TiB usable uncompressed local NVMe based disk ($\pm 10\%$). Is the 2TiB requirement on the boot drive, or are local data drives required in addition to the boot drive? If so, what is the capacity requirement on the boot drive vs local data drives?

Answer: This is the only drive, a single ~2TiB NVMe.

3. Network card requirements - 5.10.2.4.2 states Minimum of 2 ports per NIC. Just want to make sure that this applies to both the IB and 25GbE card but not the onboard?

Answer: Single on-board 1Gb/s-compatible ethernet interface is acceptable

4. What is the budget for this project?

Answer: See section 3.9.8.

5. How will you weigh the evaluation criteria?

Answer: Evaluation Criteria will be weighted in descending order as listed in the RFP document section 3.9.8.

6. Is the solution expected to be like for like with what you have today?

Answer: No, that is not expected.

7. Is there a preference toward AMD or Intel?

Answer: No.

8. Given the complexity of the solution being requested, would you consider an extension of the existing submission deadline?

Answer: No.

9. Concerning Section 5.9.6.1, Is the cost of liquid cooling included in the RFP? Will that be an additional cost?

Answer: If your system requires liquid cooling infrastructure, it must be included with your proposal and will be treated as part of the solution cost.

10. Are there future compute or storage growth requirements available that each response should include?

Answer: No.

11. What is the supporting storage and networking needed to complement this solution and are these components open for consideration in this RFP?

Answer: They are outside of the scope of this RFP.

12. 4.51.2.3 Patch Management, what is the timeframe for deploying updates and patch management using an auditable process that can be reviewed by the University?

Answer: boiler plate, not applicable.

13. 4.51.2.3 Patch Management, when should the vendor provide an initial report of patch status to the University?

Answer: boiler plate, not applicable.

14. 4.51.2.4 Scanning and Penetration Testing, what is the timeframe for providing the initial report of the scanning and penetration testing results to the University prior to the Effective Date of the Agreement?

Answer: boiler plate, not applicable.

15. 4.51.2.3 Patch Management. Does this section refer to the delivery of the patch or is the expectation for the vendor to do the patching?

Answer: boiler plate, not applicable.

16. 5.9.2 Power. Does the university have the UPS already in place or is a new 3 Phase UPS also part of what is needed?

Answer: There is a UPS in place already.

17. Is an OPEX solution or Infrastructure as a Service an option?

Answer: No.

18. What is the distance between the EDC and the central research data center?

Answer: We are presuming EDC means "enterprise data center," that is not relevant to this RFP.

19. 5.10.2.3.6 How many H100 GPU's will be required per node?

Answer: minimum 1 H100 GPU.

20. 5.10.2.3.6 What kind of link is preferred – NVLink or PCIe?

Answer: see question 1.

21. 5.10.2.3.7 How many A100 GPU's will be required per node?

Answer: minimum 1 A100 GPU.

22. 5.10.2.3.7 What kind of link is preferred – NVLink or PCIe?

Answer: see question 1.

23. In Section 3.9.5, can you please provide addition insight / clarity on the dollar amount listed and how many nodes are required per each workload? Are you able to share all HW specs / details that you are looking for the standard compute nodes, large memory workload nodes, AI/ML GPU workload nodes, scientific GPU workload?

Answer: See section 3.9.8 for budget projections, see section 5.10 for node specifications. The number of nodes purchased will depend on node pricing.

24. How many Nvidia H100 GPUs are required per AI/ML GPU node?

Answer: See question 19.

25. How many Nvidia A100 GPUs are required per scientific GPU node?

Answer: See question 21.

26. Based on the link you provided for existing Compute System Details, will this be an upgrade to existing Puma HPC cluster and based on AMD Processors? 236 Compute Nodes? 8 AI/ML & Scientific GPU nodes and 2 High Memory Nodes? Please provide as much details as possible.

Answer: This will be a separate and independent cluster. The breakdown of node types will be dependent on budget and pricing of the winning response. See section 3.9.8 for budget projections.

27. Are you able to share an existing high level Visio diagram of the current HPC cluster, connectivity, and all?

Answer: The existing cluster is not relevant to the new cluster configuration.

28. How are your existing file system(s) interacting with the HPC? Do you require scratch or persistent file system(s)?

Answer: Not relevant to the proposal. We do not require any storage bids for any use case, only the internal drive in each node as specified in section 5.10.2.5.

29. As storage is an important part of the overall design – are you open to starting a conversation as part of this RFP (or outside of it) around an end-to-end solution for your storage requirements?

Answer: Storage is outside the scope of this RFP.

30. 5.10.3 Can the BERT benchmark be submitted for 99.00 accuracy?

Answer: No, compute system benchmarks as specified in section 5.10.3 are required.

31. Regarding section 4.51.2.5 Encryption - ("All systems and devices that store, process and/or transmit Confidential Information must use an industry standard encryption protocol for data in transit and at rest") - is the quoted statement standard RFP template language or does it apply to this specific RFP award?

Answer: Boiler plate language, does not apply to this specific RFP.

32. How many GPUs per node are required/desired?

Answer: Minimum 1, maximum whatever you feel is the best price/performance/density solution within our stated environmental parameters as described in section 5.9, see questions 19-22

33. What would be the maximum number of GPUs per node allowed?

Answer: We are not specifying a maximum, noting that all of performance, power usage, heat output and node density are criteria for evaluation.

34. Can the University provide a sample configuration of the proposed networking solution planned for this cluster refresh (Including planned code/firmware versions)?

Answer: No, not relevant to the requirements of this RFP

35. As the RFP indicates that all networking switches will be purchased/acquired separately, is it expected that the winning vendor will support the switches (and cables, see question #7) required for the cluster as part of the 5-year support offering? If not, who will support the switches (and cabling)?

Answer: No, the winning vendor is expected to provide services for the installation and cabling of the fabric equipment in their compute racks. The switches and cabling will be supported by their manufacturer and/or the provider that sells them to us.

36. Will the University also be providing all necessary cabling along with the switching infrastructure, or should RFP responses include InfiniBand and Ethernet cabling?

Answer: The University will be providing the cabling.

37. Would you consider shipping the switches (and cables, if provided) to the HPE HPC integration center for integration and testing? Also, who is responsible for potential switch or cable failures during benchmarking runs?

Answer: We would consider shipping the switching and cables to the winner, but switching should not be required to perform any testing relevant to our RFP response.

Answer: Switching and cabling will not be relevant to any benchmarking runs. Please refer to section 5.10.3, specifically "All benchmarks must be run on a single node of the proposed type."

38. If shipping switches is not feasible, has the University considered how third-party hardware could affect testing/benchmarking? (It could be difficult or impossible to match the University's specific networking hardware choices, leading to differing results).

Answer: Please refer to section 5.10.3, specifically: "All benchmarks must be run on a single node of the proposed type." Our switching choices should have no salient impact on benchmark results.

39. Regarding section 5.9.2 (Power), what voltages can the RDC support – e.g., 208V, 277V, 480V?

Answer: 208V.

40. On Section 5.5 (Method of Payment) – does the University desire to make payments on this award via credit card? Or is this boilerplate language and would not apply to this RFP award?

Answer: This is boilerplate language. Payment will be net terms.

41. Please indicate how many of the following could participate in training: Systems administrators, operations staff, and researchers (graduate level students, staff, and faculty).

Answer: Systems administrators and operational staff: Minimum 2, Maximum 5.

Answer: Researchers and support staff: Minimum 5, Maximum 20.

42. Also, regarding training, does the University have a preference for any particular form for the training – in-person/online standard classes, self-paced online, small group training customized for the Universities particular solution and needs?

Answer: Our preference is for guided online, we will accept self-paced online and can accommodate small groups in-person, if necessary.

43. Over what period of time would the training need to be available after system installation and acceptance (e.g., 6 months? The duration of the 5y support contract term?)

Answer: Training credits/vouchers/program availability may expire no earlier than one year after the completion of acceptance testing.

44. Can we offer an alternative proposal for our cloud HPC services? The benefit is the hardware does not depreciate, it is no longer a Capital Expenditure and moves it to an OPX and you only pay for what you use we have hundreds of compute instance types to accommodate research workloads.

Answer: Refer to question 17, these points are not relevant to our RFP.

45. Would you be so kind as to confirm the desired networking interfaces per node?

Answer: Clarifying section 5.10.2.4, each node must include:

- i. 1x 25Gb/s SFP28 ethernet NIC with at least two ports, supporting RoCE.
- ii. 1x 1Gb/s RJ-45 ethernet NIC with at least one port, supporting PXE.
 1. This may be a NIC supporting 2.5Gb/s, 5Gb/s, or 10Gb/s RJ-45, but must be capable of operating at 1Gb/s.
- iii. 1x 100Gb/s HDR Infiniband HCA with at least one port.
- iv. Out-of-band management may include a dedicated NIC but must also support NC-SI or an equivalent method to provide connectivity via the above-listed 1Gb/s NIC.

46. Is self-insurance an acceptable substitute to the performance/payment bonds?

Answer: At this time, we are unable to accept a substitute for the performance/payout bonds.

47. If you accept self-insurance as a substitute, would this be considered taking an exception to the RFP?

Answer: At this time, we are unable to accept a substitute for the performance/payout bonds.

48. 5.10.2.3.6 - What kind of H100—GPU is preferred – NVLink or PCIe?

Answer: Refer to question 1.

49. 5.10.2.3.7 - What kind of A100—GPU link is preferred – NVLink or PCIe?

Answer: Refer to question 1.

50. 3.98 Evaluation Criteria - Please clarify Cost per Performance for the node configuration. What does performance constitute?

Answer: refer to section 5.10.3 that outlines our required benchmarks, those results constitute "performance."

51. 5.7 Background - How will the new system communicate with the existing clusters?

Answer: Not relevant/out of scope.

52. 5.10 Compute Systems Specifications - Are any additional administration nodes viz. Head nodes, Login nodes, Scheduler nodes etc required for the solution?

Answer: No.

53. 5.10.2.2 - Please clarify on the generation of the processors required for the solution?

Answer: Bids should include the most current generation of processor that respondents can reasonably expect to provide in full quantity by projected delivery dates as provided in Addendum 4.

54. 5.10.2.4 Network Interfaces - Should transceivers be included for the 25Gb/s network ports?

Answer: No.

55. 5.10.5.2 - Can you share the BOM for the switches and cables to be included in installation services? This would help scope and price the network SOW.

Answer: We cannot provide the specific BOM. The winning respondent should expect to rack no more than four 1RU networking switches and cable one port of 25Gb/s ethernet via 4-way copper breakout cable, one port of 100Gb/s HDR IB via 2-way copper breakout cable and one port of 1Gb/s ethernet via standard RJ-45 copper cable per node to one switch of each type located within the same rack. Vendors will not be responsible for any cabling external to the rack.

56. 5.10.3.1 - Can vendor optimized linpack binary be used for HPL benchmarking or is it mandatory to compile from netlib?

Answer: Vendor optimized HPL binaries are allowed.

57. 5.10.2.2 - Utilizing the previous generation processors might allow for higher density of nodes per rack due to low power consumption. Is this acceptable?

Answer: We understand the trade-off, but still require the most current generation of processor that respondents can reasonably expect to provide in full quantity by projected delivery dates as provided in Addendum 4.

58. 5.10.3 - Can similar systems be used for benchmarks and results extrapolated for submission and reproduced during acceptance testing?

Answer: Vendors may generate their benchmark projections in whatever way they deem appropriate. Do note that benchmark projections provided with proposals must be met or exceeded onsite during acceptance testing as specified in section 5.10.4. Penalties as specified in section 5.10.4.3 will be assessed as required.

59. 5.9.2 - Is 68kVA when derated to each system and PSU?

Answer. 68kVA is a mathematical calculation of the power available from 4x 3-phase 60A 208V PDUs. That arrangement of PDUs is the maximum we expect to support in a single rack. You should plan for rack loading that fits within that power envelope based on your own calculations of node power consumption relative to that arrangement.

60. 5.9.2 - Is redundancy required for PSUs of hardware?

Answer: No.

61. 5.10.2.2.3 - The Processor will have primary channel and secondary channels for memory. Do we interpret this as all primary memory channels must be populated with DIMMS no less than 32GB?

Answer: For a processor that features eight (8) memory channels, no less than 8 32GB DIMMs may be used on that processor in a 1 DIMM per channel configuration. More than 8 DIMMs may be used as required.

62. 5.10.2.2.2 - Is using different CPU manufacturer for CPU nodes and GPU nodes allowed?

Answer: It is allowed if your hardware portfolio requires it, but it is preferred that all node types use the same CPU manufacturer.

63. Is there no benchmark required for the large memory node? (We are planning to use the same hardware as the compute nodes for this just with more memory included)

Answer: Correct, there is no benchmark for the high-memory node.

64. AI/ML benchmark MLPerf – this is to be run against the H100 node (section 5.10.2.3.6) correct?

Answer: Correct.

65. GPU benchmark HPCG – this is to be run against the A100 node (section 5.10.2.3.7) correct?

Answer: Correct.

66. Symmetric multithreading – this is to be disabled for LINKPACK on compute node – is it also disabled both for MLPerf and HPCG on the units with GPUs included? The note about multithreading is specifically listed in section 5.10.3.1 however it says, “all benchmarks must be run without symmetric multithreading”. Want to confirm this applies to all three benchmarks requested.

Answer: Symmetric multithreading should be disabled on all node types for all benchmarks.

67. Section 5.9.2.3 says to specify the power plug type. The example given is a single-phase plug (L6-30), but you indicate above that you support 3-phase power. What plugs do you prefer or support for 3-phase PDUs?

Answer: Our existing 3-phase PDUs use IEC 60309-2 3P+E/460P9 plugs. This is the connector we are most familiar with and have existing connectivity for, so it would be most expedient. However, our room power distribution is an overhead modular bus-bar system, so if the 460P9 is not appropriate for your purposes, please include details of the alternative you require so we can assess our ability to support it.

68. Is it acceptable to bid a configuration that is mixed form factor (OCP, EIA) for different node types? (e.g., Compute nodes OCP, AI/ML and Scientific Computing EIA).

Answer: It is acceptable. Vendors choosing to do so should indicate explicitly why the form factor difference is advantageous and whether that vendor would consider providing an alternative unified form factor proposal if requested.

69. At the bidder's conference, it was mentioned that many of the points written in Section 4.51.2 Concepts was boilerplate language and was not applicable. Can you confirm that the entire section is not applicable, and/or list which sections specifically are not applicable?

Answer: section 4.51.2 is intended for hosted services RFPs and is not in whole or in part applicable to responses for this RFP.

70. Is section 4.45 also boiler plate data and not applicable?

Answer: We do not expect section 4.45 to be relevant, but it is nominally applicable in the event that the University is required for any reason to provide any data to respondents. If you have explicit concerns about some specific types of data, we might provide that you would not like to handle under these terms, please provide examples.

71. In section 4.45, the RFP states “if.” Can you say for certain as to whether we will indeed have access to UA Data as stated? If not, and we will not have access to UA Data, there is no concern, on the other hand, if we are we will need to know what PII, or security data is involved with “UA Data”.

Answer: As stated above, we do not expect to provide any relevant data to respondents. If, in the course of the execution of this process we are required to do so, this section would apply. If you have explicit concerns about some specific types of data, we might provide that you would not like to handle under these terms, please provide examples.

72. What is your primary supply water temperature?

Answer: See 5.9.6.

73. What is the maximum CDU pressure drop the facility can support?

Answer: As mentioned in 5.9.6.2, we do not currently have any CDUs. If your proposal requires a CDU, provide specifications and we will determine how best to support the solution, if it can be supported.

74. What is the maximum flow rate the facility can support?

Answer: This is complex to answer. Our facility is fed by the University of Arizona chilled water utility loop, we do not control the pump infrastructure. Our current fittings for racks that directly connect to our in-room manifolds are $\frac{3}{4}$," our room is serviced by 6" supply/return plumbing, a typical total room flow rate is approximately 230 gal/min, we do not have a figure for maximum flow rate.

75. Do you have a raised floor or slab? If raised, how deep is it?

Answer: See section 5.9.1.

76. Do you have residual air-cooling capacity via CRAHs or IRCs, etc.? If so, please specify what you have.

Answer: See section 5.9.6.

77. Would the University consider an alternative GPU other than the A100 for the Scientific GPU workloads?

Answer: For section 5.10.2.3.7 specifying the Scientific GPU constraints, at this time we are only interested in a node configuration with A100 GPUs.

78. Do you have any weight restrictions?

Answer: See section 5.9.1.

79. Is it possible to extend the RFP submission date to 8/25 instead of 7/25? We have not yet received answers to the technical questions asked and 7/25 submission date is too short of a timeline to work up all the services proposals and get the benchmarks run.

Answer: As indicated in question 8, no, we are not considering an extension at this time. Respondents may include supporting documentation with their proposals that indicate factors potentially impacting the services proposals. As specified in section 5.10.3, benchmarks must be run on a single node of each proposed type, not on any collection of nodes as a group. We believe this should make real world benchmarking sufficiently approachable if the respondent deems it necessary. Our process also allows estimation and extrapolation of single-node performance if the respondent chooses. See question 57 for further info.

End of addendum, all else remains the same.