Office of the Vice President
and Chief Financial Officer
Procurement Services
https://procurement.ufl.edu/

971 Elmore Drive
PO Box 115250
Gainesville, FL 32611-5250
(352) 392-1331 Fax 352-392-8837

April 20, 2020

ADDENDUM #1 to the University of Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0) scheduled to be opened on **May 5, 2020 3:00 PM** at the University of Florida, Elmore Hall Conference Room, Radio Road, Gainesville, Florida.

This addendum shall be considered part of the Contract Documents for the above mentioned **ITN20NH-139** as though it had been issued at the same time and incorporated integrally therewith. Where provisions of the following supplementary data differ from those of the original document, this addendum shall govern and take precedence. All other terms, conditions, and regulations will apply.

This addendum consists of:

1. **Responses to technical questions and inquires submitted prior to 5pm, April 14, 2020.**

Sincerely,

Nicola Heredia, Director
Procurement Services

Please acknowledge receipt of Addendum #1 by signing below, and returning this addendum with your proposal. Failure to include addendum with your proposal may result in rejection.

_____ _____
Signature Company Name

_____
Email Address

_____ _____
Company Address City/State/Zip

Responses to questions submitted for UF's ITN Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0)

Technical Specification Clarifications

Pg. 5, 1.2.1.1 - Technical Specifications: Data Center Footprint	Facility Water Flow Rate [L/min]: ~ 3,500 L/min estimated at this time. Facility Water Temperature [C]: ~ 10 C Facility Water Type: Chilled water no glycol Relative Humidity [low / high]: N/A as room control is to a dew point @ less than 10C Ambient Air Temperature [low / high deg C]: Air is supplied to racks from subfloor plenum @ 19 C mixing with exhaust to yield return ~6.7 meters AFF of ~29 C. No other supply air is necessary. No additional ambient control is needed in this environment. Size of primary piping: 8 Inches Insulation planned for primary lines [yes / no]: Yes Size of filtration planned for primary side piping [micron / mesh]: We currently anticipate a separate mesh filtration before each CDU. Max pressure that CIT heat exchanger might see: Exprected max 125 psi Min or max temperature that is required to be returned to the primary water source, if applicable [deg C]: N/A as return temperature to primary loop is not a concern.
Pg. 5, 1.2.1.1 - Technical Specifications: Data Center Footprint	CDU (Cooling Distribution Unit) Please specify what type of remote monitoring that will be required [Modbus / IP/ web server / other]: IP and Siemens Building Automation (Bacnet or Modbus). Please specify if a building level ATS (Automatic Transfer Switch), or a CDU level ATS option is required: CDU Level Height of raised floor, if applicable: 0.91 meters Weight restrictions for the floor: Standard raised access floor guidelines apply – concentrated load limit is 1,500 pounds (~680 kg).

Q1. Is UFL looking to maximize node count, core count, or GFLOPs within the 300kW power limit?

A1. We see aggregate general purpose research computing capability – some of this workload is GFLOP intensive, for other workloads a greater number of cores is more motivated. So we strive for a balance primarily between core count and GFLOPs. Node count is probably less important to us. Our users like their jobs to start quickly, which favors core count.

Q2. Can this refresh "mix and match" nodes? For instance, CLX-AP / traditional SP / GPU nodes and potentially ICX nodes? Can your scheduler recognize and establish queues for the different nodes in the cluster?

A2. The large memory, standard compute, and SLURM controller node type do not all need to match. For example, it is permissible for the large memory nodes to use different processor SKUs/types than

standard compute. All nodes within a class, however, should be the same. Our scheduler can recognize and establish queues for the different nodes in the cluster.

Q3. Does the University of Florida have a "Total CPU core count" target number?

A3. We desire to have proposals with 20,000 cores or more.

Q4. 2.3 GHz is listed as the preferred CPU clock rate for Standard Compute Nodes. In some builds, 2.0 GHz can be more advantageous. What is most valuable, maximizing performance per Watt or cost of the solution?

A4. The ITN states "Base clock rates greater than or equal to 2.3 GHz are preferred". We cannot make a statement regarding which feature would be more valuable without knowing specific figures for both.

Q5. Is there a preference on Core counts - definitely has an overall impact on cost and memory.

A5. The preferences is for at least 20,000 in total.

Q6. Do core switches need to be included in the BMC network or will the BMC edge switches be connected directly to UFL's existing networks?

A6. Core switches do not need to be included in the proposals. If port count requirements from proposals exceed available existing capacity, this would be a topic during the negotiation phase of the ITN.

Q7. Do spine switches need to be included in the Infiniband network, or will the BMC edge switches be connected directly to UFL's existing Infiniband network? If the latter, are UFL's existing Infiniband switches EDR or HDR?

A7. InfiniBand spine switches do not need to be included in the proposals. BMC edge switches will be connected to existing UFL switch ports. Integration with UFL's existing InfiniBand spine switches will be a topic during the negotiation phase of the ITN.

Q8. About this requirement in the ITN: "Switches should have sufficient port count and type to support a 2x100 Gigabit Ethernet LACP LAG for uplink to HiPerGator's existing networks", is this from each individual Ethernet switch? If so, is there an oversubscription ratio that we needs to be considered? Or is it desired to have aggregation switch/switches?

A8. Yes, the 2x100 uplink figure is per switch. No oversubscription figure needs to be considered.

Q9. Should we assume Headnodes communication to/from the nodes are through HiPerGator network?

A9. Yes.

Q10. Should we assume that nodes communication to/from storage are through 2 x 25GE (on nodes) and then via HiperGator's network?

A10. NFS traffic will use the 25GbE interfaces and then via HiPerGator's network, and Lustre traffic will use the InfiniBand interface(s).

Q11. Is it acceptable to centralize a large density of coper cabling for centralizing the management switch(es)?

A11. Yes.

Q12. What size of HDDs does UFL prefer for the 4 local disks in the SLURM controller nodes?

A12. 300GB at a minimum

Q13. Does the solution need unused InfiniBand ports for connecting to storage?

A13. No. Our current IB fabric is used for both message passing and storage access.

Q14. Is there a power limit per rack?

A14. No, however, there are practical limits, constrained by the number of receptacles available for use from the in rack power distribution units. Utilizing 60A Three Phase 208V 0U PDUs (APC 8866 and Eaton EMO328 are representative) a pair would provide for up to ~34 kW of **non-redundant** power to utilize, provided all power phases are perfectly balanced. Since that is seldom the case, consider that those same two units should be able to comfortably support 30 kW **non-redundant**, and clearly more if multiple pairs or alternate configurations are utilized. If your proposal requires more than two rack PDU per rack, please note the size and type planned so that we may vet appropriate power connectivity.

Q15. What is the maximum number of PDU whips UFL can provide per rack?

A15. We will utilize a pair of 600A Starline Track Busways, with whips supplied from power drops that connect to the busways, two whips per drop. The maximum number of drops would be limited by how many of these drops could physically fit onto the busway. However, you may safely presume that we are able to reasonably supply up to 4 input feeds per rack. This limitation is a physical limit of nearly all 0U mounting systems for racks of two PDU on each side of a rack.

Q16. What is the heat load limit per rack?

A16. You are currently planned for a non-enclosed air cooled environment. The planned supply air temperature is 66F or 19C. We currently are capable of delivering in excess of 4K CFM across the face of the rack. This has successfully cooled racks up to 34 kW in our environment. For practical purposes this establishes an upper limit, though we have enough flow to cool up to 38 kW

Q17. Data Center Footprint. — "The peak power not to exceed 300 kW." What is the max kW per rack that the data center can support?

A17. You are currently planned for a non-enclosed air cooled environment. The planned supply air temperature is 66F or 19C. We currently are capable of delivering in excess of 4K CFM across the face of the rack. This has successfully cooled racks up to 34 kW in our environment. For practical purposes this establishes an upper limit, though we have enough flow to cool up to 38 kW.

Q18. What is the kW per rack of current HiPerGator?

A18. Peak power draw per rack is approximately 30kW.

Q19. "The system must fit in 30 linear feet and not exceed 300 kW." Would the university prefer to have denser populated racks (> 20Kw per rack) such that they system would fit in less than 30 linear feet or would the university prefer less dense racks that would take up 30 linear feet?

A19. All things being equal, the University would prefer denser racks. However, if the solution suggested less dense racks for technical reasons or advantages, this is totally acceptable. There is no requirement for the solution to consume all 30 linear feet.

Q20. "If a system makes use of water cooling…," — an air cooled solution could be provided today based on the 30 linear feet and 300 kW limits. What would the university prefer, an air cooled solution that provided more cores or a water cooled solution that would provide more densely populated racks and less space but provide fewer cores?

Responses to questions submitted for UF's ITN Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0)

A20. If the systems were identical in both cases, the University would prefer an air cooled solution.

Q21. The University of Florida has required that the Data Center Footprint peak power of the system not to exceed 300kW. Is the 300kW total power with full system redundancy? (For example: 150kW per side)

A21. The 300kW total power figure is not with full system redundancy. We do not expect such from our compute nodes.

Q22. The University of Florida has defined the system to fit in 30 linear feet of floor space (rack width). Upon reviewing the width of each rack, we have determined that 12 racks would fit into the 30 linear feet of floor space allocated. By dividing the total peak power by the number of racks, 25kW would be supplied to each rack. (This power calculation may change depending on the answers from the previous 300kW questions.) Does the University of Florida have a preferred brand or model of rack to use in your Data Center?

A22: Most of our existing racks are APC brand, some Chatsworth, some "other". The University does not have a brand preference for its racks.

Q23. Are there any depth limitations for the racks to populate the 30 linear feet of floor space?

A23. Rack depth should not exceed 4 feet.

Q24. What is the highest deliverable power to a single cabinet?

A24. There are practical limits, constrained by the number of receptacles available for use from the in rack power distribution units. Utilizing 60A Three Phase 208V 0U PDUs (APC 8866 and Eaton EMO328 are representative) a pair would provide for up to ~34 kW of **non-redundant** power to utilize, provided all power phases are perfectly balanced. Since that is seldom the case, consider that those same two units should be able to comfortably support 30 kW **non-redundant**, and clearly more if multiple pairs or alternate configurations are utilized. If your proposal requires more than two rack PDU per rack, please note the size and type planned so that we may vet appropriate power connectivity.

Q25. What is the operating temperature for the University of Florida's Data Center where this HiPerGator 3.0 expansion will be installed?

A25. Currently planned for a non-enclosed air cooled environment. The planned supply air temperature is 66F or 19C. We currently are capable of delivering in excess of 4K CFM across the face of the rack. Convection carries return air to intakes ~22 Ft above the finished floor.

Q26. Is there a preferred hot and cold aisle? This would help determine if there should be port-side intake on switches or port-side exhaust.

A26. Cold aisle is in front of racks. Switches whose ports face the rear of the rack, should have "power-to-port" airflow.

Q27. What is max KW per rack. What is Voltage for 3-phase PDU –Totals amps per rack?

A27. Utilizing 60A Three Phase 208V 0U PDUs (APC 8866 and Eaton EMO328 are representative) a pair would provide for up to ~34 kW of **non-redundant** power to utilize, provided all power phases are perfectly balanced. Since that is seldom the case, consider that those same two units should be able to comfortably support 30 kW **non-redundant**, and clearly more if multiple pairs or alternate configurations are utilized. If your proposal requires more than two rack PDU per rack, please note the size and type planned so that we may vet appropriate power connectivity.

Responses to questions submitted for UF's ITN Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0)

Q28. Can both 25Gb Ethernet NICs on the compute nodes be cabled to the same switch, or is there a requirement for cabling to redundant switches?

A28. There is no **requirement** for cabling to redundant switches. However, we do prefer to have the option to cable each NIC to separate switches in a redundant fashion (e.g. LACP LAG).

Q29. Can both 25Gb Ethernet NICs on the compute nodes be cabled to the same switch, or is there a requirement for cabling to redundant switches?

A29. There is no **requirement** for cabling to redundant switches. However, we do prefer to have the option to cable each NIC to separate switches in a redundant fashion (e.g. LACP LAG).

Q30. About the Ethernet switches connections to rest of the Hipergator network, what technologies will be required? For example, for loop avoidance?

A30. Top of Rack (ToR) switches should support, at a minimum, 802.3ad link aggregation and Rapid Spanning Tree Protocol (RSTP, 802.1w). Any aggregation (i.e. core) switches included in the proposed solution should include, in addition to the preceding, multi-chassis trunk (MLAG/CLAG) support.

Q31. Are optics/transceivers required for all connections? If so, for connections to the Hipergator network, what distance optics are needed?

A31. Optics/transceivers are not required for all connections. Details regarding integration with HiPerGator's existing networks will be a topic for the negotiation phase of the ITN.

Q32. Is the vendor to provide the RHEL operating system?

A32. No, UFL already has the necessary subscriptions.

Q33. Do you have a preference for Intel vs AMD CPU?

A33. No.

Q34. What is the level of hardware and software support the University is requesting is the university requesting to have local support team assigned to UF to manage support tickets?

A34. Minimum acceptable support level on the hardware is 9-5, next-business day parts replacement. There is no requirement for any kind of premium support levels or 4-hour response times. We prefer advanced hardware replacement – that is, if a system or component fails, the University would report the problem and the vendor would send a replacement without having to wait for us to send the failed component in first. The University is not specifically requesting a dedicated local support team assigned to UF. It is highly preferred for on-site spares to be provided and replenished for the duration of the warranty in order to enable self-service of failed components. For any software provided, which includes switch and system firmware, the University requires access to updates, a support system of some type (e.g. web portal) with email and telephone options during business hours. The duration of the support for both hardware and software must be for 5 years.

Q35. Does the University require a break fix parts locker to be included in the ITN?

A35. A break fix parts locker is not specifically required by the ITN, but can certainly be proposed, and we would prefer it if it is an option.

Responses to questions submitted for UF's ITN Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0)

Q36. Attachment A: "Year 1 License and Maintenance costs" is it the intent of the University of purchase upfront only 1 yr of support with a renewal each year for the balance of 5 years support as noted in Recurring Costs (Years 2 through N)?

A36. An option should be included that would allow UF to pay for all costs initially.

Q37. About this requirement in the ITN: "It is highly preferred for on-site spares to be provided and replenished for the duration of the warranty in order to enable self-service of failed components", is the request to add spare components (switches, compute nodes, etc.) as part of the overall solution?

A37. Yes, at a minimum.

Q38. When does the new system have to be delivered to the University in order for the university to meet any operational requirement for its users?

A38. The system must be operational in Jan 2021.

Q39. Would the University prefer a Consumption as a Service model based on a costs per core? If so what is the expected growth over 5yrs?

A39. No.

Q40. Attachment A: A. Initial Costs (One-time) "Training Costs" can you describe the type of training the University would require?

A40. Training where relevant can be proposed, but there is no formal requirement of any.

Q41. Attachment A: A. Initial Costs (One-time) "Training Costs" can you describe the type of training the University would require?

A41. Training where relevant can be proposed, but there is no formal requirement of any.

Q42. The ITN does not provide a budget amount or minimum number of cores. Can the University provide guidance on these items?

A42. We desire to at least exceed the 20,000 cores we are replacing. There has been no budget set as this point, UF is looking for a solution.

Q43. Is there a overall monetary budget?

A43. There has been no budget set as this point, UF is looking for a solution.

Q44. Is there a preference for Removal and Trade-in of old equipment.

A44. There is no preference for Removal and Trade-in of old equipment. Please include in the proposal if possible.

Q45. Are there any preferred Acceptance Test Program (ATP) requirements.

A45. No, a set of programs to be tested has not yet been decided.

Q46. Is it acceptable/desired to provide multiple responses to cover different potential options?

A46. Yes, proposals including multiple options are completely acceptable. Each solution should be identified in the proposal as such.

Responses to questions submitted for UF's ITN Florida ITN20NH-139 Hardware for Supercomputer (HiPerGator 3.0)

Q47. Due to COVID-19, schools and business across the country have closed, and services have slowed down considerably to help reduce the spread of infection. Furthermore, we are expecting issues with our current courier deliveries. Would you be willing to accept an email submission of our response on the due date as opposed to a hard copy delivered by courier? Additionally, many schools and Universities are closing their doors due to the virus. If we do if we do need to submit a hard copy, will there be someone at the University to receive our submission?

A47. At the current time, UF expects to be open at the time that this solicitation opens. If that changes, an addendum will be distributed to share modified submission requirements/procedures.

Q48. The RFP asks for a manual signature. Due to circumstances, would the University accept a manual, yet digital signature in order to avoid having someone make a trip to our office?

A48. Yes, a digital signature is acceptable.

Q49. Under Tab 3 it asks to "Please complete Attachment B when submitting your responses". Can we also attach specification sheets along with Attachment B?

A49. Yes, please include any information that you feel important.

Q50. Is there any possibility that would could submit an all-electronic ITN response. Given that we are under a Stay in Place order we can't go into our office and probably shouldn't. Our factory is open but not the office area to do binders and print on high speed printers.

A50. At the current time, UF expects to be open at the time that this solicitation opens, if UF is open we will expect to receive hard copies of the proposals. If the nationwide situation changes, an addendum will be distributed to share modified submission requirements/procedures.