MISSISSIPPI STATE UNIVERSITY

Request for Proposals (RFP) 18-51
High Performance Computing System

ISSUE DATE: May 2, 2018

ISSUING AGENCY: Office of Procurement and Contracts
Mississippi State University
610 McArthur Hall
245 Barr Avenue
Mississippi State, MS 39762

Sealed Proposals, subject to the conditions made a part hereof, will be received **May 22, 2018 at 2:00 PM in the MSU Office of Procurement and Contracts, same address above**, for furnishing services and potentially, optional services as described herein.

IMPORTANT NOTE: Indicate firm name, and RFP number on the front of each sealed proposal envelope or package.

All inquiries concerning this RFP should be directed to:

Jennifer Mayfield
Office of Procurement and Contracts (Same address above)
jmayfield@procurement.msstate.edu
662-325-2550

A. Any addendum associated with this RFP will be posted at http://www.procurement.msstate.edu/procurement/bids/index.php located under RFP 18-51.
B. **Note 2**: It is the respondent's responsibility to assure that all addenda have been reviewed and if applicable, signed and returned.

1) UNIVERSITY OVERVIEW

a) Mississippi State University (MSU) is a comprehensive land grant university with extensive research activities associated with high performance computing. MSU has a total enrollment of nearly 22,000 students. The main campus is located adjacent to the community of Starkville in northeast Mississippi.

b) The MSU High Performance Computing Collaboratory (HPC2), located in the Thad Cochran Research, Technology and Economic Development Park adjacent to the university campus, is a coalition of eight affiliated research centers and institutes sharing the common characteristic a of multi-disciplinary team-oriented effort that is strategically involved in the application and advancement of computation science and

engineering through the utilization of high performance computing. Additionally the HPC2 manages the university's Computational Engineering graduate program, offering both Master of Science and Doctor of Philosophy degrees.

c) Additional information about MSU can be found at our website www.msstate.edu.

2) INVITATION TO SUBMIT PROPOSAL ON RFP

Mississippi State University (MSU) seeks a high performance computing cluster based on x86-64 processors, interconnected with a high-bandwidth low-latency network, maximizing performance and minimizing power consumption. The proposed system should be a fully integrated system include all compute nodes, support nodes (management, login, data transfer), networking systems, storage systems, cooling distribution units, and power distribution units. The proposed system design should be one that is scalable to support future growth of the system to a 3x size. The University desires to utilize the proposed solution in support of research, educational, and support activities at the High Performance Computing Collaboratory. The total cost of the fully integrated system, including all items specified in this document, must not exceed **$5,500,000**. In their response, Suppliers must describe the system architecture of their proposed solutions (vendors may submit up to 2 responses to this solicitation).

3) SCOPE OF SERVICES REQUIRED

a) General Requirements

i) All proposed hardware, software (released commercially or as open source), and infrastructure must be supported by the Supplier at the time of system acceptance and throughout the subsequent 48 months of production system use.

ii) The proposed system must be a balanced, commercially-available, production-grade HPC system that contains an appropriate combination of processor, memory, interconnect, disk input/output (I/O), and operating system (OS) capabilities in order to execute complex, tightly-coupled, large-scale, scientific calculations; more specifically, the system must be able to successfully execute a variety of workloads, including jobs which stress all subsystems and which require the simultaneous, tightly-coupled use of the full number of compute nodes within the system.

iii) All equipment must be new; no refurbished or used equipment is allowed.

iv) All nodes and operating systems that are provided must include any required fixes for Spectre/Meltdown vulnerabilities.

v) The entire system, including all compute nodes, support nodes (login, data transfer, system management), networking, and storage must fit into no more than 14 racks

(1) The maximum linear floor space available for contiguous racks (maximum row-length) is 20 feet.

(2) All racks containing compute nodes shall be configured in a contiguous manner.

(3) The storage systems racks will not be located adjacent to the compute node racks, but instead will be elsewhere within the data center. The storage system racks will be within 100 feet of the compute node racks; the exact location of these racks will be determined post-award.

(4) If a separate rack is required for the support nodes, this rack will be located near the storage system racks.
(5) If necessary, a dedicated network rack (e.g., for a director-class switch or tier-2+ aggregation switches) may be located in an adjacent row to the compute nodes.

vi) Once award is complete, the Supplier must provide a project manager for the duration of the project to include:
(1) Final configuration
(2) Site review
(3) System build and off-site integration
(4) Factory test
(5) Installation at MSU facility
(6) Acceptance testing

b) <u>Technical Requirements</u>

i) **Compute Nodes**

The system must contain two distinct (i.e., mutually exclusive) sets of compute nodes: standard and large-memory

(1) Standard Compute Node
(a) Must contain two processors
(i) The processor must be an Intel Xeon Gold 6148 or AMD EPYC 7451
(b) Must contain a minimum of 16GB of memory per channel for all available memory channels, utilizing the highest available memory speed supported by the proposed processor, must be error-correcting code (ECC) memory, and must be a balanced memory configuration.
(c) Must contain a mixed-use internal SSD with a minimum size of 240GB.
(d) Must support the latest stable release of RedHat Enterprise Linux (RHEL) and CentoOS (licenses for RHEL are not required)
(e) Must include connectivity to the High Speed Network (HSN)
(f) Must include connectivity to the Management Network (MN)
(g) Must support IPMI2.0 and/or Redfish

(2) Large-memory Compute Node
(a) The proposal must include at least four (4) large-memory nodes.
(b) Must be the same exact configuration of the proposed Standard Compute Node, but with a minimum of 2x the memory capacity.

ii) **Login Nodes**

(1) The proposal must include a minimum of six (6) login nodes, each of which:
(a) Must contain two processors, using the same processor as the Standard Compute Node
(b) Must contain a minimum of 16GB of memory per channel for all available memory channels, utilizing the highest available memory speed supported by the proposed processor, must be error-correcting code (ECC) memory, and must be a balanced memory configuration
(c) Must contain two (2) internal SSD's for mirrored OS, each of at least 400GB; and at least six (6) mixed-use SSD's, each of at least 800GB, configured in a hardware-based RAID6 set
(d) Must contain redundant power supplies

(e) Same OS as the Standard Compute Node
(f) Must include connectivity to the High Speed Network (HSN)
(g) Must include connectivity to the Management Network (MN)
(h) Must support IPMI2.0 and/or Redfish

iii) **Data Transfer Nodes**

(1) The proposal must include a minimum of four (4) data transfer nodes, each of which:
(a) Must contain two processors, using the same processor as the Standard Compute Node
(b) Must contain a minimum of 8GB of memory per channel for all available memory channels, utilizing the highest available memory speed supported by the proposed processor, must be error-correcting code (ECC) memory, and must be a balanced memory configuration
(c) Must contain two (2) internal SSD's for mirrored OS, each a minimum of 400GB
(d) Must contain redundant power supplies
(e) Same OS as the Standard Compute Node
(f) Must include connectivity to the High Speed Network (HSN)
(g) Must include connectivity to the Management Network (MN)
(h) Must include connectivity to the facility LAN (10GigE/40GigE)
(i) Must support IPMI2.0 and/or Redfish

iv) **System Management Nodes**

(1) The proposal must include a minimum of two (2) system management nodes, each of which:
(a) Must contain two processors, using the same processor as the Standard Compute Node
(b) Must contain a minimum of 8GB of memory per channel for all available memory channels, utilizing the highest available memory speed supported by the proposed processor, must be error-correcting code (ECC) memory, and must be a balanced memory configuration
(c) Must contain two (2) internal SSD's for mirrored OS
(d) Must contain redundant power supplies
(e) Same OS as the Standard Compute Node
(f) Must include connectivity to the High Speed Network (HSN)
(g) Must include connectivity to the Management Network (MN) at 10GigE or better)
(h) Must support IPMI2.0 and/or Redfish

v) **High Speed Network (HSN)**

(1) A high-throughput low-latency interconnect network, with a minimum of 100 Gb/s per node injection rate
(2) Fat-tree topology with a maximum of 2:1 bandwidth over-subscription using no more than two-tiers
(3) Redundant power supplies for any switch that would impact more than 24 nodes
(4) If a director class switch is proposed, this must come with a redundant management system
(5) The fabric manager must run on the System Management Nodes

vi) **Management Network (MN)**
   (1) One GigE connection to all equipment, including but not limited to all nodes, storage system, switches, smart/switched PDU's, RDHx's, etc.
   (2) A core switch will be provided that has a minimum of 10GigE connections to all edge switches and system management nodes
   (3) Minimum of 10GigE connectivity to the facility LAN

vii) **High-Capacity Storage System (/work)**
   (1) Must be simultaneously mounted on all compute nodes, login nodes, and data transfer nodes
   (2) Must be a Lustre-based high-performance parallel file system
      (a) Support for Lustre 2.7 or newer
      (b) Must have redundant metadata servers (MDS) configured for failover
      (c) The MDS nodes should have a minimum of 128 GB of memory to allow for maximized metadata caching
      (d) Object Storage Targets (OSTs) must be configured with RAID 6 or equivalent drive failure protection
   (3) All HDDs shall be self-encrypting drives
   (4) Must have at least 2 Petabytes (PB) of usable capacity with a minimum usable aggregate bandwidth of 25 GB/s as initially deployed, but scalable to 50 GB/s without adding additional controllers
   (5) The storage system must be expandable to at least 4 Petabytes of usable space by adding additional drives, with no other hardware changes
   (6) Must include connectivity to the High Speed Network (HSN)
   (7) Must include connectivity to the Management Network (MN) for management functions

viii) **Home Directory and Application Storage System (/home & /apps)**
   (1) Must be simultaneously mounted on all compute nodes, login nodes, data transfer nodes, and system management nodes
   (2) Must be NFS-based with a total usable capacity of at least 50 TB, to be sub-divided into 2 distinct file systems (/home & /apps), in a hardware-based RAID6 or equivalent drive failure protection configuration
   (3) Must support NFSv3 and NFSv4
   (4) All HDDs shall be self-encrypting drives
   (5) Must contain redundant power supplies
   (6) Must include connectivity to the High Speed Network (HSN)
   (7) Must include connectivity to the Management Network (MN)

ix) **Other**
   (1) The proposal must include one networked KVM with monitor.

c) <u>Software and Applications</u>

i) Cluster deployment, configuration, and software management
   (1) OpenHPC using the latest CentOS release at the time of initial deployment

ii) Programming Environment
   (1) Any software licenses for non-open source HPC tools (compilers, libraries, etc.) use to perform acceptance testing and benchmarking (described elsewhere in this document) must be included for the 48 months of production use of the system.

MSU HPC2 maintains licenses for both the Intel and PGI toolkits, so these are not necessary to be included with the proposal.

iii) Workload Manager/Scheduler
   (1) Moab HPC Suite with Accounting functionality

d) <u>Physical Infrastructure Requirements</u>

i) **Racks**
   (1) All equipment will be rack mounted in standard 42U or 48U racks; all racks with compute and support nodes must be the same size
   (2) Each rack should include a sufficient number of internal PDUs to feed all equipment mounted within the rack. No racks shall have more than a maximum of four (4) whips per rack.
   (3) All power connections shall be bottom fed. Network interconnect cables for adjacent compute racks may be top or bottom fed.

ii) **Rear Door Heat Exchanger (RDHx)**
   (1) The proposal must provide an active RDHx for each rack containing compute nodes.
   (2) The RDHx must neutralize 100% of the heat generated by the rack.
   (3) All cooling and power connections shall be bottom fed.
   (4) The RDHx must support network remote monitoring and control.
   (5) The RDHx must provide leak prevention and detection capabilities.
   (6) All input and output connections shall be bottom fed.

iii) **Cooling Distribution Unit (CDU)**
   (1) The facility will provide a single 6" chilled water supply to support cooling for the compute node racks. The proposal must provide a CDU to be connected to the facility chilled water loop, and in turn to supply chilled water to the RDHx for each rack containing compute nodes.
   (2) The CDU must support neutralization of 100% of the heat generated by all racks containing compute nodes.
   (3) The CDU must be able to automatically adjust to operation using cool water above the ambient dew point in the data center.
   (4) The proposal must provide a distribution manifold and all necessary piping/hoses to connect the CDU to RDHx units.
   (5) All input and output connections shall be bottom fed.

iv) **Power Distribution Unit (PDU)/Remote Power Panel (RPP)**
   (1) The facility will provide a single 480-volt 400-amp power feed to support the operation of the **all** compute nodes, network equipment, CDU, and RDHx units. The proposal must provide a free-standing PDU/RPP to be connected to the facility power feed, and in turn to supply power to each of the compute node racks, networking rack (if required), and CDU. The Supplier must size and configure the PDU/RPP to match the voltage and amperage required for their proposed compute racks, network racks, CDU, and RDHx units.
   (2) The proposal must provide a sufficient number of appropriately sized breakers to support their proposed compute racks, network racks, CDU, and RDHx units.

(3) The PDU/RPP must include remote network monitoring capabilities, including but not limited to input voltage and current, output voltage and current for each branch breaker, and alarms.
(4) All input and output connections shall be bottom fed
(5) The PDU/RPP will be installed by MSU personnel

e) R&D, Operational, and Outreach Partnerships

i) Long term, sustainable partnerships have always been an important part of MSU's strategy for advanced information technology and cyberinfrastructure. The outcome of this RFP will play a strong role in defining MSU's partnerships in cyberinfrastructure in the years to come.
ii) The awardee must help MSU personnel with the tuning and benchmarking efforts of the proposed system
iii) The awardee is encouraged to promote MSU interests associated with the proposed system.
iv) The awardee is encouraged to identify any ongoing partnership opportunities available in the form of collaborative research and development programs, internship programs, hardware grants relevant to R&D, education and outreach, or economic development.

4) SCOPE OF WORK

a) General

i) The awardee is responsible for installation and deployment of the fully integrated system at MSU, with appropriate help from MSU personnel.

b) Maintenance and Support

i) The purchase price of the proposed system and support/infrastructure components is to include 48 months of 8x5 hardware and software support with next business day parts and a comprehensive spare kit (e.g., parts locker) on site.

c) Acceptance Testing and Benchmarking

i) MSU personnel must have access to the system, both physical and remote, during the benchmark and acceptance phases of the installation.
ii) Fabric validation must be conducted
iii) The Supplier must specify the Rmax of the proposed system (compute nodes only). The proposed system must meet or exceed the specified Rmax utilizing the latest version of the HPL benchmark.
iv) Upon award, additional acceptance criteria may be negotiated with the Supplier to confirm that the proposed system meets the specified performance and system operability objectives.

5) INQUIRIES ABOUT RFP

a) Prospective respondents may make written inquiries concerning this request for proposal to obtain clarification of requirements. Responses to these inquiries may be made by addendum to the Request for Proposal (RFP). Please send your inquiries to

Jennifer Mayfield via electronic mail at jmayfield@procurement.msstate.edu

b) All inquiries should be marked **"URGENT INQUIRY. MSU RFP 18-51"**

6) ADDENDUM OR SUPPLEMENT TO RFP

a) In the event it becomes necessary to revise any part of this RFP, an addendum to this RFP will be provided to each respondent who received the original RFP. Respondents shall not rely on any other interpretations, changes or corrections.

7) ADMINISTRATIVE INFORMATION

a) Issuing Office

i) This RFP is issued by the following office:

Office of Procurement and Contracts
Mississippi State University
245 Barr Avenue, 610 McArthur Hall
Mississippi State, MS 39762

b) Schedule of Critical Dates

i) The following dates are for planning purposes only unless otherwise stated in this RFP progress towards their completion is at the sole discretion of the university.

(1) RFP Posted	May 2, 2018
(2) Questions from Vendors Due	May 10, 2018
(3) MSU Q&A Response Due	May 14, 2018
(4) Proposal Submission Deadline – 2:00 p.m.	May 22, 2018
(5) Award Date (Estimated Target)	July 1, 2018

8) PROPOSAL CONTENTS

a) At a minimum, the following items should be included in the contents of the Proposal:

i) Cover letter, indicating the scope of the proposal. The letter should include an overview of the services being offered. The letter should include a statement of

exceptions to any of the terms and conditions outlined in this RFP. (Cover letter should be no more than 3 pages in length.)

ii) Corporate Structure and Credentials

(1) Number of years of experience in HPC
(2) Staffing levels and support proposed
(3) Examples of similar previous work.

iii) Operations and Ability To Perform

(1) Provide operation plan. This should include, but not be limited to, acknowledgement and agreement with all requirements as well as explanations, where applicable, of the intended plan to achieve the requirements.
(2) Timeline of planned events, to include delivery, installation, and the beginning acceptance testing

iv) Cost: should include costs breakdowns for the major components of the proposal (e.g., computing equipment, network systems, storage systems, physical infrastructure, maintenance/support, installation, software, etc.)

9) DISCUSSIONS/EVALUATION CRITERIA/AWARD PROCESS

a) MSU reserves the right to conduct discussions with any or all respondents, or to make an award without such discussions based only on evaluation of the written proposals. MSU reserves the right to contact and interview anyone connected with any past or present projects with which the respondent has been associated. MSU likewise reserves the right to designate a review committee to evaluate the proposals according to the criteria set forth under this section. MSU may make a written determination showing the basis upon which the award was made and such determination shall be included in the procurement file.

b) MSU reserves the right to award this project in whole or in part, depending on what is in the best interest of MSU with MSU being the sole judge thereof.

c) The evaluation factors set forth in this section are described as follows:

i) The Vendor's ability to deliver a solution meeting the overall objective and functions described in the RFP
ii) Competitive cost
iii) Availability and access to technical support
iv) Delivery schedule
v) Vendor's HPC experience in higher education institutions and federal government
vi) Compliance with applicable State and Federal laws and regulations

d) Failure to attend a requested interview presentation before the committee may result in a proposal not being considered.

e) Upon award, the successful respondent may be asked to provide a deployment plan, milestone timelines, confirmation of acceptance testing criteria, and obtain MSU's input and concurrence before moving forward.

f) Proposals will be scored based on the following weights (100 points total):
i) Corporate Structure/Years of Experience in HPC/References – 20 pts
ii) Overall System Design/Technical Approach – 20 pts
iii) System Performance (compute/networking/storage) – 20 pts
iv) Expandability and/or Scalability of Overall System – 20 pts
v) Delivery Time/Time to Production – 10 pts
vi) Added Value Provided to MSU/Partnership Value – 10 pts

10) PROPOSAL SUBMISSION

a) Proposals shall be submitted in one sealed package (envelope or box). **Please make sure that the RFP number is clearly visible on the outside of the package.**

b) Please include at least one (1) hard copy and one (1) electronic copy, preferably on a flash drive, of the complete proposal. If you are submitting more than one hard copy, the original shall be marked on the first page "Original".

c) The proposal package must be received on or before **2:00 p.m. on May 22, 2018**. It is the responsibility of the respondent to ensure that the proposal package arrives in the Procurement and Contracts office on-time. The proposal package should be delivered or sent by mail to:

Office of Procurement and Contracts
Mississippi State University
610 McArthur Hall
245 Barr Avenue
Mississippi State, MS 39762

d) Your response must include the signature page included in this RFP (See Appendix A) and contain the signature of an authorized representative of the respondent's organization. The signature on the "Original" signature page should be in **blue** ink

e) MSU reserves the right to reject any and all proposals and to waive informalities and minor irregularities in proposals received and to accept any portion of a proposal or all items bid if deemed in the best interest of the University to do so.

f) Proposals received after the stated due date and time will be returned unopened. Submission via facsimile or other electronic means will not be accepted.

11) ACCEPTANCE TIME

a) Proposal shall be valid for one-hundred and eighty (180) days following the proposal due date.

12) RFP CANCELLATION

a) This RFP in no manner obligates MSU to the eventual purchase of any services described, implied or which may be proposed until confirmed by a written contract. Progress towards this end is solely at the discretion of MSU and may be terminated without penalty or obligations at any time prior to the signing of a contract. MSU reserves the right to cancel this RFP at any time, for any reason, and to reject any or all proposals or any parts thereof.

13) INDEPENDENT CONTRACTOR CLAUSE

a) The contractor shall acknowledge that an independent contractor relationship is established and that the employees of the contractor are not, nor shall they be deemed employees of MSU and that employees of MSU are not, nor shall they be deemed employees of the contractor.

14) OTHER REQUIREMENTS

a) **Award Terms:** This project shall be awarded at the discretion of the University based on the capabilities and overall reputation of the Supplier, as well as the cost. Acceptance shall be confirmed by the issuance of a purchase order or contract from the University.

b) **The Procurement Process:** The following is a general description of the process by which a firm will be selected to fulfill this Request for Proposal.

i) Request for Proposals (RFP) is issued to prospective Suppliers.

ii) A deadline for written questions is set.

iii) Proposals will be received as set forth in Section 10.

iv) Unsigned proposals will not be considered.

v) All proposals must be received by MSU no later than the date and time specified on the cover sheet of this RFP.

vi) At that date and time the package containing the proposals from each responding firm will be opened publicly and the name of each respondent will be announced.

vii) Proposal evaluation: The University will review each proposal.

viii) At their option, the evaluators may request oral presentations or discussions for the purpose of clarification or to amplify the materials presented in the proposal

ix) Respondents are cautioned that this is a request for proposals, not a request to contract, and the MSU reserves the unqualified right to reject any and all proposals when such rejection is deemed to be in the best interest of the University.

x) The proposals will be evaluated according to the criteria set forth in Section 9f.

APPENDIX A: SIGNATURE PAGE

Provide information requested, affix signature and return this page with your proposal:

NAME OF FIRM: _____

COMPLETE ADDRESS: _____

_____

TELEPHONE NUMBER: _____

AREA CODE/NUMBER

FACSIMILE NUMBER: _____

AREA CODE/NUMBER

E-MAIL ADDRESS: _____

AUTHORIZED

SIGNATURE: _____

PRINTED NAME: _____

TITLE: _____