# RFP Q & A

## R2 - Following are all the round 2 questions received by the deadline and the associated responses.

1. We are concerned that the file size of our response will be too large to email. Would we be able to ship or hand deliver a response on a USB key by the due date and time?

   *Yes. If you feel your response will be too large to email, you may ship or hand deliver your response by the due date on USB key to the following address:*

   *Southern Methodist University*
   *Purchasing Department – c/o Abby Kinney*
   *6116 N. Central Expressway, Suite 205A*
   *Dallas, TX 75206*

   *Note: SMU email system allows for 25MB message size.*

2. Do you <have> goals for doing business with WEBE or HUB zone?

   *SMU has not established a supplier diversity goal for this project.*

3. Cooling Questions from 3.7.1
   a) Will drawings of the room and mechanical and electrical infrastructure be provided? We need this to determine pipe runs, drain lines, electrical sources and where the responsibility starts and stops.

      *MEP drawings of the room will not be provided.*

      *Specifications for electrical busway tap-off box is provided for delivery of power to each cabinet. If this specification is not adequate for the power requirement in the optional/alternative proposal, then please document the requirement and provide alternative tap-off box & PDU specifications.*

      *If chilled water is required for cooling optional/alternative equipment proposed, then please document the requirement and assume SMU will deliver chilled water to the required location.*

   b) Will a site walkthrough be available?

      *Yes, a site walkthrough can be scheduled at time of on-site proposal presentation if requested.*

c) The RFP references air cooled row based cooling is needed, but the last sentence refers to piping for chilled water based in row cabinet is available. Is air cooled DX the preferred method?

*Industry standard air cooled server, storage and network equipment with front-to-back air flow is expected in the primary proposal.*

*Chilled water is only mentioned as an available capability in the Data Center if required by optional/alternative proposal.*

*SMU provides all Data Center mechanical equipment for cooling.*

d) What BTU or size of in row units should the design be based? What level of capacity growth is expected?

*Non-issue. SMU provides all Data Center mechanical equipment for cooling.*

e) Do we need to provide redundant units?

*Non-issue. SMU provides all Data Center mechanical equipment for cooling.*

f) Who is the manufacturer of the hot aisle containment, and may we provide an alternate solution?

No. Hot aisle containment system has already been accounted for so alternative solution is not appropriate.

g) Is a turnkey mechanical/electrical installation required?

*Non-issue. SMU provides all Data Center MEP equipment. Exception is the tap-off boxes and in-cabinet PDU per section 3.9.*

*If the question is regarding installation of tap-off box and connecting PDU to the tap-off box, then no. SMU will handle actual connection of the tap-off box to the electrical busway and connection of the tap-off box's whip to the cabinet PDU.*

*Server scope still requires delivery of the tap-off box and installation of PDU in the cabinet.*

4. Rack Question:
   a) Is the 750mm minimum, or can the rack be wider?

      *No, SMU expectation is to stick to the standard cabinet size in the primary proposal.*

      *Other cabinet dimensions would be considered in the optional/alternative proposal but only if required to accommodate the creative solution.*

   b) Are there specific needs on the cable management (cable fingers, d rings, etc…)?

      *There are no specific needs but open to the addition if helpful for how you intend to install equipment and cleanly manage the cabling in your proposal. Do not include enclosures that allow messy cabling to be hidden inside.*

   c) Are the racks to be placed on the floor or just delivered to the dock?

      *Turn-key installation is expected for cabinets within the scope of server proposal per section 3.2 and attachment 6.*

5. General Question
   Would SMU consider extending the due date by at least 2 business weeks? Considering that answers to the last round of questions will not be released until the 31st, that will only allow vendors a little less than 2 weeks to build a solution and craft a response to a very complex RFP. The additional time would allow vendors to put together the most thorough and comprehensive solution possible.

   *No. The RFP was released on August 2nd and initial proposal is expected to be submitted on September 12. Vendors will have had six weeks to review and prepare a response. SMU was timely in responding to the first round of questions and has met the deadline for this round with no amendments to the original specifications.*

6. Regarding the current 1.4PB of useable Lustra storage that is being moved from the current ManeFrame environment:
   a. 1) What is the type of host interface that is currently being used.

      *Lustre servers use a SAS interface in the current environment, but this should not be considered a limitation for the new cluster.*

      *The only issue that would be relevant to the RFP is the OSS specification for the additional servers noted in section 3.2.1.i.*

b. 2) Is the intention to install this asset in the new HPC environment, transfer the data to the new 1.4PB useable Lustra storage and then de-commission, or continue to use both 1.4PB assets in the new HPC environment?

*The 1.4PB of Lustre disk storage on the current cluster will be moved to the new cluster at some point in the future, but all of this work is outside the scope of this RFP.*

*The only part that would be relevant to the RFP is the additional OSS servers (see 6.a above) and the Lustre design to accommodate the doubling of Lustre storage fairly soon after the new cluster is up and running.*

7. Are there any federal contract vehicles available to SMU? Are any federal funds utilized to purchase this solution?

*No and No. Current funding does not include any government assistance. If you have other outside funding sources, please include specifics in your proposal.*

8. Does implementation include policy considerations other than Bright Cluster Manager? For example, User Facing Policies, Usage Policy, Sensitive Information Policy, Scheduler Policies and SLA's.

*Basic SLURM integration that enables submission of jobs to different type of nodes using Bright Cluster Manager is a requirement. Although tuning of queues or SLURM policies/parameters is not a requirement.*

*The SLURM policies all currently exist and therefore only involve transferring to the new cluster with adjustments for new node counts and some new node types such as NVIDIA GPU and Intel KNL nodes.*

*Any other policy recommendations (or policy templates) that could add value should be noted in the services section of the proposal (Attachment 7).*

9. Does delivery include documentation of HPC resources at the operational level?

*Documentation of HPC resources at the configuration level, that includes details on the decisions or choices made with appropriate rationale at the time of the decision is expected.*

10. If there is a communication portal for end-users as part of the solution, are you interested in this?

*The current cluster does not have a web base end-user portal and specifications for one has not been included in the scope of this RFP.*

*However, any solution that enhances user experience will be considered a plus. For example, NICE software add-on for Bright Cluster Management solution would be a strong consideration.*

*Reference for NICE software:*
https://www.nice-software.com/products/enginframe

11. Are you open to the idea of consolidating or virtualizing some of the management nodes e.g. Nagios, and Management?

    *No. Any consolidation, virtualization or elimination of potential duplicate services are outside the scope of this RFP and will not be considered until after the cluster is fully operational and SMU has had time to fully understand and evaluate new capabilities.*

12. Please define the exact function of the nodes listed as Management.

    *Management nodes are primarily for Nagios services and contingency for any unknowns that may arise.*

13. Bright Cluster Manager performs numerous monitoring functions, given this fact could Nagios be eliminated?

    *No. Any consolidation or elimination of potential duplicate services are outside the scope of this RFP and will not be considered until after the cluster is fully operational and SMU has had time to fully understand and evaluate new capabilities.*

14. In 3.9.3 Power is specified using NEMA receptacle. Can an alternate PDI tap off (IEC Receptacle) box that meets the same electrical specifications be used that would allow for an alternate rack mount PDU connection?

    *No. SMU's expectation is to stick to the standard receptacles specified so existing in-stock spares will be interchangeable.*

15. In Section 3.2.1 the Support Nodes are specified as 1U form factor for every Support node. However consistent with the specification of the node or server function it may be that the physical space in the 1U chassis is insufficient, or not sufficient for necessary cooling and cabling, for the necessary adapters specified or necessary for the function of the node or server. In those cases, may a 2U form factor be proposed for that node or server in the main or primary proposal?

    *Yes, 2u form factor is acceptable if necessary to accommodate the required components.*

## R1 - Following are all the round 1 questions received by the deadline and the associated responses.

1. Paragraph 3.4 Lustre Storage Hardware: Please describe the planned use for the Lustre storage system. Will it be used only for high speed scratch storage or for permanent file system space?
    - Lustre is the high performance parallel file system for the cluster. All the Lustre space is considered scratch space in-so-much-as it is not backed up. All data stored in Lustre space is expected to be recoverable as reproducible output (rerun jobs) or as input data recoverable from other local or internet sources.
      NFS storage is only included for situations where Lustre doesn't (currently) perform well and to hold traditional home directories that are backed up.

2. Paragraph 3.4 Lustre Storage Hardware: The RFP specifies a capacity of 1.4PB and a performance requirement of 30GB/s. Are both of these figures mandatory? E.g., if a proposed solution could achieve the 30GB/s at a lower capacity point and at a lower cost, would that solution be considered compliant? We're interested in fully understanding the rationale for both figures.
    - Both the capacity (1.4PB) and throughput (30GB/s) are minimal requirements and both are considered mandatory. Exceeding either or both of these minima would be considered a plus given competitive pricing.

3. Paragraph 3.4 Lustre Storage Hardware: The RFP specification appears to detail requirements for a hardware-only proposal with implementation and integration to be done by the University. Would SMU entertain proposals for pre-integrated tested Lustre solutions that meet the requirements of the RFP but come pre-assembled and tested with a warranted performance and support mechanism?
    - In general, the RFP is designed to be awarded based on best of breed component solutions. In order to provide pre-assembled and tested solutions the Company would need to win the bid for all components involved.
      Regardless, SMU will not accept the cluster without completing the work described in Attachment 8 to some degree.

4. Paragraph 3.4.3, OSS and MDS Servers. The specification indicates that these are scoped as part of the server portion of the RFP and should not be scoped as part of this section. Would SMU entertain a proposal that consisted of high density storage enclosures with embedded servers that could function as OSS and MDS servers?
    - The RFP allows for creative solutions in a secondary alternative proposal. The primary proposal expects all servers will be from the same manufacturer and included in section 3.2 for pricing, delivery, installation and support.
    - That said, a true storage appliance with embedded controllers (not just servers shoved into the storage unit) would be acceptable in the primary proposal if the proposal can meet all other requirements and demonstrate lower total cost of ownership, flexible growth in both capacity and performance and maintenance/support easily managed by HPC SysAdmin.

5. Paragraph 3.4.5 RAID support. Does the University accept a software RAID solution that includes a declustered RAID6 implementation that provided significantly faster reconstruction and less downtime as being compliant with this paragraph?
    - The RFP allows for creative solutions in a secondary alternative proposal. The primary proposal requires hardware RAID including RAID6.

6. Paragraph 3.4.6 ZFS support. Please elaborate on which features of ZFS are of particular interest to the University. Is ZFS mandatory or would another solution that offered high performance and data integrity be acceptable?
    - SMU is open to considering alternative solutions but needs to understand the pro/con issues from Company's perspective. Please provide pro/con issues for going with a) ZFS and b) alternative option.

7. Paragraph 3.4.8. 12GB SAS interface. This technology is still fairly new and not yet widely deployed. If the requested performance levels or better can be achieved with the more mainstream 6GB SAS enclosures, would that be acceptable to the University.
    - Use of 6GB SAS would be considered acceptable but not preferred. Performance and throughput requirements would not be relaxed.

8. Paragraph 3.4.17 DNE and ZFS. Would a solution that fully supports ~~DNS~~ DNE on a non-ZFS underlying file system be conserved acceptable? (Reference question above on ZFS.)
    - SMU is open to considering alternative solutions but needs to understand the pro/con issues from Company's perspective. Please provide pro/con issues for going with a) ZFS and b) alternative option.

9. Paragraph 3.4.23 MDS Servers. Would a storage enclosure that included competitive embedded controllers to function as MDS servers be considered acceptable?
    - The RFP allows for creative solutions in a secondary alternative proposal. The primary proposal expects all servers will be from the same manufacturer and included in section 3.2 for pricing, delivery, installation and support.
    - That said, a true storage appliance with embedded controllers (not just servers shoved into the storage unit) would be acceptable in the primary proposal if the proposal can meet all other requirements and demonstrate lower total cost of ownership, flexible growth in both capacity and performance and maintenance/support easily managed by HPC SysAdmin.

10. Paragraph 3.4.25 MDT RAID. We normally use RAID10 for MDT data integrity rather than RAID6. Would this be considered acceptable?
    - RAID10 is acceptable for MDT pending component failure analysis and cost delta.

11. Paragraph 3.4.27 12GB SAS. As noted above, 12GB SAS is not widely available as yet. If the performance requirements can be achieved with the more commonly available 6GB SAS enclosures, would that be deemed acceptable?
    - Use of 6GB SAS would be considered acceptable but not preferred. Performance and throughput requirements would not be relaxed.

12. Paragraph 4.4 Long Term Archive Solution. This paragraph requests only 20TB/year of archive data, even though the capacities of storage requested with this RFP seem to be much larger (36TB for the NFS storage appliance per paragraph 3.3, and 1.4PB for the Lustre Storage system per paragraph 3.4). Please clarify the data that will be stored in the Archive solution and confirm the requested capacity.
    - The estimate of 20TB per year with 10-year retention is correct for the Long Term Archive requirement. Lustre capacity and performance metrics are unrelated.
      Data format could include text, programs, images, documents and other forms of binary data produced by different applications. Please note any data format limitations of archive solution proposed.
    - To further clarify, total archive capacity delivered should be 100TB derived from 20TB per year accumulated or 5-year equipment life expectancy.
    - Please note proposed solution's maximum ability to scale beyond capacity of initial deployment.

13. Paragraph 3.4.1 of the RFP states, “Support up to 432 clients” whereas Paragraph 3.5.1 states, “432-port (or larger) EDR managed InfiniBand switch capacity in leaf/spine topology with 2:1 oversubscription.
    >Regarding scalability of the cluster does it need to scale beyond 432 nodes for future growth?
    - The 432 client capacity for network connectivity is based on Mellanox InfiniBand technology constraint as the reference capability at 2:1 over subscription. However, as stated often in the RFP, other options such as Intel Omni Path Architecture would be an acceptable substitution if it can be shown to have equivalent or better capabilities. Hence the 'or larger' reference.
    - It should also be pointed out that maximum client capacity of 432 nodes is considered good enough since current node count expected is in the 330 range leaving reasonable headroom for growth.

14. Can you send me the Master Services Agreement application?
    - SMU has provided standard terms and conditions as Attachment 11. SMU is willing to consider Company's standard master services agreement if submitted in conjunction with the RFP document to be reviewed by SMU.

15. Section 2.1.6 who will have the responsibility to move the 1.5 petabytes from the current ManeFrame to the new HPC system?
    - Data migration from existing HPC cluster (aka ManeFrame) is out of scope. SMU will handle user and data migration from ManeFrame to the new cluster.

16. Section 2.95 does the facility at 5555 N Central Expressway have a raised dock and a doc leveler.
    - The University Data Center does have a raised dock built to handle large truck/trailer rigs. Manual ramps are used to adjust height for easy unloading of equipment.

17. Section 3.9.3 please provide a complete manufacturer part number for the Busway Tap-Off Boxes referenced
    - A51-0339-1-70 – SO 24823

18. In section 1.9, there is mention of NVIDIA GPUs. Are these to be integrated with the servers, or are they stand-alone GPUs that need to be priced?
    - Expectation is for PCIe card integrated with 2u server.
      Note: To help with delivery schedule, server could be delivered and installed with other servers and the NVIDIA P100 PCIe card could be delivered and installed later. Delivery schedule expectation would need to be clearly bounded.

19. On page # 47 and 48, what applications on the list are IO bound today and require ultra high performance. (Example, what applications on the list would benefit from the shortest wall clock run time)?

    - The types of research applications that run on SMU's HPC resources vary widely and only portions of an application or computational workflow may be specifically IO bound. Generally, IO bound applications include some standard commercial and open source packages as well as custom programs developed by research groups for their specific needs. Commonly used commercial and open source packages that frequently require fast IO are quantum chemical codes such as Gaussian, CFOUR, and NWChem and molecular dynamics codes such as CHARMM and NAMD. Researchers developing their own applications that can be IO bound are frequently doing so using binary files, such as netCDF/HDF or application-specific binary formats, while working in languages such Python, C++, or Fortran to take advantage of existing packages, libraries, and workflows.

      Due to the wide breadth of research done on SMU's HPC resources, no specific applications would benefit most from a minimized wall clock run time. That is to say that a reduction of the wall clock run time for all applications is expected.

20. Can you share current wall clock completion times for the most challenging of applications?

    - Due to the nature of computational research that constantly strives to push the boundaries of limited resources, the current class of the most challenging of applications have run-times on the order of several months to not feasible with current resources. A sampling of commonly run applications with moderate run times will be used as benchmarks to verify increased system performance and stability relative to current resources.

21. What applications listed are highly parallelized?

    - Highly parallelized applications include common commercial and open source packages such as CHARMM, NAMD, Qbox, CFOUR, and LAMMPS. SMU has a number of research groups doing work developing tools, libraries, and algorithms for extreme parallelism. These research efforts commonly use standard tools such as OpenMP, MPI, CUDA, and OpenCL. Currently SMU uses: GCC, PGI, and Intel compilers for OpenMP; MPICH, OpenMPI, MVAPICH, and Intel MPI for MPI; PGI and the CUDA toolkit for CUDA and OpenCL. Suites of libraries uses for parallelization of specific functions

include ATLAS, OpenBLAS, Intel Math Kernel Library, FFTW, netCDF/HDF, HYPRE, and PETSc.

22. What is the working data set size for these types of applications?

- Currently, the most IO bound applications have working data set sizes on the order 5 TB.

23. Benchmark Questions **– Attachment 8**

**1.2 Roles and Responsibilities**
SMU will develop and execute the Benchmark jobs. Both internal and industry standard benchmarks. SMU will dedicate two (2) FTE to this phase with expected duration of 3 week
**Confirm:** Want to confirm that none of the responding vendors need to run or provide any benchmarks with our response.

- Correct. SMU will provide and execute the benchmark jobs.

**2.2 Measures of Success**
The measures of success describe how the cluster as a whole and/or each component get a passing grade. This is not designed or expected to be an onerous process but it is expected to be thorough enough to ensure the cluster is ready to rapidly accept and transition the existing ManeFrame data and user community.

General measures of success for running the SMU software noted in Table 2.3.1 is:

i. 2.2.1 Jobs run to completion
ii. 2.2.2 Run time is as good or better than ManeFrame (existing cluster).
iii. 2.2.3 Results are expected and reasonable when compared to ManeFrame results.

**Question:** Can SMU share any current benchmarks from their current ManeFrame cluster and define what specific benchmark values SMU considers “expected and reasonable?

- Not at this time. While SMU has run and will continue to run benchmark jobs on the existing cluster (ManeFrame) in preparation for the acceptance testing process, it doesn't seem a productive use of SMU's time to package them or for each Company to design a system that is optimized for their execution.

  SMU needs a general purpose cluster capable of supporting a diverse research community that will grow and evolve over the life of this new cluster. The primary purpose of the benchmark jobs is to insure hardware is properly installed and all software is properly configured to support existing workloads.

  Any cluster configuration designed to maximize synthetic benchmarks and sacrifice performance for the diverse SMU research community of the real world would be a mistake.

- **2.2.6** Job throughput should be improved on the new cluster, which is a function of not only performance, but also of general system stability demonstrated in consistent and predictable (wall clock) run times.

  Passing grade for cluster performance of SMU benchmark jobs will be a comparison of wall clock time for identical calculations. Further optimization of the job and software is ~~expected~~ to provide additional performance improvement. Throughput can be measured as wall time for a group of calculations. Group calculations will also span various calculation and computational resource requirements and the number of jobs.
  **Question:** Will SMU do this optimization of the job and software or is SMU expecting vendor to do this?
  - The word 'expected' is removed if that's the concern and 'optimization' is a relative term.
    Rewording the issue: Significant performance improvement is 'highly anticipated' as a byproduct of the work to properly configure the cluster hardware and software. Frankly, if the new cluster cannot easily outperform the existing 7+ year old cluster there will be a lot of questions we will all need to answer.

- **3.2 Measures of Success**
  General measures of success for running the Benchmark jobs are:

  3.2.1 Jobs run to completion.
  3.2.2 Run time is reasonably close to the theoretical design of the component or the cluster as a whole.
  3.2.3 Results are expected and reasonable per industry standard output.

  Success will be achieved when actual benchmark results come reasonably close to approximating the actual performance of the theoretical and/or advertised design of the component(s) and the complete cluster to the theoretical design of the full configuration.

  **Question:** Can SMU share any current industry standard benchmarks from their current ManeFrame cluster and define what specific benchmark values SMU considers “expected and reasonable?
  - Benchmark jobs are listed in Attachment 8.
  - While SMU has run and will continue to run benchmark jobs on the existing cluster (ManeFrame) in preparation for the acceptance testing process, it doesn't seem a productive use of SMU's time to package them or for each Company to design a system that is optimized for their execution.

    SMU needs a general purpose cluster capable of supporting a diverse research community that will grow and evolve over the life of this new cluster. The primary purpose of the benchmark jobs is to insure hardware is properly installed and all software is properly configured to support existing workloads.

    That said, when the current cluster (ManeFrame) was originally set up, its Lustre storage system was designed to be a 13GB/s file system. SMU benchmarked this using the end-to-end ost-survey

benchmark tool. As a result, SMU witnessed up to 12GB/s. Readings of these numbers were obtained by running ost-survey on the cluster and observed by running "ltop" tool. SMU considers this a good end-to-end test of the Lustre system that accounts for every component in the data path.

Also, SMU ran HPL benchmarks and obtained 75% of its theoretical FLOPS rating of the system. This observation was obtained with no tuning of the runs or the system. For the new cluster SMU considers 80% of the theoretical FLOPS on the system as a reasonable expectation with little effort to tune the jobs and certainly no effort to bend the configuration to suit a single benchmark.

24. Is an alternative solution acceptable in the primary response if SMU's specifications are not supported exactly as requested?
    - Question is too vague.
25. Most server specifications call for 1U or 2U servers. Are denser solutions acceptable?
    - The primary proposal expects to include the form factor specified. The RFP allows for creative solutions in a secondary alternative proposal.
26. What build and version of Linux are you planning to load on the compute nodes?
    - CentOS, version 7.2 but willing to consider 7.x as a common denominator if better support is available for all components.
27. Can you supply an OS image for load in the factory?
    - Yes, assuming delivery procedures can be worked out.
28. Please confirm that SMU will provide software and licenses for internal applications?
    - Yes, SMU will provide software and licenses for internal applications in Attachment 9.
29. Can you please clarify the instructions in section 2.4 and define "discrete bid?"
    - Each bid for the components listed in section 2.4 must be able to stand alone and be awarded as an isolated (discrete), best of breed, solution. For example, Company A could win the bid for Servers (section 3.2 spec) and Company B could win the bid for NFS Storage (section 3.3 spec) and so on for all ten (10) components listed.
    - See question 30 for additional clarification.
    - Note: SMU doesn't expect it is worthwhile or even appropriate for anyone to provide pricing for the 'Various Software' listed in Section 4.3 (Attachment 9) but it is listed for reference and completeness.

30. Section 2.4 states that we must submit individual bids for each bullet, however section 7.6 states we may provide no more than 2 different proposals. Can you clarify your definition of bid verses proposal?

    - Section 2.4 identifies the breakdown required of the primary bid and secondary bid if offered.

      Attachment 2 provides the summary format to price each major component of the primary bid. The same format would apply to alternative/optional bid if offered.

      In other words, the primary bid can be seen as a collection component bids. SMU reserves the right to award best of breed solutions based on the components identified in section 2.4, further defined in section 3, 4 & 5 with pricing format defined in Attachment 2.

      Attachment 2 also allows the Company to enter a discount that only applies to 'all or nothing' turn-key award.

      See question 29 for additional clarification.

31. Section 3.2, Item K: Must include support for booting over NVMe or M.2 device.
    **Question**: Is booting over NVMe or M.2 a mandatory requirement or is booting over regular SSD device sufficient?
    - For NVMe, the intention is to ensure future options are held open. If only a regular SSD boot device is supported it will be considered acceptable but not preferred. Value proposition should be clearly outlined for this situation.
    - For M.2, the assumption is it would be the most economical solution for solid state bootable disk device with no redundancy. Regular SSD is a suitable option to replace M.2 for boot device assuming price is competitive.

32. Section 3.2.2: Base Compute Nodes
    **a. Question**: What is driving the decision behind the 1u and 2u form factor for servers? Would the university accept an alternate denser form factor that does not meet the standard 1u server requirements for the Base Compute Nodes?
    - The primary proposal is expected to include the form factor specified. The RFP allows for creative solutions in a secondary proposal.

    **b. Question:** Is booting over M.2 a mandatory requirement or is booting over regular SSD device an acceptable alternative?
    - For M.2, the assumption is it would be the most economical solution for fast bootable device with no redundancy. Regular SSD is a suitable option to replace M.2 for boot device assuming price is competitive.

33. Section 3.2.5: **Compute Nodes with Intel Xeon Phi (KNL) Processor**

**Question:** Is booting over M.2 a mandatory requirement or is booting over regular SSD device an acceptable alternative?
- For M.2, the assumption is it would be the most economical solution for fast bootable device with no redundancy. Regular SSD is a suitable option to replace M.2 for boot device assuming price is competitive.

34. Section 3.2.6: Compute nodes with NVIDIA Tesla P100 GPU

**a. Question**: Would the university accepts an alternate denser form factor that does not meet the standard 2u server requirements for the Compute nodes with NVIDIA Tesla P100 GPU?

- The primary proposal is expected to include the form factor specified. The RFP allows for creative solutions in a secondary proposal.

**b. Question:** Is booting over M.2 a mandatory requirement or is booting over regular SSD device an acceptable alternative?

- For M.2, the assumption is it would be the most economical solution for fast bootable device with no redundancy. Regular SSD is a suitable option to replace M.2 for boot device assuming price is competitive.

35. The University is requesting support for technology that is not currently available in the requested timelines outlined and requested in the RFP. Will respondents be excluded from consideration if they plan on offering secondary responses to provide information on future technologies the University may not be considering at this time?

- Not sure what the concern is with this question. The RFP allows for creative solutions in a secondary proposal. The Company is free to offer any technology considered of value to SMU in the secondary proposal.
- There is a delivery schedule requirement. SMU is willing to make accommodations in the schedule in order to procure the three (3) technologies noted: KNL processors, NVIDIA P100 GPU and Mellanox ConnectX-5 (should InfiniBand be the selected solution).
- Other future technologies will be considered but there are limits to how far SMU is willing go to make accommodations in the schedule.