

RFP 10-12
HIGH PERFORMANCE COMPUTER CLUSTER
QUESTIONS AND ANSWERS

INQUIRIES RECEIVED NLT OCTOBER 7, 2011
RESPONSES TO BE SENT NLT OCTOBER 18, 2011

Q1. Section 3.2 - Option One shows a compute cluster, can you tell us more about this?

We are looking for a high performance distributed memory supercomputer capable of running parallel jobs using MPI.

Q2. Do you have a budget for this order? This will determine the speed of the processor we quote.

For our initial purchase we anticipate between \$200K and \$250K for compute nodes and between 0 and \$50K on disk storage. However, this initial purchase is anticipated to be approximately 25% of the final machine. Bids that offer high scalability and a low cost per GFLOP will be ranked highly. Alternatively, this must be balanced against the lifetime of the machine. Bids containing processors that are likely to be discontinued within the next 2-3 years will be rated as less scalable. In order to fine tune the initial order we will use the incremental cost to add (or delete) nodes until our initial budget is exhausted.

Q3. Should the cloud servers be identical to the compute nodes in terms of processor speed and RAM?

They can be, but do not need to be. It depends on what you have to offer that would differentiate the nodes. The cloud nodes should not contain less RAM than the compute nodes, but may contain more.

Q4. What storage/hard drive requirements do you have for the nodes?

The local disks should be sufficient for booting and local scratch space, larger is better, however, desired amount of disk space locally is pricing dependent.

Q5. What are your power requirements (15A, 20A, or 30A; 110V or 220V; single-phase or three phase)?

We can adapt to any of the above. Bidders are not required to conform to any particular power arrangement, but should specify what is required.

Q6. What type of power, power connector, and how much power and cooling are available for each rack, and for this project?

Our power capacity is sufficient to power virtually any configuration. Cooling is by hot/cold row and also is sufficient so as not to be a deciding factor. Power needs, cooling needs and power connector should be detailed, but there is no particular specification in the RFP.

Q7. Would you be willing to sign a non-disclosure agreement to obtain pricing for the next generation AMD processors scheduled to be released later this year?

We have signed a NDA with AMD. If other NDA's are required we would be willing to sign.

Q8. Does the University wish to have an on-site install and training?

We are open to any and all options. Specify the cost associated with installation and training and we will use this information in our decision. For proposals that do not include install and training, we will factor in our estimates for our own time to install and come up to speed.

Q9. What type of warranty is expected (on-site, advance cross-ship, 3-year, longer)?

At least three years with advance over-night cross-ship and 8x5 support.

Q10. Section 2.2 states that IF a separate contract is not written, the RFP, signed proposal, etc, become the contract documents. Since it is unclear whether a separate document will be written, how are changes to the terms and conditions contained within the RFP document to be proposed? For instance, we need to suggest limitation of liability language. We'd also like the ability to terminate for non-payment by the University. How are we to include these terms and conditions?

Section 1.11 of the RFP reads: "Proposal Understanding: By submitting a proposal, the bidder agrees and assures that the specifications are adequate, and the bidder accepts the terms and conditions herein. Any exceptions should be noted in your response." Note your exceptions and recommended changes to the terms and conditions in response to Section 1.11.

Q11. The RFP has two options, Option 1 being the HPC platform, and Option 2 being the storage target. These are two different H/W entities, and potentially two different vendors. Can I submit quotes for both, or can I submit for just 1 part.

You can submit as many quotes as desired for either or both parts. As you say, these can potentially be separate vendors, but they absolutely could also be the same. Additionally, one proposal can address both parts.

Q12. Please characterize the 800MBs + workload throughput requirement. What percentage of the IO is small random read, small random write, large sequential read, and large sequential write?

It is impossible to state with certainty what the future workload will be. Certainly a reasonably large fraction will be small random reads and writes.

- What are the approx Data Set Size(s)?

These vary depending on application and can range from small megabytes to large terabytes. The largest fraction is tens of gigabytes, but this will inevitably grow with increased capability.

- Would you consider an alternative to NFS, such as a special purpose HPC file system based solution?

We are open to any options that can meet our goals. Whatever is chosen will need to work with both clustered and cloud systems.

Q13. The goal is a system having 512 cores, with half or more suitable for cloud usage and half or more suitable for HPC. Interprocess bandwidth should be no less than the equivalent of one full speed connection for each 32 compute cores. Each compute node 10 should have no less than 4GB of RAM per core. Can you elaborate exactly how many cores you are looking for right from the get go on this RFP?

The RFP states 512.

Q14. The compute capability is specified in the RFP as 512 cores. Is there an associated GFLOPS expectation or any specific benchmark criteria?

We want the highest performance we can afford, subject to our other criteria.

Q15. Is interconnect latency or bandwidth more important?

While both are important, latency is more important for HPC.

Q16. We can offer either Platform or Bright as a commercial cluster management system. Is there a preference? Platform defaults to LSF job scheduler and Bright defaults to SLURM job scheduler.

There is no preference. We use Torque and Moab right now. Some sort of provision for rolling reservations and job preemption needs to be available if a cluster management system is offered.

Q17. Is there any specific cloud management system required? If not, any specific cloud management requirements, e.g. virtualization support or bare metal OS provisioning, etc.?

There is no specific cloud management system required. The cost and benefits of any proposed cloud management system will be evaluated against our estimates of our costs to install and maintain such a system.

Q18. Should the response include PDUs for the racks?

Whether PDU's are supplied or not will be weighed into our decision. In general, it depends on the architecture that you present.

Q19. What are the power/cooling constraints per rack? i.e. are there specific limits to how much power can be provided to a rack and are there limits to how much cooling capacity is available per rack? Number and type of power feeds per rack?

See Q6

Q20. Please provide information on other physical constraints for the system such as where inter-rack cables will be routed (above or below the rack); height of ceiling above the floor; limitations for rack configuration; data center floor weight limits; door clearance to data center as well as ramp and elevator access as necessary; etc

There is loading dock access. Ability to fit through a 32X83 inch door is preferred. Cable routing should be above cabinet. Ceiling height is 9 feet or higher. Weight limits are not a concern.

Q21. How will the cluster be connected to the campus network?

The vendor should provide a system that provides adequate firewall security as part of their proposal. Currently we do this by having redundant firewalls, one running on the interface machine and one as a custom built appliance. However, the vendor is not required to follow this practice. The vendor is should provide pricing for the interface machine(s).

Q22. The RFP states: "NAS system should be capable of serving NFS data over 10 GbE and QDR InfiniBand (40 Gb/sec) using RDMA". Is it necessary for the NAS system to serve data over 10GbE and Infiniband simultaneously? Or just require 10GbE and IB interfaces?

It will be looked at favorably for the system to be able to serve data over both at the same time. If this is not possible please specify what is possible.

Q23. How will performance of the solution be assessed? Will it be necessary to reproduce performance results for acceptance?

Performance evaluation and acceptance criteria will be determined during contract negotiation.

Cloud Services Software:

Q24. Can the University describe the use case(s) for the private cloud services described in this RFP including:

a. Anticipated cloud models – IAAS, PAAS, SAAS (Which applications? Is self service required?)

A solution that does not preclude any of the above is preferred. Current cloud services emphasizes IAAS. PAAS is anticipated in the future. Self service is preferred.

b. Is SLA and chargeback capability desired/required as part of these cloud services?

They are not required, desirability depends on pricing

c. Is policy based automation required for application lifecycle management, e.g. the ability to define criteria for the reclaiming of VM's?

It is not required, desirability depends on pricing

d. Is notification and alarms based on triggers required?

It is not required, desirability depends on pricing

- e. Is workflow approval process required for submission of job and/or adding additional resources to a job?

It is not required, desirability depends on pricing

- f. Is it anticipated that this cloud will be deployed to more than one datacenter?

Not currently

- g. Is network isolation required for different groups/lines of business?

It is not required, desirability depends on pricing

- h. Would you like to cap your line of businesses with quotas for different resources?

It is not required, desirability depends on pricing

- i. Would you like to provide clouds within clouds for your customers so that they feel that they have their own resources to manage?

It is not required, desirability depends on pricing

- j. The RFP describes the percentage use of the physical resources as “half or more suitable for cloud usages and half or more suitable for HPC” and “note that dual use cores should be counted as both HPC and Cloud”. Is this to say that some percentage of the nodes should have the ability to be re-provisioned as virtual infrastructure supporting the cloud or as a dedicated compute node for HPC? Or is the intent to provide some cloud resources to the HPC as needed (aka cloud bursting)?

Our expectation is that there will always be some base cloud load and some base HPC load. However, at any one time, the best ratio of HPC to cloud usage is not likely to be 50%-50% so the ability to move some nodes from one use to the other through some means is desirable.

- Q25. What does the University expect to be the most widely deployed virtual machine in this private cloud environment? (VMware, KVM, etc)

The most widely deployed virtual machine will be rather dependent on what system and supporting software is chosen. We do not have an *a priori* preference.

- Q26. What is the current infrastructure? Would this be integrated with existing systems, e.g. ticketing, DNS, asset management, LDAP/AD?

This is a new system; there is no requirement for backward compatibility with existing infrastructure. The University currently uses LDAP, Shibboleth, and CAS. A solution that interfaces with one or more of these is desirable but not required.

HPC Software:

Q27. Can the University describe the use case(s) for the HPC described in this RFP including:

a. Anticipated applications

We anticipate a wide range of applications including custom written programs.

b. Is there a desire to provide a web portal front end to the HPC to help drive broader adoption and minimize job submission errors?

It is not required, desirability depends on pricing

Q28. What are the power/cooling constraints per rack, i.e. are there specific limits to how much power can be provided to a rack and are there limits to how much cooling capacity is available per rack? Number and type of power feeds per rack?

See Q19

Q29. Please provide information on other physical constraints for the system such as where inter-rack cables will be routed (above or below the rack); height of ceiling above the floor; limitations for rack configuration; data center floor weight limits; door clearance to data center as well as ramp and elevator access as necessary; etc.

See Q20

Q30. How will the cluster be connected to the campus network?

See Q21

Q31. The RFP states: "NAS system should be capable of serving NFS data over 10 GbE and QDR InfiniBand (40 Gb/sec) using RDMA". Is it necessary for the NAS system to serve data over 10GbE and Infiniband simultaneously? Or just require 10GbE and IB interfaces?

See Q22

Q32. How will performance of the solution be assessed? Will it be necessary to reproduce performance results for acceptance?

See Q23

Q33. The goal is a system having 512 cores, with half or more suitable for cloud usage and half or more suitable for HPC. Interprocess bandwidth should be no less than the equivalent of one full speed connection for each 32 compute cores. Each compute node 10 should have no less than 4GB of RAM per core. Can you elaborate exactly how many cores you are looking for right from the get go on this RFP?

See Q13