WHITE PAPER COMPENDIUM:

Best Practices for Securing Grant Funding and Publishing an RFP for Your Acquisition of HPC Hardware Equipment

A Step-by-Step Guide

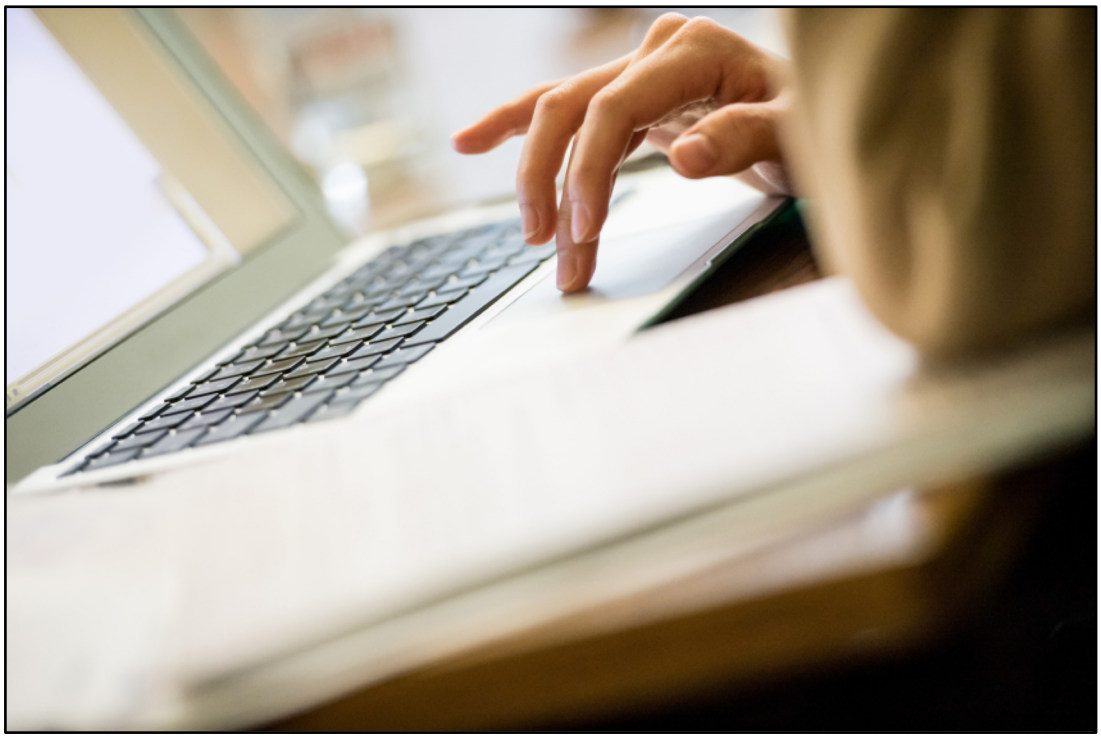

advanced clustering technologies, inc.

WHITE PAPER

Best Practices for Writing a Successful Grant Proposal for HPC Cluster Funding

Prepared by

advanced clustering technologies, inc.

SIXTH EDITION, April 2022

Introduction

Why should I trust advice about grant proposal writing that is being provided by a vendor who builds the very same high performance computing solutions that I am proposing to buy?

It's a fair question. And the answer is quite simple: Experience. For more than 20 years, Advanced Clustering Technologies has been providing quotes and guidance to customers who are embarking on the same grant proposal journey on which you now find yourself.

We have seen it all: the good, the bad and the ugly among proposals; winning proposals and losing proposals; grant writing successes and failures; institutions that get critical equipment funding and those that do not.

From that first-hand experience, we have consulted with our customers to learn what works and what doesn't when it comes to grant proposal writing so that we can advise future grant seekers.

This white paper is a result of that years-long effort, and we provide it to you now in the hope that through this white paper, you will find a blueprint for the kind of grant proposal that secures your needed funding.

The principles outlined in this white paper were designed specifically for two grant programs offered by the National Science Foundation: the Major Research Instrumentation (MRI) program and the Campus Cyberinfrastructure (CC) program, but the principles apply to best practices for most of the grant programs that exist for the purpose of funding HPC projects.

Grants for HPC projects remain competitive and difficult to obtain. In 2011, it was reported that 30% of MRI grant proposals were funded. In Nov. 2020, budget limitations saw that number dwindle down to 20%.

This white paper is intended to give you a comprehensive list of steps you can take to give your grant proposal the greatest chance for success.

We wish you all success in your pursuit of grant funding, and we are here to help any way that we can.

- Kyle Sheumaker
President and CTO
Advanced Clustering Technologies

Getting Started

If you are considering the possibility of applying for a grant from the NSF, it is a good idea to take a look at recent projects that received funding. You can do a search at **https://nsf.gov/awardsearch/** and look for keywords that would be part of your proposal. That way you might find successful proposals similar to what you would write.

Here, as an example, are some of the high performance computing projects that were funded recently by the NSF MRI grant program:

- Acquisition of a Massive Database to Accelerate Data Science Discovery
- Development of a Next-Generation Modular Ion-Trap Quantum Computer
- Acquisition of High Performance Computing System for Interdisciplinary Research and Teaching
- Acquisition of an HPC Cluster for Science and Engineering Research
- Acquisition of a Hybrid HPC Cluster for Computational Modeling
- Acquisition of Cutting-Edge GPU and MPI Nodes

We also recommend that you check within your own institution to find out if you can obtain grant templates or samples of previous grant applications that have been submitted.

These can be helpful to you but please be aware that MRI grant proposals are much more complex than the standard grant application. Often a template will not be a good starting point for your MRI proposal. You will need to follow the MRI solicitation guidelines to the letter.

Taking these steps will get you started down the right path, but before you get too involved in the writing of a proposal, you first need to take the steps outlined on the next few pages of this white paper, beginning with a determination as to whether or not your project meets the eligibility requirements of the grant program you are planning to undertake.

Learn about the Grant Programs

The NSF offers two grant programs that offer funding for high performance computing equipment. They are:

Campus Cyberinfrastructure (CC) program

Supports two program areas for up to 2 years:

- Data Storage awards supported up to $500,000
- Regional Computing awards supported up to $1,000,000
  Read the full solicitation.
- Review examples of Campus CI plans provided by CC* program awardees here to see examples of winning proposals.

Major Research Instrumentation (MRI) program

This program has two tracks:

- Track 1: request funds greater than or equal to $100,0001 and less than $1,000,000.
- Track 2: request funds greater than or equal to $1,000,000 up to and including $4,000,000
  Make sure you read the full solicitation.
- You might find it helpful to review the NSF site here to see MRI proposals that received funding in the past.

You should read the solicitations for each of these programs to gain a full understanding of the programs, including eligibility requirements for each.

You will find the links for each program's solicitation on this page.

Your greatest chance for success comes in complying with all of the requirements of the program's guidelines for your proposal.

Checklist for Your MRI Proposal

1. Cover sheet:

__ Title that is concise and conveys purpose of the proposal

__ Identify a single PI and up to four co-PIs

2. Project summary (no more than one page in length):

__ Overview

__ Statement on intellectual merit of proposed activity

__ Statement on broader impacts of proposed activity

3. Project description (maximum length is 15 pages with charts):

3a. __ Instrument location and type

__ Justification for submission as Development proposal (only required for development proposals) address these:

- __ How will the end result of the effort be a stable shared-use research instrument, rather than technology development, a device, a product or a technique/protocol?
- __ What significant new capabilities, not available in an instrument provided by a vendor, will the new instrument provide?
- __ Does the instrument development effort build capacity for instrument development activities within an MRI submission-eligible organization(s)? In what way does the instrument development require design and development work that must be undertaken or has been undertaken in- house, rather than through readily available/published designs found in the literature?
- __ To what extent does the instrument development require/benefit from a team of scientists. Engineers, and technicians that bring a variety of skills to the project?

Checklist for Your MRI Proposal

__ For what activities does the instrument development require a significant number of person-hours, more so than simple "assembly" of purchased parts? To what extent does the instrument development require timeframes for completion that are longer than are required for plug-and-play or assembled instruments?

__ Does the instrument development require the use of a machine shop or a testbed to fabricate/test unique components?

__ Does the instrument development effort involve risks in achieving the required specifications, and what is the risk mitigation plan?

3b. __ **Research activities to be enabled.**

The degree to which the planned uses of the proposed instrumentation constitute exciting, ground-breaking and/or transformative research is a significant factor in the merit review evaluation of MRI proposals. In this section describe the specific research program(s) and research training activities that will be enabled and that drive the request for the desired equipment.

3c. __ **Description of the Research Instrument and Needs**

(Suggested length: up to 2 pages for instrument acquisition; up to 5 pages for instrument development).

3d. __ **Broader impacts** (including impact on research and training infrastructure)

3e. __ **Management Plan** (suggested length: up to 2 pages for instrument acquisition; up to 5 pages for instrument development)

Checklist for Your MRI Proposal

4. __ **References Cited** (following standard PAPPG guidelines)

5. __ **Biographical Sketches** (biographical sketches of the PI and any Co-PI(s) as well as any designated Senior Personnel)

6. __ **Budget and Budget Justification** (Total project cost should be clearly stated in the budget justification, which must not exceed five pages, and must be itemized in table form using the following template.)

__ **Budgets for Acquisition Proposals** (at least 70% of the Total Project Cost must consist of items that can be included on the Equipment line of the NSF budget form)

7. __ **Current and Pending Support for PIs**

8. __ **Facilities, Equipment, and other Resources**

9. __ **Special Information and Supplementary Documents**

9a. __ For each organization receiving funds, provide on institutional letterhead from each sponsored projects office, the following statement classifying the organization(s) as either non-Ph.D.-granting, Ph.D.-granting, or non-degree-granting (as defined in Section IV). Statements must follow only the format indicated at the top of the next page:

Checklist for Your MRI Proposal

To: NSF MRI Coordinator
By signing below I certify that _____(organization)_____ is classified as _____(select one: non-Ph.D.-granting /Ph.D.-granting /non-degree- granting)_____ as defined in Section IV of the MRI solicitation.
Signed: _____ Print Name:_____
Title of Signatory:_____
_____
Date: _____

10. __ Data Management plan (no more than two pages)

11. __ Postdoctoral mentoring plan (when applicable)

12. __ Single Copy Documents
Required:
a. Collaborators & Other Affiliations (COA) Information.
Encouraged:
b. List of suggested reviewers

Additionally, each proposal must contain this statement(s). No other letter(s)/statement(s) classifying or describing the institution type(s) will be permitted:
b. __ For all proposals: Include a letter documenting the performing institution's commitment to ensuring successful operations and maintenance over the expected lifetime of the instrument. This letter (two-page maximum) must also list the MRI awards made to the organization with a start date in the previous five calendar years

and briefly describe the status of the instrumentation obtained from each award.

c. __ When applicable: A letter (one-page maximum) documenting the organization's commitment for cost sharing, if applicable, must be included.

d. __ When applicable: Proposals that include subawards (except for development proposals with subawards to institutions that do not exceed 20% of the total amount requested from NSF) must include a statement from each subawardee's sponsored projects office acknowledging that this proposal is included as part of the subawardee institution's submission limit. Otherwise, an organization may exceed its submission limit, with the result that the proposal including the subaward may be returned without review.

e. __ When applicable: If a proposed effort involves a private sector partner or other organization serving as a partner (as opposed to an individual(s)), or a large formalized collaboration (e.g., through a memorandum of understanding or other legal document), a letter (one page maximum) confirming their 12 participation must be included. In particular, proposals involving large formalized collaborations are encouraged to have the collaboration utilize this letter to document the role, importance and priority of the requested instrument in the overall efforts being undertaken by the collaboration.

f. __ When applicable: If the proposal involves organizations other than the performing organization, list all partners.

g. __ When applicable: Proposals for the acquisition or development of

an instrument to be located at an organization other than the performing organization must provide a (one-page maximum) supplementary document stating the host organization's commitment to house the instrument. For the purposes of this solicitation, use of instruments at NSF's Antarctic facilities is considered to be field deployment and a supplementary document from the host facility is not required.

h. __ For all proposals. Inclusion of representative, itemized vendor quotes is required for all MRI proposals. Although a proposal might reference and have a quote(s) for a specific make and model, the proposer is reminded that his/her organization's approved procurement processes must be utilized in the event of an award to establish the appropriate item(s) to be purchased and that applicable procurement standards for institutions of higher education and other non-profit organizations are described in 2 CFR 215.40-48.

You are also encouraged, but not required, to include:
a. Statements from individuals, on institutional letterhead, confirming substantive collaboration efforts and/or usage of the instrument may be submitted, but they must follow only the format indicated below:

By signing below I acknowledge that I am listed as a collaborator and/or major user of the instrument on this MRI proposal, entitled "_____(proposal title)_____," with _____(PI name)_____ as the Principal Investigator. I agree to undertake the tasks assigned to me, as described in the proposal, and I commit to provide or make available the resources therein designated to me.

Checklist for Your MRI Proposal

Signed: _____

Print Name:_____

Date: _____

Institution:_____

The proposal body itself should describe the nature and need for a collaboration, and/or describe the major users and their need for the instrument. Statements of collaboration beyond that specified above, including letters of support/endorsement, are not allowed. Each statement must be signed by the designated collaborator or user. PI requests to collaborators for these statements should be made well in advance of the proposal submission deadline since, if they are to be included, they must be included at the time of the proposal submission.

This checklist does not contain all instructions but exists to ensure you include all required statements and information. Please refer to the full solicitation for complete details for each section.

Checklist for Your CC Proposal

For Program Area One: Data Storage

If you are applying for the Data Storage program, your proposal is required to:

- __ address scientific and engineering projects and their research and education storage needs, describing project-specific scenarios for scientific data generation, storage, and management.
- __ address features, capabilities, and software platforms representing the proposed storage resources and services.
- __ address plans to manage the resource, data sets, and usage while ensuring adherence to FAIR principles and equitable access.
- __ include a summary table of the science drivers and their data storage environments.
- __ describe the platform architecture and open-source software/platform.
- __ describe an open source-based approach to storage system monitoring, measurement, management, and instrumentation.
- __ include a sustainability plan addressing the institution's commitment to providing an ongoing level of sustained access to storage resources.
- __ describe how the data storage system is connected via high performance network connectivity, including a network topology showing how the system is connected to the campus network.
- __ include complete itemized vendor quotes with the budget.
- __ describe the storage system as a shared resource intra-campus and inter-campus via interoperability with a national and federated data sharing fabric, with 20% or more of the disk/storage space committed to extramural scientific uses.

Checklist for Your CC Proposal

- __ deploy the system "on premise" which may include placement at an off-campus data center with a pre-existing role in campus research computing and storage.
- __ document a data management approach and plan, included as a supplementary document of no more than 3 pages.
- __ include a Campus CI plan as a Supplementary Document, limited to no more than 5 pages (see Section II. Program-wide Criteria above for more information).
- __ have titles that begin with "CC* Data Storage:" followed by the title of the project.
- __ be submitted by an Institution of Higher Education.

This checklist does not contain all instructions but exists to ensure you include all required statements and information. Please refer to the full solicitation for complete details for each section.

Checklist for Your CC Proposal

For Program Area Two: Regional Computing

If you are applying for the Regional Computing program, your proposal is required to address:

- __ scientific and engineering projects and their research and education computing needs, describing project-specific scenarios for scientific computing tied to the proposed computing resources;
- __ features, capabilities, and software platforms representing the proposed computing resources; and
- __ scientific computing codes expected to run on the resources.

Your project description must include:

- __ A summary table of the science drivers and their computing environments—these requirements should be specified in clear terms reflecting a specific understanding of the required computing resources and environment, for example, CPU/GPU type, compute job profile parameter ranges, core count ranges per job, times to completion or as part of a composition or scientific workflow profile;
- __ The platform architecture specifying cluster components, including compute node type and count, per-node memory, interconnect fabric, storage, and open-source software/platform;
- __ An open source-based approach to cluster monitoring, measurement, management, and instrumentation;
- __ A sustainability plan addressing the institution's commitment to providing an ongoing level of sustained access to computational resources; A High-Performance Network Connectivity and Specification—see below for more details; and
- __ A description of the cluster as an Inter-campus - see below for more details.

Checklist for Your CC Proposal

Your proposal for regional computing is required to:

- __ address scientific and engineering projects and their research and education computing needs, describing project-specific scenarios for scientific computing tied to the proposed computing resources.
- __ address features, capabilities, and software platforms representing the proposed computing resources.
- __ address scientific computing codes expected to run on the resources.
- __ include a summary table of the science drivers and their computing environments. These requirements should be specified in clear terms reflecting a specific understanding of the required computing resources and environment (for example, CPU/GPU type, compute job profile parameter ranges, core count ranges per job, times to completion), or as part of a composition or scientific workflow profile.
- __ describe the platform architecture specifying cluster components, including compute node type and count, per-node memory, interconnect fabric, storage, and open-source software/platform.
- __ describe an open source-based approach to cluster monitoring, measurement, management, and instrumentation.
- __ describe a sustainability plan addressing the institution's commitment to providing an ongoing level of sustained access to computational resources. describe how the cluster is connected via high-performance network connectivity, including a network topology showing how the cluster is connected to the campus network.
- __ provide a description of the cluster as a shared resource intra-campus and inter-campus, with 20% or more of the cycles

Checklist for Your CC Proposal

committed to extramural scientific uses.

- __ include complete itemized vendor quotes with the budget.
- __ include a Campus CI plan as a Supplementary Document, limited to no more than 5 pages (see Section II. Program-wide Criteria above for more information).
- __ be submitted by Institutions of Higher Education or Non--profit, Non-academic Organizations.
- __ have titles that begin with "CC* Regional Computing:" followed by the title of the project.

This checklist does not contain all instructions but exists to ensure you include all required statements and information. Please refer to the full solicitation for complete details for each section.

One Step at a Time

Understanding the grant funding programs that are available is the first step, but there are many others that must be undertaken before your proposal will be ready for submission. On the following pages, we offer advice for each step in the process as gleaned from our customers who have succeeded in winning grant funding in the past.

Determine Your Eligibility

Now that you have familiarized yourself with the two main vehicles for NSF grant funding and the specific requirements of each, it's time you determine your eligibility to apply. There are three types of MRI-eligible organizations:

1) Institutions of Higher Education

- - Ph.D.-granting institutions (awarded more than 20 Ph.D or D.Sc. degrees in combined two previous academic years)
- - Non-Ph.D.-granting institutions (awarded fewer than 20 or fewer Ph.D./D.Sc. degrees during the combined previous two academic years)
- - Non-degree-granting organizations (do not award Associate's, Bachelor's, Master's degrees or Ph.Ds. or Ds.Sc. or institutions that award all degrees outside of NSF-supported fields)

2) Not-for-Profit, Non-Degree-Granting Domestic U.S. Organizations

- 501(c)(3) tax status
- Must have an independent, permanent administrative organization (e.g., an office of sponsored research) in the United States
- May include science centers, museums, research labs, observatories or similar organizations

3) Legally Incorporated, Not-for-Profit Consortia

- Includes two or more eligible organizations
- 501(c)(3) tax status
- Must have an independent administrative structure (e.g., an office of sponsored research) in the United States

Make sure you are reading the solicitation for the most current year. Rules and guidelines may change from year to year.

Make Use of NSF Resources

You are able to submit proposals for review in advance of the final submission deadline, and our customers have found this to be a very helpful part of the process. Rather than submitting a proposal and later learning that key points are missing, it is much better to get that information in time to act on it. We highly recommend that you make use of these opportunities to get input from NSF representatives early on in the grant writing process.

Understand Cost Sharing Requirements

Cost-sharing requirements vary depending on your institution's classification.

Ph.D.-granting institutions and non-degree-granting organizations are required to cost-share 30% of the total project cost (which is different than 30% of the amount being requested in your proposal). Non-degree-granting organizations are exempt from cost sharing requirements.

You must have your institution's commitment to cost sharing in writing. It's important to note that any manufacturer or vendor discounts you are offered may not be designated as cost sharing.

Do Not Exceed Submission Limits

After you determine your eligibility, you need to find out how many applications your institution will be submitting in the MRI competition you want to enter. Why is this necessary? A single organization or institution may only submit a total of three applications during any given MRI competition – at least one of which must be for instrument acquisition.

You want to make sure your institution doesn't exceed the limit. If you were to send three MRI proposals, the NSF can return all of the without review or comment.

Cover your bases by checking with your research office to ask about policies for handling limited competitions. Then check the NSF site's search function (https://nsf.gov/awardsearch/) for active and expired awards from your institution.

Determining How Much Funding to Request

Never ask for more funding than is required. Nor should you lowball your request. It is highly recommended that after you have taken all of the steps described in this white paper, you take the quote from your HPC vendor and request the amount of funding that is needed to purchase your new HPC cluster. Nothing more, nothing less.

Consider limiting your grant proposal to an amount of less than $1 million. Proposals that exceed $1 million face increased competition from the many applications for expensive lab equipment. It is highly recommended that you limit your MRI funding request, especially if this is your first, to less than $1 million.

Casting a Vision for Your Proposed HPC Cluster

Your most important job as grant writer is to cast a vision for your project that inspires the NSF to award you the funds you seek. At the outset of your proposal, say why your project is important. Be able to back it up. You should also give the worst-case scenario: what happens if you don't win the grant? What happens to the project and those who would rely on the cluster?

Make sure the science being supported by your cluster is compelling. Describing good science is the backbone of your proposal. An interdisciplinary panel will be reviewing your proposal. Write accordingly.

Defining Your Instrument

- What is the role of the HPC cluster?
- Why is the cluster needed?
- What similar clusters are already available?
- How will the cluster attract researchers?
- Who will benefit from using this HPC cluster?
- How will researchers access the cluster?

As with any academic writing, include a strong list of supporting references and cite publications. Share data about the strong current funding for the scientific research you are proposing. If you can't show that, your proposal will be found wanting.

Broad Impact: What is the Big Picture Benefit?

At the heart of the NSF's mission is to fund research projects that will have a broad impact on the world. Your proposal needs to include that big picture view. Make sure you discuss the societal benefits of your proposed HPC cluster.

Making the case for the benefits of the many research projects that will be supported by your cluster is a great start. You can also demonstrate a broader impact by working with campus diversity programs to ensure everyone has access to your cluster.

How will your HPC cluster affect training education on campus and in the community? Talk about how the cluster will attract researchers, students and minorities from around the world.

Show how the cluster will improve research. What effect will the cluster have on STEM research?

We understand it can be challenging to conjure up a broader impact. Check out these examples on the NSF site for inspiration: https://www.nsf.gov/od/oia/special/broaderimpacts/

Give Your Cluster a Name (and a Few Heroes)

Give your cluster a name, and use it often. Make the name something that relates to your institution or organization for added effect. For example, if your university is in Dallas, Texas, consider naming your HPC cluster "Tex" or "Big D."

Provide a list of research projects that will make use of your HPC cluster. Highlight the biggest "hero" projects first – and include the most interesting details. You can devote an entire page to each of the hero projects if it's necessary to make the case for each. List all of the other research projects as well along with a brief summary of what makes each important. Talk about how much time/memory/bandwidth each project will require of the resource.

You're going to need to prove that the HPC cluster will be full all of the time – or even oversubscribed. Otherwise, it won't appear to be necessary.

Don't Forget to Share Your Proposed New Toys

NSF wants you to share. In fact, it requires that you do. Shared use of your cluster is vital to the success of your proposal.

Make sure you talk to researchers on your campus and at neighboring institutions to build a list of potential users for your cluster. Don't stop there. Graduate and undergraduate students also make nice additions to your cluster user base. The more, the merrier.

Get Institutional Support – and Get It In Writing

You'll need a letter from your institution stating that it will operate and maintain the cluster during and after the grant period.

Make sure the letter lists all MRI awards your institution has won in the last five years along with a summary of each and the status of the instrumentation from each award. Past successes, and continued use and value, go a long way toward demonstrating the validity of your current proposal.

Avoid mentioning any dollar figures here, which should be part of the proposal where you talk about the physical facilities that will be home to your proposed cluster.

Your HPC Equipment Proposal

You will need to be specific about the HPC equipment you will need, and you will also need to justify the technology you want to acquire. Many grant writers at the university level begin by consulting with on-campus IT staff, but we recommend you include at least one specific vendor quote. This is a service that Advanced Clustering is happy to provide.

Questions for Your HPC Vendor

- How many CPU core hours will you need?
- How much storage will you need?
- How much bandwidth will you need?
- Have you benchmarked your code?
- How will performance be enhanced?
- Will software be optimized?

Vendors deal with many HPC users and customers and use that experience to create the most common and useful configurations. They also work with other grant applicants and have learned what kinds of proposals get funded. They also have access to current market pricing and discounts and can also help you tailor a system that meets the specific needs of your project. Allow several weeks in your grant writing timeline to work with the vendor on your HPC cluster quote.

The NSF expects you to justify the need for the technologies you select. You can't just say you want a GPU cluster without making a clear case for why it is needed. By the same token, you're not going to be granted funding for 600TB of storage without explaining why so much storage is needed.

When including information about the HPC cluster you're proposing, remember that details matter. Core hours need to come from real data. The committee needs to see that you know the specifics of the cluster solution being sought and the practical uses for it.

If the proposal is for a new type of technology (for example, the latest NVIDIA GPUs or Intel Xeon processors), be able to talk about who will port the code to the new platform. What kind of speedup is expected on the new platform?

Be able to address the storage requirements of your proposed HPC cluster. How much storage will you need? If you need live storage, what is the maximum amount of storage at a time that will be needed for this project? If you need archival storage, what is the total amount of storage needed over the lifetime of the cluster?

Ask for help from your HPC vendor in making these determinations.

More Questions for Your HPC Vendor

There are other topics to cover with your HPC vendor as well:
What is the typical size of each dataset to be transferred?
Where will such datasets originate?
To where are you transferring them?
Why do such datasets need to be transferred between these endpoints?
What is the time window for transferring each such dataset?
Why does each such dataset need to be transferred during that specific time window?
How often do you expect to have such a data transfer need?

Get Physical Requirements on Paper, Too

One of the reasons we recommend you spend a good deal of time talking with your HPC vendor is to ensure your HPC consultants have all of the details about your proposed cluster. That includes the physical details. Where will your cluster be located? How much floor space, A/C, power and UPS is currently available.

Your HPC vendor will compare this against what will be needed so you can be aware of any shortfalls. You will need to talk about how you will meet any of these shortfalls, which cannot be covered by MRI funding.

This is where you need institutional support. You will need a letter of support from your institution that addresses their commitment to provide anything you need (space, A/C, power, UPS, etc.) to support the HPC cluster.

Data Matters, So How Will You Manage It?

Settling on your instrument management plan will take some time. You will also be required to share your plan for the management of all the data that will be produced by your proposed instrument.

You will need to talk about the types of data, samples, physical collections, software, curriculum materials, etc. that will be involved. What standards will you set for data formatting and content?

How will you protect privacy, confidentiality, security and intellectual property rights? How will data be archived? How will you preserve access to the archived data?

Define Your Management Plan

You are required to include a detailed management plan for your instrument as part of your proposal. Start by describing the facility where the instrument will reside. Describe how the instrument fits into the overall campus infrastructure.

Designate a project leader who will oversee all aspects of instrument management. A tenured faculty member is preferred. You should also identify a faculty advisory committee and talk about how long they are willing to serve (preferably for the duration of the project). Discuss who will help with day-to-day operations of the instrument.

Who will have access to the cluster, and who will manage all operators? When decisions need to be made about the equipment and its management, describe who will be involved and what procedures will be followed to determine the best courses of action. We also recommend you include a plan for sustainability. In other words, when the grant ends, what will you do with the equipment involved?

Checklist: Your Management Plan

- Where will the instrument reside?
- Who will have access to the cluster?
- Who will operate the equipment?
- How will the instrument be maintained?
- How will time be allocated for users?
- What will cluster operations cost?
- What will cluster maintenance cost?
- How much downtime do you anticipate?

Warranties and Service

Also make sure your HPC vendor includes warranty and maintenance costs and considerations for a minimum of three years.

Work with the vendor to get the most reasonable price possible. Ask a technology expert to review all cluster and HPC equipment descriptions to ensure you haven't missed anything important.

Keep in mind that proposals for smaller clusters at undergraduate-serving institutions are well-received by the NSF and fall into a different category of competition.

Institutions (especially smaller ones) often lack IT support and have a greater need of vendor support even after the cluster is delivered. These clusters are going to need consultative, high-touch vendor support versus a vendor that solely provides hardware with no ongoing support.

Learning From Others

We are sharing this feedback that one of our customers received upon submitting an MRI proposal so that you can benefit from reading the reviewers' comments about the strengths and weaknesses of the proposal.

Strengths

- Much detail about how the new infrastructure will be integrated into the existing system and the process of integrating its use into the university.
- Great description of cyber-infrastructure plan.
- Nice description of how the use of tools can reduce barriers to entry for novice users. Open source tools will be important to project success.
- Broad impacts are clearly defined for this project. Exposing undergraduate and high school students to computational research and tools develops the necessary pipeline for future research scientists as well as computing professionals.
- The proposed system represents a novel approach to HPC leveraging on demand for interaction and positions it alongside familiar campus technology resources such as email.
- The project leverages many external advisors, including local, regional, and national contacts, to ensure smooth deployment.

Strengths (continued)

- Project includes a high degree of engagement with activities to be disseminated at the campus level via posters at an established research and scholarly activity fair.

Weaknesses

- Platform architecture specifying HPC cluster components needs to be clearly described.
- Proposal lacks a statement about how the university will handle hardware upgrades once manufacturer warranties vanish.
- The work of talking to research computing providers in the state to ensure the proposed system addresses the actual need has not been done in advance, and this proposal is made weak by this failure to define the actual need.
- Proposer needs more focus on scientific use cases, the number of which will underscore the size of the instrument being requested. The types of equipment needed should be clearly identified in the use case summary. Is the plan to teach the courses using GPUs or CPUs? How much time is needed?

Weaknesses (continued)

- Where is the project plan with milestones and a plan to assess success needs? The proposal needs more information about how the system will be set up over the installation period provided by the vendor. Will the vendor be working directly with the datacenter team? Are there letters of collaboration to support this work?

- The proposal states that the PI will have only a few hours per year release time for research. This in no way equals the amount of time that must be devoted to the the maintenance and support of an HPC cluster that supports multiple courses and students. You need to identify the person or persons who will be providing backup technical support when the PI is unavailable.

- This proposal lacks references that would have added weight to the validity of what is being requested here.

- More detail is needed about the physical location where this resource would be placed. Where would it be? Is there adequate power? What about cooling?

- The proposal makes several references to a vendor quote, but the quote was not provided.

Produced By Advanced Clustering Technologies

Advanced Clustering offers customized, turn-key high performance computing clusters, servers, storage solutions and workstations to customers in the fields of aerospace, climate, defense, education, energy, engineering, finance, government, life sciences, etc. With more than 16 years of experience developing HPC solutions for universities, government agencies and industry, Advanced Clustering has advised many customers on writing successful grant proposals for HPC purchases. Learn more at advancedclustering.com.

advanced clustering
technologies, inc.

WHITE PAPER

Best Practices for Writing an RFP for the Acquisition of High Performance Computing Equipment

Prepared by

advanced clustering technologies, inc.

FIRST EDITION, March 2018

Download this white paper at http://actlink.co/rfp

Introduction

In our 17+ years in business as a provider of high performance computing clusters, servers, storage solutions and workstations, we at Advanced Clustering Technologies have seen hundreds of Requests for Proposals (RFPs).

We have seen successful RFPs that were completed with timelines intact. We have also seen RFPs that were incomplete, resulting in broken timelines and/or re-bids. That is a painful road, one we hope to help you avoid with this white paper.

This white paper is possible through the support of our customers, many of whom agreed to share their RFP writing experiences. We want to thank the following recent RFP writers who contributed mightily to this white paper:

William Garrick
Manager, Research Computing
Portland State University

Doug Jennewein
Director of Research Computing
University of South Dakota

Dr. Evan Lemley
Assistant Dean, College of Math and Science
University of Central Oklahoma

Dr. Gopinath Subramanian
Assistant Professor, Polymer Science and Engineering
University of Southern Mississippi

First Things First: Information Gathering

Information gathering is the first step toward writing your RFP, and there are a number of people who have information to help you through the process. They include:

- your purchasing department
- your datacenter personnel
- outside vendors
- your peers at other universities
- the other PIs (assuming you won a grant)

Your Purchasing Department

One of your first meetings should be with your purchasing department to discuss the details of the equipment you plan to purchase and get answers to key questions, including:

Is an RFP the right way to facilitate this purchase?

The are other kinds of proposal request documents that you want want to use. Ask your purchasing department which format you should use. Other options include:

- Request for Quote (RFQ)
- Invitation for Bid (IFB)
- Invitation to Negotiate (ITN)

Has purchasing dealt with ordering HPC hardware before?

This is important to ask because if they have done this before, they may have tools or advice that can help you. If there is no experience with RFPs for HPC hardware, you will be collaborating together to successfully navigate the process.

Is Board of Regents approval required for this purchase?

In some states, grant winners are required to get approval from the state Board of Regents prior to starting the purchasing process. If you need such approval, you will need to get on the board's agenda, which could affect your RFP timeline.

Are there state laws/university requirements to follow?

You will want to make sure your RFP is in full compliance with all relevant laws and requirements. It's very important to find out if you are allowed to include your budget amount in your RFP. If not, you will need to make sure you're including enough detail about your new cluster to get the desired responses.

Is there a template you can use for writing your RFP?

Most RFP templates are designed for equipment purchases other than HPC. It's common for templates to only ask for Product Name,

Manufacturer Name and Part Number. It's important to discuss with your purchasing department the ways in which HPC hardware is more complex than other kinds of equipment. The template will likely only be a starting point for you. You will need to expand the information you include in the RFP to ensure vendors provide complete responses.

Are there examples of successful RFPs I can review?

Because RFP templates are typically not designed for HPC equipment purchases, it's helpful to take a look at RFPs that have been successful in the past. If your purchasing department cannot provide sample RFPs, reach out to your peers at other universities or ask outside vendors for suggestions about where to find examples of strong RFPs.

Are there scoring guidelines I need to follow?

Your scoring system is critical to the success of your RFP. It's not only going to be the thing that guides you through the evaluation process. It's going to be the means by which you select the winning proposal. It's important that you understand how your RFP will be scored before you begin writing it. Here are some questions to ask your purchasing department:

- Am I required to accept the lowest bid or can I make a decision based on the best fit?
- Does the university have a scoring system in place that I need to follow?

Your Datacenter Personnel

It is equally important that you speak with the people who manage your datacenter or the facility where your new equipment will be installed. You need to learn about the physical properties of the site You will need answers to the following questions:

- How much available space do we have in our datacenter?
- Will my new cluster fit in the datacenter?
- How much power is available in the datacenter?
- What type of power connectors do we have?
- What is the cooling capacity of our datacenter?
- What is the current rack configuration?
- Is there room for new racks and, if so, how much?

Knowing the answers to these questions will help you avoid the disaster of purchasing a cluster that cannot be accommodated by your datacenter.

Outside Vendors

Before you can adequately describe the HPC system you need in your RFP, you must have an understanding of what is available in the marketplace. It's important that you seek information and quotes from outside vendors. They have expertise working with all current technologies, and they will be able to help you plan for any emerging technologies that are available as well. Failure to seek vendor input will likely result in an RFP that fails to deliver the responses you seek or the best equipment you can buy.

Your Peers at Other Universities

Your purchasing department may have examples of other technology RFPs on file to guide you, but it will also be very helpful to review RFPs written by your peers at other institutions. Reach out to colleagues elsewhere and ask if they would be willing to share their RFPs.

The Other Principal Investigators and Co-PIs

Do not forget to stay in communication with your other PIs as you work through the RFP writing process. Their input and advice will be helpful as well. Together you should compile a complete list for the user types who will be relying on your new equipment. What departments will be using the cluster? What kinds of research will they be conducting? What software applications will be required? All of this information will help you determine the ideal configuration for your new cluster.

The Next Step: Writing the RFP

Once you have finished collecting information, you are ready to start writing your RFP. If your purchasing department was able to provide you with a template, that is a good starting point. Realize, however, that most RFP templates were designed for simple purchases of a simple piece of equipment and ask merely for Product Name, Manufacture Name and Model Number. You must provide more product details in your RFP to get the desired responses.

Essential Elements of Your RFP

Q&A Period

Look at the Question and Answer period as an opportunity to fail-proof your RFP by using vendor questions to clarify what you need.

Make sure you allow at least a week between RFP publication date and the deadline for vendor questions. Also make sure you allow yourself enough time (at least one week) to answer the questions.

Consider offering two Q&A periods. Why? Because the first round of answers will inevitably lead to more questions. Allowing vendors a chance to further clarify makes it all the more likely that you will have a successful RFP. It also helps you avoid mistakes or omissions that could cause the whole process to fail. You want to avoid re-bids and vendor appeals.

Your Overall Timeline

You need to allow a minimum of three weeks between the end of your QandA period(s) and the RFP deadline. Vendors need time to negotiate deals on equipment. If you don't give them enough time, you won't be getting the best deal.

Technical Requirements

This is the most important part of your RFP because this section outlines the configuration for the cluster being requested. The technical requirements need to be as detailed as possible. Without precise requirements, the responses you receive will not present you with the best options for your hardware needs.

Do your homework. Have a strong knowledge of the equipment you want to buy. Consult with vendors to understand the current technologies. Do not become too attached to a specific technology, manufacturer, processor, etc. unless you understand fully what it means for your cluster.

If you are not able to publish a budget amount in your RFP, you need to give vendors other limiting factors to ensure you get responses that meet your need. Without a budget amount, at a minimum vendors will need to know a total core count or the total number of nodes you require.

Without this, some vendors will refuse to participate in your RFP because they won't have a clear understanding of what you want. The vendors who do participate will only be guessing what you need. Either way, this is a situation worth avoiding.

If you are required to accept the lowest bid, it's going to be important for you to ask for, as an example, the best cost for X number of cores.

Understanding the number and types of jobs being run should dictate the quantity and types of nodes to get. Some machines (such as high memory nodes) can cost many times more than a standard compute node, but if only one user needs it, it doesn't make sense to purchase multiples of these as it will cut into the number of other types of nodes you can purchase.

For your overall system, answer the following questions:

- What is your preferred core count?
- How many total compute servers do you need?
- You can request Best Price for X number of cores or X number of servers that fit your specifications

Avoid using "or better" in your RFP. Example: "Please provide 50 cores or better." This is asking for something for free. Better to ask for what you really need.

For storage solutions, define requirements clearly and try to include redundancy, volume, file system and use cases.

To demonstrate the importance of being verbose in the technical section of your RFP, we are sharing some common questions that get asked by vendors during the RFP question. Providing enough details to answer these kinds of question will enable vendors to provide with the best possible solutions from which to choose.

Sample Vendor Questions from Recent RFPs:

Compilers

Are there any specific compiler requirements?

If you want an Intel Compiler Suite, can you specify which version (there are several)?

CPUs

Does the university have any requirements for which Skylake CPU they'd like to receive?

You're requesting 128GB of RAM. However, the new Skylake CPUs have six memory channels and thus, we'd need to go to 192GB or 96GB or Ram (or more). Do you have a preference?

GPU Computing

Is there a steady growth in the interest in GPU computing? If so, in what applications or fields?

Concerning GPU nodes, will you accept non-Skylake GPU nodes?

How much memory do you need on your NVIDIA Tesla P100 GPUs? 12 or 16GB?

Interconnects

Is there a high speed interconnect other than Ethernet required? If so, what kind of performance (bandwidth and latency) is needed?

Can EDR be substituted for FDR InfiniBand?

I/O

How should the I/O solution function when a failure is exposed beyond the RAID level?

Nodes and Node Counts

How are specific nodes to be configured? In other words, what CPUs, how much memory, how much storage, etc.?

Is there a preferred core count, clock frequency or model number?

Do you want the head/login node to have high availability?

Concerning the compute nodes, is there a target quantity or performance level that you are trying to reach?

Parallel File Systems

Will scratch space be allocated in the Parallel File System (PFS)?

What is the desired throughput for the PFS?

Do you have an existing PFS? If so, what is it?

How many files do you currently have on your PFS?

What is the average file size on the existing PFS?

Where do you intend to mount the PFS (compute only, compute + management, etc.)?

Is there a preferred Parallel File System or should we recommend (GPFS, Lustre, etc.)?

Racks, Cables, Cooling and Power

Are we supposed to supply the racks and PDUs? If so, please provide makes and models.

Is there a room diagram or will the racks be adjacent to each other?

Will a floor map with environmental and connectivity configurations be made available?

Are there Vizio drawings of existing rack layouts?

Are there cable trays for networking either above or below the racks? If so, what is the distance from the top or bottom of the racks to the cable tray?

Is there a per-rack limitation on cooling?

Is the power single phase or 3 phase?

How many receptacles can the University provide per rack?

Should the server racks come pre-cabled for all compute nodes or just for the parallel file system and associated components?

Software

For commercial software, how many years of maintenance is required?

Storage

Is there an overall storage speed requirement?

Do you need to support CIFS and NFS on storage system?

When you say battery backup on the storage solution, do you mean batteries on the storage controllers or batteries for the entire storage array, as in a UPS?

Training

How many people will be attending training? And on what software and hardware products?

Allow for Innovation

You must allow for innovation in vendor responses. Vendors are the subject matter experts when it comes to HPC equipment, and you want to receive the value adds that they bring to your project. You should welcome options in your RFP process because more options mean more choices for you and, ultimately, the best HPC cluster for the job.

If the goal, for example, is to get as many nodes as you can, let vendors present you with that option. One word of caution: You need to strike a balance between how much leeway you offer vendors and how you will score such innovations when you are evaluating the responses. If vendors have too much leeway, you may not be comparing apples to apples when scoring responses.

Delivery

HPC Equipment is big and heavy. Shipping it is costly. Delivering it is painstaking work that requires a lot of pre-planning. There are countless things that can go wrong with delivery. In this, the devil is in the details. Providing delivery details will help the vendor pack, ship and deliver your equipment without issue.

Here are questions your vendors will have. You should seek answers in consultation with your datacenter personnel:

- Where will the cluster be delivered?
- Is there a loading dock?
- What are the dimensions of the loading dock?
- Any overhead limitations?
- What are the hours of operation for your loading dock?
- Will someone be there to receive the shipment?
- Are carts available for use? Size/weigh limitations?
- Are pallet jacks available?
- What is the cluster's final destination?
- How far is the final destination from the loading dock?
- Can you provide a map from loading dock to destination?
- Are there doors along the route from loading dock to destination? Dimensions of doors?
- Do we know keys/access codes? Will someone be there to assist?
- Are there elevators along the route from loading dock to destination?
- If so, what are the door dimensions and size/weight limitations?

Installation/Training

This is an often overlooked part of any RFP. Do not forget to provide your requirements for installation and training services as part of your HPC purchase. If you are going to require on-site installation, provide vendors with details about the location. This would be a good place to include details about datacenter power and cooling capabilities.

Scoring

Use all requirements or guidelines provided by your purchasing department as the basis for your scoring system. You will include each factor and its possible score.

Make sure you ask for at least three references from each submitter. Preferably references will be for customers who bought systems of similar size or configuration. Contact all references and let what they say be a factor in your selection process.

Here is an example:

Sample: Scoring Your RFP

Total Cost for Required Items	150
Technical Merit	50
Vendor Experience/References	50
Installation/Support	50
Subtotal	300
Errors/Missing Information	-20 each
Total	

The Last Steps

After you have written your RFP, have all Primary Investigators and Co-PIs review the content and offer any suggestions for changes before you arrive at what you consider to be the finished draft.

Review the Deadline

Make sure your deadline is realistic. Are you allowing enough time for vendors to write their responses? Remember, they will often be collecting quotes from suppliers and manufacturers, which can be a time-consuming process. Also make sure your time between response deadline and your decision is allowing enough time for you to evaluate and score the responses. Finally, are your deadlines in keeping with any requirements you have to spend the money by a specified date?

Be Open to Post-Award Recommendations

You should consider adding a statement within your RFP that states you are open to discussions with the winning vendor after the contract has been awarded to ensure that you have been given the best possible configuration with the most appropriate technologies and components. In the end, you want the best HPC hardware and software solutions that you can get within your budget. Allowing vendors the opportunity to present you with recommendations that might not have been possible within the constraints of the RFP is a worthwhile exercise.

Keep Your Team in the Loop

Hopefully you have kept your purchasing department and datacenter team in the loop throughout the course of your RFP information gathering and writing process. Don't forget to keep them informed as you publish the document and collect and score responses.

Share Your RFP experiences

This white paper is a work in progess. We will be updating it as we hear about the RFP writing experiences of others. Please share your experiences with us (info@advancedclustering.com) and also let us know what you thought of this white paper and its tips for RFP writers.

Produced By Advanced Clustering Technologies

Advanced Clustering offers customized, turn-key high performance computing clusters, servers, storage solutions and workstations to customers in the fields of aerospace, climate, defense, education, energy, engineering, finance, government, life sciences, etc. With more than 17 years of experience developing HPC solutions for universities, government agencies and industry, Advanced Clustering has advised many customers on writing successful grant proposals and RFPs for HPC purchases. Learn more at **advancedclustering.com.**

Download this white paper at **http://actlink.co/rfp.**

advanced clustering
technologies, inc.