CLUSTER LIFECYCLE MANAGEMENT

BY DEEPAK KHOSLA

Assessing Your HPC Needs

For many organizations, an HPC cluster ranks among its most expensive and complex IT investments. From acquisition to replacement, these complex systems must be proactively designed, managed and operated within the context of a lifecycle, one that is unique from most other IT assets. To maximize ROI within a rapidly changing market, a pragmatic approach is necessary. This column series will walk through Cluster Lifecycle Management from an initial Needs Assessment, matching needs to design, Vendor Selection, Deployment, Maintenance and finally to the stage of Replacement.

Congratulations! You have decided to join the supercomputing club, now what? First order of business, perform a Needs Assessment to determine exactly what you require in terms of HPC technology.

Put simply, understand the intended use of the cluster and expected return to the business. Understanding the cluster's intended use cases will allow you select a cluster that's the right size with the right architecture and components. Buy a cluster that is too small to meet demands, and you have a money pit requiring unanticipated expenditures for years to come. Buy one that's bigger than you need, and you've wasted money on an investment with a poor ROI. The management you worked so hard to convince that HPC was just the solution to your problems isn't likely to forget either scenario.

A Needs Assessment starts with an understanding of the application(s) and types of computations to be executed on the cluster. If you have HPC experts on staff, have them perform an analysis of how personnel will be using the cluster, what the workflow will be, and what software applications will be deployed. In these cases the HPC experts will work with the application vendor or developers to determine the ideal number and type of nodes and most appropriate operating system needed. If you don't have HPC experts on staff you will need to rely on the hardware and software vendors to work with your cluster users and business units to understand the user requirements and system demands.

More specifically, your HPC experts or the vendors will want to know the typical usage, or performance, profile of each application within your organization.

- How many users access the application at any given time?
- Is the application part of a larger workflow?
- How will they launch the application from their workstations?

These questions are paramount in helping to scope the design of the compute cluster (types of nodes and interconnects) to meet the intended usage, performance and operational objectives, and to maximize overall effectiveness of the cluster in parallel processing.

Next consider data storage capacity and I/O performance, which will be determined in part by how the users access data for input and output of the computations.

- Where is the input data, and how will they get output data and visualize results?
- How do they plan on sharing work with others?
- How big are the input and output files?

In addition to initial cluster demand requirements, considerations should be given in the Needs Assessment for future use or growth of the cluster. It is important to look at the business and technology roadmap associated with the cluster to see how its use may be expanded over time. The organization may have twice as many cluster users; there may be new applications added to the cluster, each requiring additional cluster resources and capabilities such as compute capacity, memory, storage, bandwidth or high-availability.

Once the cluster hardware and software decisions have been made, the more practical consideration of cluster management must be addressed. Infrastructure comes first. Your organization must ask whether it has the space with the proper cooling and electrical networks to install the system. Similar to compute capacity expansions, space expansion must be evaluated to ensure easy expansion in the future.

Will the cluster require High-Availability (HA)? HA is all about ensuring the systems remain available for users even if a major component of the cluster fails. This means redundant components, and duplicating critical components such as head or master nodes, switches, storage and network fabrics. Depending on the redundancy desired, implementing HA components can have a significant impact on cost, design, configuration, management complexity and infrastructure.

Now we need to consider an array of cluster specific tools, applications and services that will be required on the cluster. Somehow all these hardware and software components that make up the entire cluster must operate synergistically. Some services are required for the cluster to operate, others are optional. Similarly, some cluster tools are a must, others a convenience. Services and tools for the cluster vary and there are many options to choose from. Deciding on preferences and capabilities are important decisions for users and managers of the cluster.

An organization with typical IT staff may be disappointed to learn that those individuals lack the expertise to care for and maintain an HPC cluster. The organization must consider hiring new personnel or outsourcing management activities which will likely involve continuous monitoring, proactive cluster management, analyzing, and reporting on the system's operations. Staffing and services for the management of your cluster must be considered in your Needs Assessment given the painful alternatives.

Matching Needs to Cluster Design

As explained in the first column, much of the Needs Assessment must focus on the software applications that will be run on the cluster because their operational parameters dictate selection of many critical system components, compute nodes being first on the list. Each application has an optimal memory bandwidth, expressed in terms of gigabits per second (Gbps), which it needs to run effectively for each CPU in the node. The key is choosing / designing a node that can most effectively support the Gbps rating of the application across the number of CPUs within the node

For the typical organization, even one implementing a small HPC cluster, there are often multiple applications that will be run on the system. In this situation, the Gbps rating of the most memory bandwidth-intensive application should be taken into account when deciding on compute nodes. So if the most demanding application requires 64 Gbps of memory bandwidth, a node with that throughput will accommodate not just the one application but also all those requiring less memory bandwidth. The size of memory in the node is the other important factor. The required memory per CPU or overall for each application will dictate how much memory needs to be designed per node.

There are, however, implementations in which one application is a memory hog compared to all of the others combined. In that case, it is often less expensive to install one or more large memory nodes (often called 'fat nodes') for that particular application, and 'thin nodes' to handle the other less demanding applications.

The number of nodes installed in the cluster will depend on many variables, but one of the most important is a business consideration. Multiple nodes accelerate the speed with which the cluster performs its processing and produces the answers desired by the organization. Driven by business requirements, some organizations may need their applications to generate answers in two hours, while eight hours is sufficient for others. Multiple nodes provide results faster and will be worth the extra investment by some organizations.

Another important element in cluster design is the need for local disk space on the node. This is dependent on data input/output (I/O) requirements of the applications that will run on the system. Each application is rated for a certain frequency and volume of data reads and writes. Some applications will require the local disk to provide temporary 'scratch' space to hold data generated during processing. If scratch is needed, the I/O bandwidth of the application will dictate how many drives to put in the local node.

Reading and writing data to and from local disks takes time. For some applications, the speed with which they can read/write data can be more important than the volume of data involved. These applications are considered 'latency sensitive' (needing low latency) with respect to disk I/O.

Bandwidth and latency both play roles in selection of disk drives for the node and can significantly impact the overall cost of the implementation. Generally speaking, applications with high I/O and low latency ratings require SSD drives, which are the most expensive. High I/O and medium latency applications can get by with SAS drives, middle of the road in price. If latency is not an issue for the application, SATA drives are sufficient and also the least costly.

In addition to local disks, the cluster will need to have access to data on some central storage to facilitate getting input data and storing output data which needs to be shared among all nodes, as well as with external users. These external storage devices will move data to and from the nodes through an I/O-interconnect. Typically 1G or 10G connection per node may be sufficient for accessing shared storage. The reason is that since this is shared storage access, the throughput of the network or the storage itself is usually the limiting factor. For high throughput, low latency I/O needs, it is best to use local storage.

Another key implementation factor dictated by the applications is the need for an application-interconnect that allows nodes to communicate with each other. Some applications require low-latency communication among the nodes in the cluster, and some do not. For those that require latency in the single digit microseconds or throughput over 10Gbps, for example, Infiniband (IB) is the best fit today.

At the other extreme, if the application only requires high volume data transfers at the beginning and end of the jobs, and has very light inter-node communication needs, then a single 1G or 10G interconnect could handle both I/O and application traffic. If Ethernet is determined to be sufficient, then even if the need is 1G today, it is still recommended to specify systems with the 10G built-in port to handle future growth as it is becoming the default configuration.

The final decision in the cluster design is the operating environment. Again, the applications drive this selection. Certain applications run most efficiently in specific operating environments, and this requirement is well documented for each application. For an organization running multiple applications on its cluster, the challenge is finding one operating environment that satisfies the specifications of all applications. This is usually the least expensive option.

Often, however, multiple applications will require more than one operating environment. One option at this point is to install separate sets of nodes, each running the desired operating system. But there is another option – set up the nodes to run in multiple operating environments by establishing a provisioning system to dynamically switch from one environment to the other when particular applications are running. While complicated, this provides flexibility to match the resources to the workload.

Choosing a Vendor to Build the Cluster You've Designed

The first step in the vendor selection process is writing down all of the cluster design requirements – both hardware and software – that you derived from the needs assessment. You don't have to do a full-fledged RFP, but at the very least, the hardware requirements should include number of nodes, CPU type, memory size, local disk space, interconnect type, and storage size. For this phase of development, it's fine to define these requirements with some acceptable range of values.

Software specifications also have to be included in the requirements document. Preferred operating systems, libraries, compilers and workload managers should be identified. Again, if your application can run on a couple of different operating systems with a variety of schedulers, these options should be written down so your vendor search has some flexibility - a necessity when shopping for the best value.

Once the hardware and software requirements have been outlined, the facilities and IT department should also be contacted to determine how much power, cooling and physical space is available. Sharing the specifications will help them provide some preliminary options for these and can help highlight any IT software or systems constraints.

There are two critical risks this minimizes. Many a cluster deployment has been delayed at the last minute when it was discovered the facility lacked the needed power or air conditioning to run the system. If power and cooling upgrades will be necessary for an HPC cluster, they should be included in the planning process to prevent surprises. Further, if there are any software or network related guidelines (e.g. OS, security, access, etc.), these can also be addressed early in the process. These IT specifications should then be added to the overall requirements. Again, providing flexibility where possible will allow for more options by vendors to help achieve the maximum value.

As the requirements document is being prepared, the organization should also create a formal vendor selection committee comprised of personnel representing all of the stakeholders in the HPC development process. Business operations, end users and the internal IT staff should be involved. For some larger organizations, the IT department is divided into three teams handling storage, network and server operations. All three must participate in the selection committee.

A vendor-neutral outside consultant may also be a welcome addition to the team. Consultants with experience on multiple HPC cluster deployments can usually review the requirements and assist the organization in calculating a workable budget for the overall project, and provide risk-benefit advice on various options. They can also help during final vendor selection by giving thumbs down to unnecessary 'upsells' that inevitably creep into many proposals.

The final step before approaching potential vendors is to assess internal delivery capabilities. Treat your IT staff as a vendor and query them to determine how much of the deployment, if any, they can handle. This can save a lot of money. However, from experience, I have found deployments proceed more smoothly when the outside hardware or software vendor is also responsible for installing the clusters and getting them up and running. A single point of accountability is often easier to manage when the schedule is tight.

Internal capabilities also must be examined to decide who will manage and support the cluster once it's operational. Few traditional IT staffs have experienced HPC professionals onboard. Hiring new personnel may be necessary. Otherwise, some vendors offer management contracts as optional services, which can be upsells well worth considering. These management needs must be addressed in advance.

Finding Potential Vendors

Once the requirements document is complete, it's ready to send to select vendors. Organizations entirely new to the HPC arena might not know where to start. I recommend asking the existing IT department whether your preferred computer system vendor offers an HPC solution. If so, continuing to do business with a trusted vendor is an excellent option, assuming their solution meets your design.

A second place to look for guidance is the developers or users of your primary applications. They may know which HPC hardware and software systems have run their application successfully. Your application developers may even provide you with references of other clients who are running the software in the HPC environment. Seeking advice by posting questions on relevant LinkedIn groups has emerged as another fast and easy way to find (or avoid!) certain potential vendors.

HPC vendors can be loosely divided into three 'tiers'. The first group is comprised of major providers such as Dell, HP and IBM. Specializing in hardware and software, they are well-known providers of large HPC clusters. The next tier includes specialists who provide highly customized HPC systems, such as Cray, SGI, and Aspen. Vendors in either of these groups can offer end-to-end solutions, including hardware, software and deployment services.

The third tier of providers is system integrators with HPC experience such as PCPC Direct or Dasher Technologies. They have well developed channels they will rely on to piece together the right hardware and software components to meet your requirements. They may also be able to handle end-to-end deployment. Aside from their ability to build a highly customized solution, a significant advantage here is their ability to present you with options to compare among their own channel suppliers providing price-competitiveness.

Choose at least three vendors and send your requirements document to each one. Let them know that you are looking for options and are happy to answer any questions or provide clarifications. If the vendor has a good HPC team, very likely, they will come back to you with questions asking for details on why you made certain specifications. Expect an iterative process. Have your entire committee involved at every step. Some of this discussion may take place by phone or online meeting. Take the questions and clarifications from the first round and re-submit a revised document to the vendors.

In the meantime, request and call references furnished by each vendor. Be sure to ask these references how the vendors performed in each phase of the project – procurement, deployment, and post-deployment.

Next, move forward with personal interviews. Invite representatives from the vendors into your facility to meet with the entire selection committee. At this point, the most minute details of the design document should be nailed down. If the committee is confident the vendors have workable solutions, it's time for them to submit a written price proposal. The committee must review the proposals to ensure that each contains all of the desired HPC cluster components and that comparisons between vendors are apples-to-apples.

The committee should then vote on the vendor that offers the HPC cluster solution that most closely meets your design criteria and provides the best value. And remember, best value is not always the lowest price. The best value takes into account not only the price, but also the speed and quality with which the users can get productive with the applications running on the cluster, the reliability of the system, and the quality of the ongoing support.

And have fun with the process!

Deploy and Validate Your Cluster

As part of the vendor selection process, I recommended you ask the candidates to include deployment and validation services in their price proposals. Most, if not all, HPC system providers are prepared to install the hardware and software they sell and make sure everything runs correctly. In some cases, they may subcontract some of this work out to experienced HPC professionals.

If you have an HPC-savvy IT department, performing the deployment yourself can save money, but you should give this option considerable thought. In too many situations, I've seen an expensive cluster sit idle while an unexpected problem is worked out. This downtime can wipe out any cost savings you hoped to incur by deploying with internal staff.

Once the deployment responsibilities have been settled and the final purchase contract has been signed, your HPC selection team has not completed its work. The same group of internal stakeholders must now turn its attention to the deployment and validation phases.

The first order of business is for the team to prepare the facility that will soon house the cluster. Your team must make sure the chosen site has adequate space to accommodate the hardware. There must also be sufficient and reliable power to run the system and air conditioning systems to keep it cool. It may turn out that the local facility is not adequate; if so a co-lo facility will also have to be selected. While looking at power, pay special attention to provide the exact voltage/phase/plug specs to the vendor. Providing the right power connection cables is usually the responsibility of the vendor, but they need the correct specifications in advance. If the site is ready for installation, deployment can begin.

For the larger cluster, HPC vendors will typically 'rack and stack' the clusters offsite before shipping, which means the racks can easily be rolled into your facility and put in their chosen spots. The racks will then be connected to each other and to the power source. An important step during installation is to label equipment and cable connections so they can be readily identified.

The vendor then powers up the hardware to make sure each component is functioning and performs a burn-in. Either onsite or prior to shipping, the vendor confirms the equipment has the most up-to-date supported bios and firmware for the various system components.

The next crucial phase of deployment is loading the operating system and software. Most vendors will use a cluster management system package to get the operating system and the HPC software stack deployed on the nodes. Specifically, this specialized software ensures compute nodes are set up consistently so they start properly from the operating system and all contain an identical software stack. If nodes are set up with consistent images and have connectivity to the head node, they all look the same to the applications that will run on them. It also significantly reduces the time to redeploy a node if needed or to deploy changes to the software.

The deployment may also involve setting up any external storage that is needed by the cluster. Finally, the appropriate scheduler and applications need to be properly configured and deployed.

The cluster is now ready for basic validation. The vendor runs a suite of software specifically designed to test the nodes and the cluster, usually High Performance LINPACK (HPL). Alternatively these suites may be manufacturer-specific and shipped with the HPC stacks. For example, Intel offers its own cluster-specific requirements program called Intel Cluster

Ready. Online validation applications are also available. In addition, tests may be requested by you and designed specifically for your particular use cases and applications.

Typically, the test first validates the nodes are functioning individually and then confirms they are operating together as a cluster. Some of the problems that may be identified during basic validation are memory issues within specific nodes or interconnect errors between nodes. Tools may be included in the suite to test each interconnect and data storage drive.

Often vendors complete their validation tests at this point. But I recommend additional validation and even benchmarking as part of the process. Once the basic cluster validation tests above have been completed, it is critical that the application setup be tested by having it submit jobs through the scheduler to make sure it will run on the cluster from end to end. Only at this point can you be confident your new cluster is fully up and running.

Some vendors will go the extra mile and perform benchmark tests to determine how efficiently the cluster is operating. An HPL benchmark, for example, measures the speed with which the HPC system arrives at solutions while performing real calculations. The results serve as a performance baseline for the cluster, and some vendors take advantage of this information to tweak the system – modify various settings – to squeeze greater power and speed from it.

Assuming issues were resolved during the validation and benchmark scores are acceptable, your HPC cluster is now ready for operation.

The total time required for deployment and validation will vary with the size of the system. A smaller cluster comprised of 16 or 32 nodes, for instance, may take a week to become operational, while a 200- to 300-node system may require a month or two depending on the complexity of the overall configuration and acceptance test requirements. These times may be shorter if the vendor performs much of the work offsite before delivery and installation at your facility.

Members of your internal IT group should be on hand during deployment and validation for several reasons. Ensuring the process proceeds as planned is important, and the team will understand how equipment is connected if they observe without getting in the way of the vendors. Most vendors are willing to have a 'knowledge transfer' session, but this is best left for the completion of deployment and validation so as not to slow down progress while work is underway.

Now it's time to put your new HPC cluster to use and make sure it continues operating efficiently.

Care and Feeding of Your Cluster

Once the deployment and validation of your new HPC cluster are completed, it is time for the HPC systems management functions to begin. I am assuming the advice from the previous columns was followed and the primary HPC system administrator was identified and in place in the deployment phase. This is no time to discover you do not have an HPC expert on your staff or at your disposal. Just because the hardware and software are humming now doesn't mean they will stay that way. Like any other complex system, the HPC cluster needs to be continuously monitored, analyzed, and maintained to keep it running efficiently.

The mistake I've seen made too often, especially by larger organizations, is the assumption that someone on the existing IT staff can probably figure out the HPC system, perhaps with some minor training. Unfortunately, this rarely works out. Although HPC is a niche within the larger Information Technology space, even the best IT generalist will have little or no experience in supercomputing. It is NOT just a collection of Linux or Windows servers stacked together. HPC is a specialization unto itself.

You must have HPC expertise available to you if you want the new system to perform as expected. There are two options – hire one or more full-time HPC administrators or contract for ongoing HPC system support. Budget will likely dictate which works best for your organization. For several scenarios, contract support may be a better option due to the difficulty of finding and retaining HPC experts on staff due to intense market demand or because you may not need a full-time person. Check with your system vendor or integrator to see if they offer contracted management services.

Now that your cluster is operational and you have a skilled HPC administrator(s) on staff or under contract, the first job is to configure the cluster so that it works well operationally. The two major aspects of this responsibility are that the cluster must be configured to work optimally from both an end user usability perspective and from a systems operation perspective.

The administrator must first set up proper security access for the end users. There are two major components to a successful security design. The first addresses connectivity to the appropriate authentication system that makes sure users can securely log in. Often the cluster has to be configured to tie into an already established enterprise system such as LDAP, Windows, etc. It is critical that this authentication performs with speed and reliability. HPC jobs running in parallel will fail often if the authentication system is unreliable. The second component to success addresses the authorization requirements. The administrator must validate that the file systems and directory permissions follow the authorization policies. This is critical so that users can work smoothly all the way from submitting the jobs to reviewing the results from their workstation. These must then be set up, configured, and tested across both the compute and storage components for the unique user groups.

Additionally, policies may need to be set up on the scheduler to allocate for various user groups and application profiles, as well as on storage to meet the varying space requirements. When security, computer, and storage are configured, users can safely log into the system and know where to securely put their data.

If your cluster is brand new, the users are most likely first-time users of HPC technology. This means they will need training and instruction on how to run their applications on the system. The applications they ran on a desktop or mainframe will not perform the same way on the cluster. Users will likely need application-specific training. Depending on the scheduler, there will be different ways to submit jobs from various applications.

It will be the administrator's responsibility to begin building a written knowledge base pertaining to the cluster and each application. This hardcopy or web-based document will serve as a guide for users to understand how to submit and track jobs and what to do if a problem occurs. Depending on the level or size of the user base, it may also make sense to look at some portals that can make job management easier for the end users.

For the cluster itself, the administrator should set up monitoring and alerting tools as soon as the system becomes operational. Monitoring, reporting, and alerting of storage, network, and compute services on a continuous or periodic basis are critical to identify signs of trouble before they turn into major malfunctions. Minor usage problems could simply mean disk space is filling up, but soft memory errors could be signs of impending node failure.

Such monitoring and analysis tools are readily available. Many HPC clusters come equipped with system-specific tools, while other more robust technical and business analysis packages are commercially available. Whatever their source, these tools should be set up to identify and predict routine maintenance issues, such as disk cleanup and error log review, as well as actual malfunctions that must be repaired.

In my experience, however, pinpointing the cause of several problems in the HPC domain requires looking for clues in multiple components. When things are going wrong with an HPC cluster, alarms may be triggered in several places at once. The skilled administrator will review all of the flagged performance issues and figure out what the underlying cause actually is. Few software tools can take the place of a human in this regard.

Proper care of the cluster also requires the administrator to be proactive. Every three to six months, I recommend running a standard set of diagnostics and benchmarks to see if the cluster has some systemic issues or has fallen below baselines established during deployment. If so, further scrutiny is in order. Last, but not least, the HPC administrator must find the right way to make changes so that all applications keep working well on the cluster. Patches and changes for applications, or libraries, or OS/hardware must be carefully considered and tested if possible, before implementing. I have seen quite a few expensive outages where a simple change for one application has caused failures in other co-existing applications.

Finally, a viable back-up plan must be enacted so the system can be brought back online quickly in the event of failure. The most important things to back up are the configurations of the scheduler, head node, key software, applications and user data. While intermediate data does not often need to be backed up, user input and output data should be, especially if the time to regenerate results is high. The organization should also establish data retention policies determining when data should be backed up from the cluster to offsite storage.

Capacity Planning and Reporting

I wrapped up the Care and Feeding column by noting how critical it is to monitor the HPC system as a routine part of daily operations to detect small problems before they become big ones. This concept of gathering system data dovetails with Capacity Planning and Reporting because the information you collect each day will paint an overview of larger operational trends over time that help you plan for the future.

An HPC system is not static. In fact, most will need to undergo major upgrades, expansions or refreshes after two to three years. And in all likelihood, these changes will be prompted by new demands put on the system related to growth in your business. New users and new projects will require changes and upgrades. Or in some cases, new or upgraded applications may require more processing capacity.

Perhaps most often, the trigger that prompts a capacity upgrade is related to data. In today's world, the problems being solved by the HPC cluster are getting bigger and more complicated, requiring the crunching of larger and more complex data sets. Adding capacity is sometimes the only way to keep up.

Monitoring and reporting can tell you how efficiently the processes and applications are running, but the information must be analyzed to determine how busy the system is overall. These details can help you make better decisions on upgrading and changing the system. Specifically, you need to anticipate when to implement capacity upgrades and which components of the system should be changed. Armed with this data, you are more likely to spend your money wisely.

The HPC reporting system should help provide the information needed to decide when to upgrade or add capacity as well as what type of resources to add. Some of the typical analyses needed are:

- What are the most commonly run projects and applications?
- How do they rank by CPU time?

- How do they rank by CPU and memory usage?
- What is the throughput of various architectures (ideally in business metrics such as 'widgets built')?
- Who are the heaviest users of the system?
- How many resources are used for their jobs?
- What are the cost allocations of compute and storage by users and projects?

It's critical to understand that the answers to these questions will vary by location. With the distributed architecture that is so common with HPC implementations at large organizations, clusters, end users, and data centers may be spread around the world. Their hardware and software may not be the same from one location to the next. In addition, local variables, such as labor and electricity expenses, will impact system operating costs.

The real challenge, therefore, is for the reporting system to be able to provide this information consolidated from multiple locations so that it can be analyzed. This analysis must be conducted with a return on the investment (ROI) in mind both regionally and centrally. This analysis must examine HPC usage data from the perspective of the business. Usage metrics must be monetized so capacity expansion decisions can be weighed against an ROI. Capacity upgrades have to be planned and implemented in a way that maximizes service to end users and their projects while still providing a positive return on investment for the whole organization.

So, how is this reporting and analysis best conducted?

Smart planning for capacity changes obviously can benefit from solid data reporting. Ideally, data is collected from systems and schedulers continuously and made available for analysis as needed, rather than spending time trying to find the right data when a decision has to be made.

There are several types of tools available to provide the above information. For example, open source charting tools like Ganglia, Cacti, Zabbix and others collect performance data from the systems and the network. Several of these can be extended to add custom metrics. Some cluster managers, also come with reporting tools that give insights into cluster health and performance. Most of these solutions work across heterogeneous architectures.

At the next level are job specific reporting tools from various commercial scheduler vendors. They are able to provide basic user and job information with varying level of sophistication. In general these are proprietary to each scheduler.

At a higher level, there are generic data analytics tools like Splunk that can provide insight from various types of the above sources. These require significant expertise, customization and upkeep to provide effective results. Finally there are a few platform-independent, HPC specific, analytics tools such as DecisionHPC that provide globally consolidated, single pane-of-glass system and job reporting for heterogeneous clusters and schedulers.

On the other side of the spectrum, some HPC operators can choose to build their own custom reporting tools. This also requires significant HPC knowledge as well as development expertise to ensure that the solution can scale, can meet ever-changing user needs, and is supportable and maintainable long-term.

As adaption into commercial uses for HPC increases, the importance of having good reporting and analytics also increases, thus leading to more solutions becoming available. Ideally, you should have a long-term strategy for a reporting and analytics solution that is independent of the various operational tools that tend to change over time, that can be easily supported, and that can be easily customized to your business needs.

At the end of the day, the system manager needs a solid reporting system to meet the evolving needs of the end users and the business. Knowing in advance when you need to increase resources and capacity is a critical part of that, as getting into the budget cycle and procuring the upgrades usually require long lead times. When it comes to upgrading resources, last minute decisions and investment, without insight into current and historical trends, are not practical.

Recycle and Rebirth

I focused quite a bit in this column series on the steps that need to be taken to keep your system running efficiently so that end users have productive experiences and your organization gains a positive ROI on its investment in HPC technology. But even with your clusters running at peak efficiency, the day will come – could be within a few months or could be four years – when you will start seeing the need to upgrade the system.

In the *Capacity Planning and Reporting* column, I stressed the need for a solid reporting system that keeps the HPC manager up to date regarding how the clusters are being used, and how they are performing. This is important because a system refresh, whether big or small, can't be a snap decision. Budgeting for and procuring the new resources require a long lead time and must follow a thorough examination of how the system has been used to date. Having the usage and performance history is critical to making good decisions regarding upgrades and changes.

Symptoms that indicate change is needed will come in various forms and severities. Most commonly, your business has grown and users are asking for more speed, or more throughput, or additional storage to complete their jobs. Or perhaps the software applications have evolved and need additional compute capacity or even a different software stack to run effectively. After a couple of years, hardware can simply get out of date and can't keep up with required throughput or can become significantly more expensive to operate and maintain. Associated symptoms include more frequent hardware failures, more unexplained intermittent job failures, etc.

Having the system and job usage and performance reporting system, in combination with a trouble-ticketing and change management system, will provide you the data you need to make data driven analytical decisions as to which options are the best ones to take. As always, very seldom is there only one path!

System upgrades and refreshes come in varying degrees. Just like making a decision on buying a car – do you replace the old one, or just upgrade/repair some parts, or buy a new one and keep the old one – cluster refreshes are no different. However, the decision is not to be taken lightly due to the cost involved with each of the alternatives. Generally speaking, the options are to upgrade existing system, or to expand the existing cluster by adding additional compute and/or storage, or to replace the entire system. In my experience with scores of customers or upgrades, rolling upgrades seem the most common where customers are able to keep certain production volume going while upgrading some parts of the system.

To determine which path to take, you'll need to refer to the historical reporting information I recommended you maintain from the start of operations. Knowing what hardware to upgrade or replace will be dictated by the metrics of how your system has been used. These metrics will help you make a decision that provides the best ROI, as was your original decision to deploy HPC technology in the first place.

Some important information, or metrics, that you'll have to examine include the following:

- What is the average throughput of your system?
- What is the usage profile?

- What applications run most efficiently on which architectures?
- Which nodes are used the most? The least?
- What problems are your users reporting?
- What are the costs associated with maintenance, administration, power, and software licenses?

Your job is to perform a cost-benefit analysis using these data points so you can determine which option will yield the best return to the organization. And most importantly, your enhanced system should serve the end users more effectively and reliably than ever before. Just as you did before purchasing your first cluster, you must assess your upgrade needs based on your current and expected usage profile.

Once the decision has been made on what to replace or acquire, you must crank up the RFP (Request For Proposal) process that you used to buy the original system. If possible, put the same team together – system administrators, data center operators, end users, IT staff, and executives, and engage them in the analysis and acquisition process again. If you weren't happy with the architecture or performance of your original system, you have the opportunity now to enact changes that will improve your cluster in its second incarnation. One major consideration with a new system RFP is to ensure compatibility of the new system with existing one is taken into account. This includes management systems, schedulers, software and your workflows and policies and procedures.

Even if you had a positive experience with your original HPC purchase and want to buy from the same vendor again, I still recommend sending at least two or three other RFPs out. In today's rapidly changing technological landscape, you might be surprised to find a different vendor now offers exactly what you need at a better value. Notice I said value not price! A good vendor with knowledge of your existing system can provide valuable insight into what options may or may not be a good choice.

Just as your original RFP took into account deployment of the new system, your upgrade must be installed, deployed and validated by a seasoned professional. Upgrades, especially major, require a unique skill set because they must be accomplished with minimal disruption to the existing system. If the old system will be replaced entirely, logistics can become a problem and must be handled in the RFP. All of the implementation activities add 'soft' dollars to the cost of the upgrade and must be taken into account in the cost-benefits analysis.

Finally, what will you do with the old hardware or entire system you are replacing? Depending on the age and condition, you might want to consider donating to a local college or even selling the system to another organization.

Once the upgrade or replacement is up and running, the rebirth of your HPC system is complete, for now. You'll know the refresh has been a success if you see improvements in key metrics – throughput has increased, jobs running to completion and overall operations costs coming down. And most importantly, you should see greater satisfaction and fewer complaints from your end user.

Celebrate when you are done as this process is not easy – it often feels like replacing the tires on a moving car! But is that not why you enjoy the HPC world? It is not for the faint of heart. Nice work! Now lather, rinse and repeat. Happy clustering.

If you have any questions on this topic, please contact me directly at dkhosla@x-iss.com. For information about X-ISS and our products and services, contact sales@x-iss.com.

Deepak Khosla is president and CEO of X-ISS Inc. (www.x-iss.com).