Running Scientific Data Centers
Five Underappreciated Challenges

Addison Snell

White paper sponsored by: SGI

OVERVIEW

The common hallmark of the scientific data center is the drive to accelerate research and discovery. At the same time, these sites may also be measured on a wide range of factors from carbon footprint and energy consumption to total cost of ownership (TCO), which includes not only capital costs but also the cost of space, energy, personnel and so on.

As data centers strive to add computational resources to meet the demands of their High Productivity Computing (HPC) applications, they encounter challenges that go well beyond their basic constraints. Tabor Research believes there are five primary challenges that are frequently underappreciated at the time of system acquisition, leading to inefficiencies or barriers to expansion.

1. *Space limitations:* Lack of computer room space prohibits future expansion.
2. *Power and cooling limitations:* New systems would exceed the data center's capacity to power or cool them.
3. *System reliability:* Interruptions or failures keep users from realizing the full potential of HPC resources.
4. *Configuring the right amount of storage:* Bottlenecks in data movement result in sub-optimal performance.
5. *Selecting the right architecture:* The difference between theoretical performance and realized performance can lie in choosing an architecture that matches the needs of the application.

In this white paper Tabor Research examines each of these challenges and suggests frameworks for addressing these concerns in advance, before they become limiting factors for scientific data centers.

CHALLENGES IN SCIENTIFIC DATA CENTERS

HPC is a critical tool for accelerating the scientific discovery process in a wide range of fields. Regardless of the specific discipline, managers of scientific data centers are driven to acquire increasingly more powerful systems to support advancements in research.

The drive to acquire systems with the most computational horsepower for the least amount of money has certainly led to many advancements; however, it is a line of thinking that can also lead to inefficient solutions. Managers of business enterprise data centers frequently judge solutions in terms of total cost of ownership, reliability, or their suitability to the problem at hand. Even when "time to solution" rules, there are a number of lessons HPC data center managers can learn from general IT, especially when it comes to these specific challenges that are frequently underappreciated at the time of system acquisition.

Challenge #1: Space Limitations

The price/performance of commodity components has led to an increasing proliferation of nodes in many sites. High-throughput applications can result in the collocation of large numbers of compute nodes within a single data center. As a result data center managers may find themselves up against space limitations, either with their current installations or those coming in the near future.

Expanding beyond the current compute room space is prohibitive, both in direct cost and in the time siphoned away from the data center's research. If a compute lab is out of space, there may not be other available space that can be converted, and acquiring or constructing a new building is naturally a major undertaking.

Making better use of existing data center space means looking at increasing density wherever possible, beginning at the computing platform level. Some blade platforms offer improvements in computational density over rack-mounted components. Another consideration is rack density, with the goal being more high-density individual rack units per rack.

Another way to achieve higher density is to consider which components are necessary for optimal application performance. Diskless nodes or blades can improve density when scratch space is not required; conversely, on-node disk can obviate the need for attached storage for some applications.

A final consideration for density is the choice of architectures (Challenge #5). Improving the amount of work completed per node can lead to a reduction in the total number of nodes installed.

Challenge #2: Power and Cooling Limitations

Additional nodes stress not only the space limitations of a data center, but also the capacity to power and cool the systems. In the simplest cases, power and cooling for an HPC datacenter are merely expensive, with costs potentially rising over time. In more complex cases, new system acquisitions would exceed the capacity of the datacenter's facilities. These scenarios can be effectively the same as running out of physical space. The new system cannot be placed.

Rising electricity costs and consumption rates are becoming a major concern for data center facilities. In addition to the electricity required to run the servers themselves, other devices such as backup power systems, electrical wiring, and cooling solutions drain significant amounts of power.

Facility planning is therefore crucial for dealing with electrical power issues, to anticipate potential problems with growing electrical costs and the increasing chances of disruption to energy supplies. In many cases it is important to eliminate redundant components and to avoid the temptation to overbuild the data center infrastructure. A better approach may be to choose equipment that is designed for scalability; i.e., one that eliminates components that are superfluous for scale-out HPC applications and that can expand in the future as needed – but not before they are needed – to handle growing performance requirements.

The type of power supply can also make a significant difference across a data center. High-efficiency power supplies can cut electrical power consumption by up to a third, which in many environments translates into tens of thousands of dollars saved annually in energy costs.

High efficiency is as important for cooling solutions as it is for power. Data center managers should consider up front what cooling approaches will be most efficient for their installations.

Although air cooling is still the most common approach, water cooling can be a good choice for some data centers. That said, not all water-cooled systems are equally effective. Some systems cool water in large chillers located outside the data center. The cooled water is pumped to air-conditioning units within the computer room, thus lowering the temperature of the entire room. The challenge with this approach is that it typically does not address hot spots or other variations in temperature within the data center.

A more sophisticated version of water cooling is to associate the cooling approach with the specific pieces of equipment that are generating the heat. For instance, water-cooled rear server doors are able to deliver cooling efficiency while creating minimal impact on the ambient data center temperature. Targeted water-cooling approaches control hot spots, require less fan power to move the cold air, and can enable higher rack densities than whole-room cooling.

With high-density air-cooled systems, especially blade-based architectures, the data center manager should be aware of how the cooling design can affect reliability. For example, when there is airflow through the back of a system, the data center manager should make sure all cables are properly channeled to the sides of the racks. Improper cable management can create blocked airflow, leading to inefficient cooling and more frequent failures. Additionally, configurations that contain in complex arrangements of baffles and airflow management structures that require small fans or blowers offer more points of failure, which translates to an overall loss of efficiency. (See Challenge #3.)

Challenge #3: System Reliability

Compute nodes and other components inevitably malfunction over time, and cluster installations often lack sufficient redundancy to deal effectively and quickly with component failures. Additionally, the networks that tie the cluster nodes together can also suffer from reliability problems that grow significantly as clusters scale.

Reliability is achieved through the cumulative results of myriad details, such as:

- *Cooling:* Devices running at high temperatures are more prone to failure, which lowers long-term reliability.
- *Non-redundant components:* Disks, fans, and other components with moving parts wear out. These components should be designed with redundancy and high availability.
- *Single point of failure:* Adding redundancy where possible can add to the system cost, but the result is increased reliability. Managers should consider when this investment is worthwhile.
- *Cables:* Cables that are improperly installed or maintained can bring down an entire data center.

Different systems offer varying levels of redundancy or reliability features. Managers of scientific data centers should consider the cost of system failures and downtime. Where appropriate, it is worthwhile to consider technologies that improve system availability.

Challenge #4: Configuring the Right Amount of Storage

Data bottlenecks can keep a system from running efficiently. If a data center manager's approach is to maximize computational resources with the available budget, he or she may find that the acquired systems are not sufficiently provisioned in terms of I/O capabilities or storage capacity.

The following are some points to keep in mind when configuring storage and I/O for a scientific data center:

- Many scientific applications have significant I/O load and store time at various points in the workflow. Lags in these areas prevent the processing elements from being used effectively.
- Some organizations require high availability of data, but without the cost of full replication. Technologies to consider include file virtualization or hierarchical storage management (HSM) with data migration.
- Network-attached storage (NAS) frequently works well with high-speed networks and high-performance data centers, making storage easily accessible to users on the network while also boosting manageability and scalability.
- Storage solutions with high file directory scalability reduce management complexity by enabling a large single directory that maintains scalable performance.
- High efficiency of usable storage of application data reduces the hidden cost of converting raw storage capacity.

The consideration of I/O and storage priorities can alleviate data bottlenecks and lead to greater overall system efficiency and utilization.

Challenge #5: Selecting the Right Architecture

It is a fact of data center life that not all applications run equally well on all systems. Some require large memory per CPU but a small number of CPUs, while others require many CPUs but small memory per CPU. Some applications require low-latency, high-bandwidth interconnects to scale efficiently, while others do not.

The challenge for the scientific data center manager is to select an architecture – or perhaps a combination of architectures – that will best suit the expected workflow. Parts of the system may be tightly or loosely coupled, and some nodes may have more or less memory installed. In addition, the manager should consider when specialty components, such as accelerators, may be appropriate for accelerating an end-to-end workflow.

Benchmarking can be a useful process during acquisition. It takes time and effort, but testing an application workload with actual data sets, optimized as appropriate, can tell the true story of which system will perform the best under typical scientific data center workloads.

TABOR RESEARCH ANALYSIS

There will always be pressure to compute faster, with the goal of improved speed of scientific discovery. An analysis of that goal suggests that investment in computation alone is not always the best answer. Data center managers should consider a system's reliability, its configuration, and its facilities footprint as part of the overall evaluation process in order to avoid inefficiencies with systems that are more powerful in theory than they are in practice.

When acquiring a new HPC system, scientific data center managers should consider how proper management of these challenges can lead to improvements in top-line metrics like time to solution or TCO.

- Improved system density eases the burden of system expansion and ultimately the cost of owning a system. Having more room to upgrade systems benefits long-term research goals.
- Power and cooling costs are important TCO components, and like system density, they can also be barriers to expansion benefiting future research.

- Improving reliability directly increases the computational yield of each system, because jobs complete more efficiently. Workers within the data center will also become more efficient.
- The removal of data management bottlenecks improves jobs' times to completion. In addition, architecting the proper storage solution can lead to a reduction in the amount of disks being managed or can improve data reliability, thereby reducing TCO.
- The proper system architecture not only improves time to solution, but it may also lead to a reduction in the total number of nodes, processors, or cores in the system. A reduction in the number of nodes has the direct benefit of reduced cost, and it can also lead to better computational density, reduced power and cooling costs, and improved reliability.

These challenges are often underappreciated until they become day-to-day problems that are impeding scientific progress. By planning for these challenges in advance, scientific data center managers will find that they avoid common pitfalls, and in the end, they help their users meet their overall goals in a more efficient manner.