

# Solving Big Data Problems with Private Cloud Storage

**Addison Snell**

White paper

October 2011

## EXECUTIVE SUMMARY

Data management challenges for high performance computing (HPC) applications in science, engineering, and business aren't getting any easier. In fact, the challenges are becoming more complex, as they are fueled by more data sources, coupled with greater reliance on simulation, increasing complexity of models, and an unrelenting pursuit of increasing realism in those models. Across HPC, Intersect360 Research sees an ever-present need for storage solutions that address these challenges through scalability, data throughput, and simultaneous access to the data sets.

Orthogonal to the data explosion in HPC, zeal for cloud computing has taken root across the IT landscape, where much of the emphasis has been on cloud computing, with little focus on the underlying storage requirements. But cloud-for-cloud's-sake is not the watchword within HPC, where performance, scalability, and security are paramount. Intersect360 Research studies have shown that HPC users prefer private clouds over public clouds for their applications. The differences are substantial, and storage requirements often determine whether the workload is better handled with a private rather than a public cloud. Big data applications with high data-to-compute ratios may be better suited to private cloud deployments.

For private storage clouds to be effective, they must address several essential elements, including:

*Aggregation:* Resources are pooled to leverage economies of scale in capacity and performance.

*Capacity on demand:* Cloud solutions need to be able to scale dynamically to meet sudden increases in demand without negatively affecting performance, transparency, or ease of use for all other users.

*Resource allocation:* Cloud resources need to be elastic to reallocate resources among users and groups according to demand, preference, and priority.

*Accounting:* Organizations need tracking and prioritization mechanisms to handle the sharing of resources.

Panasas, an established provider of high-performance storage solutions for demanding HPC environments, has technology and expertise for designing solutions that address the essentials of these private storage clouds. Panasas ActiveStor parallel storage appliances are specifically designed for scalability and performance, and the PanFS parallel file system provides essential capabilities for provisioning and allocation.

In search of a high-performance cloud solution, the University of Leicester (U.K.) selected Panasas over other solutions because of ActiveStor performance and its ability to satisfy these essential elements. Intersect360 Research studies show other HPC users coming forward with similar requirements.

For those seeking to implement a private storage cloud, Panasas provides scalability, throughput, and provisioning features that are aimed specifically at satisfying private cloud requirements. For those who simply want HPC storage, these same features address the increasing complexity of everyday data challenges.

## MARKET DYNAMICS

### The Data Challenge Gets More Challenging

Like successful small businesses, storage infrastructures grow in complexity as they scale. A bread maker may begin with two ovens, a couple of recipes, and a handful of ingredients, but as more and more hungry customers queue up at the door, things get more complicated. First of all, the baker needs more of the things she already had – more ovens, more ingredients, and perhaps more recipes to satisfy increasing demand. But with the increase in scale also comes increasing complexity in operations, including inventory, throughput, cost, and the utilization of resources. Many entrepreneurs soon find that managing a growing business – a bakery, a floral shop, a hair salon – is more about general business management than it is about baking, arranging flowers, or styling hair.

The scalability of storage solutions isn't so different. With organizational success comes continual data growth (the equivalent of more ingredients for bread), which implies not only the increase in raw capacity (more ovens) but also the implementation of better data movement strategies (how to efficiently make more tasty bread). Especially within High Performance Computing (HPC), growing complexity is the hallmark of expanding ranks of data-intensive computing applications.

Many HPC applications face continually accelerating data management challenges. Some of the main pain points include:

*Managing large data files:* Emerging “big data” applications present enterprises with very high data-to-compute ratios. For example, in the oil and gas industry, seismic processing clusters may need to ingest exploration data totaling hundreds of terabytes with single files often reaching multiple terabytes in size.

*Managing many files:* Applications such as genomics and protein matching tend not to have large databases, but rather they must simultaneously manage large numbers of databases and files that arise from next-gen sequencing and more complex instrumentation.

*Simultaneous access to files:* Teams of engineers may work collaboratively on a single design, requiring file locks and version control.

*Long-term data management:* Medical and financial markets face increasing regulation that includes requirements for long-term data management. Storing large amounts of data must be done in a reliable and retrievable fashion.

These challenges become ever more difficult to manage as the data-to-compute ratio grows. The nature of scientific computing is that once a problem is solved, it is no longer interesting. You don't need to design the same chemical, protect the same investment, or discover the same oil twice; you move on to the next, more difficult challenge. This fuels several data growth trends:

*Increasing complexity:* There is a continuous desire within engineering markets to strive for greater realism in models or more degrees of freedom in simulations.

*More simulation:* Various industries continue to increase their reliance on data-driven simulation over physical experimentation.

*Greater fidelity:* Oil exploration companies have incorporated wide-azimuth (WAZ) technology which captures data at higher fidelity, and new generations of instrumentation are generating more realism (and consequently more data) in a wide range of exploration applications.

All of this is enough to fan the flames of data explosion in markets that have historically experienced high data growth rates. Beyond these markets, these same challenges are rapidly emerging in expanding “big data” applications, based on the new prevalence of data sources, many of them real-time sources that provide the basis for predictive analytics. While much of the discussion is based on the software algorithms to analyze this data, it cannot be ignored that all of this big data also needs to be stored in a way that ensures its transfer can match the I/O requirements of the applications.

Intersect360 Research sees storage for HPC as a \$4 billion market in 2010, and with a 9.2% compound annual growth rate through 2015, it is forecasted to be the fastest growing product segment, despite advancements in technology that continue to reduce the cost of capacity and data management.<sup>1</sup>

The goals of a big data storage system include:

*Scalability:* Storage scalability includes both capacity (more bytes) and capability (fast access). It is essential that the storage system can scale with the ever growing storage demands. But simply adding capacity without maintaining the ability to deliver data according to users’ requirements can diminish performance, as traffic overwhelms the total I/O bandwidth. This problem is exacerbated as the data-to-compute ratio increases.

*Throughput and data access:* Efficient movement of data to and from the storage system is the key to performance. Underperforming I/O rates can easily throttle total performance, leaving compute and network resources idle. In addition, data access times must be fast and unconcerned by where the data resides in the storage infrastructure.

*Data sharing:* The ability to efficiently share data across an organization is vital to increasing productivity. Independent silos of information create barriers and inefficiencies within an organization. Conversely, the ability to prevent over-sharing, or protecting data that is shared on a limited basis, is also important.

One of the key challenges that stand in the way of many of these goals is data fragmentation. Fragmentation happens when storage infrastructure must grow to support user needs because the scope of the initial storage infrastructure was too local, too isolated, or both. Fragmentation impedes the ability to scale storage capability, reduces overall performance, impacts data access times, and often requires additional administrative overhead. Efforts to unite multiple local storage islands with *ad hoc* solutions may fail to address scalability, throughput, and data sharing, and can actually exacerbate the problem.

A typical HPC resource normally has two types of storage, basic user storage and separate high performance storage. User storage usually consists of Network File System (NFS) mounts on the login and compute nodes that are used to hold users’ /home directories. These /home directories are often local to a cluster and may not be directly accessible from the user’s workstation.

Many clusters also have some form of high performance parallel storage. Unlike a single gateway NFS storage array, high performance file systems operate in parallel and provide a much wider I/O pipe for parallel

<sup>1</sup> Intersect360 Research, “2011-2015 Total HPC Market Forecast,” April 2011.

applications (i.e., those applications which span multiple nodes in a cluster). All data in a typical NFS file system passes through a single NFS server and can cause a bottleneck for parallel applications. By contrast, a parallel file system (such as Panasas PanFS) allows simultaneous access from compute to storage nodes, delivering higher throughput and performance.

Fragmentation may continue beyond a single cluster and can result from the growth of multiple clusters within an organization. Thus, it may be possible that users and administrators must contend with multiple /home directories and multiple high-speed file systems, all under different name spaces. These multiple storage domains can further complicate administration due to data backup and data movement issues within an organization. Faced with mounting challenges, many organizations are looking for more consolidated solutions for their storage environments.

## **The Private HPC Storage Cloud**

Cloud computing, a red-hot industry buzzword, is touted as a revolutionary approach to IT solutions that has the power to free the technology user from the expense and headache of resource management. But as with any major new industry buzzword, there are different definitions that essentially represent different levels of implementation – no definition is inherently right or wrong, each represents a different perspective.

What all clouds have in common is the notion of virtualization – an abstraction layer that hides the complexity of all or part of an IT infrastructure or workflow. The end user accesses a resource through a web interface (or web-like, in the case of some specialized apps or intranets), and the user may pay for his usage in some fashion, but the user is ambivalent to or unaware of the actual manifestation of that resource – it lies elsewhere. It is somebody else's problem.

Within that loose construct there are many ways in which clouds differ, most notably in the distinction between public and private clouds. Intersect360 Research provides this analogy: Comparing private to public clouds is a bit like saying both Americans and Europeans play football. There are high-level similarities (both are team sports that involve running around with a ball on a grassy rectangular field) and some overlapping terminology (“The referee ruled the player offside.”), but there are critical differences in implementation.

The differences between public and private clouds come down to four main toggles: ownership, possession, management, and provisioning. That is to say, if the resource in question is owned by someone else, at someone else's location, managed by them, and they decide when it is available, then it is a public cloud. In the extreme, someone may view “the Cloud” as a singular, ethereal, capital-C resource, like “the Web” or “the Internet.”

By contrast, a private cloud is one in which the resource is owned, possessed, managed, and provisioned all within the end user's organization; the end user is just disconnected through a virtualization layer. Hybrid clouds blend the two models.

A recent study by Intersect360 Research revealed that private and hybrid cloud models have greater interest within the HPC community than public cloud models, particularly among sites with larger budgets and greater investment in existing infrastructure. Overall, 23% of responders are currently using some form of cloud system to support their HPC workflow, while 48% of survey sites are actively evaluating cloud computing services.<sup>2</sup> The

---

<sup>2</sup> Intersect360 Research, “Cloud Computing in HPC: Usage and Types,” June 2011.


magnitude of this response suggests that cloud computing has moved beyond its inventor phase and is now in its early adopter phase in the HPC marketplace.

The preference for private clouds is directly related to the goals and challenges cited above for data-intensive computing markets; i.e., organizations are seeking to improve performance and scale capacity while dealing with increased capital and organizational constraints. These are amplified by security concerns as regarding public clouds. This pertains to the nature of many types of HPC research, in which regulation (as in financial services or medical research) or competitive pressure (as oil exploration) or intellectual property (as in manufacturing) or national security, compels organizations to go to lengths to ensure data cannot be leaked by malice or happenstance.

Across all sectors of respondents, bandwidth and security were ranked as the top two barriers to cloud adoption for HPC. (See Figure 1.) The ability to get data in and out of the cloud, and to ensure its absolute protection while it is there, are the hallmark challenges of running HPC applications in any cloud, most of all a public one, in which the data leaves an organization's own control, limited to internet speeds.

**Figure 1. Percent of Surveyed Respondents Rating Barriers as “Significant”**  
*(“Significant”: Rating of 4 or 5 on five-point scale)*

Source: Intersect360 Research “Cloud Computing in HPC: Barriers,” September 2011



The sample from Figure 1 includes those not using, those evaluating, and those using clouds, and shows that overall 57.4% see bandwidth as a significant barrier, and 55.6% see security as a significant barrier. If those who are already using clouds are omitted, the barriers are higher – 61% for both bandwidth and security – as of course the early adopters of clouds were those who saw fewer hurdles. These barriers tend to be multiplicative. That is, each additional barrier further reduces the likelihood or ability of the organization to adopt a public cloud solution.

The Intersect360 Research report, “Cloud Computing in HPC: Barriers,” published September 2011, describes all of the cited barriers in more detail. In particular, about bandwidth and security, we report in part:

*“In multi-tenant environments (i.e., public cloud systems), security risks seem numerous, and it will take time for cloud service providers to build up trust. ... Security considerations add risk and potential loss factors into the cost benefit equation of cloud usage. In many cases, particularly in the commercial market sector, the cost of a security failure could quickly exceed the benefits of cloud by many orders of magnitude. Thus, even very high levels of data and application protection within a cloud may prove mathematically inadequate when compared to overall intellectual property cost, or loss of competitive advantage costs. ... In addition, security can be a structural issue within organizations, that is, something that will not be approved because of established organizational practices or corporate culture.”*

*“The bandwidth challenge leads to problems in two areas. [First,] cloud configurations are needed that can efficiently handle run time communications and data movement, so that jobs can complete in a timely manner, that is, the cost savings per processor hour are not lost while CPUs are idle waiting for data. ... [Second,] moving both the application input data into the cloud and the results data back from the cloud can prove impractical. Although bandwidth is a problem that can almost always be solved by throwing money at it, such approaches are often very expensive. ... Thus, when networking costs are added into the overall cloud proposal the cost advantages of cloud can dramatically shrink or disappear.”*

In summary, many HPC applications have high “data-to-compute” ratios, which require a high rate of data ingestion to execute the desired analysis or simulation. These applications are more constrained by available I/O bandwidth than by computational cycles. For applications like these, public clouds that emphasize available compute cycles can be a poor fit, since they solve the wrong problem while exacerbating the bandwidth issue. Thus bandwidth-constrained applications may see unacceptable delays in performance or increases in cost in public cloud environments, relative to other workloads.

## Essentials of a Storage Cloud

Within the definition of clouds – even private storage clouds – there is room for different calibers of aspiring solutions. However, there are several key criteria that are often seen as essential in the creation of a cloud that addresses the intended issues.

*Aggregation:* Part of the efficiency of cloud computing comes from the economies of scale in resource consolidation. For maximum efficiency, resources are pooled, so that they can then be allocated on an as-needed basis. This is essential for efficient scalability and allocation.

*Capacity on demand:* For any cloud there is generally the assumption that additional desired capacity can be accessed (at some cost). For private clouds, this implies the need for infrastructure scalability – the

ability to add capacity to the storage infrastructure without negatively affecting an individual user's performance, transparency, or ease of use.

*Resource allocation:* Cloud resources need to be elastic in order to allocate them among users and groups according to demand, preference, and priority. However the accounting or decision making is handled, rapid provisioning is seen as an essential cloud element.

*Accounting:* Allocation prioritization in many systems is reflected by dynamic cost structures, but in research environments that do not employ chargebacks, the matter is settled by prioritization heuristics, even if they are as simple as first-in, first-out. In any case there generally must be some mechanism for tracking cloud usage.

Private cloud implementations in an HPC environment have unique challenges beyond other cloud implementations because of end-user performance and scale requirements. Flexible support often needs to be provided for both high performance (i.e., parallel) and high throughput (i.e., NFS) file access. Performance must be scalable and support a global namespace so that resource pooling and rapid elasticity can be transparently supported.

Private clouds can ameliorate the bandwidth and security barriers typically seen in public cloud models. Furthermore the HPC community's interest in private clouds can be viewed as an extension of its interest in existing resource sharing architectures. That is, many organizations have already adopted some form of private grid architecture (both compute and storage), which now tends to fall under the private cloud model. In addition, modern workflow schedulers that provide on-the-fly provisioning of HPC resources to various users and departments, also behave like private clouds.

## PANASAS PARALLEL SOLUTIONS FOR HPC CLOUD STORAGE

Traditional IT infrastructure has struggled to meet the performance and scale required by HPC environments, especially with regard to the explosive growth of big data, largely driven by high performance technical computing. Cleaning up data fragmentation is not an easy task, and attempting to support and deliver various levels of storage is time-consuming and expensive. In addition, allowing end users to design local solutions to a global problem leads to compartmentalized, underutilized islands of HPC compute and storage clusters.

Moving to a private cloud architecture can allow a company to provide flexible storage models for its HPC environment. The goal of the private storage cloud is to move the administration of the storage array out of the end user's hands and into a managed environment, allowing scientists and researchers to focus on their science, while providing the flexibility to meet their specific performance and storage needs.

One approach to creating a private HPC cloud is to deploy Panasas scale-out storage technology. Panasas focuses on scalable data management solutions for HPC and has a long history of supporting high-throughput environments. Panasas technology allows the creation of private storage clouds that provide end users and administrators with many of the essential private HPC cloud features mentioned above.

Panasas ActiveStor systems are designed with HPC cloud implementations in mind, offering parallel performance and scale for technical computing applications and big data workloads. ActiveStor addresses scalability challenges with up to six petabytes of storage in a single global namespace. With its massively parallel architecture, Panasas has demonstrated performance that scales linearly to 150GB/second, providing scalable throughput within a shared user environment.<sup>3</sup> (See Figure 2.)

To deliver an effective HPC storage cloud, the Panasas PanFS parallel file system divides files into large virtual data objects and transparently presents the parallel file system through Ethernet or InfiniBand. Administrators can provision additional resources in the storage pool after adding storage capacity, without system or service disruptions. The Panasas system dynamically expands the storage pool as more storage resources are required, and additional capacity can instantly be available for use in the global namespace. In addition, the system automatically re-balances the existing data across the newly added capacity, ensuring that performance is not bottlenecked.

The Panasas HPC parallel storage system supports a variety of protocols important to private cloud implementations, including: NFS, Windows-based CIFS, and its own Direct Flow parallel client, which is being standardized in NFS v4.1 as part of Parallel NFS (pNFS), the emerging parallel I/O extension to the ubiquitous NFS standard.

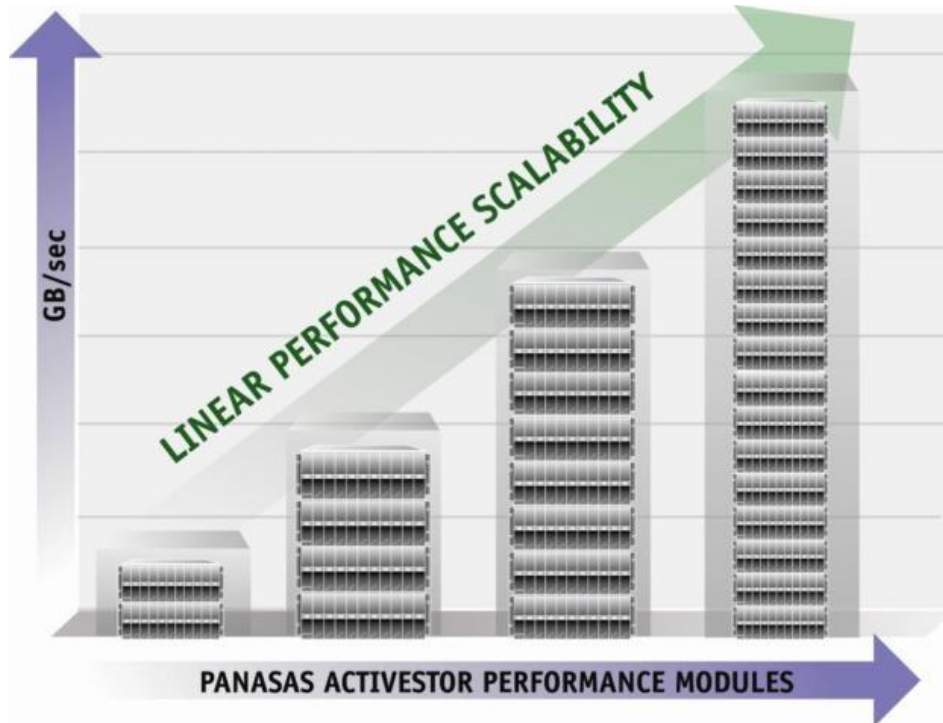
Finally, Panasas gives individual private HPC cloud consumers a clear feedback loop on resource usage. Capacity utilization and performance attainment are monitored and reported back to stakeholders and administrators. Panasas ActiveStor provides guaranteed user quotas at a departmental level that can be dynamically resized according to user needs. This allows administrators to manage budgets according to actual storage used, rather than arbitrary accounting allocations.

---

<sup>3</sup> Scalability testing performed by Enterprise Strategy Group, <http://performance.panasas.com/wp-esg-web.html>.

**Figure 2. Linear Scalability of Capacity and Performance**

Source: Panasas



Panasas designs HPC storage systems that scale in both capacity and performance, leveraging PanFS to present virtualized scale-out storage to HPC compute environments. With its scalability and management features, Panasas technology addresses the essential aspects of private HPC storage clouds.

## CASE STUDY: UNIVERSITY OF LEICESTER “ALICE” STORAGE CLOUD

The University of Leicester was faced with a common problem — fragmented HPC resources across the greater organization. Most academics were used to owning their own HPC capabilities, or sharing departmental facilities among small groups of colleagues. And a good number of academic HPC users did not work in traditional “computation heavy” disciplines such as physics, astronomy, chemistry, and engineering, and were thus not well versed in the use of clusters and parallel file systems.

To help reduce both fragmentation and the amount of overhead required for its myriad HPC systems, the university decided to deploy a centralized HPC solution and selected a Panasas ActiveStor system for its cloud storage. They named the overall solution ALICE, for Advanced Leicester Information and Computational Environment. By consolidating the university’s high performance storage and compute budget into a centralized private cloud, the university says it avoided costly over-provisioning and as a result it can better utilize its resources.

Leicester’s current facility is now far bigger than any individual department could support and provides its users with peak capacity and capability that would not otherwise be available. The system employs 256 compute nodes, two login nodes, and two management nodes, all connected by an InfiniBand network fabric. Each computational node offers a pair of quad-core 2.67GHz Intel Xeon X5550 CPUs, providing 2048 CPU cores for running jobs.

In addition to a centralized compute resource, the university wanted a centralized storage resource that provided a single global name space across all departments. In addition, the file system had to support the high performance needs of the cluster. They decided the best solution was a private storage cloud utilizing Panasas ActiveStor, a choice that was based on the results of a comprehensive suite of benchmarks developed by the team at Leicester.

The benchmarks were designed to run with varying degrees of parallelism and on different numbers of nodes, ranging from running one job on a single node to running eight jobs on each of 32 nodes. During the procurement phase two storage solutions utilizing parallel file systems were shortlisted for benchmarking, Panasas ActiveStor, a high performance parallel storage system running over 10GbE and InfiniBand, and an open source Lustre-based solution employing commodity hardware.

Leicester reported that as benchmarks were scaled across several nodes, ActiveStor demonstrated dramatically higher parallel file system performance than the Lustre-based system. This performance, coupled with the ability to support departmental NFS needs, provide rapid elasticity, and offer user quota and chargeback accounting, confirmed its choice of a private HPC storage cloud provided by Panasas. Consolidating its HPC system into a centralized cloud infrastructure allowed the University of Leicester to achieve significant cost savings while dramatically increasing the peak performance capabilities of individual departments.

## INTERSECT360 RESEARCH ANALYSIS

Cloud computing is a popular, dynamic trend across the IT landscape, but it cannot be simply assumed to be a solution to all problems. Building an HPC storage cloud means overcoming significant hurdles regarding capacity and performance scalability, along with resource allocation and data management requirements, in order to be an effective solution. Among other reasons, HPC users tend to prefer private clouds to public clouds for applications with high data-to-compute ratios, because the internet often cannot handle the data transfer necessary for their applications.

In many senses, the private cloud is the simple evolution of clusters and grids; private clouds are straightforward extensions of best-of-breed HPC technologies for allocating and managing data and computational resources. Panasas is therefore well-positioned to provide storage solutions for HPC private storage clouds. The company has established itself as a provider of high-performance storage systems at extreme scale, thereby demonstrating capabilities in scalability, throughput, and resource allocation, all of which are critical in assembling effective high-performance storage clouds.

In fact, it seems clear that had the term “CLOUD” never caught on, Panasas would still be providing these same solutions and capabilities. This could lead to a cynical observer accusing Panasas of simply slapping a cloud label on their products. But there is an important distinction between Panasas and the myriad companies that have begun to wave the cloud banner with whatever technologies they may possess. Rather than Panasas attempting to bring technology to a trend, this is more a case of a trend coming to Panasas. Intersect360 Research studies have independently confirmed end users’ desire for private clouds, along with their requirements for such solutions. These requirements match up well with Panasas technology, products, and capabilities.

This speaks well to the opportunity for Panasas. Like “grid” before it, the term “cloud” may fade in popularity, but fundamental drivers – increasing complexity of models, the desire for greater accuracy, and the increasing role of simulation – will continue to fuel data growth in HPC for years to come. Panasas will find users today who want to use their technology for private storage clouds, but the technology can continue to be successful in the future under any new buzzword. By one name or another, private clouds will be here to stay.