

## Panasas: High performance storage for huge amounts of data

July 8, 2019 (Corina Pahrman)

**Starting September 2019, “ActiveStor Ultra”, the new storage system from Panasas, will be available. Beta tests are currently underway at an airline, an energy company and a university. At ISC 2019, our author Corina Pahrman spoke with Praneetha Manthravadi, Director of Product Management and Jim Donovan, Chief Marketing Officer, about AI trends and the specifics of the new high-performance storage system.**

■ *At ISC19, Panasas introduced the new ActiveStor Ultra storage system for the first time. What does it offer and why should it be used?*

**Donovan:** Looking at the market for high-performance storage, we see three categories, each of which offers certain benefits but also requires compromise. Open source file systems like *Lustre* or *BeeGFS* are convincing due to their high performance and relatively low costs. You can combine the file system with commodity hardware and customize it to your specific requirements. This is why cutting-edge projects often rely on this environment. However, these systems are difficult to manage, and they are very fragile: they can perform a particular task very well, but if you change them a little, you have to reconfigure them.

The other extreme is enterprise-class storage in data centers: it is user-friendly, but it delivers less performance and does not work in parallel. We combine the best of both worlds with HPC (High Performance Computing) and the ease of use and reliability of the enterprise-class. We also call it Enterprise HPC or Commercial HPC. What’s special about us is that it is really is a plug-and-play file system.



Praneetha Manthravadi and Jim Donovan presented the high-performance storage system Panasas ActiveStor Ultra at the ISC (Image: Corina Pahrman).

## HPC and Plug-and-Play

■ *What exactly does plug-and-play mean in this context? How does ActiveStor Ultra fit into existing and, more importantly, future systems?*

**Manthravadi:** Our credo is the so-called Panasas Experience: from the delivery of the hardware to the work in the data center to the end of the product lifecycle, we want to ensure a high level of customer satisfaction, regardless of the size of the system. ActiveStor is designed to balance itself even when the workload changes. This way you can start small and add more tasks later. The performance is immediately available without you having to do anything: “Set it and forget it”. When you add a node, the system automatically detects it and asks if you want to integrate it.

**Donovan:** The system scales performance and compatibility independently. Currently, we have eight storage units in one rack, and if the customer needs another, they simply plug it in. The plug-and-play functionality does the rest: The IT manager just inserts the new storage and clicks “recognize it” or “detect”.



Panasas “ActiveStor” is a clustered scale-out NAS appliance designed for speed and reliability.

## Application scenarios for high performance storage

■ *Which application scenarios did you have in mind during the development phase?*

**Donovan:** On the one hand, the very traditional HPC such as research and science or for use in government agencies. On the other hand, we also target manufacturing companies, the life sciences, the energy industry and the media and entertainment industries.



Jim Donovan  
Panasas

**Manthravadi:** We are targeting customers who run large HPC simulations and have large clusters of compute nodes linking them together. However, we scale from very few to thousands of nodes.

■ ***In which areas do you see an increased need to catch up, which sectors are not yet taking advantage of the technical possibilities?***

**Donovan:** I think life sciences offer the best opportunity for using these technologies sensibly and productively. Today we talk about personalized medicine or precision medicine, and in order to develop these, it is necessary to combine several disciplines and immense amounts of data. However, the necessary IT infrastructure is far behind. So science has made a leap forward, but the technological infrastructure is not compatible. One of the biggest challenges is finding people who understand more than just technology or science. We need technicians who also understand science.

■ ***Do these research branches already collect enough data?***

**Donovan:** Yes, one of our clients is the *National Institute of Health* in the U.S. They have the largest biomedical library in the world, and thousands of researchers worldwide are working with their data. Then there are technologies like *CRISPR*, which has to do with editing genes - and which in turn generate massively complex data. This has great potential, but it is often slowed down by the lack of the ability to process all this data efficiently.

**Manthravadi:** In general, the amount of data grows faster than the methods used to process it. In the future, it's all about high-performance systems that can handle huge amounts of data.

## HPC solutions for AI projects

■ ***Why are HPC solutions so central for the development of artificial intelligence (AI)?***

**Manthravadi:** Let's look at the field of training as part of deep learning. If you want to train neural networks for image recognition, you need to feed many random images. And your system should be able to process this amount of data with very little latency in a very short time. HPC file systems can do that, providing the necessary scalability, high bandwidth and low latency.



Praneetha Manthravadi  
Panasas

■ **What trends in artificial intelligence do you see for the next three to five years?**

**Donovan:** Machine-learning, deep-learning and the training of neural networks: this is where the great potential lies. In order to acquire these technologies, companies should work on a single project - be it in the cloud, be it for a small product - and then scale it to a corporate project. Much of the AI happening today is at Amazon, Facebook, Google and Microsoft. These people have the money and the talent. The challenge is to make AI technologies available to a much broader audience.

**Manthravadi:** We are also currently working in many different frameworks. These processes will be standardized, and one or two tools such as *TensorFlow* or *Caffe* will dominate. Currently, the companies are still experimenting and trying to find out which tools they are making good progress with.

###