



To read the original German language article, [click here](#).

January 17, 2018 – StoraInsider.de –Kriemhilde Klippstätter / Rainer Graefen

Scale-out NAS for Professional Users

Panasas' HPC storage brings parallel data access to the media industry and Industry 4.0

For many years, Panasas has been serving university facilities or government research organizations, such as the Alamos National Lab, with their "Network-Attached Storage" (NAS) devices for High-Performance Computing (HPC). The people in charge at Panasas believe that the market is now ripe for commercially used HPC and have introduced a new NAS family.

Panasas was founded in 1999 by Garth Gibson, who made a name for himself in 1980 as co-author of the "Raid Papers" published by the University of California, Berkeley. In 1995, at Carnegie Mellon University, in Pittsburgh, Pennsylvania, NASA's Network Attached Storage Device (NASD) storage began, with Gibson becoming the first director of the Parallel Data Laboratory.

In 1999, together with William Courtright in Pittsburgh, he founded the startup Panasas with money from venture capitalists. A year later, the office was opened in Sunnyvale, California, today the headquarters of Panasas. It took nearly five years to deliver the first scale-out NAS product to the Los Alamos National Lab.

Parallelism When Saving and Retrieving Data

In order to adequately store the large amount of data generated by high-performance computing, Panasas has developed a parallel file system (PanFS). It allows users to access their data directly - in parallel.

The NAS head commonly used in commercial storage to determine performance has been eliminated. And thus, eliminating the bottleneck between computer systems and RAID storage. Panasas storage also separates user data and metadata, the latter are stored on fast Directors. That way, the high-performance storage units achieve the best metadata services and thus fast backup and recovery of user data.

The HPC "ActiveStor" storage system is divided into Director and Storage nodes. The Directors may be referred to as the controlling authority of the Panasas Parallel File System (PanFS). They process all metadata such as directories, file names, timestamps or access control. They also coordinate user actions to ensure the most up-to-date view of data and coherency in the data cache.

A Kingdom for a Storage Cluster

Another important task of the Directors is to look after the storage cluster - Panasas calls it "realm". These include tasks such as automatic capacity distribution, file recovery, or automatic troubleshooting in the storage cluster.

The storage nodes allow users to access or store huge amounts of data in parallel. In addition, the metadata are stored on the storage units, however, they are not processed there but in the Directors. The API is based on the definitions of the T10 SCSI-Committee for object-based storage devices, which were co-developed by Panasas and, in particular, by company founder Garth Gibson.

The heart of the parallel storage is - in addition to the file system - the "DirectFlow" protocol of the company. It supports all POSIX compliant languages. The proprietary protocol runs on every

user's computer system and provides parallel read and write cycles. It also supports Linux and MacOS systems. Panasas is proud of having co-developed the parallel file system "pNFS", which is based on the in-house DirectFlow.

Data Security and Ease of Use

For performance reasons, each file is stored in a subset on storage nodes. Each file is protected by its own distributed erasure coding. Therefore, in case of an error, the entire drive doesn't have to be restored. Snapshots allow users to access earlier versions of a file. The "always on" architecture promised by the company refers to the "Extended File System Availability". In case of a major failure, if several files are lost, the path names are displayed to provide an overview of the lost files.

Despite all the subtleties, the storage system is said to be simple to handle and ready to install in just one hour. New storage nodes are automatically detected and simplify the expansion. The entire cluster is managed by a graphical user interface. If errors occur, an automatic system for troubleshooting intervenes. Capacity and load balancing is also managed by the ActiveStor Directors.

Commercial Applications Are Now Ripe for HPC Storage

With its new storage system - ActiveStor Director 100, the matching storage node ActiveStor Hybrid 100, and the new version 7 of the PanFS file system – Panasas is ready to serve users in a commercial environment. "The growing number of small files that need to be stored present immense problems for the metadata management of traditional NAS storage," says Panasas CEO Faye Pairman. Another is the replacement of cascaded storage (multi-tier storage), whose workflow and management are rendered inefficient by the new challenges.

In addition to existing customers from the field of university research and government agencies, she sees the new ActiveStor system as particularly suitable for media and entertainment, life science and manufacturing. Industry 4.0 may well require a new memory architecture in the future as well.

###