

Dr. Jens-Peter Akelbein
IBM Deutschland Research and Development GmbH
Manager Storage Systems Architecture
Hechtsheimer Str. 2
55131 Mainz
06131 84 6064
Email: akelbein@de.ibm.com

From open storage solutions to an integrated hardware appliance

In the past years large storage solutions are implemented in more and more application areas as clusters in a scale-out architecture. Starting with high performance computing applications in the 90s the first scale-out storage products came on the market in the last decade. For file based storage or so called unstructured data solutions require a highly scalable clustered file system. Other than classical network attached storage (NAS) filers scale out NAS products provide seamless change of capacity and bandwidth according to the data growth generated by new computing workloads. Furthermore, known IP-based file protocols (NFS, CIFS, FTP, HTTP, ...) over the wire are being used for sharing the data across a very large number of clients transferring many gigabytes of data per second and providing many 100,000 I/O operations per second.

In the past years the storage software development team in Mainz developed the IBM Scale Out Network Attached Storage (SONAS) being launched as an enterprise program product in 2010. At the beginning open solutions were offered using the IBM General Parallel File System (GPFS), Linux distributions as operating systems, and open source software like the Samba project. Over time it evolved into an integrated hardware appliance consisting of up to 90 nodes and providing a capacity of 14,4 Petabyte. Heterogeneous storage can be used so different hard disk types and tape libraries. They build storage pools where data can be placed and moved based on policies.

Areas of development for an integrated scale out NAS appliance are clustering NAS protocols including node failover, user interfaces for configuring and managing the system and its functions, facilities for logging and tracing functional and non-functional characteristics, operations for changing the hardware structure, disaster recovery procedures, and facilities to provide remote support including integration into error reporting systems.

Integrated appliances are mostly known in the area of networking and storage today. Taking SONAS as an example demonstrates design principles being used for appliances and needs for enhancements.