

# STORAGE SWITZERLAND

## IS YOUR FILE SERVER CHOKING?



George Crump, Senior Analyst

Dealing with data is a top priority in 2012 for most IT organizations. The challenge in dealing with this problem is that most growth is in the area of unstructured data (data not in a database). And this isn't just user files, although those are growing too, it's also machine generated data which is quickly eclipsing that created by humans.

The era of Big Data is adding this machine-generated data such as medical records, research data, or other unstructured information in data objects of varying sizes, from 1K to 1TB or more. These files must often remain accessible for continued use, analysis, compliance or regulatory purposes. This presents a challenge for IT which now needs to scale storage for both capacity and file count, while maintaining accessibility for an indefinite amount of time.

These new requirements of almost unlimited capacity, file count and retention times may mean the traditional file server model has passed its prime. File servers can deal with moderate storage capacity growth but have limitations in file count and continued accessibility. This can result in application slow-down, reduced user productivity and more importantly, possible external customer dissatisfaction.

To combat the problem of a broken file server, IT managers have tried a variety of storage solutions including both scale-up and scale-out systems. But the fundamental problem that's often never addressed is the file system itself. Trying to address the problem with faster and more powerful legacy solutions leads, at best, to temporary and expensive fixes.

Do we really need a file system?

A traditional file system and/or NAS architecture has three layers. There is an application layer that manages login, authentication as well as group and user access. Then there is a control layer that manages access control, authentication matching, directory management, and group and user levels. Finally, the file system communicating with the attached disk via SCSI, essentially shreds a file into multiple blocks and scatters those blocks across the storage volume. The NAS or file system then creates a series of pointers called "inodes" that map the file system to the various blocks on the volume that contain file information.

As a storage system grows and the number of files stored on it increases more inodes need to be created and managed by the file system. Inode count will also increase as various storage services are enabled on the NAS, the most typical example being snapshots. This means that file systems and NAS systems will see their performance degrade as the number of files increases and the file system reaches capacity, something which is especially problematic when dealing with smaller files. As a result, performance can deteriorate quickly as file counts increase, even though there would appear to be plenty of excess capacity. This can lead to a situation that's difficult for IT to explain.

Historically, files were relatively large in size and as file counts started to reach the level where performance was noticeably impacted, the file system was more than likely running out of capacity as well. It seemed logical then that the performance impact was the result of this capacity issue, not caused by running out of inode space. However, the varying file sizes of today's corporate data and the introduction of machine-generated data are leading to performance issues long before capacity limits are reached. IT is left to explain why more capacity is needed to address performance and responsiveness issues and not a capacity problem. The inherent problems in the legacy file system make IT look bad.

Finally, many file systems simply have a limit on the number of inodes they can support. On paper this number may look so large that it's unreachable. But the combination of varying file sizes and the inefficiencies of file systems that use potentially thousands of inodes per file can reduce the effective maximum file count to a very 'reachable' number.

## Object storage

A solution is the object based storage system. Storage systems built on object based architectures have been available for several years and are commonly deployed in

cloud environments, which were among the first to incur the problems caused by high file counts. Historically, data centers have tended to shy away from object based storage since it used to require significant changes to the application, as opposed to a traditional file system storage method. However, the efficiencies object storage systems deliver, such as offering ease of scale to multiple petabytes without limitations in file count, should make data center managers take a closer look.

Companies like [Caringo](#) are making it easier for IT to implement object storage. These solutions offer an integrated application and access layer that can be accessed from a file system or directly via HTTP while leveraging an object based storage backend that performs similar functions to a typical NAS-based system. Unlike the traditional file system used on a NAS or file server, instead of shredding the file into thousands of blocks and creating a corresponding number of inodes, each file is stored as an object or group of objects with metadata and gets a Universally Unique Identifier (UUID). This UUID ensures rapid access to a file without needing to crawl inodes producing system responsiveness regardless of capacity or object count.

In addition, objects are stored with metadata that can be used to automate many management tasks including self-healing and self-optimization functionality. When combined with replication the need to backup is often eliminated. This is especially useful for organizations with large data sets that can't backup due to exceedingly long backup windows.

Performance in the object storage model is designed to be consistently responsive for massive numbers of files and users even as storage soars into the petabyte scale. This is achieved with x86 hardware delivering commodity or "cloud" economics. There is no need for specialized high transaction systems such as those used for top-end IOPS systems running databases.

The power of the cloud internal to the data center

An object-based storage backend combined with an access and protocol front end provides data centers with the best of both worlds. They now have a file-system-like environment that can grow to virtually unlimited numbers of files while leveraging other capabilities that cloud storage infrastructures are known for. These include the use of generic systems and storage hardware, a single point of management, automated storage policy engines as well as global replication capabilities. The result is a highly reliable, highly scalable environment that has a single management point with automated response capability - all while being extremely cost-effective with consistent performance and responsiveness at scale.

The reality is that the variation in file sizes will continue and file accessibility requirements will only get more demanding as just about every device you interact with on a daily basis starts to connect to the Internet. It is clear that file servers will not be able to cost effectively solve this issue and are being overcome by the current swarm of files. Integrating an object storage solution like Caringo CASTor can help solve these issues and help you fix the cause of the problem and not the symptom.

### **About Storage Switzerland**

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

*Copyright © 2011 Storage Switzerland, Inc. - All rights reserved*