



To unlock the power of Cloud Storage

- File Versioning vs. Locking

By Colm Keegan, Senior Analyst



Private or hybrid cloud file sharing provides many business benefits over traditional on-premise network attached storage. It enables organizations to consolidate infrastructure and remove duplicate silos of storage from multiple sites and office locations. Cloud file sharing reduces storage costs, simplifies storage management and reduces risk as data can be centrally managed and protected. It also helps improve user productivity and enables collaboration as data is readily accessible regardless of where the end users are located. However, ensuring that files can't be accidentally overwritten or deleted is critical to the integrity of the data and the success of a cloud based file sharing initiative.

Ensured Data Integrity

Prior to the advent of cloud storage, businesses typically relied on local NAS systems to share, manage and protect file content. NAS systems provided a critical component to ensure data integrity when files are shared out across multiple users - built-in file locking services. When using NAS file shares, only one user at a time is granted write permissions to make edits to a file. This file locking prevents file or data corruptions issues from occurring.

With multiple office locations, the cost and complexity of managing dozens of physical NAS systems can be significant. Moreover, sharing files and protecting NAS data across multiple offices requires the use of mirroring and replication technologies. These solutions also require additional "landing space" at each site for storing the replicated data, which drives up the total cost of ownership.

To expand file locking services across globally distributed NAS systems without adding the burden of managing multiple NAS platforms, global file system technology can be employed. This technology enables files to be shared across multiple geographies without risking data corruption that can occur when two users simultaneously edit the same file.

Cloud Appliance Data Access

Organizations are looking to cloud storage as a cost effective way to provide ubiquitous access to data, increase storage efficiencies and simplify global storage management. Instead of deploying islands of NAS at multiple sites across the enterprise, businesses can now deploy a single copy of their data in cloud storage, either within their own data center or by using a public cloud provider. This data can be shared and accessed simultaneously by local, branch and remote offices through "cloud integrated" storage appliances.

Cloud storage appliances provide fast local access to shared files while leveraging cloud storage for low cost and efficient data protection. The appliance essentially acts as an SSD-based file system cache and stores data in a centralized storage repository, stored either in a public cloud provider's site or in the corporate data center as a private cloud storage solution.



To facilitate rapid data access for the end users, as files are retrieved from cloud storage they are stored and served from flash SSDs on the local appliance. So while the first access to file data incurs some latency, subsequent requests will be at the speed of flash. Edits and updates to the local files will be uploaded to the cloud for offsite protection and for sharing with other appliances located at sites across the enterprise. Inactive data, on the other hand, remains in the cloud for efficient low cost storage combined with offsite data protection.

This caching technology gives users the experience of fast local NAS combined with data protection in the cloud. This simple architecture enables businesses to potentially realize significant capital and operational costs savings by collapsing silos of NAS storage infrastructure from multiple sites into a single centralized private, or public, cloud storage solution.

The File Versioning Dilemma

Cloud integrated storage is the basis for the next generation of enterprise class file sharing solutions. The challenge, however, is that some of these products don't have native file locking capabilities but instead rely on "file versioning" to create the illusion of data consistency. Rather than grant write access to a file, file versioning technology simply creates a new copy of the file each time a user requests write permissions to the original file. This creates multiple versions of a given file on the shared storage repository. As each user makes changes to a file, their unique changes are saved to a new file name; i.e. file_1, file_2, etc.

While at first glance this seems to be a usable solution to managing all the ongoing changes made to a single file, it is fraught with risk and ultimately works against productivity and wastes time. Firstly, the users must be aware of, and must follow the naming conventions used to identify the separate copies of the file. Secondly, the users are forced to merge the changes from the multiple file "copies" back together again. For example, if two users open and edit the same file, it could be very time consuming, if not impossible, for them to track and integrate all the changes made to the file. This can result in missed deadlines, lost revenue and higher production costs.

There is a better way to allow multiple users to access and edit the same file. This calls for a file locking technology that is "cloud aware".

Global File Locking

As discussed above, a global file system can span multiple sites to provide one view of an aggregated, shared file system. And the use of SSD-based cloud appliances can provide remote users with the same performance they get when accessing their files from a local NAS.

When global file locking is incorporated with the global file system, the necessary file locking services are enforced to maintain file system consistency without resorting to file versioning. When write access privileges are assigned, they will be assigned to one user at a time while other users are then limited to read access. This ensures that data will remain consistent even when hundreds of users are sharing access to the global file system across multiple sites with tens of thousands or even millions files.



Conclusion

Cloud file sharing can lower storage infrastructure costs by enabling businesses to simplify management and centralize protection of their critical business data on cost effective cloud storage. The big advantage of cloud file sharing though is that it enables a seamless end-user experience from any location geographically, without suffering performance or latency issues. This can lead to a dramatic rise in end-user productivity. But those productivity gains are lost if data integrity is not maintained through specific file locking capabilities.

Businesses need solutions which leverage the infrastructure consolidation benefits of cloud storage coupled with the file locking capabilities of a global file system. Global file system technologies that implement locking give businesses the flexibility of implementing a cloud file sharing solution on private or public cloud storage. But as importantly, this offering natively delivers the file locking services necessary to ensure that corporate data always remains in a consistent state. This enables businesses to gain all of the efficiency and operational benefits of a consolidated, high performance cloud file sharing storage solution with the assurance of data integrity.

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory based storage (Flash), big data, virtualization, cloud computing and data protection. The firm is widely recognized for its blogs, white papers, and videos on such current technologies like all-flash arrays, deduplication, software-defined storage, backup appliances, and storage networking. The “Switzerland” in the firm’s name indicates our pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor or approach.