File Level Locking vs. Volume Level Locking SAN

Jim McKenna

When designing a workflow in the modern video facility, take a look at all the options available. The promise of a “silver bullet” network is thrown around every day. The truth is, there’s no such thing as a network that does everything, or there wouldn’t be all these options available!

Let’s count the buzzwords:

1. File Sharing
2. File Level
3. Volume Level
4. Project Sharing

1. File sharing is simply that; sharing access to a file among multiple people in multiple locations. The file being shared can either be destructively shared (this means to change the contents of the file as you’re working with it) or non-destructively shared. Most every video editorial system is non-destructive. This is essential to file sharing in a collaborative environment, as destructive changes will affect every timeline or composite that uses that file.

2. File Level sharing method allows each user on the shared network to access and change files, or create new files in the same directory or volume. The popular misconception is that a video facility needs file-level sharing in order to “share files”. There are tradeoffs to this type of sharing, and some of them are substantial. The most common tradeoff is bandwidth. Native file systems (NTFS and HFS+ or Mac OS Extended) offer the greatest bandwidth due to lowest overhead and block-level access to the drives. These file systems are not file-level by nature, meaning they will only accept one writer. There are ways around this; one is by designating a “master” that will always perform the writing. This is the theory behind network-attached storage.

When you mount a volume as a network drive, you’re most likely using a NAS (not a SAN). Network Attached Storage uses a master writer/reader to access the drives, and all the clients of the master will request data writes and reads through that master. This gives the clients of the master a high level of availability, since the master sees everything. This also puts a big bottleneck in the path of the data, and can seriously affect bandwidth. On top of this, network protocol isn’t nearly as efficient as local protocols such as SCSI. This lowers the amount of data you can transmit and receive through a network connection.

Another way to allow writes from multiple stations on a SAN is the use of a custom file system. This file system is neither Mac nor Windows, and is managed by a central controller. The central controller handles all the details about the file system. This information is Metadata, or data about data. The central controller determines who can write and where they can write, to avoid two clients writing to the same section of the drives. Problems with this architecture include cost, complexity and speed. When every transaction takes a little longer due to the interaction with the controller, the entire system slows down. On top of this, there are often multiple layers of networking involved, and inevitably, costly client software licenses. There is no custom file system that does not use an installed client, and this per-seat software is another weak point, because the versions need to be closely tracked to avoid conflicts between server and clients.

A third method of sharing at a file level includes volumes that are closely managed on the clients, a master controller that designates where files can be written, and sends that information back to the clients writing the files. This method has inherent slowdown on certain I/O patterns, as each client is in constant communication with each other whenever writing to the drives. These SAN systems also suffer from fragmentation and loss of performance over time. The drives are usually not managed, nor are the volumes virtual, but physical. This offers little flexibility when partitioning space among several clients, and managing projects efficiently. Lastly, the multiple network layers can be a point of slowdown or failure when network traffic increases, or a network device like DHCP server looses contact with the clients.

3. Volume Level SAN systems have been around for a long time, and still occupy the largest percentage of the fiber channel SAN marketplace. The reason for this is simplicity, speed and cost effectiveness. A volume-level SAN is accessible from every client on the network in a read-only fashion, and each volume is available to a single client as a write volume. In this configuration, non-destructive applications can always share files, but destructive applications must have write access to any given volume.

Volume-level systems have always been simple to configure, requiring no external metadata controller. Some do still require an Ethernet network component, but those are slowly going away. The problem with some volume-level systems is in the drive and volume management. Any SAN that uses JBODS in an unmanaged architecture will have speed issues when fragmentation or full conditions arise. Also, when a SAN is partitioned into static volumes as part of a larger stripe, they lose the ability to add more clients, due to the constraints of the volume-level access permissions.

This is remedied by a virtual-volume architecture, where volumes can be created and deleted at any time, and given to any system on the network in read-only or write access. The best volume level SAN is also managed internally by use of a server chassis, not externally with a secondary client network and a standalone JBOD. It should also have various RAID protection levels that are dynamic without affecting primary data. In this way every client can have various zones of writeable and read-only storage which is high bandwidth through block-level fiber channel access without additional networking overhead, and availability of extra storage capacity at any time through additional volume creation and data protection changes.

So why not volume level? To use shared storage, you need to:

- Capture to shared storage
    - Digitize, Render Consolidate, Copy, Save
- Read back that captured material
    - Playback, View, Use
- Allow others to read back that captured material at the same time
    - Playback, View, Use
- Allow others to write to the shared storage
    - Digitize, Render Consolidate, Copy, Save

All of these things can be done with both file-level and volume-level SAN systems. Collaboration, especially in high bandwidth scenarios, can be easily achieved through simple volume-level access.

4. The project file is known as the metadata for the application, or the location of the editing timelines and footage clips that point to the video and audio data on the main storage. Project sharing does require file-level access to a project file. However, only one major manufacturer has project sharing that cannot be achieved through simple networking, if at all. Many applications simply cannot share project files. In addition, a shared project will still allow use of video and audio data on the main volume-level storage, even if the project is kept on a network share or separate SAN system entirely.

This gives the editor the opportunity to share a project file, or keep it locally to protect his changes, all while using the high-bandwidth volume-level storage to hold the video. Most applications that do not support sharing projects can still share footage and timelines by copying the project file to each station that wishes to work on the project footage. If this station also has the footage volume mounted, it will link to the footage locally in each room, and allow the editor to share the video and audio files.

To conclude, there may still be a place for file-level SAN systems at any cost, and volume-level SANs may never be able to do everything. However, the potential owner of the SAN must weigh the pros and cons carefully before deciding on a product, whether or not the marketing includes the right buzzwords.