

2017 Creative STORAGESM CONFERENCE

May 24th 2017 at the DoubleTree Hotel, Culver City, California

AN ENTERTAINMENT STORAGE ALLIANCETM EVENT 



Doug Hynes, Senior Director, Business Development, StorageDNA

TITLE: dnaLTFS: Learning File-System for Tape

ABSTRACT

Our reliance on tape continues.... In this presentation, we will highlight a new technology (dnaLTFS) for tape that can change the role tape plays in media workflows and beyond. dnaLTFS is a highly enhanced version of LTFS (Linear Tape File System) that allows LTO tape media to be mounted and used more like a hard drive. Specifically, dnaLTFS enables tape to be mounted in either WRITE mode or READ mode.

dnaLTFS: Write Mode Write mode is designed to turn tape into a high performance, streaming write device – capable of high performance write tasks. In write mode, the tape is presented over a file-system mount point (either local or over the network). Data is written directly to tape, without any disk caching. This allows for high performance offload workflows, including camera offloads, post-production (Avid, Adobe) project backups, direct camera capture and more. Tape can be used as a direct target for saving any application data for near term or long term purposes.

dnaLTFS: Read Mode Read mode is designed to allow direct access from tape without the need to perform restores. But how do we get over the fact that tape is a sequential access device? Before we explain how we enable this, let us look at some background. Tape (esp. with LTO-7) has become a high performance, sequential device – capable of over 300 MB/s (2.4 Gb/s) throughput per drive. Using multiple LTO drives enables much higher levels of bandwidth. Increased streaming performance is accompanied with consistently decreasing costs (LTO-6 is 1c/GB while LTO-7 is 2c/GB).

Additionally, it is important to note that a number of media workflow tasks are pre-dominantly sequential and conducted in batch mode. Ingesting, transcoding, QC, conforming, streaming etc. are all predominantly sequential operations. So if a number of media tasks are sequential and LTO is great at sequential bandwidth, then why can't we use LTO directly for these tasks? The challenge has largely lied in the software's inability to enable this for tape. While LTFS was the first step to solving this challenge, dnaLTFS provides the complete answer.

Talk Focus The talk will focus on explaining the technical details of dnaLTFS. We will highlight how we have achieved tasks such as offloading, transcoding and ingesting – directly from LTO – without performing any restores – using applications such as Avid, Adobe, Da Vinci Resolve and more.

We will also present a cost and workflow comparison between dnaLTFS/LTO powered Nearline and disk-based Nearline/Object storage systems. Even though LTO has become the de-facto master format for storing media content, it serves purely as a long-term, cold copy today. With the ability to now utilize LTO for direct access workflows, LTO will dramatically alter storage costs, spending, management and workflow efficiencies.

BIOGRAPHY

Doug has more than 25 years of experience working for manufacturers of integrated solutions for the M&E markets. He has held multiple leadership roles and has been a member of teams that have brought unique and game-changing products to market. Doug began his career at Avid Technology in Los Angeles, where he learned “the biz” and helped introduce editors to new editorial and collaboration tools. Prior to joining StorageDNA, Doug served as Senior Product Manager, StorNext at Quantum Corporation, where he was focused on the M&E Markets.

Doug earned a Bachelor's Degree in Communications from the Central Connecticut State University.