



A Breakthrough Storage Workflow for Tomorrow's Filmmakers

Visualize the difference between cooking for 100 people in a cafeteria compared to a restaurant. In a cafeteria, you cook in large pots and bake on large trays. The object is to cook once and serve to many. In a restaurant, higher quality demands preparing every dish individually. Of course, that's why food costs more in a restaurant than a cafeteria. However, what if you not only had to feed 100 people, but the end result was of the same quality and billed to patrons at the same amount. You'd be a fool not to prepare it cafeteria-style and save all of that extra time and cost, right?

When, then, does the video industry take the restaurant approach to storage when it can serve up data once cafeteria-style? Perhaps because she's young enough not to be saddled with a career's worth of old school assumptions, rising filmmaker Abi Corbin found herself asking this question too often, so she decided to do something about it.

"The problem is that a lot of data gets duplicated," she explains. "Department after department consistently duplicates, and it's not intuitive. The data doesn't talk to each other. For example, say we shot two minutes of a scene, but it needs to be two minutes and thirty seconds. When we go back to do it three months later, we have to grab all that information again, but it's spread across a lot of different departments, and often, when that information trickles from department to department, it gets duplicated."

Corbin and her team studied the process and found a 70% data overlap between groups, representing a tremendous amount of wasted time spread across a minimum of 16 applications.



The answer seems obvious to anyone who has ever used DropboxSM, Box[®], or a tool: store to and serve from the cloud. In reality, implementing a cloud workflow for film storage is far from obvious and riddled with towering challenges.

Corbin and many others, particularly a key group within the University of Southern California's highly influential Entertainment Technology Center[®] (ETC) decided to tackle those challenges head-on, and the results stand poised to reshape not only the film industry but all manner of content development everywhere.

Corbin and the ETC

Abi Corbin grew up in the shadows of south Boston, surrounded by poverty, barbed wire, and people mired in the lows of their lives unable to see any answers beyond the bars on their windows. Unlike most kids, Corbin's parents made no attempt to hide their daughter from reality. Instead, she says, they showed her how to change it.

Corbin graduated from high school at age 13, entered the New England Conservatory and University of Massachusetts the following year, and graduated with a Master of Arts in Performance Studies at 19. Her passion for the theater took her on a two-year journey before she landed in USC's MFA program in filmmaking, which she is just now finishing. But don't imagine her as some library-bound bookworm. Corbin's IMDb profile already shows four writing credits, four more as producer, and five in the director's seat.


The most recent of her directing credits, still in post as of this writing, is *The Suitcase*, a 20-minute film (inspired by true events tied to the 9/11 attacks) about a struggling baggage handler named Joe Franek whose life quickly turns upside down when he harmlessly steals a suitcase that turns out to contain terrorist plans, leaving him groping for the courage to place its contents into the right hands. In addition to being a compelling drama, *The Suitcase* also represents a remarkable technical collaboration spanning Corbin and her production team, large studios, technology leaders, video production software vendors and many more. Tying them all together was the ETC and its vision for a revolutionary open source workflow framework called C4. If C4 could prove itself in *The Suitcase*'s production, the major studios, many of which take technology cues from the ETC, would take immediate notice and likely change production methods across the industry.

Local Drives in a Cloud World

With his inventive hands deep in projects ranging from early Pixar to the digital face replacement technology that became standard in filmmaking, Joshua Kolden is now the CEO of visual effects firm Avalanche as well as the ETC's C4 framework lead. Whereas Abi Corbin oversees all things creative in *The Suitcase*, Kolden helms the project's technicalities, especially in regards to its voluminous storage workflow.

The trouble with mixing cloud storage and the hundreds to thousands of gigabytes typical of a film project is bandwidth. The folder that can write to a local RAID solution in minutes might take hours or even days to upload to the cloud, and no team can afford that kind of poky performance. Unfortunately, most production software assumes that one storage target is like any other, no matter the speed of the connection. C4 seeks to remedy this, and today part of that solution involves taking a hybrid approach to the problem that leverages both local storage and cloud resources.

An average day of shooting on *The Suitcase* produced roughly two terabytes of content. Data flowed into a Technicolor DIT cart containing a transcoding system and integrated RAID



storage. The challenge, as always, was to get this content off the cart and out to the people who needed it next. Assistants could plug in external drives and rush content out to editorial at the day's end, but for broader distribution, nothing could beat the cloud.

The ETC partnered with Amazon Web Services™ (AWS) and Google® to help extend the C4 framework to those companies' massive data centers, but the challenge remained of how to get data from camera media through transcoding and distributed to everyone who needed it. After weighing many factors and options, Kolden and his crew opted to use either or both of their G-Technology 24TB G-SPEED® Studio units (configured in RAID 5) as well as 8TB G-RAID® with Thunderbolt™ units (in RAID 0) to shuttle data throughout the day from the DIT cart to large partner sites, such as Technicolor® and the ETC.



Through identifiers built into the C4 framework, local software coordinated with central servers to divide the file load between them, much as larger networks will aggregate, say, multiple gigabit Ethernet feeds to achieve super-fast LAN connections. Here, though, those feeds are coming from different buildings rather than different ports. In this way, multiple locations get to leverage their big data pipes (typically using secure “dark fiber” lines) and share the job of uploading that massive data to the cloud for distribution in record time.



“With fast copying to local G-Technology drives, then having these big connections working in tandem,” Kolden says, “we were able to leverage a lot more performance. Generally, we were done transferring by the next day, but not necessarily by the first thing in the morning. That’s why we would swap out the drives, so that there was always one on-set ready to go first thing in the morning. The nice thing about having a drive be ‘too big’ was that if there was any reason why a transfer didn’t happen, more data could be dumped to the same drive — no big deal. In fact, we ended up preferring the larger, 24TB G-SPEED Studio because the unexpected happens quite expectedly. The smaller 8TB drives would get saturated pretty often.”

In all fairness, Kolden and his team ended up saturating the 24TB drives, too, just three times less often. In the controlled chaos of production, the rule most people follow is “when in doubt,

G-SPEED® Studio


delete nothing,” and so drives fill up. In future iterations of the C4 framework, easy validation routines will confirm to users when data has been successfully copied and thus eliminate this inefficient use of storage capacity.

Beyond raw capacity, Corbin and Kolden wanted to experiment with the G-RAID and G-SPEED Studio solutions as shuttle drives in part to see how size, protection, and performance trade-offs affected real world use in their application. Both products use Thunderbolt 2 connections, providing plenty of potential transfer bandwidth. Josh found that the more compact twin-drive G-RAID transferred in excess of 400 MB/s while the larger four-drive Studio (in RAID 5) reached peak speeds of about 680 MB/s.

"You end up in these situations where you have an entire crew waiting on something to be copied, and that gets very costly. Having high-end performance in your storage is really important."

Abi Corbin





Director Abi Corbin discusses the scene with cinematographer Jon Keng and actor Mojean Aria

"You end up in these situations where you have an entire crew waiting on something to be copied, and that gets very costly," says Corbin. "Having high-end performance in your storage is really important."

All told, the team's workflow entailed outputting a full day of footage from the Technicolor cart onto three external shuttle drives. If nobody was able to perform transfers during the day, making those copies before adopting the C4 framework and G-Technology solutions required about 80 minutes of transferring data to drives sequentially. (Despite having multiple transfers active simultaneously, most file systems only pass data to or from drives one at a time, making the others wait for a turn.) Due to C4's caching and bandwidth optimization features, combined with the G-Technology units' performance, that 80-minute average copying time dropped to roughly 25 minutes.

A Big Leap, More to Come

The Suitcase project proved that a proper cloud workflow can deliver tremendous benefits for film production, but it will still require capacious, convenient, and fast local storage (at least for the foreseeable future) in the critical position between set crew and cloud uplinks. If failure should strike a shuttle drive at or between those two points, the resulting delays could entail days of productivity loss and significant cost. Additionally, Corbin found that network issues and bandwidth constraints could still occasionally necessitate sending runners with drives. In the emerging world of cloud production, compact solutions with RAID protection,

Enterprise-grade dependability, and the fastest possible interfaces will remain essential.

"I have used G-Technology drives since I started as a filmmaker," says Corbin, "and I've never had a problem with performance. Actually, the only time I ever had an 'issue' with a G-Technology drive on this project was when...this is a little embarrassing...I left it on top of my car and drove away. And even then, the drive was a little worse for the wear, but the information on it was fine."

There's far more to the C4 framework than we've described here. In particular, Kolden and his many ETC colleagues have crafted a new metadata scheme that will allow file versions to remain synchronized across a scattered workflow spanning countless applications and locations. The true goal of the cloud infrastructure — and something of a technical Holy Grail for the film industry — is to allow data to flow seamlessly between applications from Adobe® to Autodesk® to AVID. JSON APIs provide the bridges over what have previously been gaping spans between silos, and the metadata attached to files via the C4 framework allows for easy synchronization of files across these software chasms, regardless of time or place.

Being the guinea pig for this revolutionary charge was no small feat. *The Suitcase* promises to be an amazing short film, and perhaps few will remember it as the proving ground that helped to reshape film workflows. But maybe not. Perhaps Abi Corbin, Joshua Kolden, and all of the others will gain some recognition for their tireless help in pushing filmmaking deeper into the 21st century. Even if no one remembers, though, Corbin will be content. The lessons learned by a little girl on the bad streets of Boston have born their fruit.

"If you want to impact the world," she says, "you have to be willing to pay a price, and you need courage. You have to sacrifice a lot to excel in any way. And if you want to be a leader, then your biggest job is to be a servant to those around you."

G-SPEED® Studio



The Suitcase
thesuitcasefilm.com

Abi Corbin
twitter.com/abicorbin

Joshua Kolden
twitter.com/joshuakolden

Abi Corbin and Joshua Kolden are leaders in their respective fields who use G-Technology products in their day-to-day work lives and are compensated for their participation.

g-technology.com

G-Technology, G-SPEED and G-RAID are trademarks of Western Digital Corporation or its affiliates in the United States and/or other countries. Thunderbolt is a trademark of Intel Corporation in the U.S. and other countries. Adobe is a trademark of Adobe Systems Incorporated in the U.S. and other countries. Avid is a trademark or registered trademark of Avid Technology, Inc. in the US and/or other countries. Autodesk is a trademark of Autodesk Incorporated in the U.S. and other countries. Dropbox is a trademark of Dropbox Incorporated in the U.S. and other countries. Box is a trademark of Box Incorporated in the U.S. and other countries. Entertainment Technology Center is a trademark of University of Southern California Incorporated in the U.S. and other countries. Google is a trademark of Google Incorporated in the U.S. and other countries. Amazon and Amazon Web Services are trademarks of Amazon Technologies Incorporated in the U.S. and other countries. Pixar is a trademark of Pixar Animation Studios in the U.S. and other countries. Other trademarks may be the property of their respective owners. ©2016 Western Digital Corporation and its affiliates. R0 8/16

