grass valley
A BELDEN BRAND

K2 System – Ready for Nonlinear Media Production: QOS, Bandwidth and File Compatibility in File-based Workflows

Karel Rasovsky, Marketing Operations Grass Valley, a Belden Brand — August 2013

WHITE PAPER

What makes an efficient file-based media workflow? This application note looks at the challenges encountered in file-based media processing, and the approaches to make file-based workflows more efficient. The K2 system was designed from the ground up for file-based operations and offers capabilities unmatched by other media system vendors.

What Is a File-Based Workflow?

Strictly speaking, a file-based or "tapeless" workflow (Figure 1) is an end-to-end workflow, from ingest to playout, where digital media is contained in files (unlike digital videotape which records media digitally, but not as a file). When working with media in this all-digital file-based domain, it is advantageous to create compressed copies of the original work that can be more efficiently moved around and processed into the desired media product. This implies that the entire file-based workflow occurs in compressed media formats. Video is encoded as early as possible, sometimes even in the camera, and that content remains compressed until it is played out to air — at which time it is decoded back to baseband (SDI/HD-SDI) and sent out. Ingest, editing, metadata management, archive, browse, content management and playout management are all performed using compressed files.

File-based Workflow in Broadcast

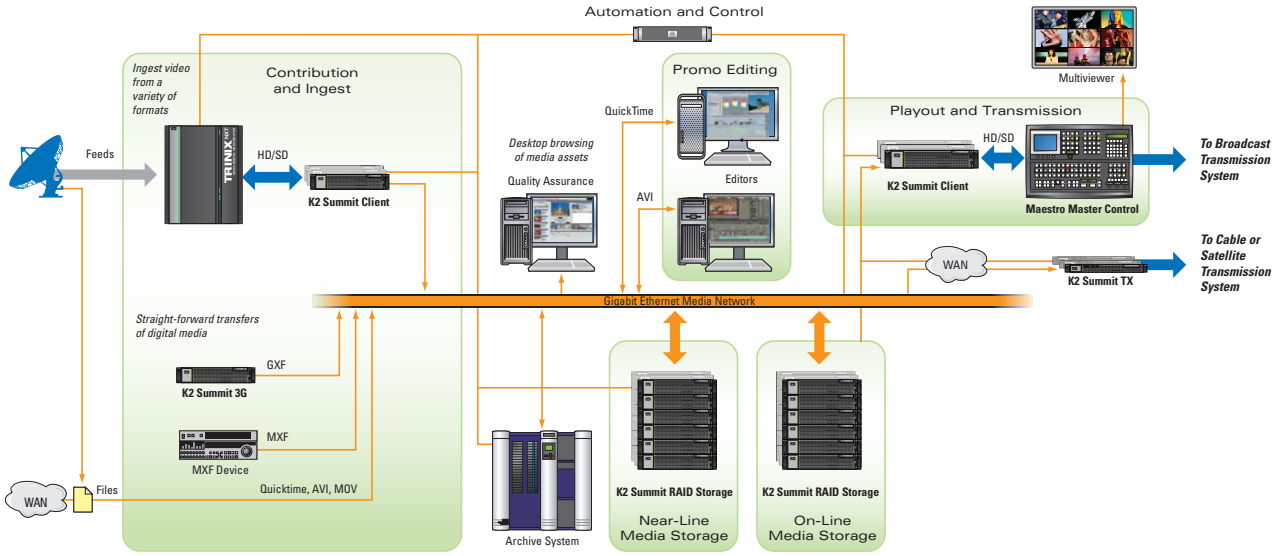

Figure 1 – In a file-based workflow, video is compressed to a file on ingest (left side) to be processed in subsequent steps of the file-based workflow.

What Makes an Efficient File-Based Workflow?

To design and deploy workflows that are efficient and best utilize creative talent, an infrastructure is needed with the following attributes:

1. Support for **a wide variety of industry-standard formats** – both compression formats and file wrappers.

2. **Ample bandwidth** to move the data around, with the ability to scale as needed.

3. Ability to manage that bandwidth with **Quality of Service (QOS)** so bandwidth is always available where it is needed most.

Industry Standard Formats

When talking about files, it is easy to get confused between compression formats and file wrappers. The compression format and the wrapper are totally independent of each other. Production and playout systems need to be compatible with both.

Compression is the way the video is reduced in size from 270 Mb/s for SD and 1.5 Gb/s for HD to a manageable 8 Mb/s to 100 Mb/s of data. Standard compression formats today are MPEG-2, DVCPRO, AVC-Intra, DNxHD, MPEG-4, as well as many others. Once compressed into smaller amounts of data, it then needs to be stored in a wrapper so it can be moved from one device to another. Standard wrappers today are MXF, GXF, .AVI and .MOV (QuickTime). For example, MXF files can contain DVCPRO or MPEG-2 video or MPEG-2 video can be wrapped as MXF or GXF. For the sake of this discussion, it is less important to understand all the compression formats and wrappers, but rather appreciate the concept and the importance of this in selecting an infrastructure for various media workflows.

The K2 system supports an extensive list of compression formats, currently including:

- SD: DVCPRO 25, DVCPRO 50, IMX 30/40/50, DV, DVCAM, XDCAM, MPEG-2 I-Frame/Long GOP
- HD: DVCPRO HD, XDCAM HD, XDCAM EX, MPEG-2 I-Frame/Long GOP, AVC-Intra 50, AVC-Intra 100, AVC LongG, DNxHD 115/120/145/175/185/220, AVCHD, H.264 MPEG-4

The wrappers supported on K2 include MXF, GXF, .AVI, .MOV (QuickTime) and MPEG. This flexibility of formats available in the K2 system minimizes the amount of transcoding and rewrapping taking place in K2-based workflows.

Ample Bandwidth

The second critical attribute is bandwidth. There are two important questions to ask:

1. How much is there?

2. How easy is it to add more as the need for bandwidth increases over time?

What is Bandwidth and Where Does it Come From?

Bandwidth is closely linked to service availability, which in turn is a crucial element of QOS. The network infrastructure must be designed to be highly available before QOS can be successfully implemented. Obviously workflow efficiency improves with more available bandwidth because users spend less time waiting for data. It's important to understand that all bandwidth comes from the disk storage system (Figure 2). If it is determined that 500 MB/s of bandwidth is needed, then the storage system must be able to reliably read or write 500 MB/s of data to or from the disks, including when a RAID disk is being rebuilt. In addition, the network infrastructure must be able to move 500 MB/s of data. Both the storage system and the network need to be optimized to meet this specification under all conditions, or the result is a poorly designed infrastructure that will eventually end up dropping frames when playing to air. Black in the middle of a commercial can cost a station hundreds of thousands of dollars — a network, millions.

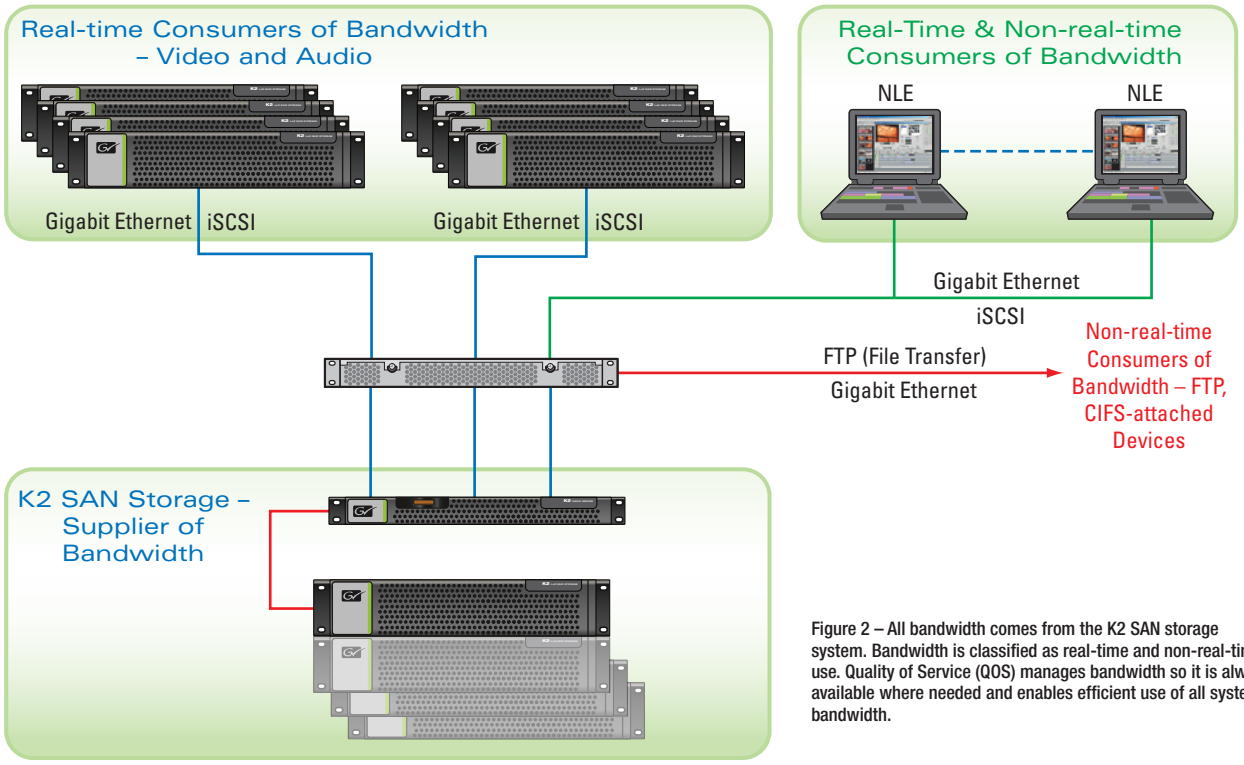

Figure 2 – All bandwidth comes from the K2 SAN storage system. Bandwidth is classified as real-time and non-real-time use. Quality of Service (QOS) manages bandwidth so it is always available where needed and enables efficient use of all system bandwidth.

What Makes an Efficient File-Based Workflow? (Cont.)

Although more available bandwidth is desirable, this has to be balanced against the associated infrastructure cost. Some media network solutions simply over-design the amount of raw bandwidth, making for a rather inefficient approach. Others come up with bandwidth allocation schemes, where a fixed share of bandwidth is allocated to video and another fixed amount allocated for file transfer. But these fixed bundles of bandwidth don't adapt to the real-world conditions, where the demand for video or file transfer is dynamic and always changing.

Another important consideration is the scalability of a media platform. The evolving needs of media production over time inevitably results in growing the storage and bandwidth requirements.

Controlling Bandwidth with Quality of Service (QOS)

Planning an infrastructure with the required bandwidth and support for the compression and file types needed is just the first step. If the bandwidth cannot be controlled to make sure it gets used for the highest priority tasks, then the workflow will fail. Given that not all users on a media network are created equal in terms of their bandwidth needs and priorities, building a "first come, first served" model for bandwidth allocation can be disastrous. Applying this scheme to the example in Figure 3, if there are 10 editors on a pool of storage with 200 MB/s of bandwidth, the reasoning would be that each editor has 20 MB/s of bandwidth. But one of the editors may log in first and launch an edit session with 10 layers of DVCPRO HD consuming up 100 MB of the available bandwidth. That leaves the other four editors with a totally unacceptable performance to get their job done.

In another scenario, a batch file transfer could be initiated to archives, when in midstream, a scheduled commercial has to play. In this situation, there may not be adequate bandwidth to play the spot without black frames. Note that a black frame is where the playout server did not get data in time to play and thus plays black — it is almost always related to a stall due to inadequate bandwidth.

K2 Quality of Service – QOS

As bandwidth comes at a cost, it is critical that media organizations have the ability to control available bandwidth in a usage-sensitive, adaptive way. The K2 system solves this challenge with a unique Quality of Service mechanism to control bandwidth usage across the system. The build-in QOS classifies media streams and users as "real-time," "non real-time" (Figure 2) and "reserved." The QOS mechanism dynamically allocates bandwidth to these activities according to the following rules:

- Record and playback are treated as real-time media streams, with guaranteed bandwidth.
- Editing clients are classified as reserved for time-critical production activity — they are treated as real-time up to a defined bandwidth limit that they cannot exceed, and so become non-real-time for any amount of bandwidth above that limit. Different amounts of bandwidth can be allocated to specific clients.
- FTP file transfers are treated as non-real-time streams. Bandwidth unused by record, play or edit clients can be used for FTP transfer.

Based on this classification, the QOS mechanism can limit a device/user to a specific amount of bandwidth (Figure 4).

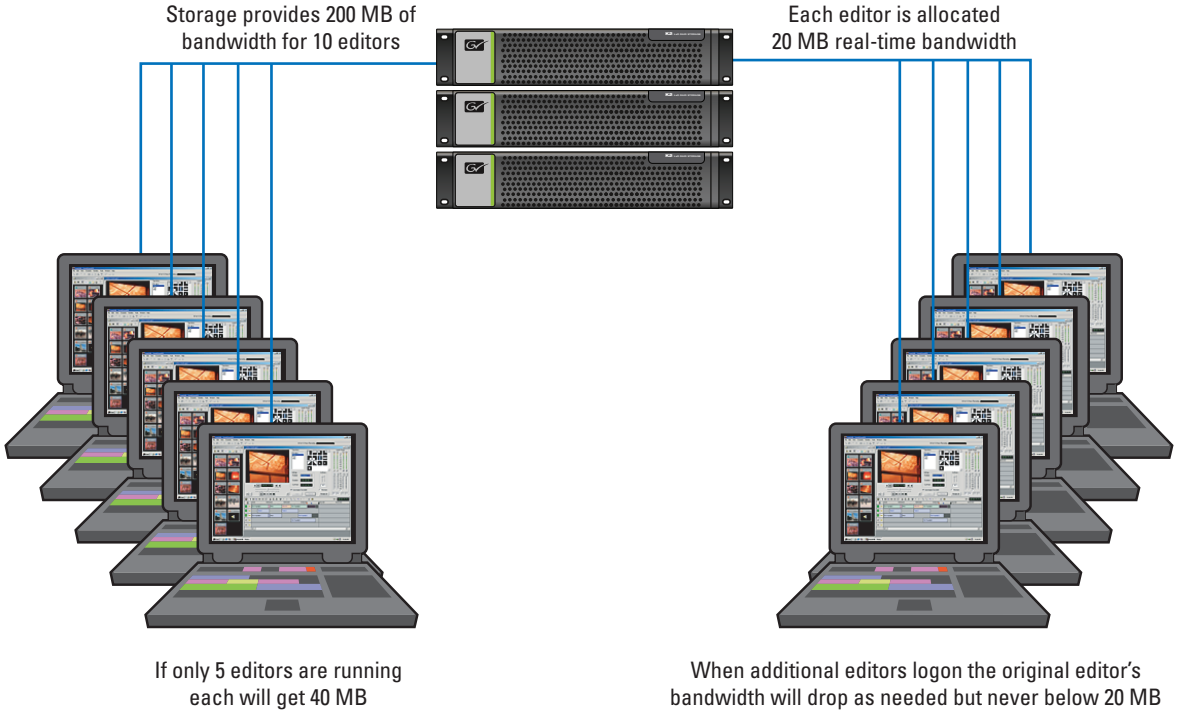

Figure 3 – In a shared editing environment, editing bandwidth is provided on a equal basis so all editors are guaranteed adequate bandwidth to do their job.

NLE Example of Shared Bandwidth

Many media platforms on the market today have no concept of dynamic allocation of bandwidth, applying rather a trivial "first come, first served" model instead, which inevitably results in inadequate performance for some users in the pool. In contrast, all the storage bandwidth in K2 system is available to the editors that are currently running. As more editors log on, they get their share of bandwidth and each editor is guaranteed an equal amount of bandwidth.

K2 Production Storage Example with QOS Bandwidth Allocation

The benefits of QOS mechanism in K2 systems:

- Bandwidth is always prioritized for recording and playing (real-time) so a frame of data is never missed.
- Non-real-time users can never take bandwidth from a real-time user, ensuring playout is never impacted.
- One NLE user cannot consume more than their allocated share of bandwidth in reserved mode, ensuring other simultaneous NLE users have bandwidth available for their work.
- Unused real-time bandwidth can dynamically be allocated to non-real-time users.

Summary

Bandwidth management of real-time, non-real-time and reserved streams and clients is fundamental to achieving file-based workflows that harmonize bandwidth usage across the entire organization, while ensuring that no frame gets dropped when it counts most — on-air. This kind of intelligent bandwidth allocation has to be built into the media platform, so as to dynamically react to changing user needs and workflows.

The K2 system delivers market-leading performance and scalability, with built-in QOS to manage bandwidth for appropriate use. The modular architecture of the K2 system can scale to 1000s of MB/s for video, editing or FTP transfer, a scale unmatched by any other media system manufacturer. The K2 storage architecture enables easy scaling for more bandwidth as needs grow. This means media organizations can deploy a media system with the current requirements and grow them gracefully over time, optimizing the cost for the initial system installation. Built on standard IT technologies optimized for media production, K2 introduces a substantial increases in bandwidth with each generation of new system components, scaling with media production needs now and into the future.

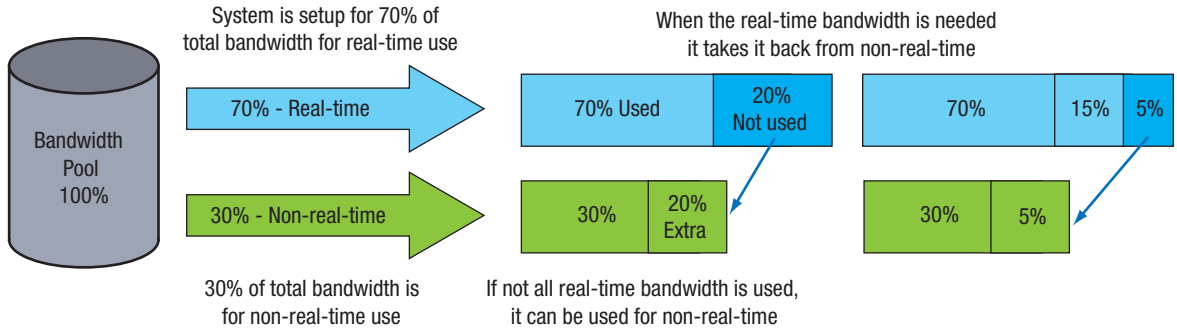

Figure 4 – K2 QOS can dynamically allocate unused bandwidth from the real-time pool to be used for non-real-time use. Bandwidth will dynamically be allocated back to real-time when needed so real-time devices are never starved for bandwidth. Most systems on the market allocate bandwidth for one use or the other and cannot use bandwidth that is not being used by the other side.