An Evolphin® Software Technology White Paper

2410 Camino Ramon #228

San Ramon, CA 94583

+1 888-386-4114

info@evolphin.com

www.evolphin.com

Transitioning to the Big Video Workflow: Key technical insights to successfully build high-performance video workflows with a MAM

Rahul Bhargava

Chief Technology Officer

Evolphin Software, Inc.

Contents

Introduction

Putting Big Video into Context

Defining Big Video Workflow

Big Video Use Cases

Use Case 1: Publishing Co. Moving to Online Videos

Use Case 2: Security Agency's Training Videos

Challenges 1 to 6

1: Big Video Footage can be Massive

2: Post-production System Architecture may Melt Down

3: Lots of Containers and Codecs

4. The Need for Hi-Res and Proxy Workflows

5. Siloes, Different Asset Types on Different Servers

6. Storage Management takes a Larger Role

Technical Solutions to Big Video Challenges

Switching to High-Performance Asset Ingest

Front-end De-duplication to Arrest Bandwidth & Storage Growth

Proxy-based Edit Workflows

Unified Information Repository Model

Flexible Transcoding Framework

Metadata-driven Video Project Namespace

Conclusion

Abstract

Be it Media & Entertainment companies producing traditional video content or internal corporate studios generating training videos for staff, it's no secret that videos are everywhere. In this climate, video production teams are presented with new challenges every year to complete their projects on time and within budget. Higher resolution formats such as 4K video demand even more resources to produce and manage.

This paper will discuss the key challenges that occur when organizations transition to big video workflows and how they can be mitigated with technological, business & operational insights gleaned from some of the leading video production teams around the world.

Issues and solutions covered include: front-end de-duplication technologies that help video editors deal with large video files; judicious use of Fiber Channel to surmount inevitable Ethernet slow-down while editing videos; flexible transcoding frameworks to rescue users from video container formats & codec tsunamis; hi-resolution proxy workflows to speed up production; unified information repository models that breakdown content siloes, and last but not least, this paper will present two use cases from organizations that have successfully transitioned to the big video workflow.

Introduction

Putting Big Video into Context

According to analysts, the usage statistics from the popular video upload and sharing service, YouTube, alone are staggering:

- Over a Billion people view 4+ Billion videos every day on YouTube
- 6 Billion hours of video are watched per month on YouTube
- 300 hours of video are uploaded to YouTube every minute

According to an initiative by Cisco to track Video[1] usage:

[1]: Cisco Visual Networking Index 2016

- Video traffic will amount to 82 percent of all Internet traffic by 2020, up from 70 percent in 2015. Internet video traffic will grow fourfold from 2015 to 2020, a CAGR of 31 percent.
- Yearly Internet data traffic will exceed the zettabyte ([ZB]; 1000 exabytes [EB]) benchmark by the end of 2016, and will hit 2.3 ZB per year by 2020. By the end of 2016, global Internet traffic will be at 1.1 ZB per year, or 88.7 EB per month, and by 2020 global Internet traffic will reach 2.3 ZB per year, or 194 EB per month.
- In other words 1.9 ZB (1900 EB or 1.9 million Petabyte) of video traffic would flow over the Internet by 2020. Of this video traffic, at least 40% would be 4K video.

This highlights that both production and consumption of videos has reached a tipping point. Video is no longer only in the realm of the Media & Entertainment industry. Video is taking a much bigger role in how companies, irrespective of industry segment or verticals they operate in, engage with their employees and customers.

In order to produce a final video that circulates on the Internet, it can take a post-production team a lot more video footage, images and other graphics elements than what an end-user sees in the final produced video. For instance, when an end-user streams a 5-minute clip on YouTube that is around 500 MB, the post-production team might have used 20 different video footages that could have an aggregated size of 5 Terabytes.

The resolution of video footages for all these videos is increasing from HD to 4K, to an upcoming 8K. Today our smart phones can easily capture 4K videos. Camera manufacturers are touting 8K resolutions. 16K video camera & 100 Megapixel DSLR cameras are already on the horizon. In other words, the video content avalanche is even more pronounced internally where it's produced than when it's on the Internet being streamed.

Defining Big Video Workflow

Simply put, the big video workflow is the manipulation of large hi-resolution video footages (HD/UHD/4K/6K/8K) used to produce content for omni-channels. Short production cycles & ever growing content volume has created major challenges to efficiently implement the big video workflow.

Big video content management places a huge amount of pressure on an organization's infrastructure including servers, network (LAN/WAN), desktop devices and the associated file systems. Managing huge video files falls into the realm of high performance computing.

This paper will look at two video production workflows; a 24x7 media outlet and secondly, a security agency. Then, the paper will look at the key technological challenges facing video production houses around the world, specifically focusing on big video workflows. For each of these challenges, the paper will concentrate on the technical, business and operational insights being used to mitigate them.

Big Video Use Cases

Use Case 1: Publishing Co. Moving to Online Videos

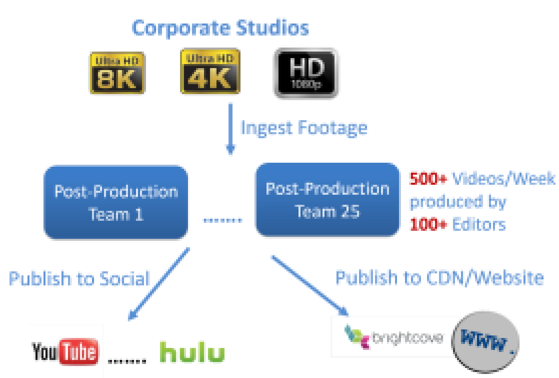

An iconic publishing company made the strategic decision to transition from a traditional print publishing business model into an online video business model in order to counter the declining print ad revenue stream. For each of their 80+ brands, they set up a recording studio that uses Canon/Sony/RED cameras to capture & ingest footage. As part of the process, 150+ video editors in each of the brand's post-production teams then edit ingested footage, post-produce and export 500+ final videos every week to a variety of video portals, social media channels, CDN (content delivery networks) and even traditional broadcast TV channels. The final videos are typically less than 10 minutes long but often link with 100+ footage files & graphics elements that collectively span over 100s of hours.

Over 3 Terra Bytes (TB) of new video footage is ingested every day to drive these post-production workflows. Traditional file and asset management systems are unable to cope with this staggering amount of file data and the associated post-production jobs were causing crippling delays and leading to dysfunctional video editing workflows. Meeting breaking event deadlines was becoming impossible. As a result, a complete re-think of the big video workflow was required.

Use Case 2: Security Agency's Training Videos

A government security agency needed to produce around 10, 1 hour long videos per week for staff training as well as for informational content for airports across the country. Even though the volume is not as high as the previous use-case, the agency struggled with big video workflows because the shared storage infrastructure over Ethernet could not sustain the high bitrates needed for editing large HD footage.

Within the context of the two use cases, the paper will now focus on the technical challenges encountered in big video workflows.

Challenge 1: Big Video Footage can be Massive

Video footage sizes are 1000s of times larger than the largest assets typically managed by traditional DAM (Digital Asset Management) systems. For example, videos produced by smartphones can generate massive video files.

Typical Smartphone UHD Camera footage statistics:

- The UHD Frame Resolution is 3840x2160 pixels => 8.3MP/frame
- An uncompressed RGB 8-bit frame could be ~ 24.9 MB (8.3 MP x 24 bits)
- An uncompressed RGB 16-bit frame could be ~ 49.8 MB (8.3 MP x 48 bits)
- An uncompressed 4K 1-min recorded clip @ 30 fps (frames/s) could require over 40 GB

As the post-production world transitions to even higher resolution cameras such as 6K, it gets worse:

- At 6K resolution, a frame would have 9x the pixels of an HD frame
- Higher resolution cameras often record at a higher frame rate, for example a 6K camera @100fps, would require over 500 GB storage for an uncompressed 5 min clip

Capturing raw video in a native format is often essential to post-produce final videos with color correction, zoom-in and other such effects applied. Hundreds of footage files from various camera card and independent video content producers might be incorporated into a final video.

Professional quality DSLR still cameras at 18 Megapixel resolutions will generate images that are barely 25-30 MB compared to the large footage files produced by even the most basic HD video cameras. Traditional Digital Asset Management (DAM) solutions are designed to handle still images, not large videos.

Video file sizes as mentioned above can be 3000x the largest still image. Traditional DAM systems are not built to accommodate such large files; the underlying database technology and the core network and storage architecture have not been designed with such large file sizes in mind.

Challenge 2: Post-production System Architecture may Melt Down

Big video workflows warrant change to the typical asset edit, post-production system setup. Ethernet-based storage access technologies often utilized by MAM/DAM solutions cannot keep up especially given the fact that big video file streaming requires a dedicated high-speed connection.

For instance, a 4K video with each frame > 8 Megapixels, would need > 800 Mbps for an uncompressed raw stream. Now imagine there were 10 video editors trying to edit 10 parallel streams; one can easily see a requirement close to 10 Gbps of dedicated bandwidth, and this does not even account for other applications that will share the available bandwidth. Further, Ethernet based protocols such as UDP/TCP have to contend with packet collision detection, packet loss, and the re-transmission of packets, all of which can slow down even a high-bandwidth connection causing the sustained transfer speed to be smaller than the theoretically available bandwidth.

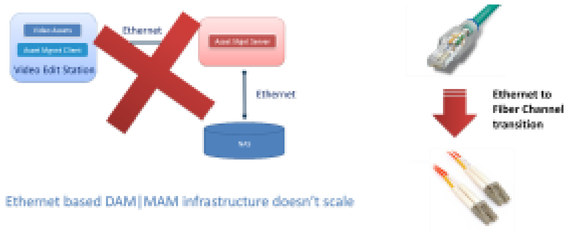

Ethernet-based connectivity within the corporate LAN often needs to be upgraded when architecting high-performance post-production workflows involving big videos. Later in this paper, we will look at fiber channel-based topologies that are better suited for certain parts of big video workflows.

Challenge 3: Lots of Containers & Codecs

Unlike images & 2D design files, videos come in a variety of containers, and with containers, there can be many different supported codecs (encoder/decoder). Each codec can have its own profile based on bit rate, frames/s (FPS), number of audit channels etc.

There is no standardization on containers or codecs: each organization will often have their own in-house codecs & containers they prefer to use. In addition, the specific camera being used brings its own flavor of codecs into the equation. Combinations are numerous and vary widely, and supporting these can be a very tedious task for the asset management solutions that need to deal with video workflows. To make it worse, traditional asset management systems are notorious for lack of coverage across this wide spectrum of containers & codecs.

This can be very frustrating for the creator or consumer of video content, especially when they have large format videos because the conversions from one format or codec to another can take hours; furthermore, at the end of the conversion, the quality may not even be what the end-user was expecting. Codecs demand too many compute resources and often require hardware offloading to a server tier optimized with GPUs for transcoding, which puts them beyond the realm of most asset management systems.

Challenge 4: The Need for Hi-Res & Proxy Workflows

Big video workflows often require Hi-Resolution native assets to be ingested from camera cards or other content channels. However, video editing in a LAN is hard, and does not even take into account that on the WAN, you may need to sustain 8-10 Gbps video stream per workstation.

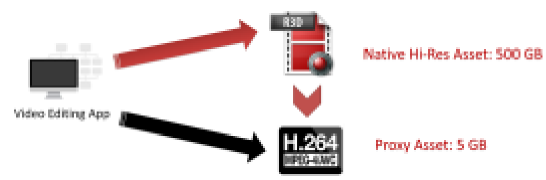

This situation often requires that teams work with video proxies. Proxies are placeholders with the same dimensions as the native videos, but with more compressed frames. They can often be 1-5% of the native video. Non-linear video editing (NLE) tools can create edit decision lists (EDL) that are first generated on a proxy, and subsequently applied to the native Hi-Res Big Video asset.

Depending upon the needs of the post-production organization, multiple levels of proxies are often required. For instance, for a web production, a low-res H.264 proxy may be good enough, but for a TV producer, an HD proxy in DNx145 codec may be needed. The editor when working remotely from home or out of the office would prefer to make edit decisions on a low-res H.264 proxy because it is easier to stream on a low bandwidth connection, but the same video editor when producing the final video, might need to zoom into a scene and may need access to a 6K native asset.

Challenge 5: Siloes, Different Asset Types on Different Servers

Post-producing videos with a combination of big hi-res assets & smaller assets, such as graphics elements, often require multiple asset types that need to be accessed by an editor. Graphics elements, 3D motion graphics, audio files & Big Video footages need to be manipulated together in a video editing application for the final export into any modern post-production workflow. Often these assets are located on different servers in the organization. This creates siloes of various asset types that live in independent asset repositories. Without a unified asset repository, it means editors and end-users have to constantly search across multiple systems, copy, transfer, and convert files from one system to another. If an organization is post-producing hundreds of videos a week, this creates an impedance mismatch in the workflow. Current asset management systems typically cannot converge separate asset type islands, leading to fragmented asset siloes.

Challenge 6: Storage Management takes a Larger Role

A lot of manual system administration labor goes into managing storage tiers with Big Videos because the hyper storage requirements warrant non-essential files be moved to a slower and cheaper storage once a video project is finished, to allow the video editors to utilize more expensive, high-performance storage for real-time post-production of new projects.

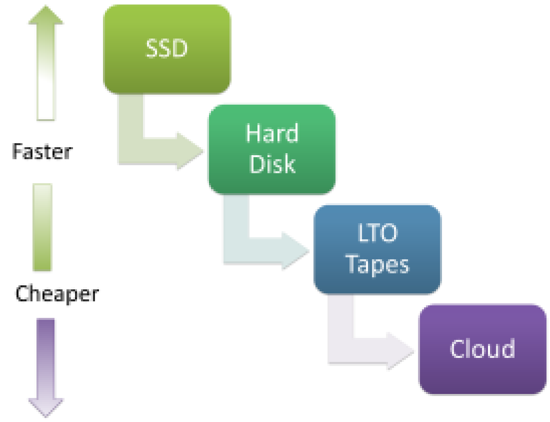

The problem is that IT administrators need to often get involved to move large video files from/to the back-end archive. In a fast-paced video production environment, this manual approach doesn't scale too well as it takes too much time to archive and restore.

Unfortunately, if the HSM (Hierarchical Storage Management) products live in a silo, then end-users need to preview assets before restoring them. In this case, it would be ideal to have metadata search enabled so users can find the assets they may need to restore. Often, the low-resolution proxy can be left online, and the hi-res and original native files can be archived to slower storage tiers. Unfortunately, this intelligence doesn't exist in storage software. In addition, storage doesn't know the video project's boundaries, for instance storage systems cannot determine at the file system tier what constitutes a video project with its dependent links. This metadata, unless extracted by the MAM system, is not exposed as file attributes. Therefore storage systems by themselves can't decide to move a file to a cheaper storage tier without breaking the video project's linkages.

A great example of this is an Adobe Premiere or Apple Final Cut project that is being edited, where hundreds of linked footages are not being modified. They are in use via links placed in a video project file and should not be archived till the entire video project is done. In another example, an editor may be working with low-res proxies, but when creating the final exports, needs to access hi-res assets that might be in an archive. HSM products would be too low-level to take into account the editor's project workflows.

Finally, each storage tier has its own costs and pros & cons. For instance, cloud storage like Amazon Glacier may be cost effective, but restores can be 1000x slower than an enterprise NAS managing offline files. To sum up, storage-tier management decisions today require too much manual intervention and are inadequate for fast paced Big Video workflows.

Technical Solutions to the Big Video Challenges

In this section, we look at technical solutions that allow us to overcome the Big Video challenges we outlined earlier.

Switching to High-Performance Asset Ingest

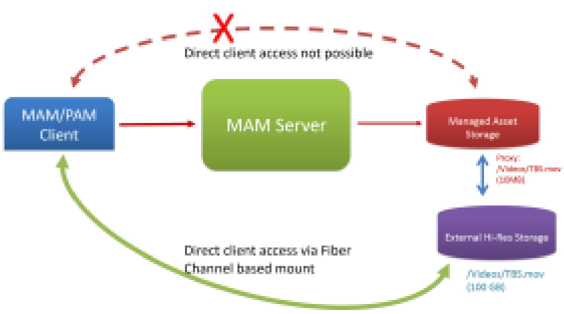

One strategy for coping with bandwidth limitations of Ethernet-based LAN is to use Fiber channel connectivity or a hybrid scheme. Fiber channel-based connection will avoid packet loss & transmission issues that can plague shared Ethernet connections, thus lowering latency and guaranteeing sustained throughputs to the storage directly.

The way this often works in practice is a user's edit workstation is set up with both a Fiber channel adapter and an Ethernet network interface card. MAM systems capable of managing multiple storage systems for regular & Hi-Res assets can then be set up to use Fiber channel mount points for Big Videos while using Ethernet-based connections for work in progress/production files such as Final Cut Pro or Adobe Premiere projects, as well as exported videos that need to be iterated. The reason for this divide between Ethernet versus Fiber channel access is because typically for read-only footage, direct Fiber channel access to storage makes more sense since the bit rate to stream them tends to be in a higher magnitude than smaller project files. For relatively smaller writeable files being edited or exported, a MAM gateway server is better suited to manage the storage in order to avoid the collaboration issues that plague shared network systems.

The MAM client capable of supporting hybrid connection models can be supplied with rules to differentiate Big Video streams from smaller video streams, and choose a connection accordingly. For instance, bit rate could be used as a parameter to allow the MAM client to use Fiber Channel for large native video files. This would allow the MAM client to inspect the bit rate metadata in the video file to decide which channel to use to copy the native file. Files with higher bit rate could be detected as big video footages that need to use Fiber channel for a faster transfer time.

Front-end De-duplication to Arrest Bandwidth & Storage Growth

In video production environments, it is not unusual for video editors to go through 15-20 iterations of a video project during an edit-review cycle between an editor and their end-clients, such as a video producers, marketing directors etc. Changes suggested by reviewers require edits on the video project file with tools such as Adobe Premiere or After Effects. The final videos then need to be exported and stored on the network for the next approval cycle. This can lead to storing multiple copies of an exported video, which can cause an excessive amount of storage consumption as depicted in the diagram below:

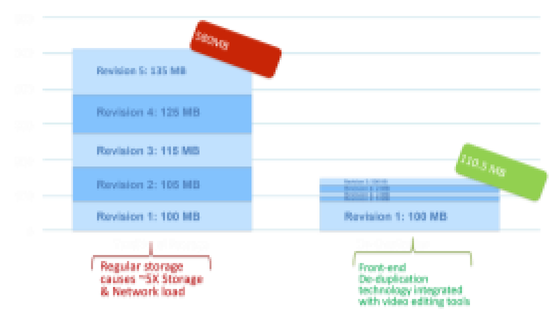

Traditional de-duplication is the general technique of identifying blocks of file data that changed in a storage volume, and only storing the changes. This allows de-duplication to be used for incremental backups and archiving solutions. De-duplication can be slow if the entire volume needs to be checked against changes to a file.

This is where front-end de-duplication technologies come into play. Using the knowledge of files that are being modified by the user on a desktop, front-end de-duplication technology can identify quickly the modified blocks on a user workstation itself, and then transmit and store the changed blocks on a back-end MAM repository. If these de-duplication technologies are integrated directly into the video editing applications, users can export videos and save incremental changes, reducing the network bandwidth needed to transmit a re-exported video as well storing various versions of an exported video. The front-end de-duplication tools are usually generic enough to work with any file type and can help with graphics elements such as Photoshop files embedded in a video project. This can be key to arresting the storage growth that can easily over-run the storage capacity in a typical video post-production organization.

Proxy-based Edit Workflows

Video Proxies are a frame accurate representation of the native Hi-Res Big Video file. They allow users to work with a small derivative of the native video asset and often can be 1-5% the size of the original file. Thus instead of working with a 100 GB MXF video file, an editor could work with a 100MB MP4 file for the initial layout, design and editing of a video project. This is a must-have for remote teams, where the quick edits need to be performed on a low-bandwidth connection, often over a WAN or over shared Ethernet LAN.

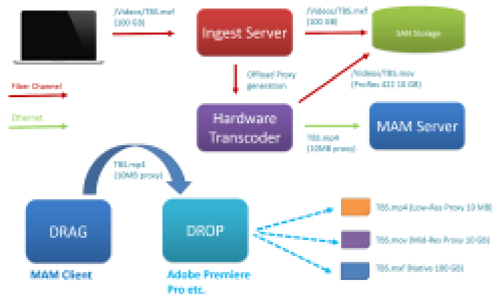

Once the video is approved, the proxies can be replaced with higher resolution derivatives before doing the final exports. Editors can switch the proxy resolution using a MAM client that has built-in proxy management support. MAM systems can integrate with the video editing tools to enable proxy switching via dragging and dropping into a video edit sequence. MAM systems that are capable of supporting auto-mounting various types of proxies can make this process seamless to the end-user.

Unified Information Repository Model

In a modern post-production environment, video editors, graphics designers, motion graphics artists, 3D modeling artists, production assistants, and producers all need to collaborate on the same video project. Each stakeholder works with their own set of tools and file types. A video project thus ends up manipulating Audio/Video, 2D, 3D & text assets. This can lead to content silos where videos are stored in a video-oriented MAM, images & 2D assets are stored in a traditional DAM, and briefs including text documents and PDF files are stored in a separate document management system. Content silos make it very difficult to manipulate and manage different file types in the same video editing workflows because files have to be moved or copied around, or transformed from one system to another.

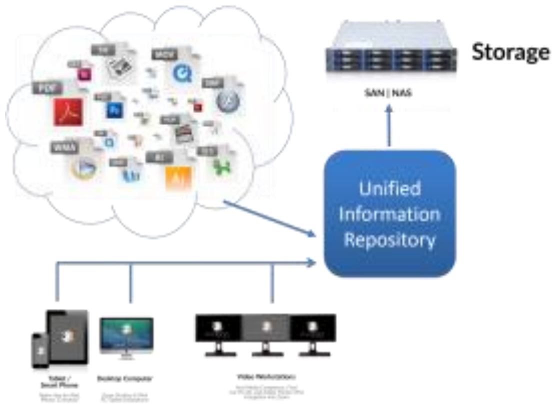

MAM systems that support unified information architectures eliminate asset type islands/silos. With these MAM systems, graphics elements & Big Videos can all be manipulated together. If an editor needs to incorporate a graphics logo into a video project, she can have the same consistent search interface for all file types, making it seamless to work with non-video assets.

Flexible Transcoding Framework

With so many video containers & codecs to manipulate in a post-production video workflow, it is often hard for one transcoding system to handle it all. Next generation MAM systems have a flexible transcoding framework that allows rule-driven algorithms to select the correct transcoder for a specific scenario.

For example, distribution to delivery end-points like VOD (Video on demand) services scenario may require cloud-based transcoders to be triggered upon copying the exported files to the VOD end-point. The MAM system could avoid wasting time transcoding prior to cloud upload if it detects the end-point to be a cloud-based VOD service.

For creating proxies during a camera card ingest, fast hardware based in-house transcoding systems might be required. For creating previews of exported or native videos for visual markups, a software transcoder like open source Ffmpeg might be enough, and could also offload expensive hardware-based transcoders to focus on more critical jobs like camera card ingest.

Metadata-driven Video Project Namespace

Conventions for file/folder names and metadata are critical to effective search in any asset management system. With Big Video workflows, studio teams can often ingest thousands of sub-clips across camera cards & live studio feeds; unless a scheme is implemented to automatically manage the file/folder names on the storage volume, it can quickly become a logistics nightmare to manage these assets in a MAM.

MAM systems capable of generating the file/folder names using metadata applied during ingest can be critical in achieving a consistent folder structure for in-place ingested hi-res assets. This ensures all the sub-clips, raw video footages, native hi-resolution videos are copied or moved in a folder structure namespace managed via metadata.

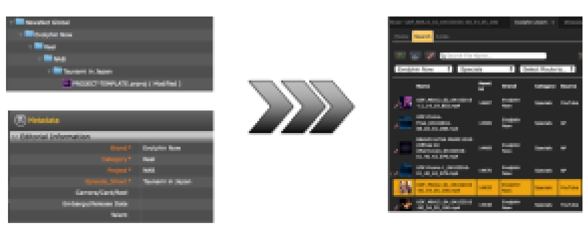

Once the namespace is under the control of the MAM system, it allows MAM clients including extensions/plugins into video editing applications to easily search hi-res footages and manage proxy switching without requiring the users to hunt down relevant files in a file system browser with a large number of assets.

Conclusion

Managing Big Video is a whole new beast. It requires orchestrating multiple systems to solve the end-to-end video workflow by ensuring the following:

- ✓ You will need a high performance, **hyper** scalable media asset management system; pay attention to the underlying asset database technology. As presented earlier, the Big Video workflows depends on a variety of asset types including graphics, 3D, and hi-resolution footage, therefore requiring a new breed of media asset management systems with an underlying database that can manage all these media types. A Unified Information Repository will simplify processing of diverse media types and will reduce the cost and support associated with multiple 'silos of content'.

- ✓ Big Video workflows require conversion between native video formats and final exported video formats, which often requires specialized hardware based transcoders. Once Big Video projects with all the dependent assets are finished, online storage space need to be reclaimed by archiving the assets. The ability of the media asset management solution to integrate with 3rd party systems such as hardware transcoders and archive systems therefore is critical.

- ✓ You will need an infrastructure designed for high performance computing; processing Big Video content requires careful orchestration between editing, transcoding, distribution and archiving systems. This in turn places great load on your network as well as application and storage servers. Fiber channel based high-speed connectivity between the workstations used for ingesting and editing large footages and the media asset management system is recommended.

- ✓ Traditional Ethernet based connectivity can be used for regular assets.

All solutions mentioned in this Big Video Whitepaper can be achieved with Evolphin Zoom.

About Evolphin Software, Inc.

Evolphin Software is the first of its kind in Rich Media Asset Management. Zoom leverages the industry's first high-performance in-memory deduplication of content, allowing users to create, share, and collect insights, and provide rich media to their clients in the fastest way possible. See Zoom in action for a different approach to media workflow management that will make you re-think conventional wisdom with a new generation of products and capabilities. More than 60 companies around the globe, covering the Advertising, Broadcasting, Media/Entertainment, Retail, and Healthcare industries, currently deploy Zoom.

For more information about Evolphin, including pricing, please visit www.evolphin.com, email Evolphin at info@evolphin.com , or call 1.888.386.4114.