WHITE PAPER

Elements of Transcoding Performance

Abstract

Accelerating media creation and delivery pipelines with BlueArc Storage platforms and Digital Rapids transcoding systems

BLUE ARC®

TABLE OF CONTENTS

BLUE ARC®

Introduction

The task of reaching audiences whenever and wherever they are has escalated in importance and complexity in today's highly connected, information- and entertainment-driven world. For media and entertainment companies to stay competitive, it's essential that they be able to distribute information and entertainment across a constantly growing number of delivery platforms as quickly as possible.

But in looking at content creation and delivery pipelines for media capture, ingest, post-production, encoding/transcoding and delivery, there are a number of factors that can dramatically affect the overall performance of the application components. Storage systems are key in this regard. As we move to a world dominated by stereoscopic 3D, 2K, 4K and higher resolution content, moving very large files onto editing systems, servers, archival and playout systems quickly and efficiently is essential—and the role of high-performance storage takes on even greater importance. Here, the profile of the data itself adds complexity to potential bottlenecks in the system. For example, performance can be affected by file sizes, the random or sequential nature of the data, the ratio of reads/writes being performed, even the number and nature of layers within video files. It's also not enough to move data across the network to the CPUs over the interconnects as well as to and from storage systems using automated tools and software suites that don't give operators exacting control over job optimization.

Effective high-volume content creation and delivery pipelines require a high-performance storage platform that is:

- Integrated with the appropriate software tools to deliver a quality product from content capture through broadcast/delivery;
- Up to the task of handling many I/O intensive applications, including such functions as meta-data extraction, keyframe extraction for storyboarding, rendering, editing and compositing, or transcoding between formats.

Without a finely tuned high-performance pipeline, repurposing content for the Web, mobile devices, and to many other platforms can become cumbersome and time consuming.

Regardless of the technical challenges, the pressure to deliver content on demand, without delay, to any device and with maximum quality is becoming pervasive. For example:

- Movie productions increasingly rely on same-day-turnaround digital dailies that directors, producers, investors and other stakeholders can review online from virtually any location.
- Broadcast and cable companies are delivering news, sports and entertainment content across platforms that range from broadcast, satellite and cable transmission to DVD, Blu-ray Disc and online downloads, as well as delivery on various devices, including iPhone, iPad, Zune, Android, PlayStation 3, Xbox 360 and more.
- Streaming services such as Hulu, Netflix, Amazon and others are delivering content to more devices than ever, including Internet-enabled TVs, Blu-ray devices, game consoles, mobile devices and set-top boxes.

In each of these cases, high-volume, fast-turnaround content production and delivery places new levels of stress on production workflows and transcoding pipelines. This document examines potential stress points in transcoding ecosystems and looks at how Digital Rapids transcoding solutions in concert with BlueArc Network Attached Storage (NAS) platforms help increase workflow efficiency, eliminate bottlenecks and improve overall return on investment.

Measuring Up

Identifying potential bottlenecks in your transcoding system architecture helps maximize transcoding speed and efficiency, but transcoding performance is more than a measure of raw speed. First among the other factors that need to be considered are image and audio quality.

Output quality starts with the quality of your source files—the old adage, garbage in, garbage out still applies in the early 21st Century. Regardless of whether you're encoding professional quality source files or user generated content shot with low-resolution cell phone or web cams, it's safe to say that not all encoding and transcoding systems are created equal. Output quality also depends on the particular codec implementation used by the transcoding system. Different transcoding solutions often use different codec implementations, even for the same formats. Even within a particular system, there may be multiple choices of codec for the same output format—for example, Digital Rapids solutions offer multiple codec implementations for the H.264 format, with some optimized for particular applications or target devices.

While the transcoding offerings from many vendors have similar lists of processes that can be applied to video, the quality of the results can vary considerably between them. This is true even when the underlying technology powering competing systems is from the same codec vendor, because there's more to creating great quality than just the compression algorithms.

Balancing Speed and Quality

Some transcoding systems by design or by default emphasize encoding speed over encoding quality. Others emphasize quality over speed. And the tradeoff between quality and speed isn't necessarily proportional—getting a 10- to 15-percent boost in quality could have a 50-percent or greater impact on speed, because the encoding features associated with boosting quality are typically computationally intensive.

Transcoding systems expose various parameters to the end user for adjusting the balance between speed and quality. Such parameters are usually format- and codec-specific. For example:

- The "Quality" setting of the codec—it could be labeled "Encoding Complexity," "Quality/Performance Balance," or simply "Quality." To the user, the choices are fairly high-level (in some cases, the choice is a scale from 1 to 50; in others, it's a "Good/Very Good/Best/Maximum" selection). These settings affect under-the-hood processing, such as the precision/strength/depth/choice of algorithm for, say, motion estimation or in-loop deblocking filtering. Higher settings usually take longer to encode.
- The choice between CABAC and CAVLC entropy encoding—specific to H.264 Main and High Profile encoding, CABAC (Context Adaptive Binary Arithmetic Coding) requires a feedback mechanism that can add latency to the encoding process, but CABAC can be up to 10-20 percent more efficient than CAVLC (Context Adaptive Variable Length Coding).
- The choice of 1-pass vs. 2-pass VBR—while many high-end applications require tight constant bitrate (CBR) encoding, variable bitrate (VBR) encoding is still performed for less stringent applications like web video, and performing two-pass VBR encoding takes more time than single-pass VBR.
- The use of B-frames, in addition to I-frames and P-frames—increases compression efficiency, but slows down performance.
- Reference frames—this setting adjusts the distance the encoder will look backward or forward when creating B-frames or P-frames—in H.264, up to 16 frames can be references for any one frame. Again, higher settings slow things down.
- Scene change detection—improves quality where there's a significant visual change in the scene, but slows things down because the content is being analyzed to detect that change.

As a rule of thumb, increased compression efficiency, that is, achieving higher quality at a given bitrate is computationally expensive. Turning on all of the bells and whistles of a given codec means it will take longer to encode a given file.

Preprocessing (operations performed before compression begins) can also impact encoding speed—image scaling, color space conversion, graphic overlays, forensic watermarking and inverse telecine are some examples.

Compliance, the Other Performance Factor

Beyond image and audio quality, output compliance plays a major role in whether your content will display properly on devices that require precise adherence to standards or protocols such as mobile phones, IPTV systems, set-top boxes, video servers and streaming systems. Such compliance could be based on specific subsets of an overall format specification, particular metadata needed in the output package or even bitrate accuracy to the target data rates. For example, VOD files created with any method that doesn't result in a true, tight CBR stream are almost guaranteed to encounter compatibility issues with many devices, so many cable operators, broadcasters and telcos simply won't allow such files on their servers.

The Role of Storage in Performance

Of course, if your storage system can't keep up with the demands imposed on it by a high-volume transcoding operation, both the speed and the quality of your final output will suffer. Data throughput performance will be stressed, for example, in enterprise-scale transcoding operations that utilize tens or hundreds of transcoding clients hosted on modern multi-threaded, multi-core CPUs. Converting a number of large (~300GB or greater) source files to a much larger number of smaller target files requires a storage platform with capacity for high input/output operations per second (IOPS), as well as the capacity to perform many multiple concurrent read/write operations. Large-scale one-to-many encoding, for example, might take place when a media company encodes each of a high volume of source files into several outputs for delivery across multiple platforms (e.g. mobile phone, PCs, game consoles, etc.) or adaptive bit rate delivery.

Sometimes, it's advantageous to "split" an encoding job across multiple transcoding hosts. Some storage solutions don't provide the ability to accommodate the kinds of load balancing this shift in workload requires to get optimal performance. Intelligent NAS solutions such as BlueArc Network Storage Systems provide comprehensive virtualization tools that simplify the administration of file system functions, for example, allowing storage read blocks to be resized and fine-tuned for optimal transcoding performance.

In addition, BlueArc Network Storage Systems are able to provide optimal data throughput running over a 10GbE network, even as more transcoding clients are added to the system, thanks to a unique system architecture and Dynamic Read Caching capabilities that instantly copy active files from lower performance storage tiers to a high-performance cache storage tier—SAS, Fibre Channel or even SSD, for example—for use across physical or virtual servers. This aggregates bandwidth and improves response time to prevent quality of service issues during spikes in demand.

QoS is another area where storage performance plays a critical role across the entire content creation and delivery ecosystem. This is especially true when modern codecs that rely on reference frames to achieve efficiency are being read or written—a dropped reference frame will result in loss of picture until the next reference frame is decoded.

Room to Grow

Scalability is another key consideration for high-volume transcoding operations. As audience appetite for on-demand content grows, content providers need to scale their media ingest, encoding and transcoding operations to accommodate growth at every stage of their production and delivery workflows. At the production level, for example:

- File-based production formats for feature films are getting increasingly ambitious—filmmakers are embracing 2K, 4K and higher resolution digital formats as well as stereoscopic 3D, which has a voracious appetite for storage—double the amount required for traditional, 2D planar films.
- Studios still need to encode analog, film-based source footage to digital formats like DPX that are suitable for sharing between post-production and special effects houses, and the need to move huge files between facilities over closed networks quickly with maximum data fidelity is growing, as is the need to digitally archive finished footage.

At the delivery level, content providers need to be poised to leverage such things as:

- Increasing delivery network bandwidth, which is continuing to grow across delivery channels as Internet service providers build out their high-speed network infrastructures, the 4G mobile broadband network rolls out and so on.
- Emerging delivery formats and channels.
- New processes that can be applied simultaneously with transcoding like digital watermarking and fingerprinting.

BlueArc and Digital Rapids Solutions

Employing highly scalable transcoding and storage solutions that work hand-in-hand, such as Digital Rapids Transcode Manager software and BlueArc Network Attached Storage (NAS) solutions, allows content creators and providers to focus on their core business rather than constantly having to worry about building out and supporting additional infrastructure.

Digital Rapids Transcode Manager

Digital Rapids Transcode Manager software provides enterprise-level workflow management for transcoding high volumes of media files to multiple output formats. Transcode Manager works with multiple Transcode Engines, combining intelligent load balancing and efficient use of resources to boost throughput and streamline operations—resulting in more profitable media publishing.

Digital Rapids Transcode Manager goes beyond standard load balancing mechanisms with intelligent job assignment, predicting which node can most efficiently process a given job. As transcoding operations grow, newer host systems will feature improved processor capabilities, making certain host systems better for particular transcoding jobs than others. The Transcode Manager server tracks the performance of all Transcode Engines for each type of encode, "learning" which systems are best for each and allocating jobs accordingly to optimize utilization and throughput. This sophisticated scoring system delivers greater efficiency than simple queuing or round-robin task allocation could provide.

Digital Rapids Transcode Manager also allocates the creation of multiple outputs from the same source file to minimize the impact on decoding, storage and network resources. Say you have four source files that need to be transcoded to 10 different formats. With four available Transcode Engines on the network, Transcode Manager will assign one job (and thus one source media file) to each of the four available Engines. Each Engine will then generate the 10 output formats from its source file. In doing so, the 10 encodes can "share" the decoding of the source file; all 10 types of output media can be encoded from a common pool of available in-memory decoded frames. This not only decreases source file decoding (each source file is decoded once, instead of

Distributed Transcoding: Digital Rapids Transcode Manager

Digital Rapids Transcode Manager distributed transcoding systems offer enterprise-level workflow management for high-volume transcoding.

- A highly efficient, scalable thread and stream management architecture optimized to take advantage of systems with multiple processor cores ensures efficient use of all available processing power in multi-core, multi-threaded systems.
- Codecs are deeply configurable and finely tuned for superior quality while maintaining exceptional performance.
- Transcode Manager can scale to meet the needs of even the most demanding enterprise-level transcoding operations running up to an unlimited number of Transcode Engines.
- Advanced fault tolerance and workflow features increase reliability and automation while integrating seamlessly into existing enterprise operations.

four times), but also is very beneficial for storage and network resources. Had each output format been encoded on a separate Transcode Engine, each source file would have been read across the network 10 times.

There are specific cases where this may not be the ideal approach. For example, if a source file is being encoded into 10 extremely complex and processor-intensive outputs, the decoding time and network/storage benefits of performing all of the encoding on a single Transcode Engine may be insignificant compared to the advantage of encoding each output on a separate high-performance system. Transcode Manager provides the flexibility needed to handle such scenarios, and lets you choose to have Transcode Manager split the multiple outputs into separate jobs, each with one output format from the source file. Alternatively, the transcoding settings can be crafted to split up the output formats in an operator-preferred fashion. In both of these scenarios, a high-performance, intelligent storage system such as BlueArc Mercury Server can make short work of the increased data throughput.

All Digital Rapids encoding solutions, including Transcode Manager, feature a highly efficient, scalable thread and stream management architecture optimized to take advantage of systems with multiple processor cores.

Digital Rapids Transcode Manager also incorporates deeply configurable codecs optimized and finely tuned for superior quality while still maintaining exceptional performance.

Transcode Manager features automated failure recovery for automated reliability. The Transcode Manager server monitors the status of its networked Engines, attempting to restart any failed transcodes. If a failure is detected and the transcode cannot be successfully restarted on that system, the job is resubmitted to another Engine, ensuring that the task is completed (job-level failover). Server-level failover automatically transfers control to a backup server if the primary Transcode Manager FE server fails, ensuring that transcoding operations continue uninterrupted.

Digital Rapids capture/encoding and transcoding solutions also feature advanced software plug-ins that deliver superior quality, performance and/or functionality. The Digital Rapids Scaler, for example, provides better scaling quality and reduced CPU overhead compared to the built-in scaling typical in popular codecs. Standard and optional plug-ins including video and graphic overlay, adaptive Inverse Telecine, color space conversion, forensic watermarking and more complement features such as automated clip trimming and adding head/tail clips.

In addition to being able to create files that conform to industry specifications (such as compatibility with the CableLabs® content encoding specification), Transcode Manager output files adhere to the exacting specifications, stream and bit rate conformance required by professional applications and devices such as set-top boxes.

When you need to increase transcoding capacity, new systems running Transcode Engine software-based transcoding instances are automatically discovered on the network and can be added to a Transcode Manager installation without interruption to existing transcoding nodes and jobs. Lower deployment costs, faster ramp up time for high volume jobs and greater end to end throughput add up to higher-volume media encoding and faster turnaround.

Digital Rapids Transcode Manager's automated capabilities significantly reduce the operational overhead and costs associated with manual management and monitoring of high-volume media transcoding, while Transcode Manager's cost of ownership (both initial software investment and recurring maintenance fees) is more affordable than most other enterprise-level transcoding solutions.

BlueArc Storage Platform

Designed for high-performance, mid-range customers and up, the BlueArc Storage Platforms employ field programmable gate arrays (FPGAs) and traditional multi-core processors to optimize and separate data movement and management processes that normally compete for system resources—and impact transcoding performance. Rather than using shared buses and shared memory that require arbitration and can cause significant performance fluctuations and latency, data is transferred between logical blocks in a point-to-point fashion over dedicated data transfer paths. As a result, the platform's Hybrid-Core Architecture supports advanced file system functions while preventing conflicts or bottlenecks and ensuring consistently high performance required by transcoding applications that rely on high throughput and great numbers of concurrent read/write operations.

With capabilities that go well beyond conventional PC or processor-based servers, Mercury's unique Hardware Accelerated File system delivers no-compromise system performance and scalability, while multiple levels of virtualization overcome the complexities of large-scale file system management and data availability.

The Cluster Namespace feature makes it easy to assign and reassign file systems to virtual servers and physical nodes as usage and performance requirements change, without impacting user access to files and data. With their ability to virtualize up to 2 petabytes in a single namespace of usable, tiered storage, BlueArc Mercury servers can scale with the growing data requirements associated with digital content creation and distribution.

Mercury also logically organizes RAID storage into shared, virtualized storage pools from which multiple file systems, physical servers and virtual servers can be provisioned. Dynamic Storage Balancing stripes the data set across the maximum number of available drive spindles, optimizing performance and resource utilization while eliminating hot spots. As capacity is added, file blocks are redistributed across the new spindles to immediately improve performance while maintaining a high level of resource utilization. This is especially helpful in large-scale transcoding operations, when source files can be tens or hundreds of gigabytes. Dynamic Storage Balancing minimizes read/write head seek times and improves performance.

Mercury's Intelligent Tiered Storage allows organizations to reduce infrastructure costs by automatically placing data on the most cost-effective or highest utility tier of storage without compromising accessibility. For example, high-performance drives can be used to hold high-priority content for on-demand transcoding, moderate-performance drives can be used to hold secondary content and archival footage can be stored on lower-performance systems. The Mercury Storage Ecosystem supports best-in-class storage arrays from multiple vendors that include SSD, FC, SAS and SATA drives, as well as leading tape libraries.

To scale capacity and performance, most storage systems require an expensive, "forklift" upgrade. BlueArc's scalable, modular solution eliminates these painful hardware replacement costs and allows enterprise and media companies to upgrade server modules and capacity without replacing the entire system.

In short, the BlueArc architecture is designed to overcome the data network bottlenecks that plague lesser storage systems in transcoding operations and delivers superior performance without sacrificing scalability, storage bandwidth or overall capacity.

Network Storage: The BlueArc solution

BlueArc intelligent high-performance storage helps deal with the major pain points in a transcoding facility.

- Intelligent high-performance NAS allows transcoding system storage to avoid the bottlenecks and latency issues that are common with less intelligent storage systems.
- I/O performance scales to meet the needs of concurrent users running multiple encoding engines.
- High performance combined with data migration software allows BlueArc network storage solutions to support processing and copying large video and source files required in enterprise-scale transcoding workflows.
- State-of-the-art system stability and reliability increase the efficiency and effectiveness of transcoding workflows.

Performance Test

To measure the performance of BlueArc Storage in a transcoding pipeline running Digital Rapids transcoding technology a test environment based on real-world, enterprise-scale workloads was set up.

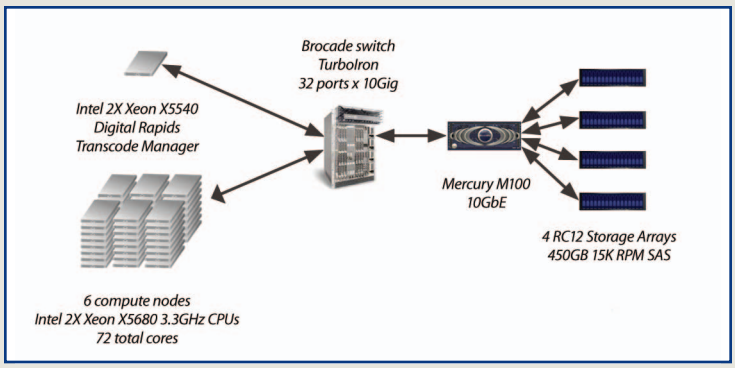

Six systems based on Intel 2X Xeon X5680 processor series, multi-threaded, multi-core 3.3GHz CPUs with 24GB RAM were used to host 12 Digital Rapids Transcode Engines, and a single system based on an Intel 2X 5540 processor series 2.5GHz CPU with 24GB RAM ran Digital Rapids Transcode Manager software.

A 10Gb Ethernet network connection provided end-to-end connectivity between the hosts, switch and a BlueArc Mercury 100 server. The Mercury 100 consisted of 48 15,000-rpm SAS drives configured as a 4TB file system. Additionally, the BlueArc Mercury "read-ahead" function was enabled and the WFS-2 file system utilized 32K blocks.

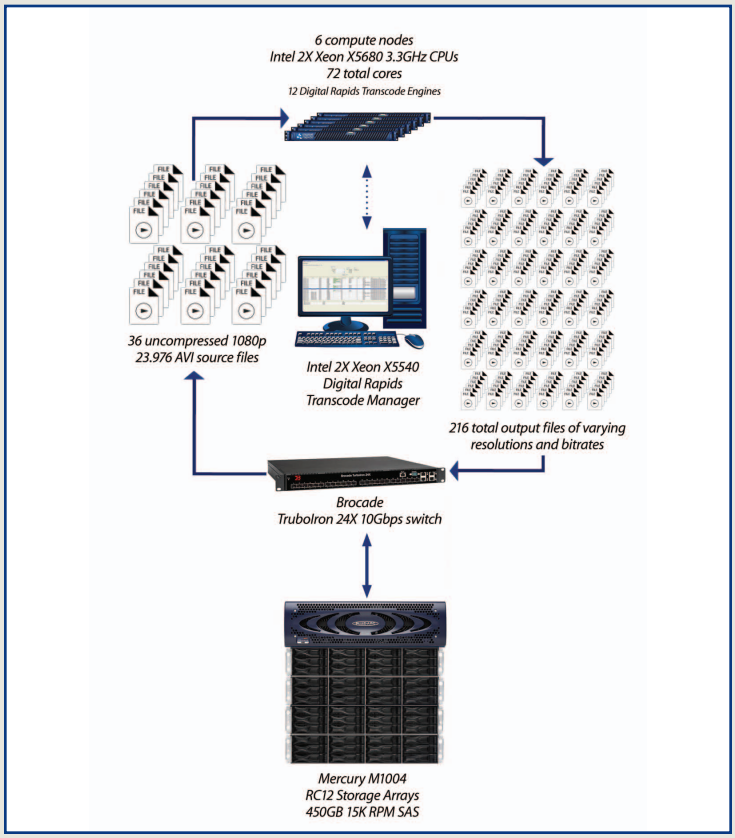

The transcoding workload consisted of 36 uncompressed 1080p23.976 AVI source files, each of which was transcoded to six different resolutions/bitrates using H.264 video and AAC audio compression in an MP4 container: one 1920p, three 720p at varying bit rates, and two 480p at varying bit rates. The encoding parameters used were the actual profiles used by a major studio to create premium deliverables for audience distribution.

Target resolutions, bitrates and formats used in the test:

1920x1080p23 8Mbps
1280x720p23 8Mbps
1280x720p23 6Mbps
1280x720p23 3Mbps
720x480p23 1.5Mbps
720x480p23 1Mbps

The source files totaled 2.04TB and 6 hours, 16 minutes and 12 seconds in duration. After transcoding, the 216 output files totaled 68.1GB.

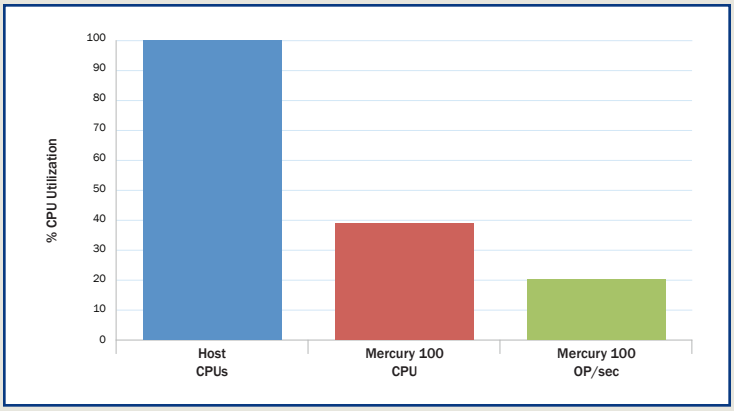

Total time to completion of all transcoded outputs: 97 minutes, 8 seconds
Average CPU utilization of Transcode Engines hosts: 100%
Mercury 100 CPU utilization: 39%
Average number of IOPS: 14,179
Mercury 10GbE throughput: 3,896Mbps

Conclusions

These results clearly demonstrate that Digital Rapids transcoding systems benefit greatly from BlueArc NAS solutions, which offer a distinct advantage over traditional storage systems when used in large-scale transcoding pipelines. Overall system throughput was more than up to the encoding workload, and the NAS system had plenty of IOPS and CPU horsepower to spare—proof that the highly scalable storage system can easily accommodate a larger number of host CPUs, Encoding Engines and heavier workloads.

Summary

High-volume, enterprise-level media transcoding requires high-performance, scalable transcoding software and storage solutions. BlueArc Network Storage Systems in conjunction with Digital Rapids transcoding software help eliminate bottlenecks, leverage all of the available processing power in modern multi-core, multi-threaded CPU architectures and deliver maximum throughput with solid stability that increase the efficiency and effectiveness of transcoding workflows.

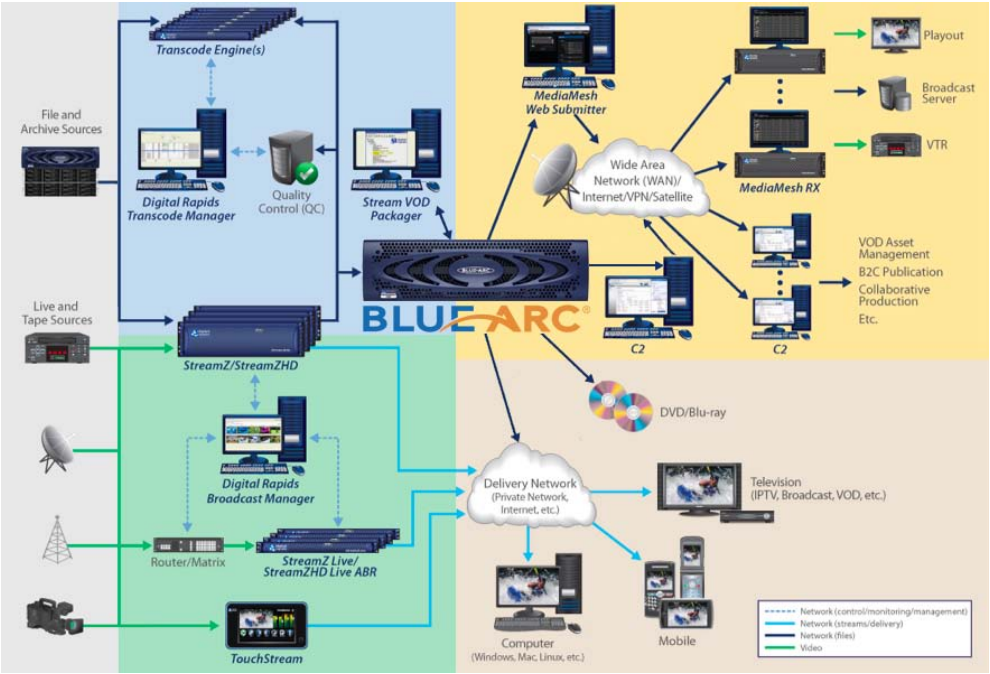

BlueArc Storage solutions like Mercury integrate seamlessly into Digital Rapids media workflows, enabling high-volume, high-throughput media ingest, encoding and transcoding that can scale to meet the needs of entertainment and media companies that need to deliver content on demand, to any device and with maximum quality.

The Usual Suspects

When building an enterprise-level ecosystem for media ingest, production, encoding, transcoding and delivery it's essential to consider the impact of each component on overall system performance. Bottlenecks, after all, can occur at just about any level.

When the amount of content moving through a transcoding farm is massive—on the order of many terabytes or petabytes—seemingly small latencies add up fast. Possible bottlenecks in a transcoding ecosystem's hardware and software include:

- The network fabric—for example, GbE connections are acceptable up to a point, but large-scale transcoding operations greatly benefit from 10GbE connections, which supply the throughput necessary for processing and copying large video source files.
- Inefficient CPU usage—this is sometimes unavoidable, say because delivering content to a particular audience requires using a legacy codec that isn't multi-threaded. Codecs that aren't multi-threaded don't take advantage of all of the processing power that the latest multi-core, multi-threaded CPUs have to offer. Additionally, not all transcoding solutions take full advantage of systems with multiple processor cores.
- The storage system's ability to keep up with high-volume simultaneous I/O operations—some storage solutions claim to be able to handle a high number of IOPS, but in reality fail when tasked with performing huge numbers of simultaneous read/write operations, a situation common in high-volume transcoding where multiple concurrent transcoding clients are used to create a large volume of small files.
- Inefficient use of heterogeneous systems—when companies choose to add to their existing transcoding system rather than swap all of their old gear for new, each new generation of hardware brings new capabilities that coexist with legacy hardware. Addressing this optimally requires intelligent load balancing—distributing encoding tasks to the nodes best suited for them to maximize throughput.
- Inadequate storage system stability, reliability and data integrity—unstable storage systems can lead to data corruption, failed writes and even complete data loss.

About BlueArc

At BlueArc, we know that lightning fast storage is the accelerator to your driving great visual results. In media and entertainment, we deliver stop watch metrics: reducing rendering times from hours to minutes, and file access times from minutes to seconds. In CGI development, quality and speed are vital. BlueArc lets artists create more revisions and dramatically speeds render farm output fueling extraordinary productions on schedule while reducing costs. BlueArc's Storage solutions scale with your digital workflow and BlueArc's tiered storage enables a more cost effective digital workflow. Scale Bigger, Store Smarter, and Accelerate Everything. Information about BlueArc solutions and services can be found at www.bluearc.com

About Digital Rapids

Digital Rapids provides the leading hardware and software solutions for transforming and delivering media, enabling the multi-platform experiences that are changing how audiences view content. Scaling from stand-alone appliances to global workflows, Digital Rapids solutions enable media professionals to maximize their productivity, quality and the value of their content. For more information see www.digital-rapids.com

BLUEARC®

BlueArc Corporation
Corporate Headquarters
50 Rio Robles
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202