



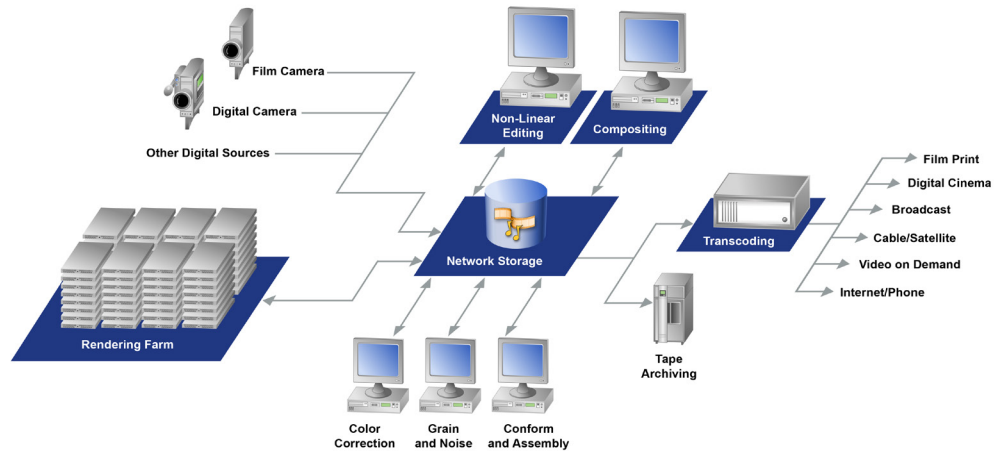
# Supercharging Media Workflows

Intelligent Storage to Enable Tomorrow's Blockbusters

Thomas Coughlin, Storage Industry Expert

## Modern Media Workflow

The increasing use of higher definition video content combined with broad use of stereoscopic content capture has driven storage capacity and bandwidth requirements for every element in the professional digital movie, TV special effects and video games workflow (see Figure 1). Between 2009 and 2015 we expect about a 10X increase in the required digital storage capacity and about a 12X growth in storage capacity shipments per year (from 4,094 PB to 47,291 PB) to support all storage requirements in the media and entertainment industry!



**Figure 1.** Professional Digital Movie Workflow

The rapid increase in required storage capacity, combined with performance requirements unique to video production, put pressures on production and post-production facilities to upgrade their facilities. This pressure for higher performance storage systems is occurring at the same time that economic factors force these same organizations to seek the most cost-effective solutions that control initial capital as well as operating costs.

Intelligent, high performance digital storage with appropriate content management software has become a requirement to handle key parts of the media production workflow. This is evident especially with a number of blockbuster movie post production processes. Similar digital storage requirement for the production of content are found in preparing content for electronic distribution. With the increasing number of distribution channels being used for an ever larger pool of content, the storage and transcoding of content to support all these distribution channels is daunting. Nevertheless efficient transcoding of finished source content into formats for different distribution channels is crucial for the realization of a return on investment.

This paper explains storage capacity demand and performance characteristics for three important areas of professional video workflows, including content distribution (dark boxes in Figure 1):

1. Rendering, 2. Non-linear Editing and Compositing, and 3. Transcoding

We will show how intelligent high performance storage solutions that meet these demands for capacity and performance can also improve the bottom line for video production facilities by increasing the efficiency and effectiveness of creative talent.

Digital storage for video post-production workflows must satisfy some important criteria that are very different from storage requirements for traditional IT operations. These requirements are:

- There can be no pauses in real-time streaming and no dropped frames
- Increasing resolution demands, particularly for the original content, drive very high data rates
- Latency requirements for data access varies in workflow and is lowest (lowest latency) where the creative process takes place—capture, editing and other post production work.



## Creating the Illusion

### Rendering

The creation of Computer Generated Imagery (CGI) played a significant part of the creation of many of today's most popular feature films. Rendering is a key element in the creation of most CGI and good rendering (as well as a good plot) leads to movies and other videos that are more immersive and generate greater consumer interest leading to greater revenue by the production company. CGI and other special effects have become more sophisticated as processing power and network bandwidth have allowed more sophisticated physical rendering and modeling programs to create ever more accurate representations of the real world or a world that, although fantastic, appears real.

Rendering is the process of generating an image from a model, by means of computer programs. The model is a description of three-dimensional objects in a strictly defined language or data structure and contains geometry, viewpoint, texture, lighting, and shading information. Rendering software applications from Autodesk and Pixar create realistic seeming CGI involving large amounts of data generation, requiring ever more digital storage capacity to support rendering, shading and physical modeling capabilities. To support and sustain modern rendering operations intelligent high performance digital storage systems are required. Figure 2 shows a typical rendering facility layout.

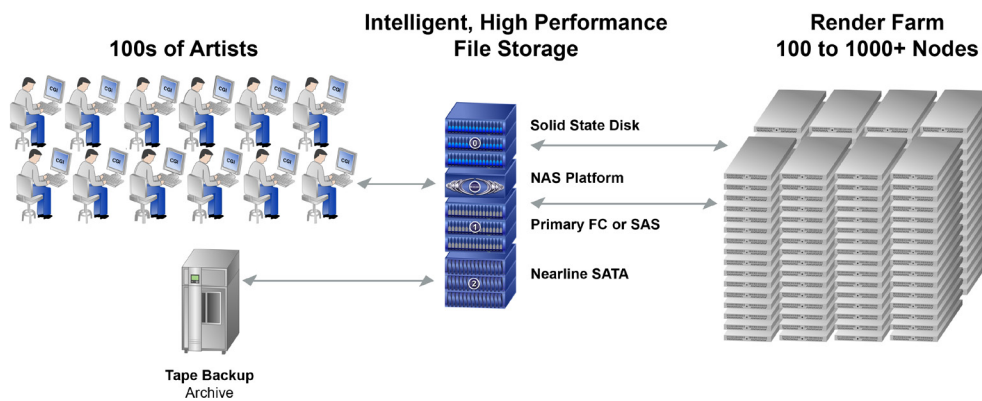


Figure 2. Layout of Rendering Facility

The advent of higher definition and 3D content has increased demand for storage and performance in CGI production. For example, although Weta Digital, a visual effects vendor, had 100 TB of storage to support production work on King Kong in 2005, the requirements for the special effects in Avatar with high resolution 3D images used in the entire film were several times greater. Avatar was produced with a “render wall” containing 35,000 CPU cores, 500 TB of BlueArc Titan high performance storage and 700 terabytes of nearline storage (1.2 PB of storage capacity total). With the popularity of 3D content, and Avatar in particular, demand for storage used for CGI and other special effects will increase considerably in the next few years.

Many networked storage systems have a hard time meeting the intermittent high intensity data access requirements for rendering and modeling. As a consequence of these issues there are a number of serious CGI pain points. The three main pain points with most networked storage systems are 1) low performance and lack of scaling relative to workflow requirements, 2) lack of automated storage tiering support to meet workflow requirements and 3) various storage and system management issues. We look at these issues a bit more in depth below:

- **Low performance and lack of scaling relative to workflow requirements**
  - Slow rendering forces nightly render scheduling times and keeping more version copies on TIER 1 (expensive high performance) disk arrays leading to higher production costs and less efficient creative work

## BlueArc Rendering Solution

BlueArc is a leading provider of scalable, high-performance intelligent network storage solutions for rendering and CGI environments with many high-profile customers. BlueArc was the choice of four-time Academy Award® winning visual effects creator Weta Digital for artistic rendering of groundbreaking animation and scene development for the major motion picture “Avatar.” This is because BlueArc Network Attached Storage (NAS) platforms can increase workflow efficiency and give a better return on investment using intelligent scalable, tiered storage. This leadership is the result of several advantages found using BlueArc digital storage and software in CGI workflow environments:

### HIGH PERFORMANCE AND SCALING

- Unique hardware (patent) architecture designed for “peak” performance protection as seen with your network switch architecture
- Built for mixed I/O types (throughput/IOPS) and data patterns thus providing “project / workflow protection
- Hardware file system approach provides highly optimized metadata capabilities for file searches as well as supporting deep directories. (support 16 million files per directory with 16 million directories)
- Provide reduced artist desktop wait times, which increase content iterations, improving quality and creative efficiency
- Resource management across shares, users and projects

### INTELLIGENT TIERED STORAGE SUPPORT

- Provides the widest range of TIER 1-3 disk/SSD based solutions for high end Studios
- Policy based migration between TIERED Storage as well as built in “Entertainment space optimized file compression/dedupe” capabilities
- Able to support “on the fly” migration of HOT content (read copy) of data from TIER 2/3 to TIER 1 disk/SSD local or cluster wide caches providing opportunities to meet sudden market demand such as providing historical performance content right after Michael Jackson’s unexpected death

- Enable 24/7 Rendering while reducing copies of versions on TIER 1 disk as well as increasing the quality and speed of creative work
- Optimized for tape backup, leveraging LAN Free backup architecture with mainstream backup software vendors and FC tape libraries

**STORAGE AND CONTENT MANAGEMENT SUPPORT**

- A single name space across TIERED storage (up to 4PB) and clustered heads (N+1)
- Supports large file system /directories (currently at 256TB) leading to less system administrator data asset management and frustrated users
- Unique “patented” object based file system provides the lowest file system utilization levels in the Entertainment space today. This leads to lower disk footprint while maintaining performance and capacity requirements
- Lowers the cost of hardware, software, support and licensing by consolidated storage
- Delivers robust mixed mode management across open systems with support for ADS, LDAP and NIS authentication environments.

**BlueArc Compositing Solution**

BlueArc scaled intelligent NAS platform provides tiered storage to support compositing work. These solutions support unique features for compositing that resolve the major pain points in compositing:

- Store and serve up millions of files across deep directories
- Provides the ability to scale not just workstation access load but also related composition software required to combine video images.
- Use of fast intelligent NAS allow system administrators to push storage for iterative operations to cost effective TIER 2 storage while maintaining a single view of the data across TIER 1 and TIER 2 storage for creative artists.

**CUSTOMER CASE STUDY**

A BlueArc customer was in the process of kicking off their own full feature film production project while continuing with their core TV animation customers projects. The customer was expected

- An inability to react to software workflow changes such as smaller files and random I/O (Maya network optimizations) leads to less efficient network traffic management and production inefficiency
- An inability to cost effectively scale existing rendering resources so a rendering house can increase the number of scene revisions impacts the volume and quality of work that can be done
- Some storage solutions are built for throughput only environments and do not scale well with hundreds of desktop/rendering nodes and this means they are inflexible to workflow and content size increases, impacting the flexibility of a rendering facility
- **Lack of automated storage tiering support to meet workflow requirements**
  - Not having LAN-free back up support can lead to poor tape backup performance bogging down network resources and reducing the efficiency of a rendering facility
  - Without a file system or user share control rendering node capacity management issues develop that reduces facility throughput and thus return on equipment investment
  - Without intelligent storage tiering and effective storage management software, as the size and complexity of the storage infrastructure increases serious data center room cooling and real estate issues can increase the costs of a rendering facility
- **Storage and content management issues**
  - Without proper storage and format management software there can be poor and/or no mixed mode capabilities across platforms (Apple, Windows, IRIX, and Linux)
  - Many storage environments lack a single name space and ability to provide a single view of the rendered data across all storage tiers impacting the ability of creative staff to find and use created content
  - Lack of storage and content management there are increased costs for hardware, software and administration in order to manage an increasing number of storage systems

Clearly new high performance storage architectures are needed for storage systems supporting rendering facilities. A storage system for these environments should be an intelligent storage system allowing scalable high performance operation, automated data migration to appropriate storage tiers as required and easy integration into rendering facilities to support the service level required in these intermittent data intensive environments and to help create tomorrow's blockbusters.

**Taming the Beast**

**Compositing and Non-Linear Editing**

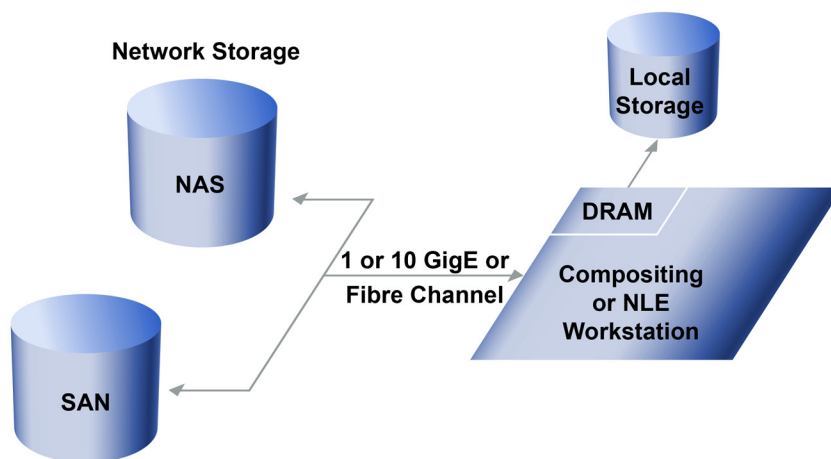
All professional video production uses digital non-linear editing (NLE) to put the content together in a final form ready for distribution. Compositing is another common step in video production to combine multiple videos to create a single “composite” video. Applications software from Apple, Adobe, Autodesk, Avid and several others are used for compositing and NLE. Compositing requires storage and serving of millions of digital files to workstations and needs the ability to scale without impact on performance. Hence a smart tiered scalable NAS platform designed to improve the efficiency and quality of compositing operations with these software packages is a necessity. It can also support content storage libraries for NLE facilities.

For NLE workstations the digital storage used may be a combination of local storage (including local caching and buffering with DRAM), a block-based SAN with shared file system such as Bright Systems which are either accessible from the video editing station and often NAS storage such as BlueArc platforms to keep less frequently accessed source content. Compositing with NAS platforms providing shared storage along with some local storage for caching and buffering provides an effective and efficient compositing work environment. Figure 3 illustrates the storage tiering that may be used to support an NLE or compositing workstation.

High performance NAS used for compositing or NLE should be an intelligent scaled NAS hardware such as that used so successfully in rendering environments. As the size of a post production facility gets larger direct attached and local storage capacity on individual workstations decreases while network storage capacity increases. This is because with more seats, the value of sharing content between users and post production operations increases, resulting in faster project throughput and more revenue as a function of the number of creative professionals in a facility. Many of the characteristics that are valuable for NAS storage used in rendering are also useful for compositing and NLE.

Dealing with the compositing pain points of many NAS storage systems allows for faster and more efficient post production processes, allowing production companies to get their products to their customers faster. The pain points impacting the iterative processes in compositing with many NAS systems are:

- Can't scale to meet the needs of compositing workflows leading to lower performance and usually additional costs to increase storage capacity or performance
- Don't provide adequate support of storage tiering behind the same NAS head leading to disjointed workflows and content locations and excessive support time to keep a workflow operating—this leads to additional costs and inefficiencies
- Slow loading times on workstations with many storage systems forces the use of multiple workstations per artist to keep their productivity at acceptable levels
- Problems with directory lookup due to metadata management issues when there are more than 100,000 files per directory
- Throughput and scaling issues limit the number of content iterations that an artist can make leading to fewer options for video content and thus lower overall finished product quality



**Figure 3.** Compositing or NLE Workstation Configuration

to scale their environment from 80 to well over 200 creative artists along with a major overhaul of their rendering capabilities. The customer expected to utilize a large array of composition software suites which included Nuke, Shake and Adobe After Effects. The customer required a high performance and scalable solution that could not just scale their composition capability but also the rendering side of the business without impacting the animation, and lighting aspects of their production workflow.

The customer's prior NAS storage was unable to both scale in size and provide sustainable performance during peak traffic events that occur near the end of a compositing project. The customer switched to a BlueArc Titan 3210 NAS storage solution which allowed them to store specific composition assets on different TIERS of storage depending on the performance and critical nature of the assets during the life cycle of each of their TV animation and/or FILM production work.

Furthermore the customer leveraged BlueArc unique Data Migrator capability of storing creative content iterations both on TIER 1 and 2 storage while providing a single file system view of the data. The data migrator capability moved data based on user defined policies such as the last access time and importance of the content during the life cycle of each of the customer's TV animation projects. This allowed the customer to leverage more of the TIER 1 disk investment for critical high performance random I/O traffic and the TIER 2 storage for assets requiring near-line storage and low I/O sequential performance loads.

The company's embrace of the BlueArc architecture allowed them to ramp up their TV workload while working on a full feature animation FILM project. Fast sustainable performance of the BlueArc NAS allowed them to produce far more content iterations in shorter time frames. Part of these improvements was due to faster access of their metadata directories that assisted them in their composition, animation, lighting and rendering workflows.

## BlueArc Transcoding Solution

BlueArc intelligent high performance storage deals 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
- Scales I/O performance to meet the needs of concurrent customers running multiple encoding applications
- High performance combined with data migration software allows it to support processing and copying large video and source files as required in transcoding workflow
- State of the art system stability and reliability to increase the efficiency and effectiveness of transcoding workflows
- Supports large file transfers using Aspera to aid in moving source and transcoded content to where it needs to reside

### CUSTOMER CASE STUDY

A BlueArc customer was launching a new streaming service for movies. The customer's encoding farm was using Digital Rapids as the transcoding software package and needed to have fast, centralized storage to perform the transcoding functions at the highest levels of storage performance. The prior storage infrastructure did not meet the performance requirements for this application. The BlueArc Titan 3210 was integrated into their workflow as the centralized storage system running on a 10GbE network. Tiered intelligent storage with the BlueArc NAS allowed the company to keep their high performance movie transcoding workflow on Serial Attached SCSI (SAS) storage and then as the transcoding was completed move the digital content using the BlueArc Data Migrator to secondary Serial ATA (SATA) storage. As a consequence the company was able to control the costs of their transcoding facility and improve their overall return on investment.

## Delivering the Goods

### Transcoding

Once digital content has been created it must be distributed to various distribution channels. Traditional distribution channels include broadcast, cable and satellite, digital cinema, video on demand (VOD) and more recently internet distribution as well as content distribution to smart mobile phones and other mobile internet devices (MIDs). Since each of these channels uses different content formats to match the transmission and other needs of the channel, transcoding of finished source content into formats for each of the distribution channels is crucial to the delivery of content and the realization of a return on investment for all the prior content production steps. Also note that transcoding generally works in a Windows CIFS environment, rather than NFS, using applications such as Digital Rapids and Rhozet.

Figure 4 shows a block diagram of a transcoding system using NAS storage with access to a content library to convert content into formats common for local editing and content distribution through internet and mobile networks.

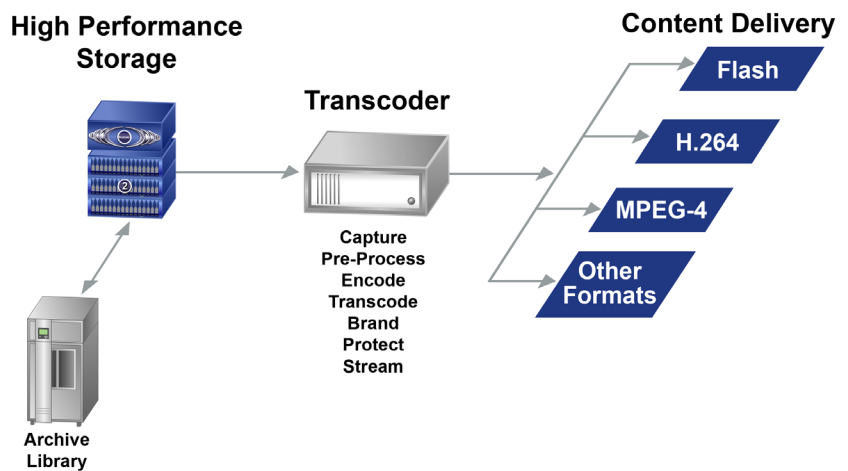
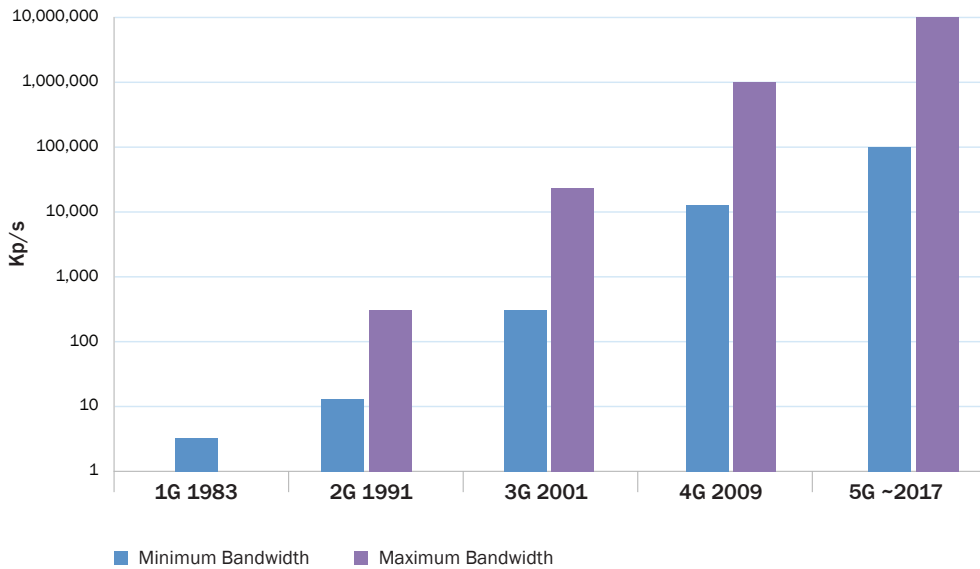


Figure 4. Transcoding for Internet and Mobile Content Delivery

Transcoded content is becoming richer to provide content for the increasing bandwidth available on mobile and internet networks. Figure 5 shows the evolution of minimum and maximum streaming bandwidth for mobile phone networks<sup>2</sup>. Increased bandwidth means larger storage capacity and an intelligent high performance NAS can be a valuable tool for transcoding content for delivery in these ever higher bandwidth networks.

Transcoding is important for the transfer of existing content libraries (in their original stored format) into a format that can be used for a new distribution channel. Transcoding of content libraries directly impacts time to market for distribution of movies to mobile phones and MIDs, game machines, home video servers and home networks. Providing more high quality legal content to these channels increases the popularity of consumer hardware and the channel offering content delivery services to these devices. This directly benefits the bottom line of the content owner and delivery companies.

Increasing the speed of transcoding requires intelligent accelerated access to content libraries. Since the video content is usually accessed as a file during the encoding operation, it is appropriate for a NAS storage system to be a key element in performing rapid transcoding.



**Figure 5. Mobile Network Minimum and Maximum Streaming Bandwidth**

Following are some major pain points for transcoding facilities:

- Inefficient workflows designed to work around issues with traditional slower storage systems leading to bottlenecks and latency issues in the network as well as the storage system
- Issues in scaling I/O performance from a number of concurrent clients using multiple encoding applications (such as from Rhozet, Inlet Fathom and Digital Rapids) and a large volume of small files
- Throughput for processing and copying large video and source files
- Storage system stability, reliability and data integrity leading to data corruption, failed writes and even complete data loss

In addition to providing rapid transcoding of existing digital content libraries, it may be possible to use intelligent high performance NAS to support analog to digital conversion and preservation projects. Transcoding of existing assets, either currently digital or still analog, is a direct path to increasing revenues through providing that content to distribution channels.

## Conclusions

Speedy intelligent NAS is a key part of today’s media workflows. NAS storage costs tend to be lower than storage costs for high performance SANs and NAS is attractive where the price/performance requirements combined with traffic management support make NAS an attractive option to media facility managers. Longer term storage of content in “active archives” is another growing application for NAS storage in media and entertainment workflows.

Getting video files into the increasing number of distribution channels has become very challenging. Each distribution channel tends to support a different content format. Making an existing content library available to a new distribution channel requires transcoding of the library content into these various formats. Since distribution is literally where the money is made in the media and entertainment market, this is an area where the return on investment for rapid transcoding is very fast. Intelligent NAS that supports rapid transcoding is a path towards greater revenue for media and entertainment companies.

High performance scalable NAS is useful by large as well as smaller companies with facilities doing rendering, compositing and transcoding. Storage tiering and content migration management with an intelligent NAS storage system helps content companies control their bottom line while increasing operational efficiency and final product quality.

Storage capacity and performance requirements are increasing as video resolution increases and with the introduction of an increasing amount of stereoscopic production. Increasing system demands and often constrained budgets limit the choices for storage systems that support operations such as rendering, compositing and transcoding. Increased performance of NAS storage using hardware acceleration combined with intelligent data migration and management software provides cost effective storage solutions to many parts of the professional media workflow. Intelligent high performance NAS is playing a pivotal role in enabling tomorrow's blockbusters.

### Footnotes

1 T. M. Coughlin, 2010 Digital Storage for Media and Entertainment Report, Coughlin Associates, <http://www.tomcoughlin.com/techpapers.htm>.

2 2010 DesignCon Keynote, Vincent Hu, Altera

### About BlueArc

BlueArc is a leading provider of high performance unified network storage systems to enterprise markets, as well as data intensive markets, such as electronic discovery, entertainment, federal government, higher education, Internet services, oil and gas and life sciences. Our products support both network attached storage, or NAS, and storage area network, or SAN, services on a converged network storage platform.

We enable companies to expand the ways they explore, discover, research, create, process and innovate in data-intensive environments. Our products replace complex and performance-limited products with high performance, scalable and easy to use systems capable of handling the most data intensive applications and environments. Further, we believe that our energy efficient design and our products' ability to consolidate legacy storage infrastructures, dramatically increases storage utilization rates and reduces our customers' total cost of ownership.

### About Thomas M. Coughlin

Tom Coughlin, President, Coughlin Associates has been working for over 30 years in the data storage industry as an engineer and senior level manager. He has over 60 publications and six patents to his credit. Tom is active with IDEMA, the IEEE, SMPTE, SNIA and other professional organizations. He is the founder and organizer of the Storage Visions Conference, a partner to the annual Consumer Electronics Show as well as the annual Creative Storage Conference. Coughlin Associates provides market and technology analysis on digital storage and applications as well as Data Storage Technical Consulting services. Publications include an annual report on data storage for Professional Media and Entertainment (2010 Digital Storage for Media and Entertainment Report). He is the author of Digital Storage in Consumer Electronics: The Essential Guide published by a division of Elsevier. He also publishes a quarterly newsletter (The Data Storage Technology Newsletter). For more information go to [www.tomcoughlin.com](http://www.tomcoughlin.com)

Thomas M. Coughlin  
408-978-8184  
[tom@tomcoughlin.com](mailto:tom@tomcoughlin.com)



**BlueArc Corporation**  
*Corporate Headquarters*  
50 Rio Robles Drive  
San Jose, CA 95134  
t 408 576 6600  
f 408 576 6601  
[www.bluearc.com](http://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