Helsinki Metropolia University of Applied Sciences
Degree Programme in Media Engineering

Jussi Uosukainen

Enhancing the Film Post Production Workflow

Bachelor's Thesis. 24 April 2010

Supervisor: Erkki Aalto, Head of Degree Programme

For providing me with the possibility to write my thesis about this wonderful subject, I would like to thank Erkki Aalto at Metropolia University of Applied Sciences and Tomi Nieminen at Generator Post.

Thanks are also in order to my teachers and fellow students at Evtek and Metropolia.

I would also like to thank my colleagues at ProAv, Humac and Generator Post for teaching me and allowing me grow to be good at my jobs.

To Eero and Joonas, thank you for putting up with a very tired dad at times.

But most of all I need to thank my wife Mirja, without whom I would have given up many years ago.

Espoo, June 2010, Jussi Uosukainen

Metropolia Ammattikorkeakoulu **Insinöörityön tiivistelma**

Tekijä Otsikko Sivumäärä Päivämäärä	Jussi Uosukainen Elokuvan korkeimman laatutason jälkituotantoprosessien hallinta 49 sivua 24.5.2010
Koulutusohjelma	Media Engineering
Tutkinto	Bachelor of Engineering
Ohjaava opettaja	yliopettaja Erkki Aalto

Insinöörityön tavoitteena oli tehostaa elokuvien jälkituotantoprosesseja suuressa jälkituotantoyrityksessä. Elokuvien tuottamisessa on viime vuosina siirrytty täysin digitaalisiin tekniikoihin. Jälkituotantotyöt on jo pitkään tehty digitaalisessa muodossa, mutta kuvausmediana on ollut filmi, samoin esitysformaattina. Nyt uudet toimijat ovat tuoneet digitaalisen kameratekniikan jokaisen elokuvantekijän ulottuville. Samalla ovat jakelutietkin digitalisoituneet, kun elokuvateatterit ovat uusineet kalustoaan yhteensopivaksi.

Uudet valmistajat ovat mullistaneet elokuvatuotannon kameratekniikan. Edullisen digitaalisen elokuvakameran suunnittelussa on otettu monta teknistä edistysaskelta, yhtenä suurimpana uusien pakkaustekniikoiden käyttäminen. Nämä pakkaustekniikat ovat hyvin raskaita ja vaativat monta erilaista työvaihetta ja muunnosta toimiakseen perinteisessä elokuvatuotannon ympäristössä. DCI-konsortion spesifioima digitaalinen jakeluformaatti parantaa jokaisen elokuvan katselijan katselunautintoa ja mahdollistaa 3D-elokuvien toistamisen minimaalisella ylimääräisellä panostuksella.

Nykyaikainen digitaalinen työkulku on yleistynyt prosessointitehon parantuessa. Työkulku on yleensä rakennettu toimimaan 2048 x 1556 -resoluutiossa käyttäen dpx-formaatin kuvia. Jälkituotannon eri vaiheet (materiaalin digitointi ja siirto, offline-editointi, online-editointi, värimäärittely, efektityöt ja masterointi jakeluun) on optimoitu tähän työkulkuun lähes kaikkialla.

Insinöörityön tilaajan tuotantoympäristössä oli haasteena olemassa olevan tekniikan käyttäminen uusien tekniikoiden kanssa. Samalla kasvavien tietomäärien kanssa toimimiseen tarvittiin uusia ratkaisuja. Keskitetyn tallennusverkon (SAN) asennus vähensi tarvetta siirtää materiaalia työvaiheiden välillä, kun siihen liitetyt työasemat pystyivät käsittelemään materiaalia suoraan jaetulta tallennuslevyltä. Hajautetun laskennan käyttöönotto nopeutti materiaalin käsittelyä ja vapautti työsemat operaattoritöihin.

Uudet tekniset ratkaisut paransivat työkulkua, nopeuttivat projektien valmistumista ja vähensivät manuaalisen työn määrää työkuluissa. Valitut tekniikat ovat osoittautuneet luotettaviksi ja hyvin skaalautuviksi.

Hakusanat	elokuvajakelu, elokuvatuotano, hajautettu laskenta, värimärittely, tallennusverkot, elokuvatekniikka, filmitekniikka

Helsinki Metropolia University of Applied Sciences **Abstract**

Author Title Pages Date	Jussi Uosukainen Enhancing the Film Post Production Workflow 49 pages 24.05.2010
Degree Programme	Media Engineering
Degree	Bachelor of Engineering
Supervisor	Erkki Aalto, Head of Degree Programme

The purpose of the thesis was to refine and improve the workflows of the biggest post production company in Finland. There is an ongoing transformation from film-based production to completely digital productions, where everything from image acquisition to delivery is done in digital formats. This means a lot of new oppotrunities but also many new problems for companies in the field.

New camera manufacturers have made significant changes in the field of cinematography. They offer filmmakers cameras that are cheaper and offer similar image quality to the cameras they are used to working with. These new digital cinema cameras use compressed images to make use of common storage mediums. In the postproduction process this material needs to be uncompressed into formats that are more friendly for the post production processes.

In the other end of the post production pipeline, Digital Cinema Initiatives (DCI) have created a standard that defines the way movie theaters can display movies in digital format. The DCI-specification uses JPEG2000 for compression of images, and standardises the interfaces on the playout systems to the projectors.

Advancements in IT-technologies can be used to improve the workflows used in post production. In this thesis, installing a Storage Area Network (SAN) in the facility greatly improved the workflows by decreasing the need to move material between systems and units. Using a distributed rendering system improved the speed of rendering material for other steps of the workflow. Both of these advancements were installed wither using existing hardware or with a minimal need to make changes in the existing systems.

The resulting systems have been in production use for some time, with great success. Both systems enhanced the way data flows through the facility in question.

Keywords	motionpicture technology, distributed rendering, color grading, SAN, visual effects, DCI, post production

Contents

1 Introduction

The whole motion picture industry is going through huge changes. New digital film cameras, the penetration of DI (Digital Intermediate) and the fast growth of DCI (Digital Cinema Initiatives) enabled cinemas mean that we are quickly approaching a completely digital production chain, from the cameras to post production and distribution. This means that there are completely new challenges awaiting the post production industry.

The digital post production process has greatly benefited from the advancements in computer technology. Almost all motion pictures produced today have al their frames transferred into digital form and manipulated. Images that would have filled whole hard drives just 15 years ago are now used 24 times a second and new camera formats change the whole way the postproduction process integrates into the whole production.

To properly understand how digital post production works, one needs to understand the basics of film technology and how the digital techniques have derived from the film processes that preceded them. Also the basics of new technologies are important, how the new cameras and more importantly how the new compression formats work. This is especially true for the Red-camera and their compression scheme, REDCODE.

New delivery formats are also a part of the post production workflow. Digital Cinema Initiatives have provided a standard for delivering digital copies to cinemas. The DCI-spec, based on JPEG2000 compression enables cinemas to receive digital cinema packages containing the images and sound, as well as subtitles in one package that fits on a portable hard drive. These packages are also very strongly encrypted and only work with a designated system, so that the packages do not end in general distribution on the internet.

The objective of the actual work is to document a post production workflow, from ingest to editorial, effects, grading and output that supports emerging formats and standards. Where are the obvious bottlenecks, how can they be averted? How can new technologies be efficiently used in conjunction with the existing tools? The client for the thesis is Generator Post, the leading post production company in Finland.

2 Film Technology

2.1 Prehistory

The history of film production can be said to have started with the brothers Lumière in december of 1895. This is the first time a commercial film was screened for a paying audience. In the preceding decades the technology had been worked on by several people in many countries, but the brothers Lumière were the ones to make moving pictures financially viable. During its first few decades, film was used as a speciality attraction, showing footage shot in other cities and continents. These first films were crudely assembled, cutting and gluing filmstrips together and playing back the resulting strip.[1,20]

2.2 Film Gauges

Most motion pictures shot within the last 100 years have been shot on some variant of the 35mm film format. The gauge was introduced for the first time in 1892 by Thomas Alva Edison, who had preceded the brothers Lumière in some aspects of creating the techniques to be used in film production and projection. After some prolonged patent litigation, it was standardized for film production in 1909. [2] Still today, over 100 years later, most films in cinemas are projected using 35mm film.

16mm film was introduced by Eastman-Kodak in 1923, and has since become the industry standard for documentary and high-end television work.

There are several other film gauges in use: 8mm is for purely consumer use, while 65mm and 70mm film gauges are mainly used in IMAX-theaters and other “speciality” venues. Bulk of all motion picture work is done in the two aforementioned formats.

2.2.1 35mm Versions and Aspect-Ratios

There have been several 35mm film versions on the market. Although the physical film has always had the same physical properties, there have been big differences in the actual size of the frame exposed. At first, the film was shot using the aperture provided by the area on the negative, 1:1,33 (the same as SD-television, 4:3). After television got popular in American homes, the movie studios needed to do something to lure the audience into the theatres. Two of the biggest things the cinema could provide over television were colour and a wider, bigger screen.

After several colour versions where several strips of film were run simultaneously through the same projector, thus creating a colour image. The most well known of these techniques was called Technicolor, which was based on three separate strips of film, each representing one colour and running in parallel. Eastman-Kodak company presented the real colour negative, introduced in 1950, starting the boom with colour motion pictures.[3]

The bigger screen was also developed, at the same time. It was decided to make the picture much wider, supposedly because it made it easier for cinemas to convert to a wider format. The movie makers were left to figure out how to create images in the new, vastly different format. The first option was to do a “hard-matte”, and cover either the top or the bottom of the image , thus ending up with an image that is wider than it is tall. However, this process compromises picture quality as the full frame is not utilized. Cropping was used to achieve an aspect ratio between 1,75 and 2:1.[4]

Several techniques (such as Cinemascope) utilize anamorphic squeezing. The film is shot using a specialized lens, which squeezes a widescreen image onto a normal-aspect ratio frame [1, 167]. The image is then projected using a similar lens, thus reversing the effect and creating a widescreen image of greater quality than cropping.[5]

2.3 Film Postproduction

Before digital postproduction became the norm in the 21st century, film postproduction was a very manual process. Effects and colour correction were done by optical and chemical processes. Many other post production steps were also possible with film, but working with the physical medium always made the processes expensive, difficult and error prone. A good example is Colour timing. [6,43]

A good example of traditional postproduction techniques is the way colour was added to film before a colour negative was introduced. The image was sent to a prism after the lens, and split two or more ways. Each image form the prism was then filmed to a different film thru a filter. Combining the filtered images later created a full colour image. [3;7]

Digital processes became the norm in the late 1990s, when processing power and storage systems caught up with the requirements of the motion picture industry. The biggest single change came with the advent of Digital Intermediate (DI) -systems.. These systems enables the whole motion picture to be digitized and finalized digitally before printing to film. [6,61]

3 RED One

The most important product of the new digital cameras is the RED One, manufactured by RED. Especially in smaller markets RED has won a big market share, due to great image quality and a cheap price.

3.1 The Company

RED was the first manufacturer to bring a digital camera capable of creating images comparable to 35mm film images to most motion picture creators. RED One completely rewrote the pricing for professional equipment and took most of the established names by surprise.[7]

Red was founded in 2005 by Jim Jannard, an entrepreneur who had a little earlier sold Oakley, a company manufacturing sunglasses, recreational wear and other similar things for several billion dollars. A long time camera geek, he wanted to set off creating the ultimate digital motion picture camera.

RED One was first announced in 2006 at the NAB trade show in Las Vegas. It was met with a lot of doubting and labelling as vapourware, since the claims by the company were quite unheard of in the marketplace. At NAB 2007 the company had three working prototypes available. The company starting taking deposits from prospective owners, and promised to start shipping cameras before the end of the year. In a short time, the company had received more than 500 deposits of 1000 USD.[8]

3.2 The Camera

RED One camera has 12 Megapixel CMOS Mysterium sensor, designed and built in-house. The sensor is physically 24.4mm x 13.7mm in size, and has 4900 x 2580 pixels. The biggest actual frame size shot is 4096 x 2304. The camera is equipped with a standard Panavision Lens-mount, with interchangeable mounts for Nikon and Canon cameras optional. The camera supports shooting at variable frame rates, up to 113 frames per second.

The price point of the camera is unique (All prices as of January 2010): $17500 for the camera body is about 1/10th of the price of the competition (Sony F23 is priced at around $150 000), which makes it affordable. Many filmmakers can now afford to buy their own cameras instead of renting. RED does not have a resales organization, instead all cameras need to purchased directly from the company. Support is heavily peer-to-peer driven, with several active forums giving support to users in all aspects of the workflow, from shooting to postproduction.[8]

RED has promised presented new cameras with improved resolution and CMOS-technology. These Epic and Scarlet cameras will have many more options available in terms of the sensor and other preferences. The resolutions vary between 3K and 28K.

3.2.1 REDCODE

The secret of the camera is in the proprietary recording format. An uncompressed 4K data stream (around 77 GB/min) could not be handily recorded onto any conventional medium, so compression is used to record the data onto normal hard drives and flash drives. This compression is called REDCODE, and denoted by a R3D extension in the file. REDCODE records the 4K image at only 2,2 GB/min, a 97% decrease in file size. It comes in two different versions, REDCODE28 (28MB/s) and REDCODE36 (36MB/s). REDCODE is based on wavelet compression.

RED likes to refer to REDCODE as a RAW format, but this is somewhat misleading. In the still-camera world, a raw-file means that all the information from the sensor is put in the file. Red has somehow turned that around, and here “raw” means that they have discarded most of the data available.

Proprietary REDCODE recording format means trouble in the post-production pipeline: The format is not supported by most NLEs, as native REDCODE support is only in the pipeline for most companies (the API was released fall 2008). Decoding into other formats takes quite a long time (in Generator Post internal testing, a 8-Core 3,2 GHz Mac Pro took between 2:1 and 10:1 realtime in decoding, all cores running, depending on end resolution).

With the new cameras new versions of the REDCODE have been introduced, like the “mathematically uncompressed” 1:2,5 compressed REDCODE. This is a completely different proposition to the heavily compressed REDCODE introduced with Red ONE.

How this will impact the postproduction chain is still not known, but it definatly means that the storage requirements are growing constantly.

3.2.2 Workflow

Working with the REDCODE files is a different workflow when compared to normal film or video based workflows. This is the basic workflow:

- Decode REDCODE Raw-files into intermediate format for offline editing. Usually ProRes422 is used, as it is a good compromise between image quality and file size.
- Use an XML to conform raw-files into DPX-images for online editing
- Do online editing and grading, print to film or compress for distribution

Online-editing is done using DPX files. In the process visual effects-shots are added, and the image is flattened for colour grading. The image is then rendered from the grading unit to the film printer or DCI system.

3.2.3 Problems

While using the camera, finnish companies have noticed several issues that are easily reproducible with the camera: low-light photography is not quite up to the level or traditional film, and highlights are sometimes clipped. There is also significant “rolling shutter” effect due to CMOS-technology: the camera does not expose the whole frame at once to the the sensor, instead it exposes a pixel line at a time. This means that during pans vertical lines can become distorted and slanted. Depth of Field becomes and issue when shooting something else than 4K. Both of these issues have been reproduced by camera owners across the world.

One of the bigger problems is storage. Companies are used to having the original negatives and film prints in storage, incase something needed to be re-cut etc. Now everything resides only in computer data, and most companies do not have a data storage and archival system in place in order to store the huge amounts of data generated during the photography and postproduction stages of a RED project. This is something many companies are figuring out at the moment.

4 Digital Cinema Initiatives

The Digital Cinema Initiatives (DCI) is an organization formed by motion picture studios and manufacturers of motion picture equipment. They have created the group to advance the digitalization of the cinemas around the world. [9]

Digital delivery of motion pictures has advantages to all the stakeholders in the industry. Distributors save great deals of money in both creating digital copies to cinemas (instead of film copies) and shipping costs (shipping thumbdrives or USB-disks instead of film reels). The distributor can also distribute the international versions of the material over the web instead of freight, saving a lot in the process as well.

The customer has a more uniform viewing experience, since there is not wear and tear on the source. The cinema owners save money on equipment upkeep, as digital equipment is easier to maintain then film equipment. Cinema owners also get the possibility of easily upgrading their systems to show 3D-films, which they can charge more for per person.

The main issue has been money. Cinema owners have been reluctant to upgrade theatres with their money, when the bulk of the actual short-term savings (distribution costs) would would go to the studios. But now the companies are working together, and digital screens are coming up everywhere. In Finland alone it is expected that the amount of digital screens will double during 2009 from 10 to 20 screens.

The final push to digital screening around the world came with James Camerons Avatar, which created a need for 3D-screening across the globe. At the time of writing this, the average times for manufacturers to ship digital cinema equipment from date of order is about 7 months, due to an avalanche of worldwide demand for systems.

4.1 DCI Specs

The Dci specification is based on the JPEG2000 format [6,116;10]. In the following tables, the most important details of the specification are listed. [6,54]

Image properties:				
Resolution	Framerate	Compression	Colour depth	Maximum bitra
2048 x 1080	48 fps 24 fps	JPEG2000	3 x 12 bits / pixel	250 Mbits/s
4096 x 2160	24 fps			

Table 1: DCI Image properties

Audio properties:			
Sample depth	Samplerate	Compression	Channels
24 bit	96 kHz	PCM	16

Table 2: DCI Audio properties

There have been tests done comparing normal modern compression schemes to JPEG2000, and JPEG2000 has proven itself very good, especially at high bit-rates. [11]

4.2 Security and Encryption

DCI has been built on the strict security measures. The Hollywood studios that have formed the initiative have insisted on very strict measures in order to protect their material from being shown in unauthorized venues.

Each movie package carries along a key for each projector the package is going to be shown on. Without the key, playback is impossible. New security features demand that each DCI-player has to have its image processing board completely tamper-proof, for a military specification. [12]

5 Digital Post Production

5.1 Digitalization of Postproduction

Motion picture postproduction has evolved in leaps and bounds since the second half of the 1990s. Advancements in the computer-industry often find their way into post production very quickly, because of the demands the post industry puts on the hardware. A single frame of film, scanned into the industry standard file, varies from 8MB to 12MB. A film shot has 24 of these a second, which demands quite a lot of the storage system in use. Compositing or applying corrections to these images requires a lot of processing power and working memory.

The first elements to come through the digital pipeline was computer-generated imagery (later cgi). The first cgi elements were introduced in the early 70s, with the final breakthrough finally happening in the end of the 80s with films such as The Abyss and Indiana Jones and Last Crusade bringing these elements to mainstream movies, the former with a realistic cgi monster, latter being the first one

Thus digitalization of the entire process has only happened quite recently, with computer hardware finally catching up with the post production industry.

5.2 Digital Images

A digital motion picture workflow requires high quality images which can faithfully reproduce the image recorded on traditional film stock. Several file sizes and formats have been tried and discarded, with the whole industry slowly adopting the same standard formats. Uncompressed 2K post production is usually performed using DPX-files.

DPX is a SMPTE standardized file, which is based on the Kodak-developed Cineon-format. The files can be in either linear or logarithmic colour space and make extensive use of look up tables (LUTs). [6,70;6,121]

5.2.1 Frame Sizes

Figure 1 shows the relative frame sizes between different images used in post production.

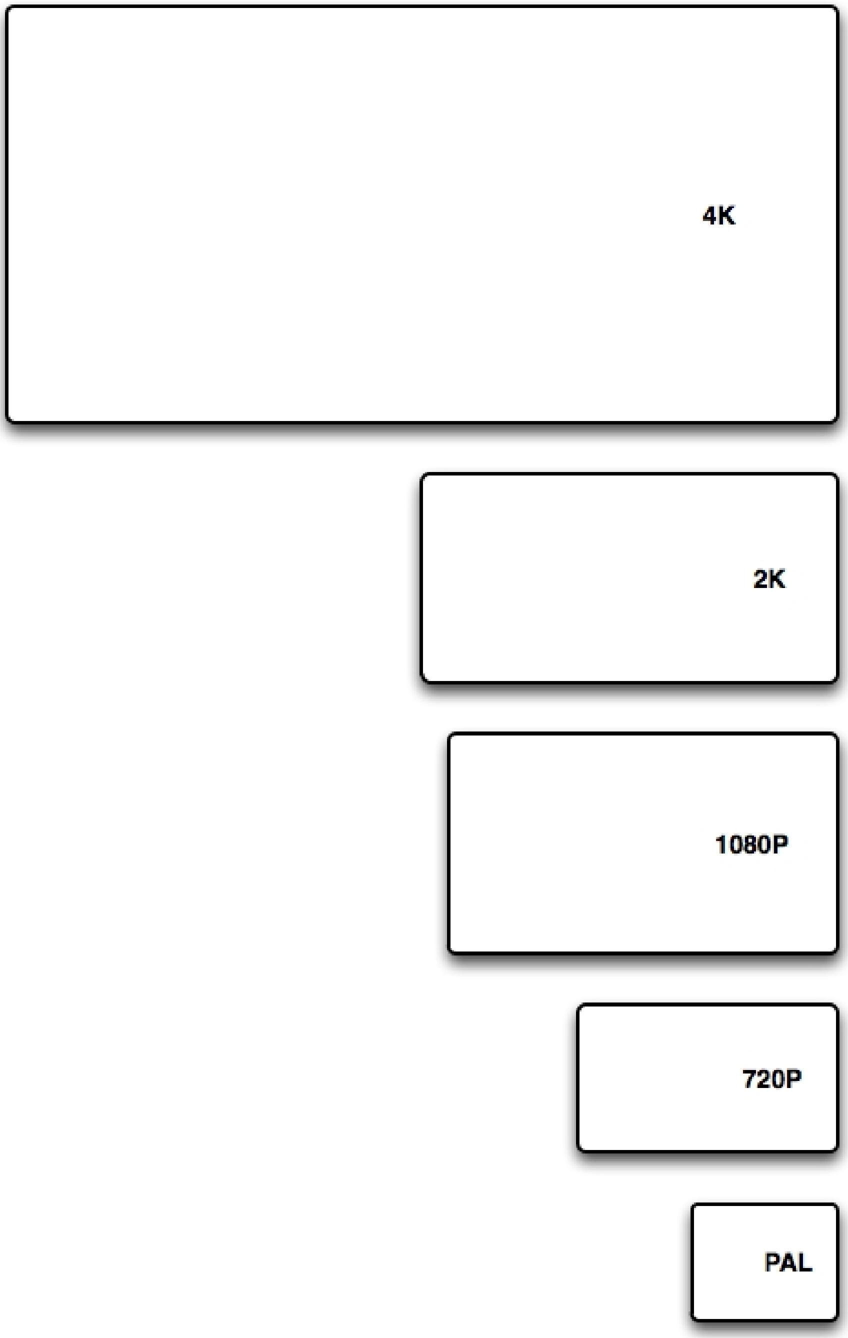

Figure 1: Relative framesizes

A normal PAL image used in standard definition broadcasts in Europe is the smallest frame in the picture, with dimensions of 720x576. There are two resolutions of HD: 1920x1080 (1080 from now on) and 1280x720 (720 from now on). HD is the standard for high resolution digital broadcasting, and has been widely adopted by NLE and camera manufacturers in professional, prosumer and consumer markets.

2K and 4K are formats used in the motion picture postproduction industry. The moniker 2K means the image is 2048 pixels wide and 4K means that the image is 4096 pixels wide [13,317]. There are several conventions used when scanning and processing frames that are cropped or anamorphic, but these are the more known and used frame sizes [12,85]:

Scanning Resolution	Full Frame	CinemaScope
4K	4,096 x 3,112	3,656 x 3,112
2K	2,048 x 1,556	1,828 x 1,556
HD	1920 x 1080	N/A
Aspect Ratio	1,316	2,35

Table 1: Frame sizes for Digital formats

5.2.2 File Sizes

Here are some examples of the sizes of files that post production works with. The File sizes are calculated by the formula H x W x Depth x 3.

Scanning Resolution	16Bit / frame	16Bit / min	10Bit / frame	10Bit / min
4K	55,27 MB	77,73 GB	36,85 MB	51,82 GB
2K	12,66 MB	17,80 GB	9,1 MB	11,87 GB
HD	11,87 MB	16,69 GB	7,91 MB	11,12 GB

Table 2: File sizes for Digital formats

The table clearly shows the big jump in file size when going from 2K production to 4K production. 4K requires approximately 4 times the capacity in both storage requirements but also in processing power.

5.3 IT-Technologies

IT services and related technologies are at the heart of a modern Postproduction facility. Every frame is at some point digitized and modified in a computer. The advances in IT-technologies have made it possible to work with better-resolution files, at a faster pace. Here I outline a couple of the technologies that have changed the way postproduction works.

5.3.1 Storage Area Networks

SANs have been a huge advancement for the post production industry. Servers have been used for sharing files between several computers for quite some time, but only SANs have enabled several workstations to access the same material in real time (which means viewing the material at the correct frame rate), concurrently.

Figure 2 (below) shows a basic diagram of a SAN and the two separate networks it is built upon.

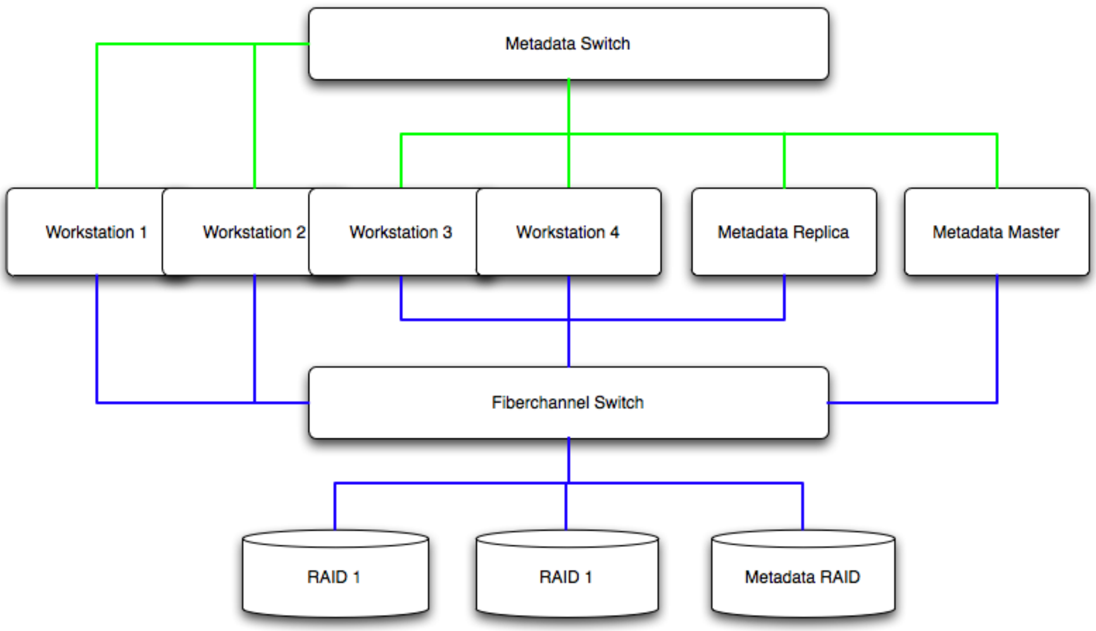

Figure 2: SAN - an overview

Before SAN technology came affordable, file sharing was done using normal file-sharing protocols, like NFS (Network File System), SMB (Server Message Block), AFP (Apple Filing Protocol) or even FTP (File Transfer Protocol). These file-sharing protocols enable users to read files from the disk of another computer or server. Depending on the hardware used, the speeds have been convenient enough to share content between workstations. But in order to work on the files in real time they have been moved onto the local storage of the workstation used.[14,8]

Usually the speed limitations of these protocols are easily noticeable. Each workstation reading material from the server uses the same network connection to connect to the server, which means that the more connections the server has, the slower each workstation receives (or uploads) material. The speed advancements in storage technology were lost in the network topology. [14,12]

Storage Area Networks changed this. A basic Fiberchannel SAN allows each workstation to connect straight to the storage system, with fiberchannel speeds. This way the network does not slow down the connections to the storage. [6,149;14,20] Now the bottleneck has been moved from the network to the actual storage subsystem performance. And this performance is easier to improve, just adding the amount of disks (or “spindles”, as they are sometimes called) usually brings around a huge increase in performance.

5.3.2 Processing Clusters

Processing clusters or render farms are clusters of computers (nodes) dedicated for rendering. Their usage has been the norm for all 3D work for several years, but with applications like Apple Qmaster there are several ways to utilize clustering to enhance rendering speeds for motion picture-work as well.

The advantage received from a render farm is not as easily calculated as adding together the number of cpus in the system. Each new node adds the amount of power lost to overhead needed for sharing the material to several nodes. A SAN system reduces the overhead drastically, eliminating the overhead that was lost in normal network applications (like NFS or SMB).

Another question is whether a dedicated farm is needed. In a dedicated farm there are several machines that are only used for rendering in the render farm. A general farm allows machines that are used for something else (office computers, workstations with less use) to be used for calculations when the machine is not used for it's primary purpose. A good renderfarm manager (like Autodesk Backburner) allows machines to have set timetables when they are available for rendering jobs, while simpler systems (like the Apple Xgrid) only have a simple option of having a machine available always or only when idle (which in Xgrids case means 15 minutes of inactivity of the mouse and keyboard), which in turn leads to interesting results if the machine accepting the remote rendering job is already rendering something without user input such as a very long Final Cut or Shake render.

6 Post Production Workflows

The following figure shows the different steps of the post production process in sequential order.

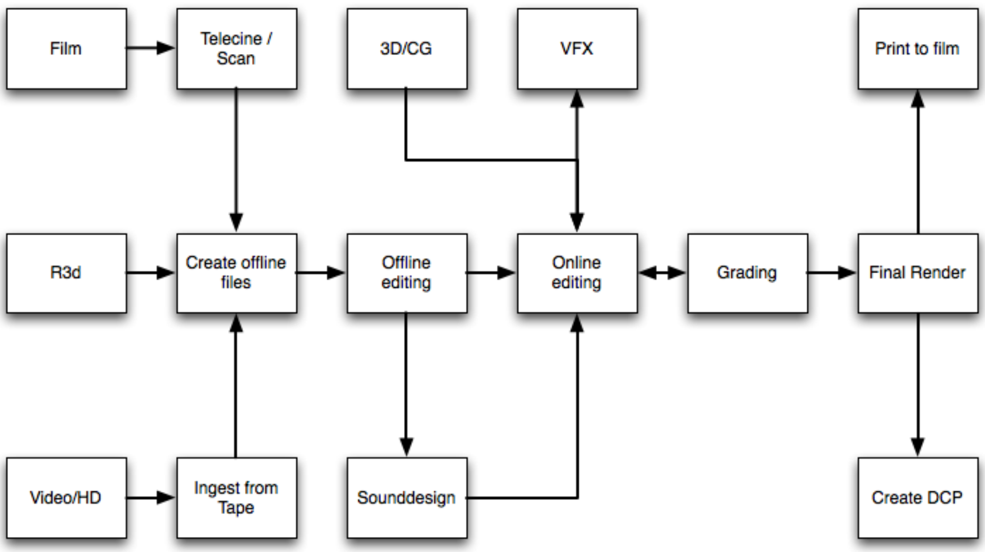

Figure 3: Digital post production diagram

On the left there are the ingest and input operations, in the middle the are the modification processes and on the right there are the delivery processes. Not all workflows go through all steps, but all steps can be categorised under one of the categories presented here.

6.1 Ingest Operations

Ingesting is the operation where the material is moved to the workflow of the post production facility. Sometimes ingest is simply recording files from tape, sometimes it is just copying files from an external disk, sometimes some operations are needed before the material is ready for use, like debayering REDCODE files.

6.2 Offline-edit

Offline editing is the artistic editorial process, where an editor chooses the material to be used in the program, and creates the structure, pacing and story of the piece, usually working closely with the director. To conserve hard disk space, the offline edit usually done using compressed frames, like DV or Apple ProRes422. These frames are usually received from a telecine scanner via videotape.

From a technical post-production point of view, offline editing is not usually “in the pipeline”. Offline editing is usually a very long process not done at the post production facility, since the compressed footage can be edited on a pretty simple machine. Only the result of the offline-edit, EDL (Edit Decision List) is handled in the postproduction workflow. [6,136]

6.3 Scanning

For online use, only the needed frames (the necessary frames are decided in the offline-process) from film are either scanned straight into files (usually DPX) using a data scanner. An other option is to use a telecine scanner and scan the frames onto a HD videotape (HDCAM or HDCAM SR), then ingest the frames from the tape to disk. [6,168] Data scanning enables resolutions greater than HD (2K and 4K), and is becoming the preferred method for scanning.

6.4 Online Editing

Online-editing is the process of arranging the full quality material into the order the offline-editor has decided. Online-editing is also the process where all the new material is added to the work (3D, visual effects, final audio track). Grading, while nominally the last phase of production, usually takes place in parallel with the online-editing. [6,136]

6.5 Visual Effects

VFX (visual effects) range from the easily noticeable to the invisible: from adding CG (computer generated) elements to live-action shots to removing wiring in effect shots. In the postproduction workflow, VFX are usually done in parallel with online-editing. In the offline-phase, the editor usually only works with rudimentary “placeholders” or even the raw material.

6.6 Grading

Grading is a relative newcomer to the field of digital postproduction. It is the process of colour correcting the whole program, each frame, thus affecting the look and feel of the whole. Digital grading has only been possible for a short time, since the processing power needed to effectively grade a whole feature film is considerable. [6,94]

6.7 Printing / Distribution

Film prints are still an important delivery medium. The final rendered frames are prepared for film printing using proprietary software created by the printer manufacturer. Then the films are printed back on to positive film. [6,141]

With digital distribution coming more important and pervasive, a DCI-master is often created concurrently with the version for film printing. Before digital mastering, the frames are rendered into suitable JPEG2000 frames using a suitable LUT. [6,121]

7 Generator Post – Tools

Generator Post has always ensured that they are at the forefront of technical advancements in the field, steadily upgrading their hardware, software and infrastructure to support the operators. This means that the operators have also had a big say in what tools they have been using. The high-end systems have been complimented with simpler desktop systems, and everything has been tied together with storage systems.

7.1 Autodesk tools

The main products from Autodesk Media&Entertainment are the IFFFS (Inferno Flame Flint Fire Smoke) range and Lustre. Inferno, Flame and Flint are high-end compositing and visual-effects workstations that have been around for almost 20 years, Fire and Smoke are online editing solutions with comparable history. Lustre is a grading solution that has been developed by Autodesk after they acquired the company that had released the product 5 years ago.

All current IFFFS programs run on HP 64bit x86 workstations with Red Hat Linux. These workstation feature proprietary drivers for the AJA-video cards installed. These units can work with any resolution material from a variety of sources. They only work with single frames, so though you can ingest several formats, the material will be changed to image frames, usually DPX frames. Realtime performance is achieved when working with 2K sources, in addition to HD and SD.

Autodesk also develop 3ds max, the worlds most used professional 3D animation program, which is used extensively at the Generator Post 3D- department.

7.1.1 IFFFS Systems

The first of the IFFFS -systems was Flame, introduced in 1993. While most of the technology running the system has changed since, the basic UI (Figure 4, below) still has many of the elements first introduced in the first half of the 90s.

Figure 4: The user interface for Autodesk Smoke 2010

IFFFS workstations usually come with a directly attached RAID-system. Previously this RAID-system was formatted with the proprietary Stone Filesystem, and is thus called Stone frame store. This disk system was not visible to the OS, so it could not be shared with other system via the OS. In more recent versions Autodesk added the ability to format the frame store as XFS, which makes it visible to the OS. [15]

The programs can also work from SAN shares or any other standard filesystems that offer enough speed (200 MB/s required for 2K work. When a normal share is enabled for work with an IFFFS workstation (enabled as a frame store), a new folder structure is needed on the filesystem. This folder system uses a proprietary way of storing information, so while the files can be viewed, the naming conventions and folder structures inside this folder do not make any sense to anyone viewing it.

7.1.2 Lustre Systems

Lustre comes in several three basic configurations: Linux, Windows and Incinerator. The Linux and Windows configurations run on a basic HP workstation, configured with an AJA video card and a powerful Nvidia GPU. The Lustre Incinerator system comprises of the Lustre frame server, Lustre Frame server and 8 incinerator nodes and a high-bandwidth Infiniband network that connects them all. The nodes enable the system to play back realtime 2K with several animated masks and grades activated. It can even playback 4K material with a basic grade at 20 frames / second.[16]

7.1.3 Workflow Tools

In the heart of the Autodesk workflow tools are wire and wiretap. Wire is the technology that allows Autodesk products to read each others stone-frame stores, and wiretap is a full-fledged API allowing the reading and writing from proprietary stone frame stores. They enable interoperability between Autodesk systems and even third-party products. Wiretap is also required for expanded format support, most notably support for red-files. [17]

7.2 Other Tools

7.2.1 Quantel

Generator Post has a Quantel iQ, the flagship product from Quantel. The iQ is a very powerful workstation for a large variety of uses, from motion graphics to online-editing and vfx-work. Quantel also offers a large variety of other products, from grading solutions (Pablo) to workflow products (Genetic Engineering). All of the products are very efficient and of high quality, but their interoperability with third-party products is very poor. In order for other systems to read data from the iQ frame store, the iQ program itself needs to be closed and a second program (SoftSAM) needs to be started.

The iQ interface is clear continuation on the interfaces seen in Henry and Editbox (Figure 5, below).

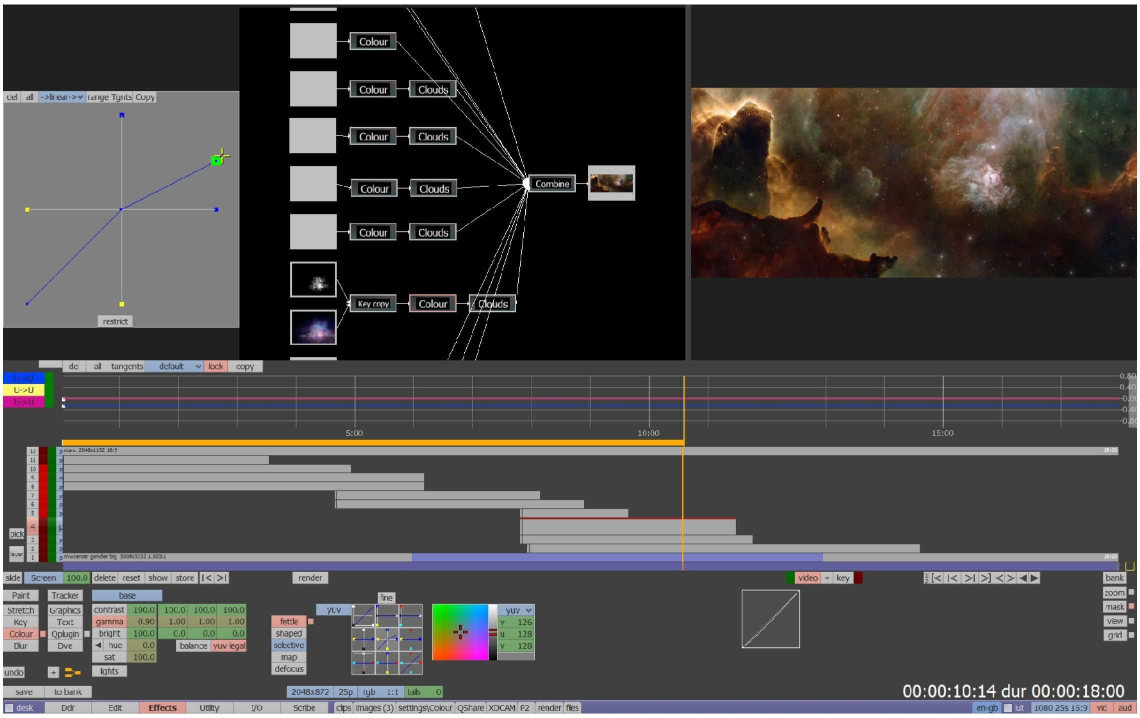

Figure 5: Quantel iQ User Interface

Quantel iQ does not work very well with a standard SAN, either. At Generator Post the decision was made to use the DLC (distributed LAN client)- technology offered by Stornext for the iQ, since Quantum claimed that the best performance we could expect from the iQ was 15 to 20 fps.

7.2.2 Apple

Apple Final Cut Studio is being widely used in a variety of roles throughout the workflow. Final Cut Pro (editing), Motion (motion graphics), Compressor (compression), Colour (grading) are an integrated production system on the desktop. Shake, though starting to show its age, is still a formidable composting tool and very useful for some very specific work (resizing with great image quality being an example).

Of the programs mentioned above, Final Cut Pro is probably the most used. It has quickly become the industry standard in editing on the desktop. The ability to interchange projects with clients easily (just moving the project files), a good set of editing tools and reasonable output options make FCP the swiss army knife of finishing projects for all uses. Colour is a promising program, but the speed when working with full 2K frames is still a question.

7.2.3 Adobe

Adobe tools are ubiquitous tools used in many phases of post production. Photoshop is the standard tool for any still-image work. Premiere is a good editing tool with a unrivalled support for different formats. Media Encoder is a useful tool for deliverables (in many cases a more user friendly program than Apple Compressor). After Effects is an industry leading desktop tool for motion graphics and compositing, and also supports a wide variety of formats.

7.2.4 DVS

DVS Clipster is an online-editing and finishing solution. The windows-based system works well with a variety of formats, and can be used to playback everything from Quicktime DV-video files to 2K DPX-sequences, in realtime to tape or for the linear DaVinci 2K realtime colour grading system. The unit comes with its own storage system, but can also use the SAN.

The DVS unit at Generator Post is an older model, and cannot be upgraded to support r3d in realtime, like the newer systems can.

7.2.5 Doremi Labs

Doremi labs is a company making disk-based video recording and playback devices. At Generator Post, the DMS-2000 system from Doremi is the tool used for creating DCI-compliant packages for distribution. This system creates VBR and CBR streams at 2K and 4K resolutions in realtime, from a number of sources (HD-SDI, files, etc.). The DMS-2000 supports all the encryption and security features demanded by the DCI-organisation.

7.2.6 Spruce DVD Maestro

This DVD-authoring solution manufactured in 1999 is still the cornerstone of DVD production. As an authoring solution the software has been surpassed by newer solutions (like the Apple DVD Studio Pro), but the encoding card in the unit (MPX3000) has no rivals in realtime encoding quality. The unit has an SDI input for capturing from a digital tape deck.

7.2.7 Celco Fury

The Celco Fury is a CRT-based film recorder. The unit is capable of creating high quality images at great speeds and high resolution. The system is controlled by an ancient SGI Octane IRIX-based computer that is probably the oldest piece of equipment in the digital production chain. The fury is fed files in it's own proprietary format through a normal NFS-share.

The printer is based on a cathode-ray tube (CRT) that creates the images. It can be used to print on negative or positive film. The printer is capable of printing images up 8K (8096x6120), and prints out normal 2K frames at 1 per second.

8 Generator Post Workflow

8.1 Generator Post Tools and Practices

With more and more of the work being done on digitized frames in a non-linear fashion, moving the data about became one of the biggest issues at Generator Post. As each unit would only work from it's own, dedicated hard drive, the frames needed to be moved throughout the workflow using a normal Gigabit Ethernet network. There were several shared storage devices, but these were not meant to be used directly for working from, but to be only used as intermediately storage. For centralized storage, a SAN was needed to cut down the amount of time spent waiting for transfers between stations.

Another big question was the growing amount of REDCODE material handled at the company. With the amount of files needed to be converted into ProRes422 for offline editing and DPX's for online work, some kind of clustering system was probably needed, so that the conversion work could be as automatized and fast as possible.

8.1.1 Archival

Another issue with the digital production pipeline is the lack of physical copies. In the traditional film workflow, several physical versions of the film are saved (the raw, filmed material, positive, negative), which can always be retrieved when needed (unless some physical harm comes to the reels). In a digital pipeline preservation of copies for perpetuity has become a major issue, one that is most conveniently tackled at the post production facility. Generator Post is a natural place for copying data on to the only sensible medium, magnetic tape data storage, preferably LTO-4.

8.2 Quantum Stornext SAN

The Stornext SAN is an advanced SAN based on a fibrechannel topology. The following figure illustrates the connections between the different parts of the system.

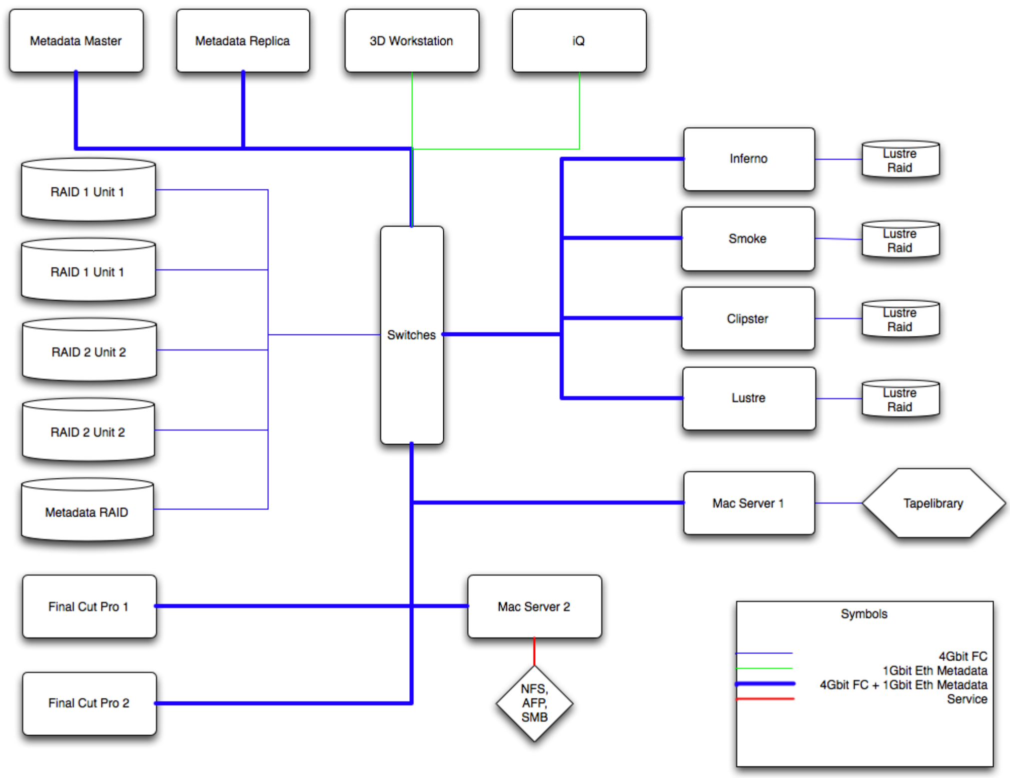

Figure 6: Generator Post SAN configuration

All the systems on the SAN are connected to the metadata switches with gigabit ethernet. Most of the systems are also connected via the fibrechannel switches. Generator Post chose Quantum Stornext as their SAN system for many reasons, of which the most important ones are mentioned below:

- Multi platform support (Mac, Linux, Windows)
- Robust technology (used in Apple Xsan, for instance)
- Very good industry reputation
- Scalable and flexible for future uses
- Support for Distributed LAN clients (access to SAN using ethernet wire speeds, without fibre channel cabling or NICs)

The Stornext system (formerly from Adic) is based on the proprietary cvfs filesystem. It can be set up on either Windows or Linux metadata servers (virtually the same product is sold for Macs as Xsan). Generator Post chose Red Hat Linux as their platform, with the server hardware coming from HP.

8.2.1 Setting Up the Stornext SAN

Before installing the SAN software itself, there were several steps needed to prepare the facility.

1. Cabling and network setup
2. Installing and setting up storage sub-units
3. Installing and setting up metadata servers

8.2.1.1 Network Setup

Setting up the network meant ensuring that each client and server system had appropriate cabling and network connections set up, while setting up the network meant installing and upgrading the metadata network switch and the fibrechannel switches.

Each client system that needed to be connected to the SAN needed both a second Ethernet connection (for the metadata information) and one or two pairs of fibrechannel cabling for the data connection. Of course each unit needed corresponding interfaces to make the connections.

At Generator Post, the following network hardware was chosen:

Make and Model	Type	Ports	Network	Amount
HP Procurve 2800	Gigabit Ethernet	48	Metadata	1
Qlogic Sanbox 5602	4Gbit Fiberchannel	16x4Gbit, 4x10 Gbit	Fiberchannel	2

Table 3: SAN Network hardware

The network setup was quite straightforward. Most units were connected to the metadata network with their secondary ethernet connection (eth1), the first interface being connected to the production network. The most important units were connected to the fibrechannel network with two pairs of fibrechannel connections, to provide both redundancy and a speed increase. The two fibrechannel switches were connected each other using the stacking possibilities and two 10Gbit interfaces.

8.2.1.2 Storage Subsystem Setup

There were five different RAID-subsystems used in the setup. They were configured in the following manner:

Make and Model	Name	Configuration	Size	Usage
Promise Vtrak	Leffa1_a	SAS/Raid6	3.7 Tb	Main online datastore
Promise Vtrak	Leffa1_b	SAS/Raid6	3.7 Tb	Main online datastore
Gatorraid	Metadata	SATA/Raid1	1 Tb	Metadata disk
Gatorraid	Tynnyri_a	SATA/Raid6	9 Tb	Near line datastore
Gatorraid	Tynnyri_b	SATA/Raid6	9Tb	Near line datastore

Table 4: Storage subsystems and configuration

The Stornext documentation recommends using a separate disk system for metadata, for performance reasons. The goal at Generator Post was to get realtime 2K performance from at least to workstations at all times (Autodesk Smoke and Lustre), so the aim was to get achieve at least 200MB/s performance from the disk array on both workstations simultaneously. Additionally, the data Generator Post is working with is 2k DPX frames, not big movie files directly.

In consultations with Stornext, it was decided that a SAS disk system is needed for the online media storage, while the metadata and near line media storage could be on SATA disks. SAS has a better seek time than SATA, which makes the difference when a lot of files are concurrently accessed. Also, it was decided that the a separate disk subsystem was needed to handle the metadata information, to enable maximum performance (as opposed to the metadata LUNs specified on the same subsystems as the media storage.

The Promise Vtrak units are RAID-units with 16 disks each and a whole lot of redundancy features: each unit has two controllers, two power supplies and a built in redundancy system, where incase of controller failure, the other controller would assume control over both parts of the disk subsystem.

LUN (Logical Unit Number) is the way disk subsystems show their raided disks to the fibrechannel hosts. The Promise units were split into two LUNs each, 1 disk for hot-spare and 7 in a RAID-6 configuration. Each Lun in this configuration has a size of 1.7 TB. The metadata LUNs were on a separate subsystem, just two 750Gb disks in a RAID1 configuration + hot spare.

8.2.1.3 Installing Stornext, Setting Up Filesystems

Installing the Stornext application was straightforward. After installing the OS with all tools needed to build programs and kernel extensions from source code, running a simple installation script using the CLI was all that was needed.

Setting up a new filesystem in Stornext is pretty straightforward:

1. Label LUNs
2. Choose filesystem location and name
3. Choose the number of stripe groups

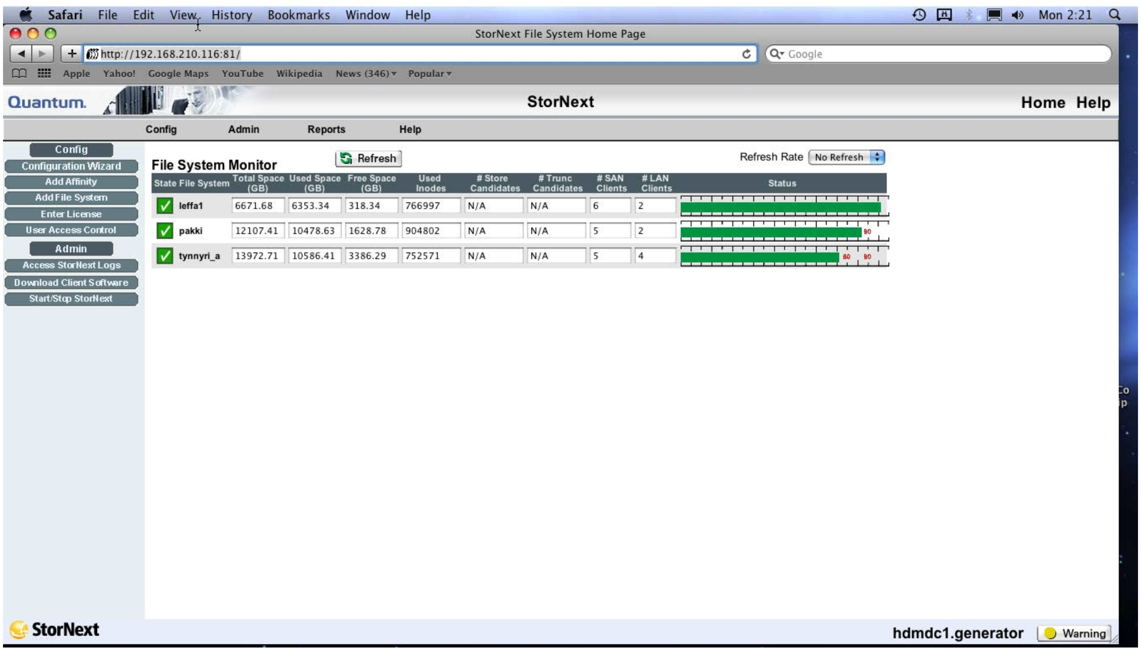

Figure 7: Stornext web-GUI after installation

The web-GUI shows all the active filesystems, how much data is stored on each filesystem, how many clients are connected to the filesystem and some properties of the filesystem.

8.2.1.4 Optimizing Performance

For Generator Post, the most important thing about performance was getting realtime 2K streams to both Lustre and Smoke from the online filesystem. This meant achieving at least 300MB/s on each of the workstations, simultaneously.

Optimizing a Stornext filesystem means changing the parameters for Stripe breadth and filesystem block. These two values can only be set when the filesystem is built, so the parameters need to be correct once the system goes into production. To find the correct values for the system, we created a two-tier testing system: first we would try out many configurations by using the standard unix dd command, then the best configurations would be tested with some more specific tools, like the AJA IO-tools and Autodesk built-in disk speed tests. The results of the AJA tests are in the table below.

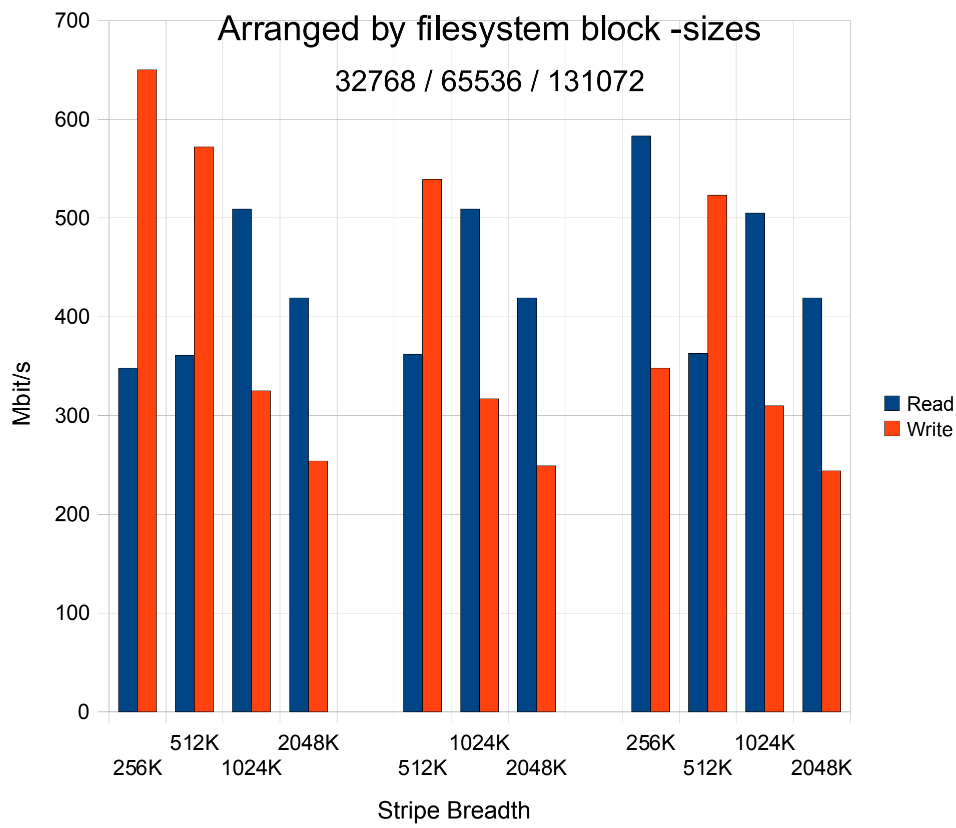

Table 3: Disk Read/Write tests using the AJA IO-Tools

In the setup at Generator Post, the most important thing (after threshold level of a 250Mb/s write for a single tape capture) was getting multiple streams of data to be read simultaneously, so the read performance was what we were focusing on. The choice was made to go with a filesystem block- size of 131072 and a stripe breadth of 256K.

8.3 Apple Xgrid for Distributed Rendering

Apple Xgrid is a system for distributed rendering. It has been created mainly for scientific calculations like genetic research, and comes as standard with each installation of Mac OS X. The system consists of a xgrid controller, a Mac OS X Server machine to which all the rendering nodes connect to. The controller is the machine that parses the commands and sends the separate jobs to the render nodes, called agents.

Setting up Xgrid for use with Redline, the command line tool from Red for transforming r3d files, required the following steps:

- setting up xgrid using single-sign-on (Kerberos).
- ensure that each machine only accepts a single job
- setting up shared folders so that all the SAN-shares are mounted identically on all systems (an issue only on machines that are not directly connected to the SAN)
- install redline on all the xgrid agents
- submit the RedRushes script as a xgrid batch

The first part is needed to ensure that the controller can connect and run commands on all the agents without separate authentication or setting passwords. This required a working Open Directory master and all the machines needed to be tied to that domain. After that it was pretty straight forward to get all the machines show up on Xgrid.

Ensuring that each machine only accepted a single job at a time was quite important, as each quad-core Mac Pro in the Xgrid accepted 8 jobs by default. The jobs in this case are heavy rendering jobs of varying lengths, and it is important that each computer only works on a single job at a time. The redline software used to do the actual conversion is well written and threaded, and a single process can use the power offered by several processors very well. Changing the amount of jobs accepted meant modifying one file on each client: /System/Library/LaunchDaemons/com.apple.xgridagentd.plist (example file on page 49).

The next image shows the Xgrid Admin GUI, with all the available processor cores showing.

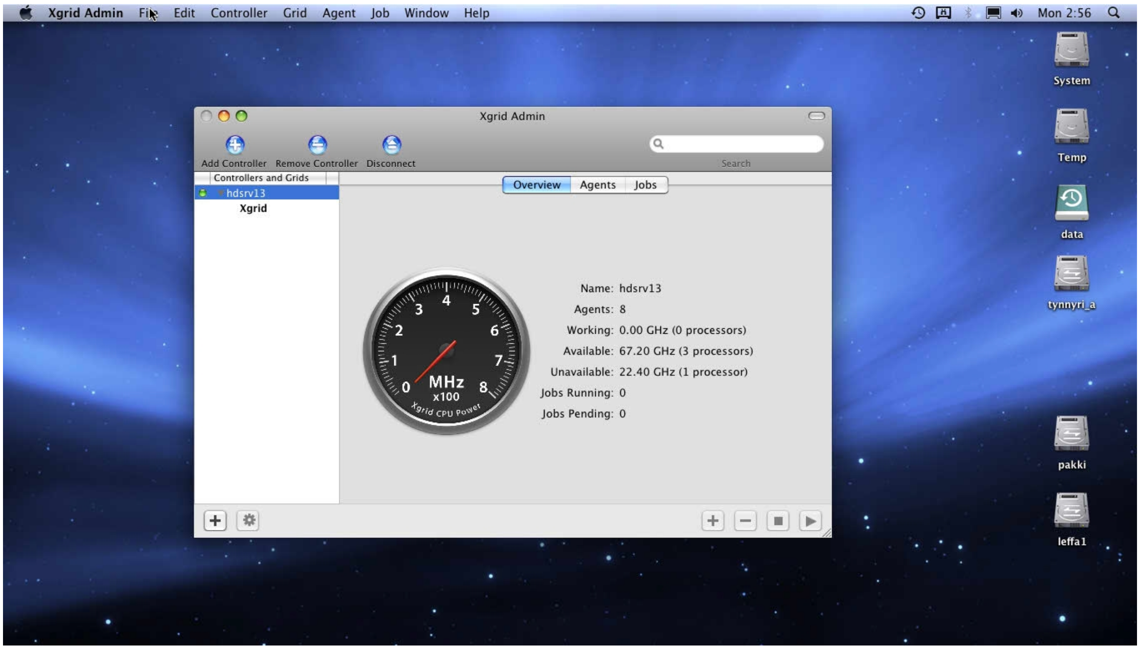

Figure 8: Xgrid Admin showing the amount of processing power available

The second part did create some problems. Using OD to automount to mount the SAN shares on all the machines tied to the SAN using NFS. This presented problems if a machine actually had the SAN installed, because then the computer was unable to mount either share (SAN or NFS), as they were trying to mount to the same folder. The solution was to manually make the machines mount the correct shares using Directory Utility.

The third step presented no issues at all.

The fourth step did present some issues. The redrushes scripts can easily have hundreds of lines, and manual work was out of question. A ruby script was written to take the name of the redrushes script, and create a xgrid batch (xml) that can be submitted to the controller.

The full workflow for the user:

1. upload the r3d files to the shared folder
2. open Redrushes, choose your settings, choose the correct folder, create the output folder, and export the script
3. run the script through the ruby script (redgrid.rb), and submit the resulting xml file to the xgrid
4. see the files render at incredible speeds

8.4 Autodesk Wiretap Central

Wiretap Central is a web-based program programmed using Adobe AIR. The UI is very similar in look and feel to the IFFFS line of Autodesk products.

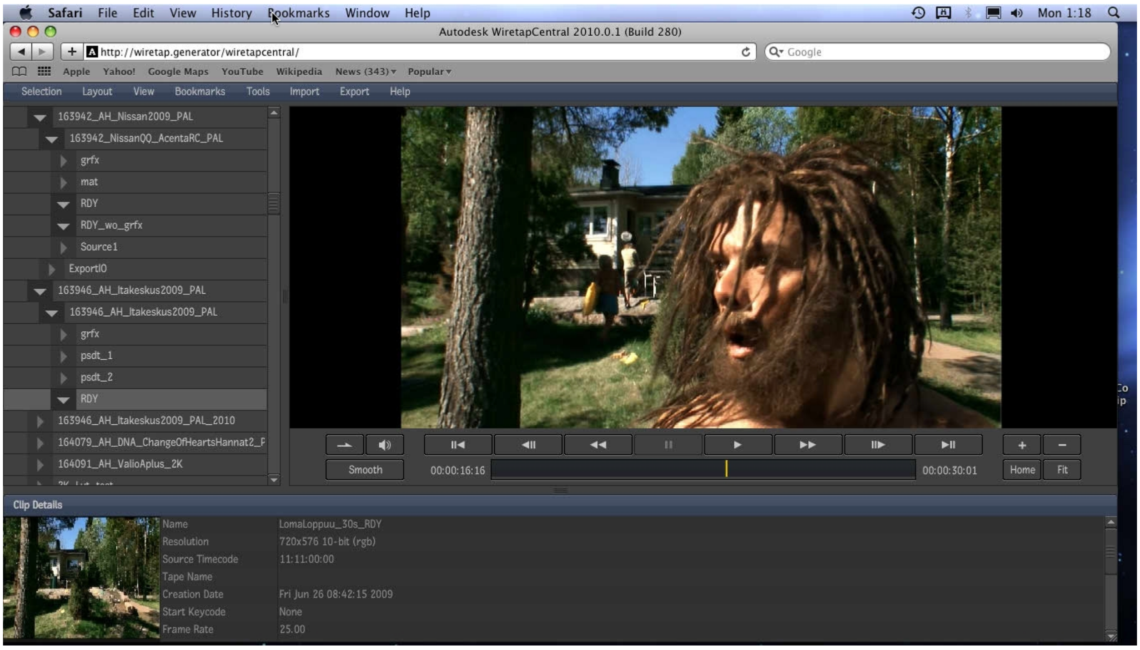

Figure 9: Wiretap Central UI

Wiretap central is web-based program that provides several services to users. It can be used to export clips directly from Autodesk systems to several distribution formats or import red-files to Autodesk-compatible systems. It consists of several components:

- Wiretap Central, the Flex-based web application
- Wiretap Gateway, the component that provides access to SAN and shared filesystems to Wiretap Central.
- Backburner, a distributed rendering system
- MI/O, the backburner component that provides support for several professional formats to backburner

The Wiretap Central UI for RED-import.

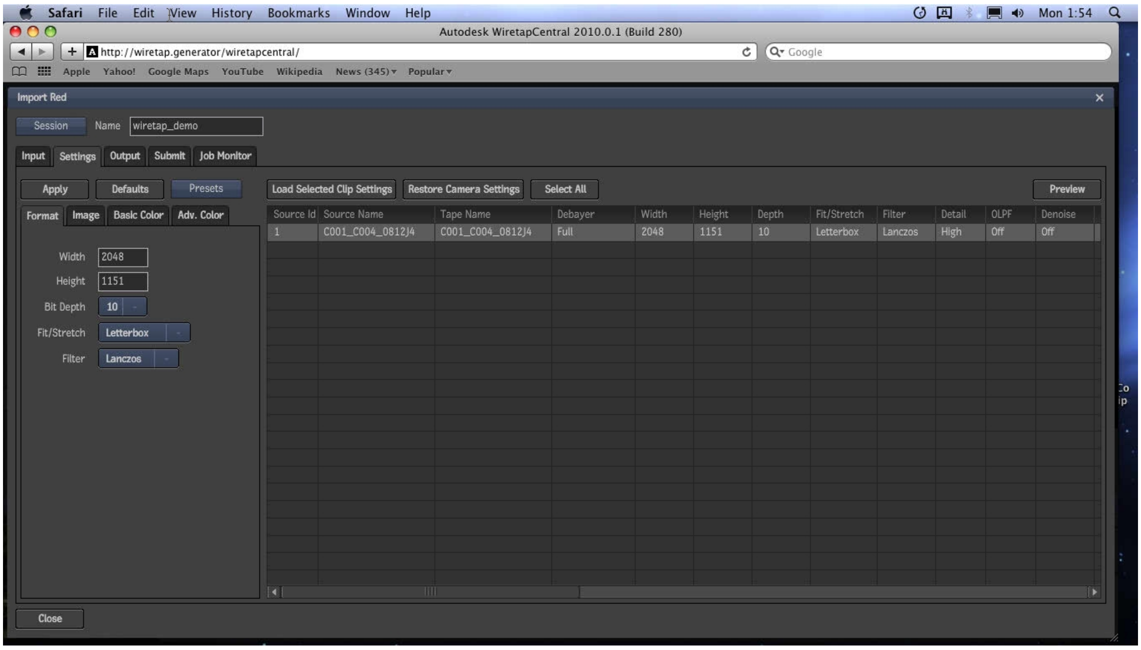

Figure 10: Wiretap Central r3d-import dialog

The biggest advantage wiretap central offers is distributed rendering of r3d files into DPX-sequences [18]. Wiretap Central allows inexpensive linux computers running the free software from Autodesk to be used in the process, which speeds up the process considerably.

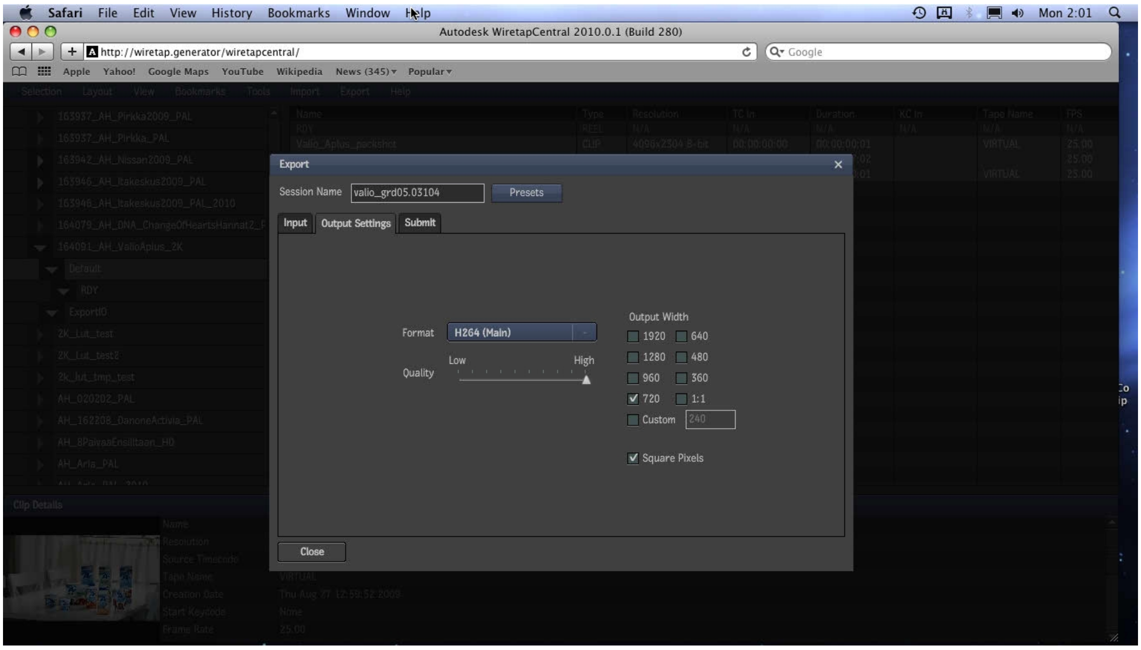

Figure 11: Wiretap Central export dialog

Wiretap Central also offers a wide range of colour correction options for the material, so that the image can be tweaked somewhat at this point. Usually a very "flat" image is optimal for colour correction work. Wiretap Central also provides web-based access to the projects on the Stones of Smoke and Inferno, allowing clips and edits to be reviewed from another computer while the workstation is being used for work. Clips and projects can also be exported into several distribution formats without using processing power from the workstation itself, thus freeing the workstation for actual work while exporting files.

8.5 Full R3D-workflow

1. Create offline-editing materials
   - 1.1. Import Red-material to SAN
   - 1.2. Create Redrushes-script, use redgrid.rb to submit job to xgrid
   - 1.3. copy finished ProRes files to external FW-disk, send to client
2. Create DPX-files for postproduction
   - 2.1. Copy EDL or Final Cut XML to SAN

2.2. modify XML / EDL to work with Wiretap Central

2.3. Use Wiretap-render farm to render out needed DPX frames from original R3D-files. Make basic colour choices here.

3. Soft-import DPX-frames to Autodesk Smoke, prepare timeline, use wiretap to read project from Autodesk Lustre
4. Render out from Lustre to SAN
5. Read finished frames to printing and digital distribution

The greatest benefits of this system are speed and efficiency. The parts of the workflow that require the most CPU-cycles are handled by render farms that offer easily expandable processing power. The media is only stored on shared storage so there is hardly any need to move the media from one system to another, eliminating transfer times.

This workflow also enables consistent archiving of all material to tape in a consistent manner. Before a project is started, a producer decides with the client which material and intermediate steps are archived on to data tape.

9 Conclusions

The changes Generator Post wanted to achieve with the new systems were clearly achieved. The new storage systems allowed work to be done on the same material on different workstations without moving the material across the network at all. This meant a faster turn-around time for projects and less idle waiting time. Apple Xgrid allowed the efficient use of idle hardware across the facility to improve the speed of rendering RED-material for offline use.

The existing systems were tightly integrated to the new SAN without the need to make big changes. The existing systems also got full of the new storage, with several workstations being able to work with 2K material off the SAN in real time. Finding out the correct settings for all the systems was time consuming, mostly due to a lack of clearly defined rules from the multiple vendors included in the project.

The biggest caveat is probably the amount of space offered by the system for real time work. With a total size of 7 TB the work disk is not very big, and the biggest reason for this is the smaller size of the SAS disks as compared to normal SATA-disks. The speed advances brought by SAS-disks could have been replaced by using more of the slower SATA-disks.

With the advances in storage technology, I believe that the Stornext SAN will be in use about the year 2012 from the writing of this thesis. At that time there will be need to support bigger data streams (multiple streams of 2K, 4K work etc.). This means that the Fibrechannel backbone in use will be outdated and either the SAN will be upgraded to 8Gb Fibrechannel or the whole SAN will be bypassed and a 10Gbit ethernet network built instead. That would offer similar performance but without a need for a separate SAN-filesystem and associated costs.

Creating the Apple Xgrid rendering system was born from the need to render out a large amount of material every day for several months during the filming of a motion picture. Using Xgrid of rendering media was somewhat unorthodox, with Apple clearly preferring clients to use Qmaster for this kind of work. However, the client and server for Xgrid were much more stable and less prone to suddenly stop working, so the decision was made to use it instead. This meant a lot of trial-and-error testing. Another big part of the job was making the interface to the system easy enough for average operators to use.

The Xgrid system worked very well for its intended use. Using several computers clearly helped to deliver the material to the clients more quickly. There were some stability issues, and some mysteriously missing clips every once in a while, but these were small issues that were easily worked around.

References

1. von Bagh P. Elokuvan Historia. Otava; 1998.

2. 35 mm film - Wikipedia, the free encyclopedia [online]. URL: http://en.wikipedia.org/wiki/35mm_film. Accessed 2010 May 8.

3. Technicolor - Wikipedia, the free encyclopedia [online]. URL: http://en.wikipedia.org/wiki/Technicolour#Three-strip_Technicolor. Accessed 2010 Apr 25.

4. Widescreen - Wikipedia, the free encyclopedia [online]. URL: http://en.wikipedia.org/wiki/Widescreen. Accessed 2010 May 8.

5. Anamorphic format - Wikipedia, the free encyclopedia [online]. URL: http://en.wikipedia.org/wiki/Anamorphic. Accessed 2010 May 8.

6. Pank B. The Digital Fact Book. 20th ed. Quantel; 2008.

7. Luff J. Digital cinema cameras. Broadcast Engineering. 49(10):130-133.

8. Behar M. Analog Meets Its Match in Red Digital Cinema's Ultrahigh-Res Camera. Wired Magazine. 2008 Aug 18;16(09).

9. Digital Cinema Initiatives (DCI) - DIGITAL CINEMA SYSTEM SPECIFICATION, VERSION 1.2 [online]. URL: http://www.dcimovies.com/press/. Accessed 2010 May 8.

10. Bilgin A, Marcellin M. JPEG2000 for digital cinema. In: Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on. 2006. p. 4 pp.-3881.

11. Boxin Shi, Lin Liu, Chao Xu. Comparison between JPEG2000 and H.264 for digital cinema. In: Multimedia and Expo, 2008 IEEE International Conference on. 2008. p. 725-728.

12. Bloom J. Digital Cinema Content Security and the DCI. In: Information Sciences and Systems, 2006 40th Annual Conference on. 2006. p. 1176-1181.

13. Brinkmann R. The Art and Science of Digital Compositing. 1st ed. Morgan Kaufmann; 1999.

14. Preston WC. Using SANs and NAS. 1st ed. O'Reilly Media, Inc.; 2002.

15. Effects&Finishing Installation and Configuration Guide 2011. 2010.

16. Lustre Installation and Configuration Guide 2011. 2010.

17. Autodesk® WiretapCentral™ 2011 and Wiretap Gateway 2011 Installation Guide. 2010.

Appendices

Appendix a: redgrid.rb

```
#!/usr/bin/env ruby -wKU
# 2009-07-05
# (c) Copyright 2009 Jussi Uosukainen (jussi.uosukainen@iki.fi). All Rights Reserved.

# A script designed to write redrushes scripts into a xgrid patch
# Asks for the redrushes script, then names all elements according to that file

# Require one argument
unless ARGV.length == 1
puts "Usage: ruby redgrid.rb redrushes_script.txt"
exit
end

# Pull filename from argument
filename = ARGV[0]

# Name project after according to file
project = filename
project = project.gsub(".txt", "")

# Change output-filename to xml from text
filename = filename.gsub("txt", "xml")

# Open the file for output
output = File.new(filename, "w+")

# Create file header
header = "{\njobSpecification = {\nname = \"#{project}\";\nnotificationEmail
= \"#{email}\";\ntaskSpecifications = {"

# write the header to the file
output.write header

# Open redrushes script
redlines = File.open(ARGV[0]).readlines

# Read number of lines
noc1 = redlines.size
```

```
# Create iterator
i = 0

# While loop for creating the xgrid commands based on each line
while i < noc1 do

# Split each line using space as separator
com0 = redlines[i].split(" ")

# Create command
command = "\nchr#{i+1} = {\narguments = (\n\"#{com0[1]}\", \n\"#{com0[2]}\",
\n\"#{com0[3]}\", \n\"#{com0[4]}\", \n\"#{com0[5]}\",\n\"#{com0[6]}\",
\n\"#{com0[7]}\", \n\"#{com0[8]}\", \n\"#{com0[9]}\", \n\"#{com0[10]}\",
\n\"#{com0[11]}\", \n\"#{com0[12]}\", \n\"#{com0[13]}\", \n\"#{com0[14]}\",
\n\"#{com0[15]}\", \n\"#{com0[16]}\", \n\"#{com0[17]}\", \n\"#{com0[18]}\",
\n\"#{com0[19]}\", \n\"#{com0[20]}\", \n\"#{com0[21]}\", \n\"#{com0[22]}\", \n);
\ncommand= \"/usr/sbin/REDline\"; \n};"

# Append command to file
output.write command
i += 1
end

# Create footer
footer = "\n};\n};\n}"

# Write footer to file
output.write footer

# Close file
output.close
```

Appendix b: com.apple.xgridagentd.plist


```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN"
"http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
	<key>BindToFirstAvailable</key>
	<false/>
	<key>MaximumTaskCount</key>
	<string>1</string>
	<key>ControllerAuthentication</key>
	<string>Kerberos</string>
	<key>ControllerName</key>
	<string>xgrid.hostname</string>
	<key>OnlyWhenIdle</key>
	<false/>
	<key>ResolveNameAsNetService</key>
	<false/>
	<key>SuspendWhenNotIdle</key>
	<false/>
</dict>
</plist>
```