Scaleable Intelligent Video Server System

Title	**Work Package 12 Publishable Summary**
Revision	**A**
Deliverable #	**Summary D12.1 and D12.2**
Author	Wolfgang Huther, Rolf Hedtke
Company	Grass Valley Germany GmbH
Date	17 April 2007
Filename	
Dissemination[2]	**PU**

REVISION	**DATE**	**DESCRIPTION**
A	17 apr 2007	Created by WH & RH

[2] **CO** = Confidential (only for members of the consortium + EC); **RE** = Restricted to a stated circulation list (+ EC) *[replace this footnote with the list]*; **PP** = Restricted to other FP6 participants (+ EC); **PU** = Public

TABLE OF CONTENTS

1 SUMMARY OF WP12 CONTRIBUTION TO SIVSS

The main function of this work package was to deliver the next generation acquisition technology for integration with a large scale video archive or multimedia platform. The complete acquisition technology system is also compatible with legacy video formats to provide a mechanism for digitisation of old and culturally significant film footage. The system is able to store the input content at the highest quality possible for any particular input medium. This requires the high data throughput of the SIVSS server architecture.

The necessary interface software to integrate a full Media Asset Management Service has been provided. This involves the creation of metadata tags associated with the video or audio content. The digitisation of video or audio content creates extremely large files. The metadata tags provide the way for efficiently search for particular types of footage from the key words or attributes within metadata file. The metafile file and actual video footage are linked within the storage management system of the software application.

2 INTRODUCTION WITH OBJECTIVES

2.1 Complete Input Acquisition System.

GRASS VALLEY has delivered the complete input acquisition system with

- a video acquisition mechanism capable of capturing more than 27 million pixels per second
- film scanning resolution greater than 3072 x 2048 pixels at very high frame rate
- compatible with the film transfer needs of SDTV, HDTV, 2K, 4K data from 70mm, 35mm,16mm and Super 8mm film formats legacy support for application server target platform recommendations

2.2 Media Asset Management interface software

Goal of the deliverable D12.2 was to develop a Media Asset Management interface software package which provides the necessary file system and management services (e.g. format conversion, proxy generation, scene detection, metadata management) which are required to integrate the SIVSS Input Acquisition System (IAS, as defined in D12.1) into (Enterprise) Media Asset Management systems. The successful integration into the Grass Valley “Contentshare.net” media management system has been shown with the demonstrator hardware.

3 DELIVERABLE D12.1 INPUT ACQUISITION SYSTEM

3.1 SIVSS Ingest Interface Investigation

With the advent of HDTV in Europe, high resolution, high speed digitisation/scanning is a basic requirement for future HDTV programs. For the SIVSS postproduction ingest part the most important Video/Audio sources are

- 2K real-time filmscanner (2048 x 1756)

- 4K real-fast filmscanner (4096 x 3512)
- Digital HDTV Video camera (1920 x 1080, up to 150 frames/sec)
- Digital film camera (2048 x 1080)

Therefore the requirement for SIVSS ingest data rate was defined as follows:

Each frame has a resolution of 2K - coded in the three components R, G and B where each component is coded with at least 10bits. The film frames are transmitted with a rate of 24 frames per second. The required bandwidth is calculated according to

BW = pixel_per_line * lines_per_image * (approx. 3*10bit) * frame_rate

BW = 2048 * 1756 * 4Byte * 24 frames/sec

BW = 346 MByte/sec

(Because of computer-systems are byte oriented we have chosen the nearest Byte number to 3x10bit which is 4Bytes).

For the SIVSS ingest system only interfaces with a data rate > 400MByte/Sec were investigated. The interfaces should be commercially available, therefore more or less proprietary interfaces like Quadrics QSNet II were not considered here.

Table 1 shows the nominal data rate of interface candidates.

Technology	Nominal Bandwidth
Fibre Channel 4G FC	400 MByte/s
GSN	800 MByte/s
10 Gigabit Ethernet	1 000 MByte/s
InfiniBand 4x	1 000 MByte/s
InfiniBand 8x	2 000 MByte/s
InfiniBand 12x	3 000 MByte/s

Table 1: Interface data rates

3.1.1 GSN (Gigabit System Network formerly known as Hippi 6400)

GSN has been used in research labs and government sites for several years. It never has been accepted by a broader community and therefore production quantities are very low with very high costs (e.g. 15k$ per NIC). Real implementations have proofed that reliable transfers are possible up to 600MByte/sec. Physical transport medium is copper or fibre. This technology is currently used as an interface for the filmscanner and allows reliable data transfer up to 400MB/sec. But recently the hardware support of this technology has been stopped by the manufactures, so for new development this kind of interface is not appropriate anymore.

3.1.2 Fibre Channel

Fibre Channel has been used in disc arrays for several years, therefore this interface was evaluated during integration of SIVSS storage devices. As SCSI- devices, the Fibre Channel-drives and

RAID-arrays allow to monitor and tune a set of functionalities, like ReadCaching, WriteCaching, CommandTagQueuing etc, this is essential for tuning the throughput of a system for high-performance real-time-playback and recording. With other technologies likes ATA or SATA this functionality is not available. Compared to other technologies the protocol overhead of a Fibre Channel transmission reduces the data rates significantly. However the flexibility with regards to number of devices of up to 127 per Fibre Channel and the possibility to extent the optical cable length to hundreds of meters rates the Fibre Channel network superior to any other technology.

For larger installations with many high performance clients working on the same set of media without the need of copying the Data from one machine to the other Fibre Channel is the only technology that supports the required Storage Area Network (SAN) technology. From the viewpoint of a local or SAN file system there is no difference between a single Fibre Channel disc in a "Just a Bunch of Disc" (JBOD) array and a RAID controller in a RAID arrays. Thus the internal array technology is of no relevance for the file system. The physical transport medium is copper or fibre.

3.1.3 10 Gigabit Ethernet

This network is the natural upgrade path of existing computer networks and is widespread available. This is a logical candidate for every high speed network. The nominal high data rate can only be achieved with unreliable connections. There is a dramatic decrease in data rate if a reliable connection is required. Reliable connections normally require a significant protocol overhead (TCP/IP stack). New interface boards became available in 2005 which implement the protocol stack in hardware as an integral part of the host channel adapter. This technology is known as **T**CP **O**ffload **E**ngine (TOE). But TOEs require a modification within the controlling operating system. This technology is not widely supported by currently available operating systems. Physical transport medium is fibre but cheaper implementations using copper will become available in the future.

3.1.4 InfiniBand

InfiniBand is a technology which became popular in building up computer clusters and will probably replace proprietary technologies like Quadrics QSNet and Myrinet. InfiniBand is also scalable to some amount. Reliable connections can be handled without large overhead. Physical transport medium is copper but optical enabled host channel adapters and optical media converters are available as well.

Several low level and high level protocols are already established. **R**emote **D**irect **M**emory **A**ccess **RDMA** is an integral part of the InfiniBand technology which gives some advantages over 10Gig Ethernet. InfiniBand was tested with a low level protocol (VERBS layer) and transfer rates of 500 MBytes/sec (server to server) could be achieved.

It was decided to select InfiniBand 4x as future platform for any scanner ingest activities. It combines low cost, reliable connection together with a minimized CPU overhead.

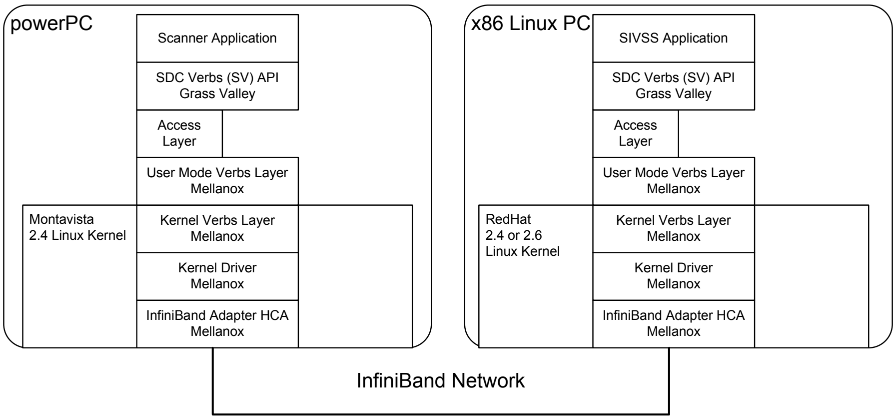

Figure 1: Layers for IB based Scanner to Ingest Application data transfer

3.2 Selection of Workstation/Server Bus System

Host channel adapters are offered with different PCI Bus technologies. Table 2 shows the nominal data rates of the different PCI implementations.

Technology	Nominal Bandwidth
PCI –X 66	512Mbyte/Sec
PCI –X 100	768Mbyte/Sec
PCI –X 133	1 GByte/s
PCIe 4 x	1 GByte/s
PCIe 8 x	2 GByte/s
PCIe 16x	4 GByte/s

Table 2: Bus Data rates

The tests during year 2004 were made with PCI-X based systems. In 2005 PCIexpress (PCIe) emerged as a new technology which was tested as well. The tests showed that the required data rates could be achieved with PCI-X and PCIe 4x systems. The nominal data rate of the PCIe bus is decayed by currently available chipsets for PC motherboards. The PCIe standard allows transactions with a burst-length up to 4KByte. The tests of the SIVSS ingest system were made with a HP8200 workstation which allows only transactions of 64 Byte. With the 4x implementation two ingest data streams can be transferred in parallel. It is assumed that the transaction size of future motherboards will go up to 256 Byte in the next 2 years. Table 3 shows the data rates which can be achieved with different PCIexpress implementations.

Transaction Payload Size	PCIe 4 x	PCIe 8 x	PCIe 16 x
64 Bytes	760 MByte/s	1446 MByte/s	2895 MByte/s
128 Bytes	862 MByte/s	1643 MByte/s	3286 MByte/s
256 Bytes	926 MByte/s	1762 MByte/s	3525 MByte/s

Table 3: Achievable Data Rates on PCIe

3.3 Pre Evaluation of Acquisition System performance

In 2004 we proofed that a system based on the hp xw8200 workstation was able to deliver the required internal bandwidth for real-time 2k streaming. The goal of the new evaluation tests was to prove that the performance of the components of a real system, including clustered file system and disk array, meets the requirements. For this test the new hp xw9300 workstation was used, which will be the successor of the hp xw8200. The xw9300 was the first 64 bit CPU we used in the SIVSS project.

The rightmost column in table 4 shows the system configuration we used for the test. After some minor problems with Red Hat Enterprise 3 (due to insufficient support of the new AMD-ET64-Architecture) have been fixed and the latest "Engineering-build" of ADIC's Cvfs file system has been installed we were able run first performance tests, but never were able to achieve more than 14 fps in ingest and 15 fps (< 200Mbyte/sec) in play mode. Using Linux kernel 2.6 improved performance by approx. 10% but was still not sufficient.

	Xw8000	**Xw8200**	**Xw9300 (64Bit)**
CPU	2 x XEON 3.0 GHz	2 x XEON 3.2 GHz	2 x Opteron 280 DP 2.6 GHz
RAM	2 GByte	2 GByte	4 GByte
GRAPHICS	nVidia FX3000	nVidia FX 3400	nVidia FX 3400
I/O	2 x PCI-X 100 MHz 1 x PCI-X 133 MHz 2 x PCI	1 x PCI-Express 4 x 2 x PCI-X 100 MHz 1 x PCI-X 133 MHz 2 x PCI	2 x PCI-Express 16x 1 x PCI-X 100 MHz 2x PCI-X 133 MHz 1 x PCI
Networking	1 x Gig-Ethernet	1 x Gig-Ethernet	1 x Gig-Ethernet
OS		RedHat EL 3.0 WSS Update 6 x86_64, kernel 2.4.21-37.ELsmp	
File system		ADIC Cvfs 2.6.3 Build 30 (Beta version)	
Storage		Xyratex JBOD RS-1600, 6 Disks Seagate 36GByte, two loops with 8 disks each.	
Controller:		QLogic qla2342: double 2gbit, connected with two optical cables.	

Table 4: Workstation platforms used in 2005

3.4 Integration Work and IBC 2005 Demonstration

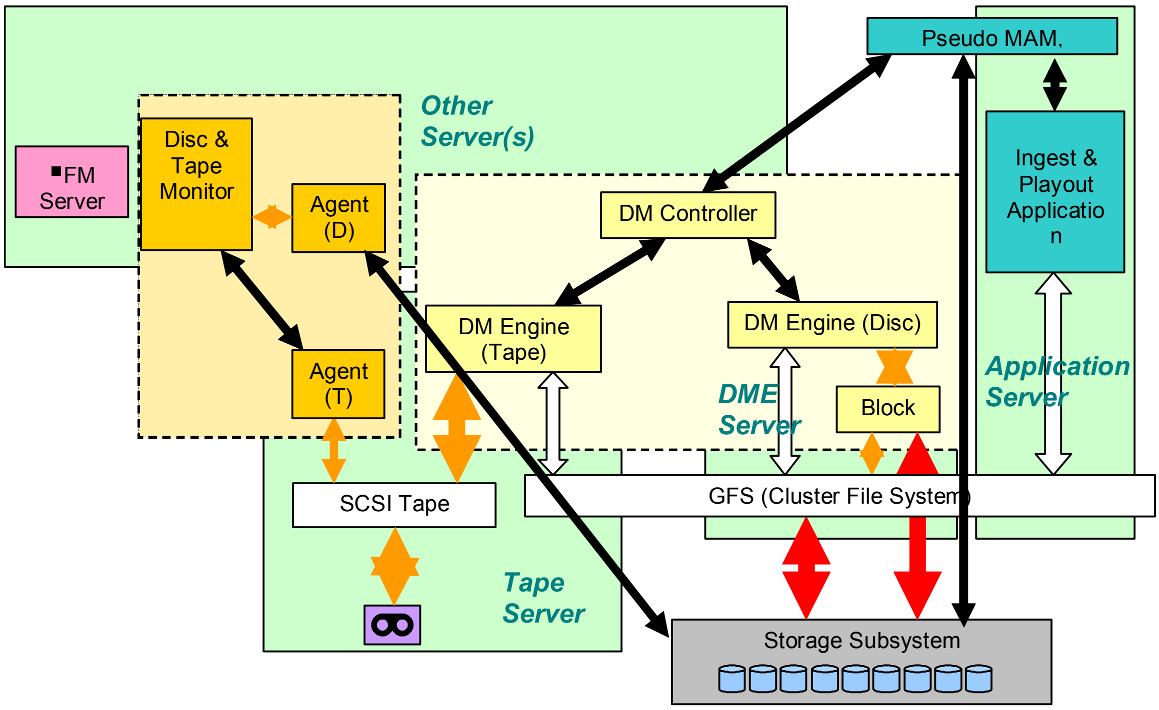

Figure 2: IBC Demonstrator Software & Interfaces

At the IBC exhibition 2005 in Amsterdam we demonstrated a prototype system (see Figure2) and achieved a sustained simultaneous ingest and playback of 1920x1080 life video from 2 independent hosts to Xyratex storage through GFS filesytem . This equals a sustained read/write performance of approx. 220 MByte, i.e. 25 file i/o transfers per second to GFS. The basic MAM functionality allowed browsing, managing and viewing of recorded material together with archiving and retrieval to/from tape by using the HISTOR data mover.

3.5 Description of Input Acquisition System

The high-level Video I/O software modules are organized as follows:

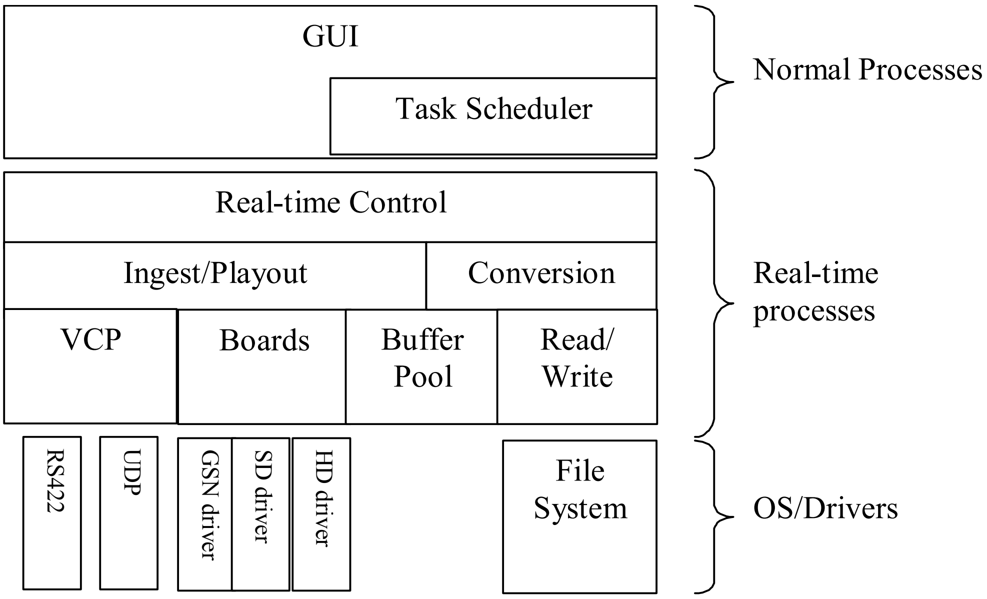

Figure 3: IAS Software Structure

For the acquisition of realtime HDTV signals a hardware interface card with dedicated HDTV and SDTV I/O functionality was used. The acquisition and playout software tools developed for the SIVSS projects were implemented on the HP workstation using LINUX operating system. The connection to the SIVSS hardware was made by a dual AS (4x) endpoint card.

3.6 Final Evaluation of Acquisition System Performance

The primary choice for SIVSS was to use GFS as the distributed filesystem for the SIVSS platform, and this was used in the pre-analysis testing and characterisation. In the final platform, however it was not possible to use GFS and so a commercial filesystem (StorNext or CVFS) was implemented.

The reason GFS could not be used was that the hardware platform required the use of RedHat Enterprise Server Version 4 at update level 4 in order to support the latest release of PCI-E incorporated in the test servers. This was only discovered after an extended period of debug and analysis on the platform, where previous versions of the hardware would not interact correctly and necessitated the deployment of new servers with RHES4.4. Unfortunately GFS was not available for this platform, only for RHES4.3 and so the decision was made to deploy CVFS as an alternative to allow the planned testing to be completed.

With this configuration the performance of more than 320MB/s could be achieved. This means with this the operationis capable of 1 Ingest or Playback Stream from 1 Host in HD with a concurrent/simultaneous Archive/Restore Operation. The bandwidth of the integrated system is adequate for the purpose of video streaming, providing ~ 300 MB/s for each streaming channel implemented. The fact that there were two channels in operation at the same time during this tests showed that the system is scalable and that the behaviour of the two sections are independent of each other.

The throughput capability of the storage subsystem has been demonstrated to be more than adequate for the SIVSS requirements.

At some point in the future, GFS development will catch up with the latest Linux Distribution Release and this problem will no longer exist.

4 DELIVERABLE D12.2 MEDIA ASSET MANAGEMENT INTERFACE

4.1 Use case definition

The postproduction workflow has been analysed and the use cases defined as shown in figure 4. Two input HDTV streams were captured in parallel and stored on the SIVSS storage subsystem. The data of this storage system was transferred automatically to the tape archive by the HISTOR data mover software.

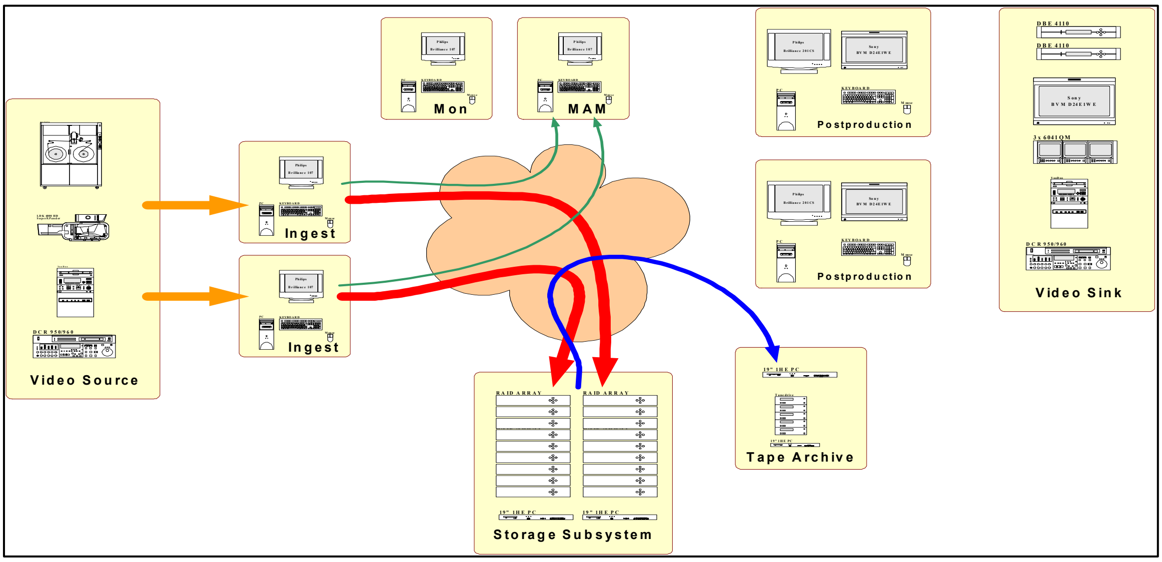

Figure 4 Postproduction Ingest Use Case

4.1.1 Ingest

Video content from several high resolution sources like filmscanners or digital cameras together with sound from multiple sources will be received (digitized if needed) by the acquisition application, metadata from the source devices added or generated (automatically and/or annotated by the operator). Content will be stored in the central storage subsystem, while the metadata will reside on a separate database. The Media Asset Management system (MAM) will be informed whenever new material has been ingested.

4.1.2 Playout

Initiated by the MAM or a local operator the IAS will read images from the storage subsystem, convert it into the required format and stream the images thru the dedicated outputs.

4.1.3 Archive

For archive/retrieval the MAM sends appropriate commands to the HISTOR data mover (D8.1) which will move data from storage subsystem to the tape archive or vice versa and informs the MAM on completion.

4.2 Requirements

The analysis of the use cases above showed the need for services like

- Audio functionalities (ingest, playout, flowgraph, mixing etc).
- Metadata Handling (create, modify, delete, export of metadata objects)
- Interoperability with a “global” media management system (MAM)
- Interoperability with several IAS stations

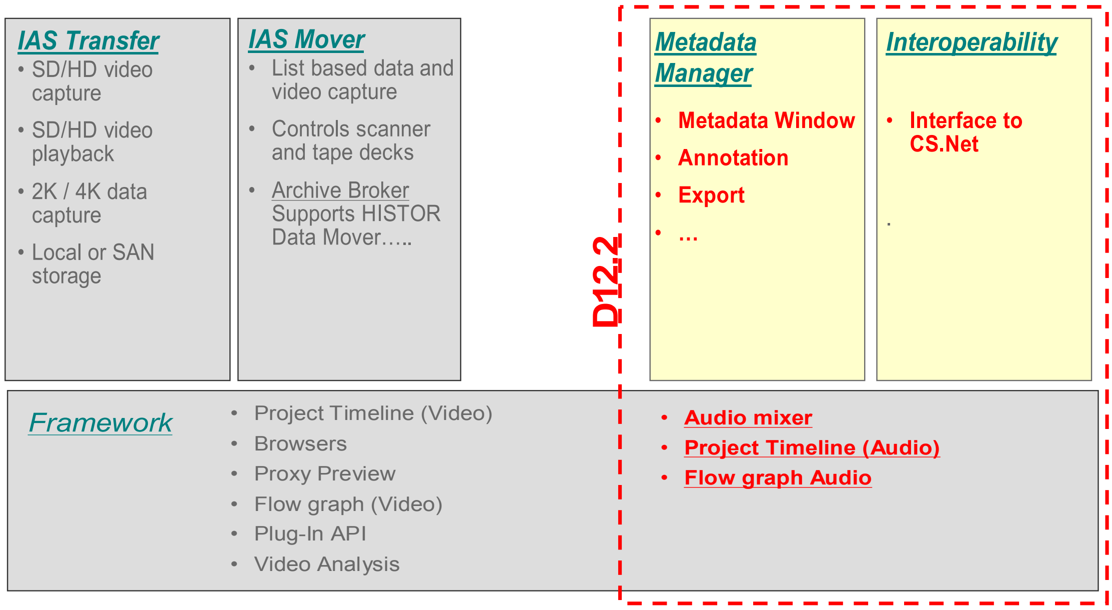

Figure 5 SIVSS IAS feature set Oct. 2006

4.3 Audio Subsystem

Adding Audio functionalities to a Video ingest system does not only require the development of audio I/O and audio processing algorithms. Moreover a framework for integrating audio and video processing is needed. The design had to deal with the following questions:

– Audio processing should be represented in a separate flowgraph structure from the image flowgraphs.
– how changes in connectivity of an image processing flowgraph shall be reflected in audio flowgraphs
– at what point in an image and audio processing chain will the synchronization occur
– how existing synchronization shall be affected by subsequent changes in connectivity of image and audio nodes

While the smallest unit of a video clip is a single frame, the smallest unit of an audio clip is a sample. IAS should provide up to 48,000 samples per second of audio (48 kHz) opposed to the respective video frequency setting (frames per second) for the video output.

In order to achieve seamless play-out in the timeline of elements which have varying frame/sample rates, the unit of measure in the timeline and players will be switched from frames to ticks. If there are for example 7200 ticks per second, then a film frame is 300 ticks, a video frame is 240 ticks and a sound sample is approximately 6.7 ticks.

4.4 Metadata

One major task for this workpackage was to develop the basic utility software package for providing the necessary management services, e.g. metadata handling, scene detection, and annotation. This is an integral part of the ingest and postproduction system with an interface to the media asset management system.

4.5 Metadata handling in SIVSS Workflow

For the SIVSS project the need for metadata management is obvious. The pieces of metadata which are created in the process of a production must be managed in such a way that it facilitates the customer's workflow as much as possible, rather than adding a lot of extra working steps. At the end of a production, the user must be able to easily collect the set of metadata and store it together with the video/audio data. Therefore a metadata solution is needed that serves all needs of the project.

In this context the word "Metadata" describes the pieces of information which are collected in the process of a production and which are independent of the current timeline / flowgraph operations.

4.6 Interoperability

In addition to the "local" metadata storage of a SIVSS system there is an interface to a "Enterprise" media management system, which will be *Content Share.net*, developed by a Thomson group.

The SIVSS system consists of a 4 station layout. Each station has a specific function (video ingest, audio ingest, video⇔audio sync, color correction, playout station). All stations will have to share the same set of metadata. Station A should automatically be informed about metadata changes made by station B and so on.

The Metadata feature is used as a supporting process and must not decrease the system performance.

4.7 Metadata Model selected

A metadata structure is the overall format for all records in the database. It can be a list of fields acting as a metadata dictionary, or a scheme that shows relationships between data elements. The broadcasting, library, and archival communities have developed different data structure standards.

The *Society of Motion Picture and Television Engineers* (SMPTE) started work on the handling of metadata with the SMPTE Metadata Dictionary (or Register of Metadata Elements, RP210) in 1998. The SMPTE Metadata Dictionary is defined by SMPTE Standard 335M. It is regularly updated by the addition of new metadata elements (deletions are not permitted in order to give some degree of future-proofing and backward compatibility). Currently, more than 1,700 elements are listed in the dictionary. Each entry is uniquely registered so that it can be unambiguously referenced. For management purposes, entries in the SMPTE Metadata Dictionary are grouped under a number of "nodes": -identifiers and locators, administration, interpretive, parametric, process, temporal, and experimental. There are also nodes for elements registered for public use by user organizations and elements registered as private. The SMPTE Groups Register contains a list of uniquely registered groups (or sets) of metadata elements used in systems and is defined by SMPTE 359M.

The SMPTE labels register is defined by SMPTE400M and Recommended Practice RP-224. This is a specialized list of enumerated or controlled terms used mainly to label audio or video data. A label is a specialized and unique identifier that is attached to the video or audio data. Although originally intended for use in key-Length-value (KLV) encoded system, the use of the registries is not restricted to KLV. They can be used also with XML encodings. Entries in the SMPTE Metadata Dictionary can be used in conjunction with groups, types, labels, and controlled vocabularies to build descriptive metadata schemes. There is also the Descriptive Metadata Scheme (DMS-l) standard (SMPTE 380M), which was intended for use with the Material eXchange Format (MXF). It contains many of the descriptive metadata elements commonly used in program making and archiving. The full scheme has three frameworks, each with many groups and each group comprising several elements. The three frameworks are for *production, scene,* and *clip metadata*: production metadata is overall metadata that applies to the whole program; scene metadata describes conceptual and editorial information (e.g., where a scene is supposed to be, not where it was actually shot); clip information describes factual information about the capture of the essence (e.g., where a scene was actually shot, not where it is supposed to be).

The *EBU P/Meta* working group has designed a standard as a metadata vocabulary for programme exchange in the professional broadcast industry. It is not intended as an internal representation of a broadcaster's system but as an exchange format for programme-related information in a business-to-business use case. P/Meta consists of a number of attributes (some of them with a controlled list of values) which are organized into sets. P /Meta has syntactical rules, which must be followed when constructing metadata groups. This scheme is technology independent; it can be represented in KLV (key, length, value) format or as XML.

The goal of the *MPEG-7* standard is to allow searching, indexing, filtering, and access of audiovisual content across separate and diverse systems by enabling interoperability among devices and applications that deal with A/V content descriptions. MPEG-7 descriptions have two possible forms: a textual XML form suitable for editing, searching, and filtering, and a binary form suitable for storage, transmission, and streaming delivery.

The MPEG-7 descriptors primarily describe low-level audio or visual features such as color, texture, motion, and so forth, as well as some of the text-based attributes of A/V content such as location, time, and quality. Both human users and automated systems that process audiovisual information are within the scope of the MPEG-7 standard, formally called the "Multimedia Content Description Interface". MPEG-7 facilitates the creation of tools to describe multimedia content in a standardized way. MPEG-7 is not used directly to create descriptions, but rather to produce the tools that create the description. MPEG-7 defines a set of description tools, called description schemes and descriptors. Descriptors represent single properties of the content description, while description schemes are containers for descriptors and other description schemes. The definition of description schemes and descriptors use the data definition language (DDL) which is largely based on XML schema.

Because of the richness and complexity of MPEG-7, it is necessary to agree on a subset for a certain application in order to simplify exchange of descriptions. Version 2 of MPEG-7 introduces the concept of Profiles which are defined subsets for certain application domains.

The *Dublin Core* metadata standard was originally developed to describe electronic text documents but has later been extended to audiovisual material. Focusing on simplicity, it contains a list of 15 attributes belonging to three groups (content, version and intellectual property). Some of these elements can be refined using qualifiers to narrow down the meaning. This is called Qualified Dublin Core. The Scandinavian Audiovisual Metadata Group has also evaluated the Dublin Core set and has proposed a metadata set based on Qualified Dublin Core. This basic set of metadata descriptors is considered optimal for an elementary common set of descriptors that can be shared by multiple archives. Moreover, Dublin Core recommends using controlled vocabularies for providing the values for these elements.

Common to all these metadata schemes is the goal for defining metadata structures which allow for worldwide international exchange of programme material. Each entry must be uniquely registered within an international organization so that it can be unambiguously referenced. This implies a large amount of administrative work to be done and the metadata codewords are long, but unique binary strings. The primary goal of the SIVSS project, however, was to develop a high speed video server system. The generation and handling metadata is inside the system and must not degrade the performance of the SIVSS operation in any way. Therefore a much simpler and faster approach was chosen here. The metadata structure was adapted to the SIVSS system to ensure that the bandwidth is not affected by the metadata handling. This metadata format is only used inside the SIVSS database and has an interface to the content management system as well. If there would be an external exchange of programme material for worldwide delivery or archiving, the essence (video and audio) and the associated metadata can be retrieved from the databases and can be encoded into a MXF format, or any other format which is requested. This is not a part of the SIVSS project but can be implemented later.

4.8 Interoperability with “Enterprise” MAM ContentShare.Net

One major requirement for the selection of a high end production system was the possibility to interoperate within larger systems either on an “Enterprise” level or in smaller workgroups. Reason for this requirement is the need to accelerate the Postproduction process by parallelisation of tasks (collaborative workflow). SIVSS IAS should provide the right features to be used in both areas. Chapter Metadata describes the development done for the workgroup application, where the featureset for the Enterprise model is described below.

In order to limit the “lack of control on the Media Asset Manager (MAM), on the intelligence or functionalities it supports”, we decided to work together with the Grass Valley internal development of a MAM middleware software called ContentShareNet (CS.Net).

4.9 Short description of Content Share.Net

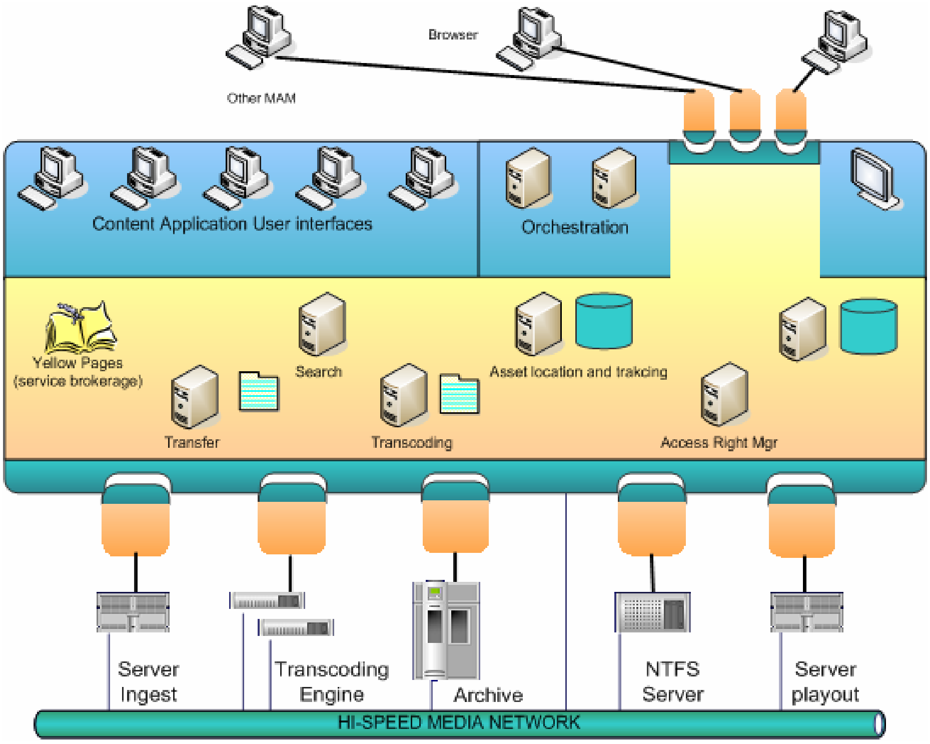

Figure 6 ContentShare (from 10000 feet)

ContentShare is a multi-layered Service Oriented Architecture based on the Microsoft.Net framework with:

- Interfaces to:
  - Automation, Traffic, Billing
  - Other Media Asset Management systems
- Applications:
  - Search, Browse, Logging, Transfer / transcode, Administration
  - Workflow automation
- Core components to administrate and control resources
- Device interfaces with simplified device management (Brokers)

Brokers are interfaces to devices that provide a level of abstraction:

- Unified control for different devices
- Provide connectivity via a unified framework
- Report device capabilities and availability
- Control basic functions to handle content
- View, edit, import and export metadata

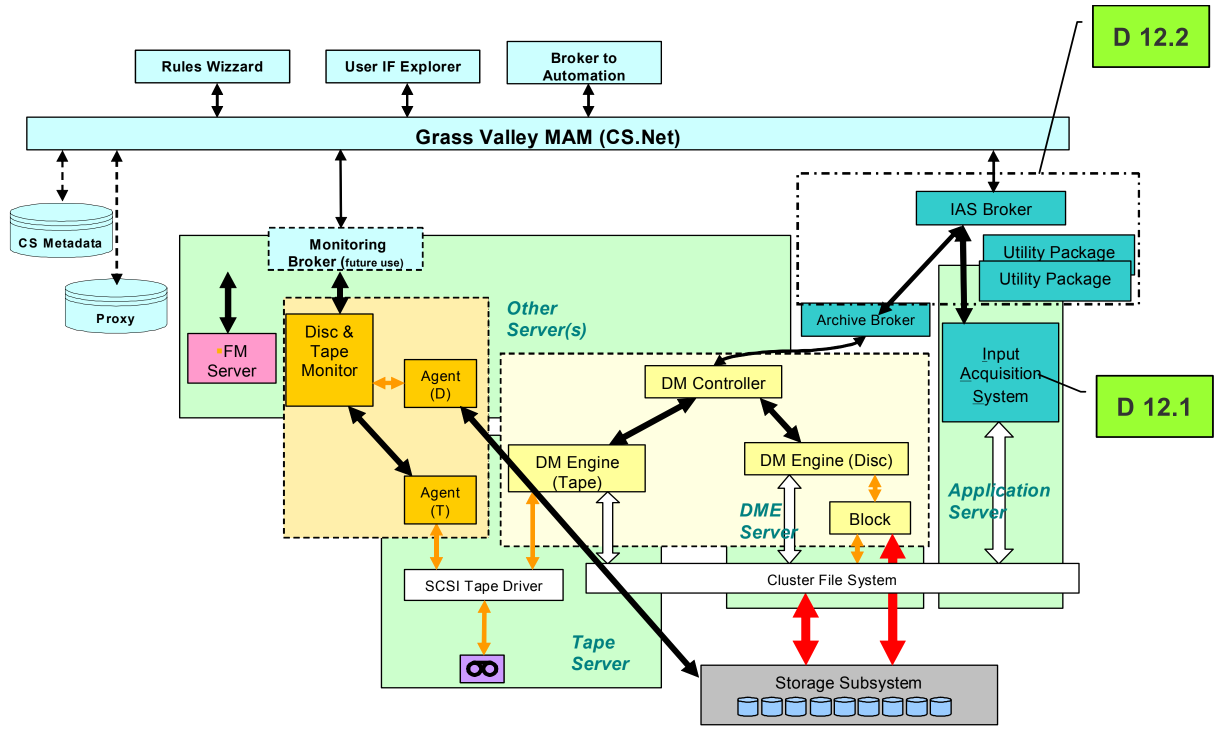

Figure 7 SIVSS Software Software component incl. CS.Net

4.10 Description of the SIVSS IAS Broker

The IAS Broker provides the following services and qualities:

1. Data management, e.g. browsing, movement of audio and video objects
2. Metadata Handling, e.g. update metadata databases between CS.Net and SIVSS IAS
3. no negative impact on ingest & playout performance
4. metadata update within 1 sec

Data Management:

We decided to use the Advanced Media Protocol (AMP). It provides the required featureset for datamanagement and is supported bei CS.Net. We developed a module which interfaces between AMP messages and the IAS Mover.

Metadata Handling:

In order to exchange metadata between CS.Net and SIVSS, CS.Net provides a so-called “watch folder”. The following modifications to the current metadata set are needed:

creation of a new metadata object
deletion of a metadata object

Parameter value changes within a metadata object are encoded into XML (following the CS.Net XML-specific syntax) and are copied into the watch folder, using ftp. This is done automatically by the SIVSS application without any user interaction. CS.Net checks the watch folder periodically and updates its database accordingly. Vice versa, CS.Net sends their metadata updates (encoded into XML files) into an outgoing watch folder.

5 CONCLUSION AND OUTLOOK

The combination of low cost, high density, high performance storage devices with COTS computer platforms proved to be able to support the requirements of even the high end postproduction industries:

- high scalability
- up to 320 MByte of bandwidth performance
- highest possible image quality
- support many legacy video and data formats

This industry is moving from a film/tape based production to a data centric one with requirements like:

- manage many different versions of content (Cinema, Airline, DVD, …, DVB-H)
- massive collaborative work (shorten time to market)
- accessing huge data volumes (don't copy, don't move)

A flexible Ingest & Playout system has been developed and integrated into a larger System managed by the GVG ContenShare.Net Media Asset Management system. The Architecture proved it's flexibility (new Data I, new HW platforms, file systems, OS, new applications). The combination of low cost, high density, high performance storage devices with COTS computer platforms proved to be able to support the requirements of even the high end postproduction industry:

- highly scalable
- up to 320 MByte of bandwidth performance
- highest possible image quality
- support many legacy video and data formats

Grass Valley will use the SIVSS IAS software components as a starting point for the product development of a multi seat postproduction tool. We already are in contact to major players in the film postproduction market, trying to understand their future needs in a, mixed 4K, 2K, HD to multiple deliverable formats, environment.