Request for Technology

Joint EBU – AMWA Task Force
on
Frameworks for Interoperable
Media Services (FIMS)

Publication Copy

Geneva
23 April 2010

[*] Page intentionally left blank. This document is paginated for two-sided printing

Contents

[*] Page intentionally left blank. This document is paginated for two-sided printing

PART 1: Request For Technology - Scope and Submission Guidelines

1. Executive summary

The Advanced Media Workflow Association (AMWA) and the European Broadcasting Union (EBU) are co-publishing this Request For Technology (RFT) as the result of joint activities in the FIMS Task Force. The Task Force has been created to address issues around interoperability in digital workflows. The joint effort between the EBU and the AMWA to seek out a common approach to the integration of software components in modern media production facilities is believed to be a fundamental need of the entire media industry.

In this RFT, we are soliciting information regarding service definitions. The scope of this detailed request is limited to some very specific functions. We are also looking for information on an overall framework which would support the business functions of the professional media industry. The focus of this RFT is on Service Oriented Architecture (SOA), and reflects a move by the media industry toward systems that use this approach. Media companies are moving towards more rapid adoption of these systems because of a need to increase agility in a market where user requirements are changing rapidly. In addition to the need for agility, these companies require the increased maintainability, scalability and extensibility that a Service Oriented approach would provide.

The Task Force's planned activities (roadmap) are as follows:

1. Requirements Collection: The Task Force has collected user requirements, use cases, and technical requirements. This information can be found in Part 2 of the RFT.

2. Publication of this RFT.

3. Intention to Respond: Potential Respondents are asked to notify their intent to respond.

4. Scope of response: Potential Respondents will be invited to provide preliminary, two-page abstract stating areas of the RFT that will be covered by their response.

5. Pre-submission clarification – Respondents will be able to request clarification of the RFT through e-mail. In addition, a face-to-face meeting with Respondents will be held for discussion of their preliminary responses and to answer any questions.

6. Submission of RFT Response by Respondents.

7. RFT Evaluation: Responses to the RFT will be evaluated against the requirements collected in the first step above. Interviews with Respondents will be organised as required to clarify information contained in the RFT response.

8. Preliminary publication: Preliminary conclusions are scheduled for release in September 2010.

9. FIMS Report: A report with the findings and conclusions will be co-published by EBU and AMWA, which may recommend the publication of technical specifications and/or or other recommendations.

10. Preparation of AMWA Specification and EBU Technical Report: It is planned that, if responses are sufficient, the AMWA and the EBU will co-author Specification(s) and Tech Report(s) for use by the industry.

11. Standardisation: We anticipate that the results of this work will likely be submitted to the appropriate due-process standards bodies for standardisation.

All interested parties are invited to respond to this RFT. Contributors need not be members of either the EBU or the AMWA. Only contributions which are free of intellectual property claims, or contributions which are made under compensation-free Reasonable And Non-Discriminatory licensing will be accepted. This RFT process and the preparation of Specifications will be carried out and conformed to the AMWA Intellectual Property Policy. See § 5.8.9 for more information.

2. What problem are we trying to solve?

Currently in the media industry, users are implementing services-based systems using existing proprietary systems with bespoke software 'glue' holding it all together. They are doing this without an open, agreed framework and without standardised interfaces as these do not exist yet. While work in the EBU has identified a number of common processes such as capture which are performed essentially the same way throughout the industry, users are currently implementing those processes in many different ways. At the same time, vendors of technology to the industry technology vendors are responding to demand for services-based products, but interoperability between different implementations is non-existent, not because this is an explicit goal of technology vendors, but because there is a lack of an agreed framework and service definitions in the media industry.

This RFT is our attempt to address this lack of standardization before bespoke implementations become so embedded that it becomes difficult to impossible for the industry to move to a more standardized, interchangeable approach. Users face tremendous pressure to create more flexible workflows and more efficient operations. We believe that without this work, many of the benefits of SOA technology will fail to be delivered to users precisely at a time when the industry is in need of these benefits the most.

3. What information are we seeking?

In this RFT (see Annex 1), we are asking Respondents to provide information in three areas:

1. An overall framework for integration of reusable components for multimedia content production.

2. Details of how this overall architecture would be applied to three interface models for three basic services - 1) Interactive Capture, 2) Transform Media and 3) Transfer Media.

3. A decomposition of services for the entirety or a subset of the professional media industry.

4. Objectives and Scope

The primary objective of the FIMS initiative is to request technology to enable the creation of standard software interfaces for services in a standard framework.

The benefit from this activity is to aid in the adoption of services-based technologies in the media industry.

To properly exploit this technology a common framework should be adopted to help ensure integration interoperability, interchangeability and reusability of services. This will drastically

reduce integration costs, allow users to more freely choose the most appropriate products on the market at any given time, improve maintainability, and aid in the adoption of new technologies.

It is recognized that the potential scope of SOA in the media industry is extremely large. Because of this, it is the intention of the FIMS project to approach our work in an iterative fashion. While this document requests information around a relatively limited scope, it is expected that we will issue additional RFTs over time, and with those, increase the scope of our work. The intended flow of the Task Force work for current and future work can be illustrated by the following diagram.

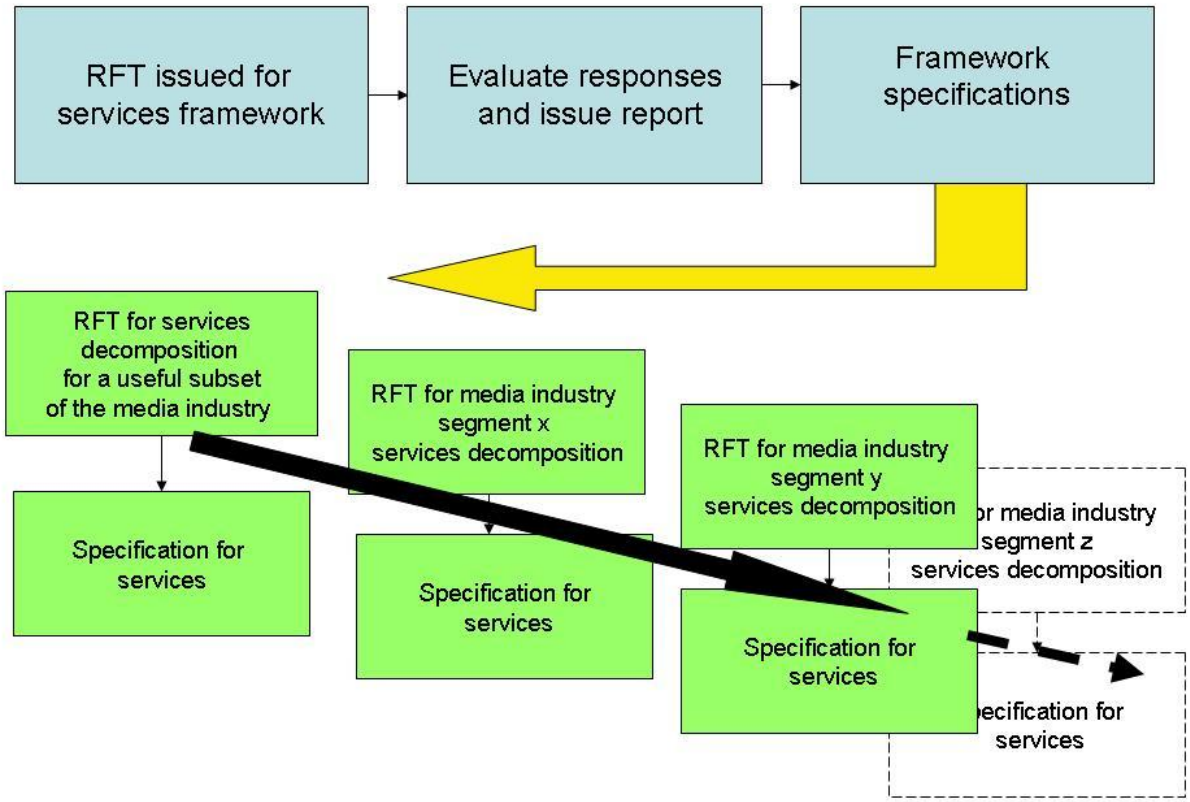

The following is the scope of this RFT:

1. **To gather information on the current availability of technologies to support the implementation of a framework for the integration of reusable components for media content production.**

The Respondents are requested to describe suitable technologies to cover the following items:

a. Service description model
b. Service interface language
c. Service directory
d. Service communication
e. Service lifecycle management
f. Service SLA management
g. Media files management
h. Time and synchronization (synch with external sources and events)

To the degree to which a Respondent feels this list is insufficient to fully describe an appropriate framework, we would welcome additional items.

2. **Respondents are requested to characterize the proposed technology for three basic services:**
    a. Media transform
    b. Media move
    c. Media capture

3. **Respondents are invited to propose their vision of decomposition of the media industry.**

Respondents are encouraged to respond to all aspects of this RFT, but we will accept partial responses. Respondents are directed to Annex 1, 'Requirements of Respondents' of this document for specific responses sought by the FIMS Task Force.

5. Submission Guidelines

5.1 General Guidelines

Vendors and developers of technological solutions, systems and products (hereafter, 'Respondents') are invited to study this document and to respond to the Task Force with firm technology proposals which meet the Requirements of Respondents.

The EBU-AMWA FIMS Task Force invites participation from all interested organisations. Responses are invited from members and non-members of the EBU and AMWA organisations.

5.2 Communications regarding this RFT

All communications regarding this RFT should be directed to the FIMS RFT management team at ebu_amwa_adm@list.ebu.ch.

5.3 Single Point of Contact Required

Respondents shall provide a single point of contact for all communications regarding the RFT. It is the responsibility of the Point of Contact to disseminate communications from the FIMS management team appropriately within his/her organization.

5.4 Intent to Respond

We ask that you notify us by 16th. May, if you intend to respond. Such notification should be by e-mail to ebu_amwa_adm@list.ebu.ch and should include the organization name and the single point of contact.

5.5 Withdrawal of Responses

Withdrawal of responses to this RFT must be made to the e-mail address above.

5.6 Rejection, Combination and Expansion of Responses

We reserve the right to reject any or all responses to this RFT at our sole discretion. We may work with Respondents to modify their responses, we may choose to accept portions of a response, and

we may merge concepts presented in multiple responses into a single final report. We may also choose to expand on concepts presented in a response to this RFT.

5.7 Visibility of Submissions

Respondents to the RFT accept that their contribution will be made available in electronic form to the FIMS members on AMWA and EBU Websites but they will not be made public. Please note that the FIMS membership includes media companies, technology providers, consultants, and other interested parties.

Portions of responses may be selected for inclusion in public reports and specifications. Working Drafts can also be made available to the general public.

5.7.1 Sharing the RFT

Receivers of this document may share it with other parties they believe would have a bona-fide interest in potentially responding to the RFT. Upon providing a copy to another party, please provide the contact details of such party to ebu_amwa_adm@list.ebu.ch for our records.

5.8 Submission Process

5.8.1 Respondent Meeting

The FIMS management team will hold a face-to-face meeting with Respondents in Turin, Italy, on 3rd. and 4th. June 2010, at which time the team will discuss the RFT and address questions which Respondents may have. Respondents should be aware that all questions and responses will be documented and may be shared with other Respondents in the interest of fairness, whether they attend the meeting or not.

Respondents are required to submit a response abstract of about 2 pages prior to the meeting, for discussion.

5.8.2 Queries

It is recognized that Respondents may have questions which come up as they review this RFT. Respondents are encouraged to contact the FIMS Task Force management team at ebu_amwa_adm@list.ebu.ch with their questions. The team will attempt to assist Respondents with background information, additional explanations, etc. Respondents should be aware that in the interest of fairness, all questions and responses which serve to clarify the RFT or provide additional information from us may be shared with other Respondents in the interests of fairness and clarity.

5.8.3 Definitions

The meaning of key terms in this document shall be in accordance with **IETF RFC 2119**.

5.8.4 Partial Responses

While Respondents are expected, at a minimum, to address the essential requirements related to the areas of their responses, it is recognised that no single technology solution can be expected to fully address all the issues that are described in the RFT. The final selection of technologies is likely to be a combination of solutions that meet the user requirements, and that a combination of submissions received may yield the best overall balance between functionality/extensibility and cost/practicability as judged by consensus within the Task Force.

5.8.5 Evaluation

The Task Force will evaluate responses against the user requirements in this RFT, and a selection of the technologies, judged as appropriate by consensus within the Task Force, will be made. The selected technology proposals may be drafted into preliminary standards documents for final standardisation work by organisations to be selected in light of the solutions retained.

5.8.6 Submission

- Respondents should contact the EBU - AMWA FIMS Management Team at ebu_amwa_adm@list.ebu.ch as described above to obtain a submission reference number. Please use this submission reference number when you include your submission and in all correspondence regarding your submission.
- This submission number should also be used when submitting a two page abstract of the proposal prior to the Respondents' Meeting on June 3rd. and June 4th. 2010. Your response must follow the following file naming convention: FIMSxxx-0.yyy where xxx is your submission reference number and yyy is the common file suffix for the document file format.
- The complete submission must follow the naming convention above, except that the -0 becomes -1. Any supplementary documents must be named FIMSxxx-z.yyy where z is incremented by 1 for each additional document.
- Upload the file to the URL you will be provided with together with your submission number.
- Announce your submission by sending an e-mail to the general e-mail reflector (ebu-amwa@list.ebu.ch). You should receive a confirmation e-mail from the Management Team acknowledging your submission. If you do not receive a confirmation e-mail, you should send an e-mail to the Management Team as described above requesting acknowledgment of receipt of your submission.
- Your submission **must** contain an IPR statement as defined in this RFT. (See § 5.8.10)

Please do not forget to indicate your contact details and address when communicating with the project management team.

NOTE: To be considered a response to this RFT, all responses must be submitted per the process described in this RFT. Verbal or written submissions which are not made per the contribution guidelines described here will NOT be considered part of the RFT response.

5.8.7 Submission Format

Respondents shall use the RFT Submission Template which is provided in Annex 2 of this RFT. The response will only be accepted in electronic form; paper submissions will not be accepted. Respondents should take special note of the associated Requirements of Respondents spreadsheet [FIMS_Respondent_Spreadsheet.xls]. Respondents shall indicate in the spreadsheet, the page and section number in their response which corresponds to the user requirement, or they should indicate that they are not responding to that user requirement, as appropriate.

Submissions should be submitted as described above not later than 5:00 PM [US Pacific Time] 13th. June 2010.

5.8.8 Schedule

The anticipated time schedule is as follows.

- Presentation and discussion of proposals with Respondents will be done at a Task Force open meeting, to be held on 3rd. and 4th. June 2010' in Turin, Italy. (Duration and timing presentations will be directly communicated to Respondents).

- Questions about the RFT sent in via e-mail will receive written responses or will be subject to teleconferences and/or web meetings until 3rd. June, 2010.
- Responses to the RFT received by the management team not later than 13th. June, 2010.
- For more details on these guidelines please see § 5.9
- Start of technology selection and preparation of a 'Request For Standardization': Target: September 2010.
- Hand-over to formal standardization work Target: December 2010.

5.8.9 IPR Policy

The FIMS project, all responses to any RFTs issued by the FIMS project, and the preparation and publication of any Specifications under the FIMS project will be governed by the Advanced Media Workflow Association IPR Policy[1]. Note that this description of AMWA IPR policy is for information only. Respondents shall review the full AMWA IPR Policy, including the ***Participation Agreement*** and the ***Implementer's Agreement***.

All necessary IPR documents may be downloaded as PDFs from the AMWA website at http://www.amwa.tv/policies .

All entities participating in FIMS activities or submitting responses under this RFT will be required to execute a ***Participation Agreement***. Execution of this agreement will be required of everyone whether they are AMWA members, EBU members, or other entities interested in participating or contributing to this process. Companies do not have to join the AMWA or the EBU to respond to this RFT.

The Scope of Use of any contributed Necessary Claims shall be limited to what is contained in any AMWA Specification which is written as a result of the FIMS project. All contributors shall have time to review any Specifications for non-contributed Necessary Claims before final approval.

5.8.10 IPR Declarations required with Contributions

All responses to the RFT shall be accompanied by an IPR Declaration as stipulated in the AMWA IPR Policy. A Respondent may make one of two IPR declarations, as set out fully in section 4.5 of the AMWA IPR Policy[1]. These declarations are summarised here for convenience.

No License Required	The Respondent declares that no license is required to implement any contributions made in their RFT response.
Compensation-Free, (RAND-Z) Reasonable and Non-Discriminatory License	The Respondent declares that they will grant a license to all implementers regarding any contributions made in their RFT response

Submissions which do not contain an IPR Declaration in accordance with one of two options listed above will be rejected.

*Note: The AMWA IPR policy allows for two additional declarations – **'RAND'** and **'None Of The Above'**. These two declarations are only allowed during final review of an AMWA Specification, and only for Necessary Claims which a contributor discovers were included in the document by a third party (non-contributed Necessary Claims).) RAND and None Of The Above are disqualifying IPR Declarations for any responses to this RFT.*

[1] http://www.amwa.tv/about/policies/AMWA_IPR_Policy2.0[adopted11_14_08].pdf

5.8.11 Gaining Access to Licenses for Contributed Intellectual Property

If the AMWA and the EBU decide to publish Specifications as a result of the FIMS project, Implementers may obtain any RAND-Z licenses for contributed Necessary Claims by executing an AMWA ***Implementers Agreement***. The Implementers Agreement is available to all members of industry, whether they are AMWA members or not. The AMWA Implementers Agreement may be downloaded from the AMWA website, noted above.

5.9 Timelines

The following is an overall project timeline for the FIMS project:

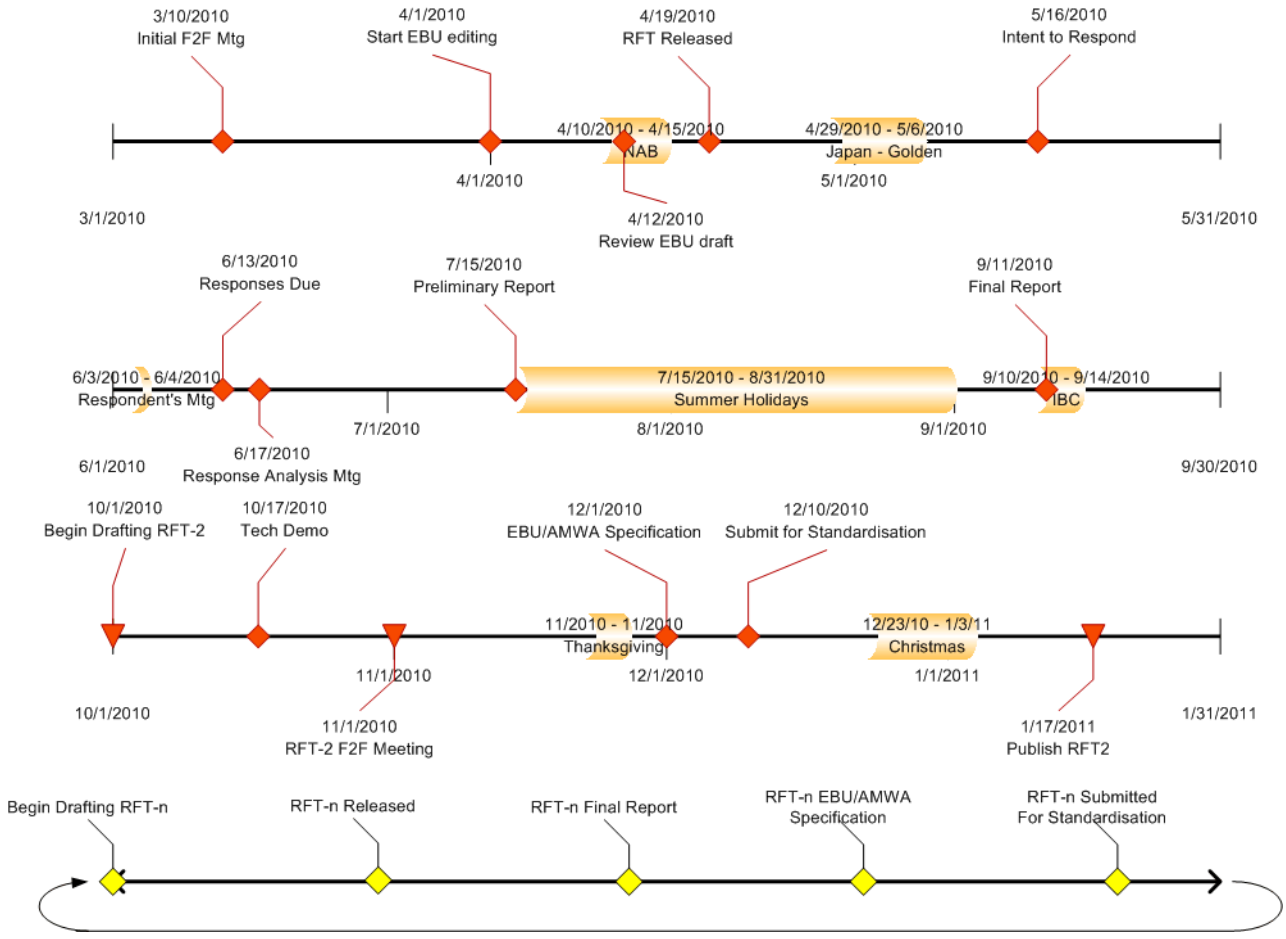

6. About the FIMS Task Force

6.1 Why did the EBU and AMWA initiate this Joint Task Force?

While the market is moving away from traditional proprietary media technologies in favour of IT based ones, the design and management of enterprise level production systems is still a challenge for most broadcasters. A main obstacle has been identified in the lack of standard interfaces between components and systems that forces system integrators to devote consistent resources in the development of custom adapters to integrate components from different vendors. This in turn generates scalability and maintenance problems as the substitution or upgrade of one component can require further adaptation expenses.

The EBU, through its ECI-NP (Expert Community on Integrated Production - Networked Production) group, is evaluating the use of Service Oriented Architecture (SOA), a system design pattern that is believed to offer the potential for greatly improved interoperability over the current system design practices based on proprietary interfaces. AMWA has set up a Media Services Architecture Group (MSAG) with a similar mandate.

Therefore, as both the EBU and AMWA have similar activities and intents, the two organizations have agreed to engage in joint efforts in order to gather a broader consensus between users and technology vendors as well as to share knowledge and join resources to better serve users and the industry as a whole.

6.2 What are the expected benefits?

The introduction of IT-based production models is considered a strategic move to fulfil the market demand for new integrated content production and distribution services. Interoperability of systems and components is needed to guarantee flexibility in the design, configuration and upgrade of production infrastructures and to keep its associated Total Cost of Ownership as low as possible.

The joint Task Force is made up of both technology vendors and users. Technology vendor members are interested in reducing their costs and risks associated with integration. User members are interested in faster time of integration, with lower cost and risk. The adoption of standard interfaces at the business level would allow these goals to be achieved.

6.3 What are the risks if the activity is not done?

The risks are that market solutions do not spontaneously converge towards interoperable models. This could lead to high costs and low efficiency in the deployment of the new IT technologies in the broadcast environment. It could also lead to a loss of competitive advantage as users are unable to implement more flexible workflow designs which are required to meet the requirements of an ever-changing business environment.

6.4 Co-chairs:

John Footen, on behalf of AMWA assisted by Brad Gilmer, AMWA

Giorgio Dimino, RAI, on behalf of the EBU, assisted by Jean-Pierre Evain, EBU.

PART 2 - Use Cases and Background Information

7. Introduction

A Reference Architecture is defined as 'a template/blueprint solution for an architecture for a particular domain that also provides a common vocabulary with which to discuss implementations, often with the aim to stress commonality.' The AMWA/EBU SOA Reference Architecture described in this document is meant to provide a technical reference view which complements the business process view.

The Reference Architecture will first be used to support the RFT process.

The Reference Architecture will then be used as the framework for the FIMS report.

The reference Architecture is meant to be:

- Technology neutral
- Encouraging the use of existing standards
- Supporting the concept of a Service Oriented Architecture (SOA)

8. Limitations of current approaches

In the framework of IT-based media production, the interoperability of vendor solutions and the design of open, scalable and upgradeable systems is still a challenge. System integrators tend to offer systems where the various components are tightly coupled. Replacing a single component with an equivalent one from a different vendor may require substantial system reworking and in-depth knowledge of the system. This results in high costs of change for the owner and a dependency on the system integrator.

Open systems can be designed only after the definition of common functional specifications and identification of base components. In the traditional media technology environment, a few standardized interfaces are sufficient to ensure interoperability in a multivendor environment, namely the audio and video signal formats and electrical interfaces. Manual operation is accepted as necessary in many trivial tasks, such as loading/unloading media in a player.

A great advantage of data centric approaches is content format independence and greater automation capabilities. Unfortunately this multiplies the number of interface points required for seamless integration and orchestration of components. While most software products on the market expose APIs (either publicly or under NDA conditions), specialized system integrators are often needed to write the software glue necessary for integration of systems. In the absence of standardized middleware and interfaces, system integration often results in complex craftsmanship where each integration requirement needs extensive testing and provides scarce reusability possibilities. Some of the most serious integration issues are described below:

8.1 Interoperability between vendors

Later replacement of a component with a different one with equivalent functionalities generally requires the revision of the software, with the associated risks, costs and testing times. Stronger

vendors tend to make bilateral agreements with complementary partners and provide built-in interfaces for specific products. Market competition is therefore somewhat limited.

8.2 Long term maintenance

Integration of components via custom interfaces poses problems of long term maintenance and vendor lock in of customers. Replacing phased out products can have serious impacts on the overall system performance that must be evaluated case by case. Generally, software revision by the original system integrator is unavoidable, adding up considerably to the replacement costs. Obsolete hardware failures and operating systems upgrades can equally become challenging operations and a source of hidden incompatibilities between applications.

8.3 System scaling

During the lifecycle of a system it is often required to extend its functionalities or provide headroom for more users and higher throughput of content. An upgrade path that was not designed from the beginning could easily hit architectural limitations and provoke unexpected bottlenecks difficult to track down and eliminate.

8.4 Dynamic reconfiguration of processes

The media business is a very dynamic one where new delivery channels and associated production models are rapidly introduced with varying success. The impact on existing systems may vary from the integration of new functionalities to the implementation of new work flows. Tightly coupled systems can require major modifications in the integration software and extensive testing to fulfil the new requirements.

8.5 Exception handling

Even when workflows are carefully designed and verified there can occasionally arise situations where a divergence from standard process is required to handle specific unforeseen needs. When the orchestration logic is not exposed on an independent layer and can be bypassed in emergency serious glitches in the effectiveness of the system could occur.

8.6 Vendor reliance

In tightly integrated systems, market stability of key components is a crucial issue. This includes bug fixing, feature evolution and API stability. Unfortunately, the ability of a vendor to fulfil these requirements depends from a number of external conditions, like market success of the product, profitability of the market segment the product is located in and the financial solidity of the vendor. A key component that cannot be updated could cause the freezing of the evolution of the entire system. Market failure of the system integrator, especially when the integration source code is not available to the customer, can undermine the maintainability of the overall system.

8.7 IT innovation cycles

Technological innovation in IT, as expressed by Moore's law, causes a product's phase-out pace to be much faster than what has been observed with traditional video equipment. This requires that customer investment policies be revised accordingly, but also that components evolution or replacement becomes an essential feature of system design. This requirement is very difficult to incorporate in systems based on custom interfaces.

8.8 Availability of technical talents

In the broadcast environment, IT still represents a major paradigm shift for technical personnel in the commissioning, design and management of operational infrastructure. Proper requirements, dimensioning of resources and analysis of workflows are often gathered on a trial and error basis starting from the experience obtained in traditional video infrastructures. Therefore, systems designed without sufficient flexibility could suffer long start up periods. Simplicity in maintenance operations and clear bug tracking is also very valuable.

9. General Media Industry Characteristics

While SOA technology is used in a number of different industries, there are a number of unique characteristics of media technology and the media business that must be considered in any response. Below is a non-comprehensive list of such characteristics.

9.1 Craft-oriented industry

Unlike manufacturing or many other industries, a certain amount of creativity and uniqueness needs to go into every asset produced by a media facility. Media production is not entirely an assembly line process. There are certainly large areas of production that can benefit from heavy business process orchestration and formalized rules. But there is also a little bit of reinventing the wheel that must be done every single time, simply because creativity is involved. No one makes the exact same movie twice.

9.2 Significant Human Workflows

There is a limit to the amount of pure automation because there cannot be too much automation without having the product suffer. Regardless of how much automation technology is available, there will still be large portions of the production process that are human-led. Creativity cannot be automated. Resource allocation is complex. For media professionals, it does matter who does colour correction, it does matter who is editing the news story, and it does matter who does the sound mixing.

9.3 Highly Collaborative

A media enterprise also features a much more collaborative environment than many other types of companies. Media is the work of many people, some only working on a few seconds of the duration of an entire finished product. A single asset will need to be accessed by many individuals, sometimes for an extended period of time and sometimes simultaneously. Each may require a different portion of that asset or a different view into a particular production.

9.4 Geographically Distributed

Participants in media business processes can be found all over the world working from home, studios, on-site, etc. Many of these participants work on a freelance or project-to-project basis, and many production companies will move from location to location with each new project. This means there is a need for a highly secure, federated infrastructure.

9.5 Large File Sizes

Media file formats are not the same thing as a company's e-mails or Word documents. Media files are usually orders of magnitude larger than traditional IT files. This means that the media Use Cases are different in terms of size and network loading compared to the typical Use Cases

encountered in SOA designs. Instead of simple messages needing to move from service to service, there is a need for entire media assets to move through the software infrastructure. With such large files come a number of unique challenges with regards to storage, transfer, and capacity.

For real-time media transfer of a single file, the available bandwidth must be guaranteed to be greater than the bit rate of the format being transferred. With so many media files needed in a functioning media enterprise, meeting storage requirements is crucial. In most traditional SOA implementations, a relational database is attached to the middleware layer and is ample to capture and safeguard important service communication. In a media application, network-attached storage (NAS) or storage-area network (SAN) solutions are required to hold the large media files of a media organization.

Regarding bandwidth and throughput of media, there are a number of areas where SOA in the media industry is a challenge. For example, while the transfer of a record from one bank to another (or one hospital to another, etc.) can be considered a trivial task if both ends of the transfer are available at the time, the transfer of media assets from one place to another is significantly more difficult. A single 4k asset can take hours to transfer over a standard Ethernet connection, and the transfer cannot just be initiated and then forgotten. Media file transfers must be managed. For example, the priority of various transfers moving through bandwidth limited connections may need to be juggled due to unforeseen circumstances such as breaking news in a news environment. Furthermore, network conditions are dynamic, and a transfer which initially was predicted to complete in fifteen minutes may suddenly require a total of 45 minutes, which may be unacceptable to the user.

9.6 Streaming

There is often a desire to begin working with or playing out media before an event is complete. For example, a news department may want to begin editing a story while the feed from the remote site is still coming in. This is sometimes called streaming. Such a need is unique to media and must be accounted for.

9.7 Quality degradation with processing/coding

The media information that moves throughout the media company is unique not only because it is large and time sensitive, but also because it is quality sensitive. Media quality is important. It is not sufficient to have the wrong media at the right place at the right time. The media must be at the right resolution and in the right format to be used effectively. When money is processed in a bank, it does not lose some of its worth at each step — the amount remains consistent throughout any given business process no matter how it is split up or distributed.

With media, however, this is not the case. Data loss may be an acceptable aspect of the normal course of business in the media industry. Each processing step in media requires a specific quality and format of media, and there are many actions such as transcoding that affect the quality of the media that pass through it. Sometimes these processing steps have the unfortunate side effect of lowering the overall possible quality of media through the introduction of compression or processing artefacts. It is important that technologies understand and account properly for these kind of ‘generational’ effects.

9.8 Multiple Resolutions

Many production services such as editing or logging are usually not quality dependent and can work just as well with a low-resolution proxy version of a media asset as they can with the full-resolution asset. In fact, some services may work better with a low-res proxy because the smaller asset size

will allow for processing on slower computers and transfers over more restricted networks. Technology should account for the potential use of proxies in services whenever appropriate.

End consumption of media is often done with different formats. So while it may be the same basic content, it may be distributed in different resolutions. Abstraction of media management issues such as this may require understanding of these details.

9.9 Large number of exception cases

Every industry has its exception cases. The media industry, especially in areas such as news production and feature production, is no exception. The whole concept of 'breaking news' feels like, in reality, one huge exception case. The technology must be able to accommodate such situations. The technology should allow users to circumvent it and, if needed, 'synchronize' afterwards to account for any disconnects.

9.10 Special time requirements

In the media industry, real-time means that the information in one frame of media is transmitted in the time it would take a viewer to watch that one frame. Real-time for SOA technologies, however, probably means that the information is delivered within an acceptable reaction time for a human user. This could be sub-second, or it could be several seconds. Some elements of media workflow are real-time in the media sense and the technology should appropriately interact with such workflows.

Media assets — regardless of their quality, completeness, or intended recipients — all have duration. This means that media assets have a temporal context: once a piece of media starts playing, the audience will get confused should it end prematurely. Media also may have temporal metadata associated with it such as scene labels, dialog and closed captioning, or mark-in and mark-out points. Unlike the stateless information that flows through other industries' infrastructures, media is time-based.

This means:

- It can only be consumed properly at a fixed rate and in a fixed amount of time.
- It often does not make sense when viewed partially, interrupted, or in the wrong order.
- It must ultimately be streamed to the user (although that streaming may occur on the client).
- It must ultimately be decrypted and converted to an 'analogue medium' to be consumed by 'analogue' human beings.

Additionally, many services in the media industry have strict time-based requirements. They may require that something be done ASAP, or may require that it be done by a specific time or it is useless. For example, if a programme is scheduled to air at noon, it will not be useful if the service processing it finishes at 12:30.

9.11 High level of reliability

In television or other broadcast media, consumer expectations regarding reliability are stringent. No consumer is happy with even small disruptions in what they view. In addition, regardless of the manner of distribution, the acquisition of media is always going to be time sensitive. If the system drops three frames of footage on acquisition of a news conference, that time will never occur again. Technology should account for extremely high levels of reliability.

9.12 High amount of variability in Organizational Size

While the industry as a whole is not particularly large compared to other industries, there is a great degree of variance in the sizes of the companies and their appetite for adopting technologies such as SOA. Organizations can be as large as a major Hollywood Studio to as small as an independent post-production company or a freelance individual.

9.13 Media is 'a Horizontal'

It is argued that in many senses, media is not a vertical. It is a horizontal that crosses all industries. All industries make use of media in different ways. This characteristic may mean that technologies and approaches adopted for media by the media vertical may also be adopted by those who use media in other industries.

10. Reference Model Overview

10.1 Precepts

Our vision is that system integration of IT-based media production systems can be obtained through a pool of production resources interconnected over a network via standardized interfaces.

Business processes are then implemented via a proper configuration of a resource orchestrator.

A conceptual model can be defined by the following layers:

Architecture model

- Application Layer
- Process Orchestration Layer
- Abstract Services Interface Layer
- Services Implementation Layer

This model is illustrated in the following Figure.

For modularity sake these four layers should be as independent as is possible.

As the complete definition of a complex production infrastructure is a titanic objective, this initiative concentrates its efforts on the identification of the middleware needed to define proper standard interfaces to interconnect components that will form a framework to implement the *Abstract Services Interface Layer*. The other 3 layers are out of the scope of this RFT.

It is envisaged that the business intelligence related to media handling and processing can form an ecosystem of SOA based components. To properly exploit this technology a common framework should be adopted to support the integration of the various components in real production environments.

The aim of this initiative is to foster the birth of a market for reusable components and the underlying infrastructure. Therefore, definition criteria and granularity of services are important concepts to be specified in the present context. A number of these criteria to be observed are defined by the OASIS [Organization for the Advancement of Structured Information Standards] model. They must be adapted to the media content production environment.

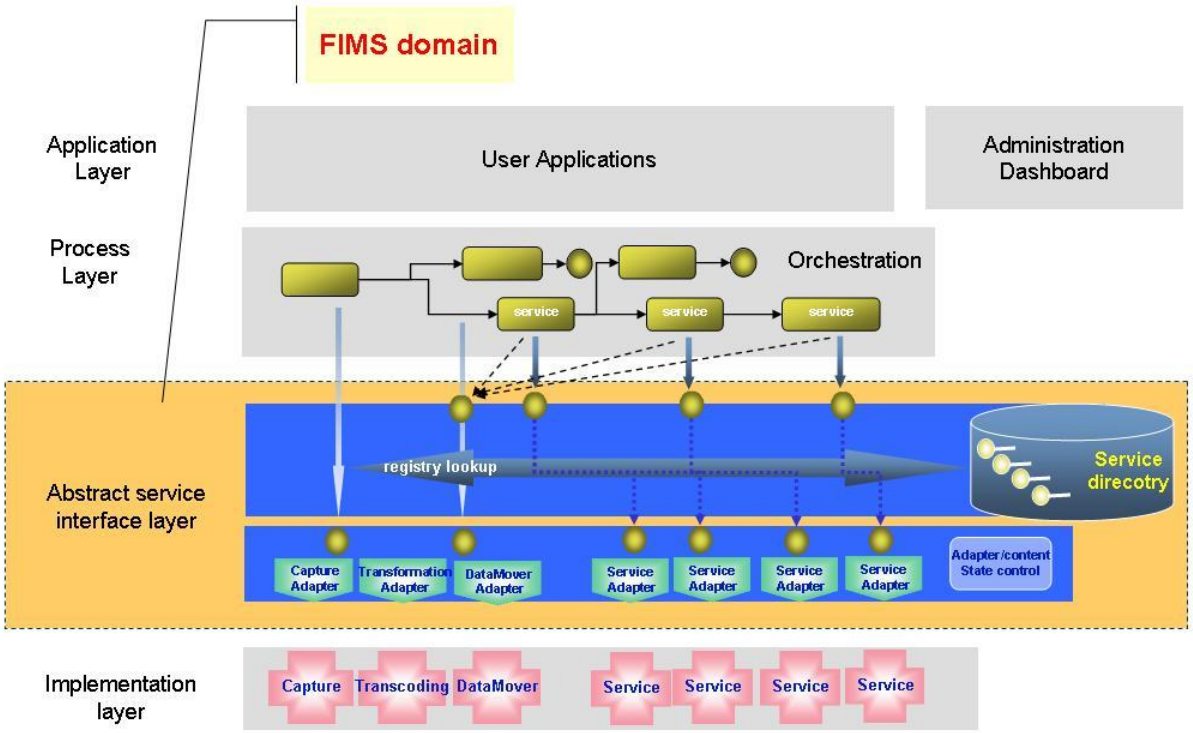

Figure: The Architecture Model and the FIMS domain

A service is a business logic artefact modelled in such a way that it assumes a pure reactive role. It is invoked by a consumer with a set of arguments, specified by the service interface, and executed according to a standardized service contract, known to both partners (service and consumer), synchronously or asynchronously. It is worth noting that in many cases an artefact can be implemented both as a standalone application and as a service. The choice is done depending on the role played: if it is active, that is the artefact can be activated spontaneously, it is an application. If it activates on demand from a consumer, then it is a service. Consumers select services according to their service contracts, communicate with services via their service interface exclusively, and don't need to know any detail about the service implementation. This principle is known as 'loose coupling'.

Basic services implement simple content manipulation operations (e.g., transcoding, rendering, play, capture) and can be assimilated to appliances that can be found on the market. Composite services are aggregates of basic (or composite) services properly orchestrated to accomplish some complex task (e.g., ingestion, editing, play out, archiving).

10.2 OASIS Reference Model

It is important in the context of this RFT to allow the fair comparison of different proposals. To assist in achieving this goal, it is considered essential that a set of common definitions be used to describe the different components of services.

OASIS has done significant work in the identification and definition of the main components characterising a service, which fits very well with this RFT [see section 4.3 of http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.pdf].

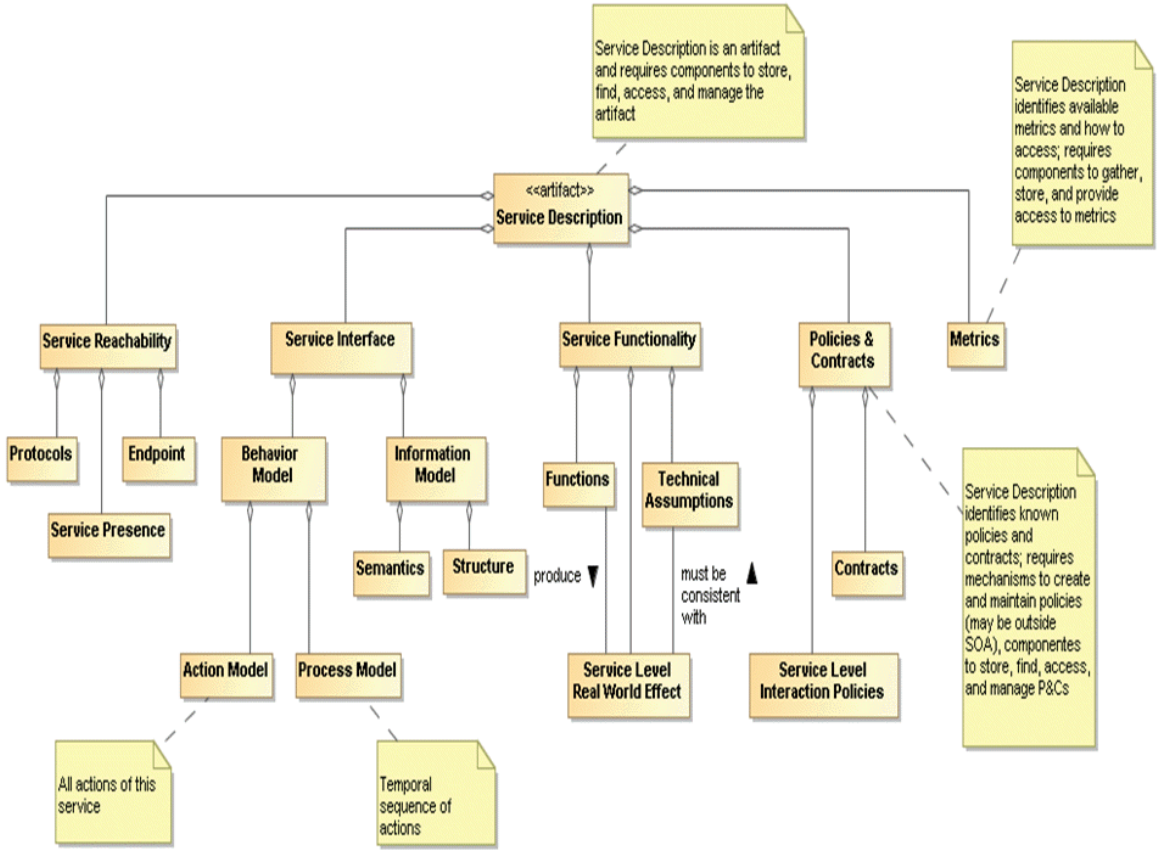

Figure: OASIS Reference Model

Respondents are therefore encouraged to use this OASIS model to shape their responses in order to maximise the understanding of their technology in comparison to others.

The use of the proposed technology shall be applied to the RFT basic services contextualised in use cases in another section of this document.

As an example, the Figure overleaf shows how the Capture Use Case components would fit within the OASIS reference model:

10.3 References and other material

Several major organizations are involved with the development and oversight of open standards for SOA:

- Organization for the Advancement of Structured Information Standards (OASIS; http://www.oasis-open.org)
- Web Services Interoperability Organization (WS-I; http://www.ws-i.org)
- World Wide Web Consortium (W3C; http://www.w3.org)
- Open Service Oriented Architecture (OSOA; http://www.osoa.org)

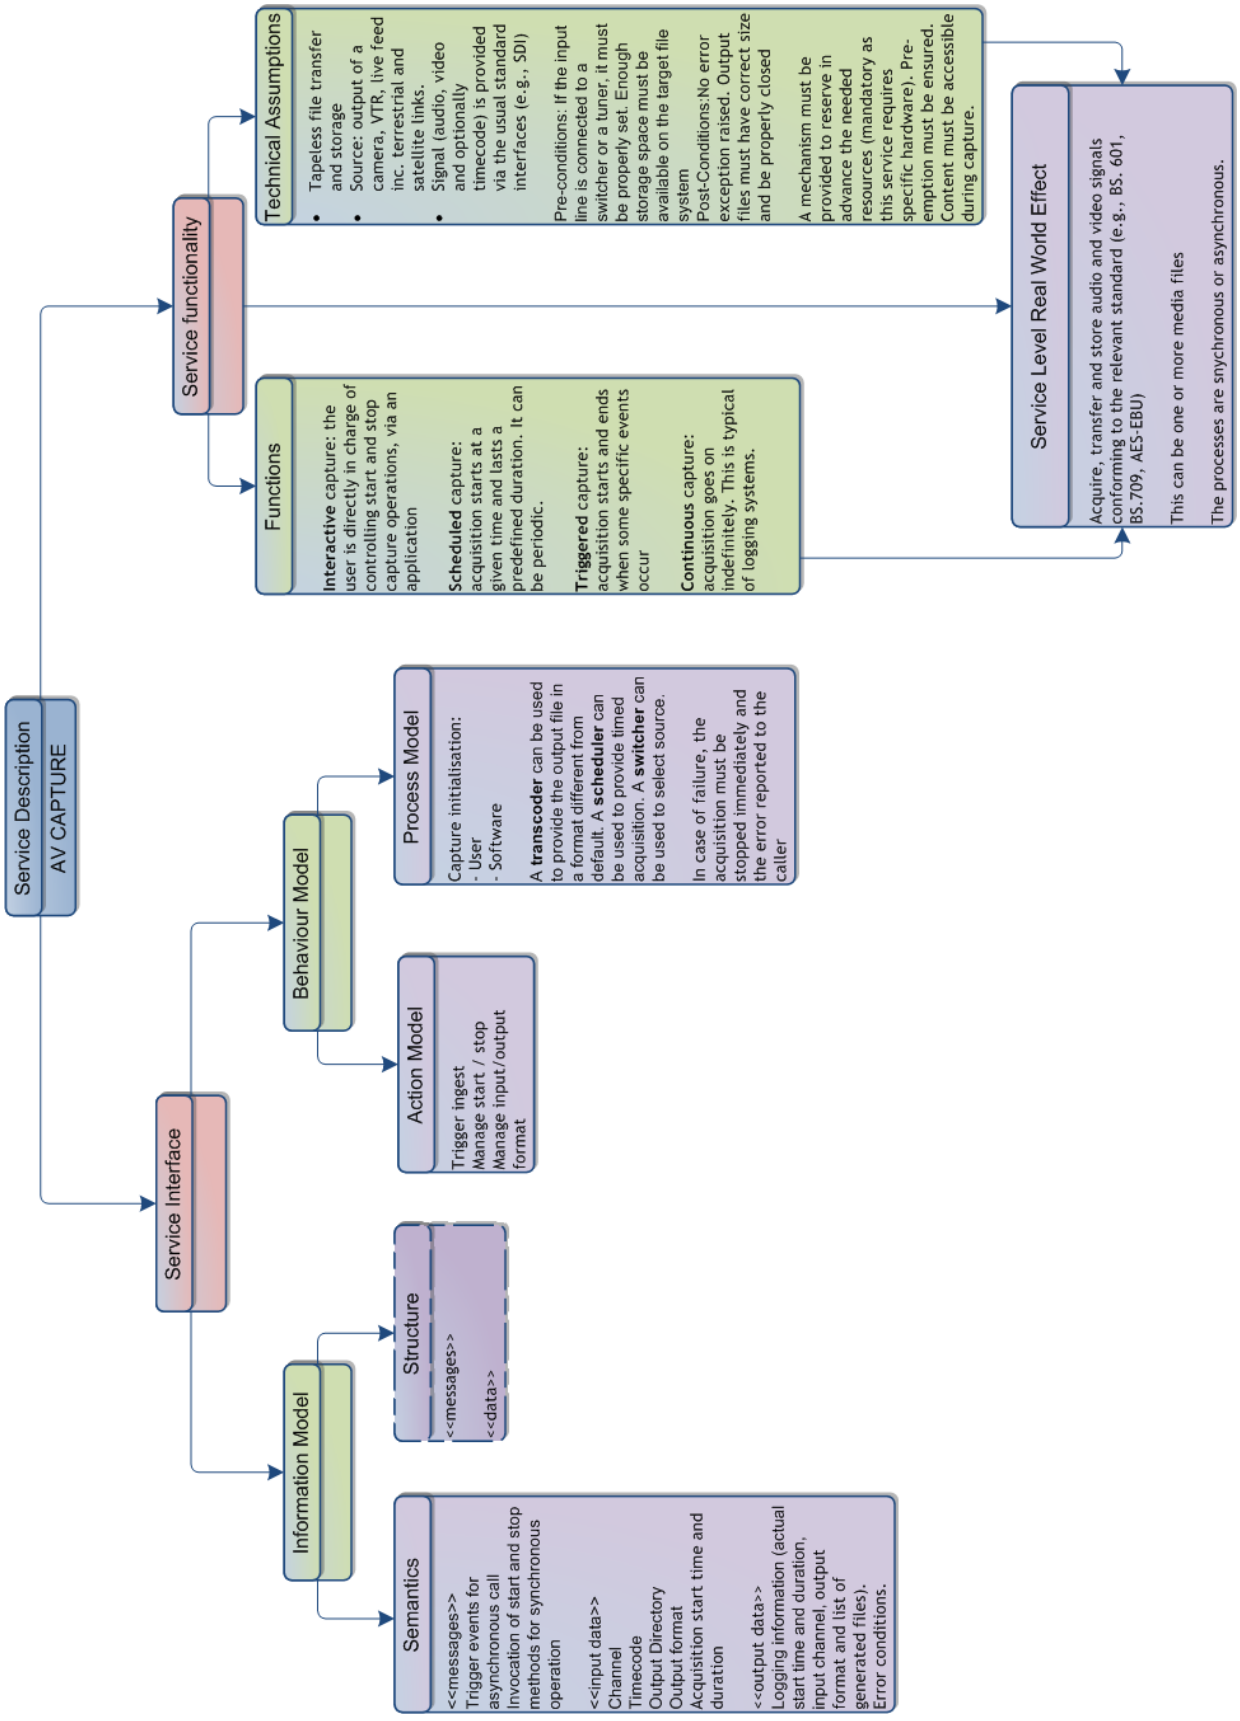

Figure: AV Content Capture Use Case within the OASIS Model

11. Use Cases

This section details some use cases of the overall AMWA/EBU Framework for Services. These use cases should be linked to the business function and technical decomposition of how and why it is related.

Respondents are invited to test their technology against these basic use cases, which will be the only ones taken into account for comparing solutions/technologies. Information on the usability of the technology to a wider range of use cases is welcomed for informational purposes but is not required.

11.1 Basic Use Cases that will be tested in this RFT

11.1.1 Use-Case – AV Content Capture

Name

Interactive capture of AV content

Summary

Content items are created from one or more related video and audio streams and associated metadata. Often a human operator starts and stops the capture operation. A typical scenario where this use case is used is in a studio using tapeless recording of camera and microphone feeds. Variants of this will apply to other types of capture (e.g. on a camera with removable media).

Description

Video and audio streams, typically provided via an interface such as HD-SDI, AES-48, etc, are captured as media files, using one or more codec and wrapper formats appropriate to the workflow. Often the capture is initiated by a human operator, who will also enter appropriate metadata about the capture (e.g. programme name and location). As part of the capture use case, it may be necessary to copy files to another location (for robustness) or to transcode to other formats (for workflow reasons). It may also be necessary to notify users about the success or otherwise of the capture.

Initiating actor

Human operator or an automation system.

Supporting actors

Depending on the workflow and robustness requirements the following optional actors may be involved:

- Transcoding Service
- Data Moving Service
- Scheduling Service
- Notification Service

Inputs

Information about the capture configuration:

- Metadata such as programme name, location, scene number.
- What sources are to be captured
- What versions are to be created (e.g. full quality, web proxy)
- What video, audio and wrapper formats are required for each version.

- When each version is required (immediately or at a later time).
- Whether recordings should be broken into smaller files, and if so on what basis (e.g. at regular intervals, on scene changes, under manual control).
- What actions should be performed when the files have been captured (for example the full quality version may be sent to a production asset management system, while the proxy is sent to a web server, and a notification sent to users that capture has completed).
- Video, audio and timecode sources. These might be local cameras/microphone, feeds from a remote source, or (for legacy purposes) VTRs. Sources might be connected via SDI, Firewire, IP or other connections.

Outputs

Content item(s):

- The completed set of AV files are available from the recording at the required locations and in the required formats
- Metadata corresponding to at least the minimum set required

Information about the capture itself, e.g.:

- resources used
- any technical issues observed

Pre-conditions

Storage space is available for the media files that will be created. If capturing from a VTR or other pre-recorded source, this is cued to the appropriate point.

Post-conditions

The content item is available

Any users have been notified

Default flow

Initiating actor provides configuration information about the capture.

Initiating actor requests start of capture.

Media files are created.

Initiating actor requests end of capture. Media files are finalised.

As required:

- additional versions are created by transcoding
- the media files are transferred to the required destination(s)
- users are notified

Information about the capture is output.

Exception handling

If the minimum metadata set is not available the capture request is rejected

If an input is missing, has the wrong standard, or some other problem, the initiating actor is informed.

If a problem develops with an input during capture, the initiating actor is notified, and a log of the problem is included in the information about the capture.

Optional actions

If the configuration requires long recordings to be broken into smaller files, then at appropriate intervals during capture each media file are finalised and a new one created.

During capture, the operator may log information about the recording (e.g. suspected technical problems)

11.1.2 Use-Case – Content Publishing

Name

Publisher (Playout)

Summary

Provisioning from source service (originator – e.g. editing) to a publishing service (destination – e.g. playout)

Description

From a source service, which could be scheduled, or a manually triggered user task, an essence shall be provisioned into a publishing service. The targeted publishing service describes a category of services, which at this point could for example be a video server for delivery to a Web or mobile channel, a broadband channel, or a playout service for delivery on air.

Initiating Actor

Initiating actor could either be a direct request for publishing by a human task service or another service in the workflow before – for example editing service, search/select from archive, scheduling service, etc. The initiating actor has to provide the corresponding properties of the appropriate method for the publishing service with the request.

Supporting Actors

Depending on content characteristics of actual content and required content characteristics of the target service the following optional actors are involved:

- Transcoding Service
- Data Moving Service

Inputs

Input depends on the method type of the request:

add:

- actual content in the proper format
- content might be required local to the service
- add properties of/for the actual content(schedule time, channel, etc)

remove:

- content for playout from schedule
- essence from server storage – might be optional

change:

can have 2 variations & needs to refer to an existing playout content

- exchange essence for playout (in this case it is similar to 'add' but needs to refer to an existing playout content)
- change schedule for playout

list:
- list schedules - options: day/time, playout-content, etc.

status:
- checks status of playout content

Outputs

No content returned by the process. Information on service processing returned. Additional information returned by the service depends on the publish service method that has been requested:

add:

feedback on operation (success, warning, etc.)

remove:

see above

change:

see above

list:

RC, plus requested information

status:

see above

Pre-conditions

- Playout Application has to be available with appropriate adapters
- Playout Application has to be registered and integrated / set-up as a service
- Initiating Actor has to be up and running
- Essence has to be available if required – this might also be handled via exceptions or compensation

Post-conditions

Post conditions depend on publisher method:

add:

publishing content successfully integrated into publishing schedules and corresponding essence is ready for publishing

remove:

content successfully removed from publishing, schedule of essence optionally removed if present and desired.

change:

updates / changes successful according the request

list:
requested information successfully delivered

status:
requested information successfully delivered
For all of the above methods there exist exceptions and, optionally, compensations. Thoughts on these from Respondents are welcome.

Non functional Requirements
Securing the transmission of data – especially if it goes to another site – e.g. remote site with data encryption might be a requirement. Performance, especially for the transmission of huge essence data, has to be considered. Availability and reliability of services and infrastructure has to be considered on various conceptual levels.

Default flow
When publishing is initiated (manually, scheduled), based on the request, the appropriate distribution channel is selected with the related metadata and or attributes. If necessary, a transcoding and data moving services might be requested in preparation for the publishing call.

Exception handling
Exception handling should be considered based on the 'type' of failure:

- Timeouts waiting for responses from services – this might have various reasons and might be handled by appropriate actions (retry, self-healing, notification, etc.)
- Concrete service failures – those exceptions could for example be handled by appropriate policies, this could even mean to invoke additional services to rollback the entire transactional process.
- System failure – this could interrupt the process in the middle of operation and might requires specific status information for recovery.

Optional actions
Besides the optional sub-processing of transcoding and data moving there might be considerations to combine transcoding with data protection technologies such as Watermarking and Fingerprinting that could be combined with transcoding to optimise performance.

11.2 Other Informative Use Cases

11.2.1 Use-Case – AV Content Playback

Name
Playback

Summary
Some content stored as a file must be played back

Description
This use case describes basic playback of a media file.
A first case is upon a user request, thus the usual transport controls (Play, Stop, Pause, Fast forward, Rewind) must be provided. The display can be the user pc screen or a monitor connected

to a dedicated video output (SDI) either from the user pc if a proper hardware is installed or from a remote video switcher.
A second case is automatic play out, where the operation is driven by a scheduler.

Initiating Actor
The initiating actor can be an orchestrator for automatic functioning or a user application.

Supporting Actors
A scheduler can be used to provide timed play out. A switcher can be used to select destination.

Inputs
- One or more media files
- Play out channel(s) (dedicated output, PC display, etc.).
- Play out format (it can be different from native content format).
- Play, Stop, Pause, FF, Rewind, Go to controls

Outputs
- The media content is displayed or streamed
- Logging information about the operation
- Current stream position.
- Status of the player.

Pre-conditions
The input media file must be available for play and the transfer bandwidth must be sufficient to sustain continuous play. The requested display must be available and its format must be compatible with that requested.

Post-conditions
No play discontinuity has been recorded. No premature end of media has been encountered.

Non-functional requirements
A mechanism must be provided to reserve in advance the needed resources (mandatory when specific hardware is used). Pre-emption must be ensured.

Default flow
In case of playback initiated by a user application (of which this service can be a component), the transport controls are rendered on screen, the user selects (if not defaulted) an output channel and the reproduction can be started and stopped at any time. The user can seek to any point in the stream and request to view the content on a display whose format is different from that of the content. The current stream position must be available to other components (e.g., to capture mark points).
In case of automatic play out the player must be ready to start playing at a predetermined time or upon a given event. The operation will stop at the end of the media file or immediately upon request of the caller.

Exception Handling
In case of failure, the playback must be stopped immediately and the error reported to the caller In case of automatic play out the error handling can be more complex (e.g., switch to a fallback

source)

11.2.2 Use-Case – Interaction with Metadata

Name
Interaction with asset metadata

Summary
Allowing users to view, add, remove or edit asset metadata contained within a clip or database.

Description
A programme asset exists within a system. The method by which it was ingested into the system is (for this use case) irrelevant. The user, or a process wishes to add, change or delete metadata associated with the asset. The following processes are expected:

A request is made to the service to retrieve the metadata associated with a specified asset. Requested metadata could be structural, descriptive or both. The service returns the requested information to the requestor for display.

A user or process requests that the service change, add or delete a metadata field. The service receives the change request – in the form of a key and value – from the initiating process and modifies, inserts, or deletes the metadata within the asset. This operation is disallowed for structural metadata

Initiating Actor
A 'view clip metadata' process or an 'edit clip metadata' process

Supporting Actors
An operator (at some form of GUI), an external control process

Inputs
File-based programme assets
Information flows:

- Asset ID
- Type of request – query, change
- Metadata key
- Modification to be enacted
    - Change existing data
    - Append to existing data
    - Add new metadata field
    - Delete Metadata value (fill with null)
    - Delete key and value (purge)

Outputs
Metadata is modified, added or deleted within the asset

Information

- Process succeeded
- Process failed

- Process pending – asset not yet complete enough

Pre-conditions

1. The asset is ready for use
2. The asset type is sufficiently complete for metadata modification to be a valid operation

Post-conditions

1. Process complete status is set to TRUE

Non-functional requirements

- The service can access files on the file system in question (security, access rights)
- The initiating process has sufficient access rights to enact such metadata changes

Default flow

1. A programme asset is added to a system
2. A request is made to the Metadata service to display or modify metadata within the asset
3. A business rule is used to ensure that the request is validated for access/modification rights
4. Metadata is extracted from the asset and returned to the requestor.
5. Any (authenticated) requested change is made to the metadata within the asset
6. The newly written metadata is re-extracted from the asset and returned to the requestor (confirmation step)
7. The pass/fail status of the service is returned to the requesting process

Exception Handling

Exceptions include:

- The process is not authenticated for access to the file system in question
- The requesting service is not authorized to modify metadata in the asset

11.2.3 Use-Case – Content Approval and Quality Control

Name

Quality Check with approval stages.

Summary

The quality check and approval of a programme asset employing manual and automatic systems.

Description

A programme asset has been added to a system using a 'digitise from tape' (ingest) or 'file import' process. The technical quality of the asset is then checked to verify if it is suitable for subsequent use. The following processes are expected:

A request is made to the service to check a specified asset. Identification of the asset in a digital asset management system (DAM) is provided together with the metadata required for the QC process.

Using the metadata provided, a business rule is used to select a quality check process that is relevant to this particular asset (Asset types will typically be Programme part, Promotion, or Commercial). If the asset type can be managed by the process it requests a file from the DAM

service, if it cannot manage the asset type it is rejected. If it can be managed the DAM service delivers the file to an appropriate location from where it is processed by an automatic system that makes a number of measurements. The number and type of measurements are determined by the asset type. The results of the measurements are compared to a set of 'pass/fail' parameters and a set of results created which are available as metadata for use by other processes.

The pass/fail data is used to update the status of the asset. If all tests are passed the asset status is set to 'available', making it suitable for later use. If any tests fail the asset is added to job list for a subsequent manual check by an operator. The metadata produced by the automatic check is available to the operator and used to determine the manual checks required.

The operator selects the job from a queue and checks the asset using a manual toolset (typically including quality sound and vision monitoring, waveform and vector-style displays etc). Every 'fail' point in the asset identified by the automatic process is checked; data is manually appended to the asset record indicating whether remedial action is required. If remedial action is required the asset status remains 'not available', otherwise the asset is approved and the status is set to 'available'.

If the asset requires a rectification process it is placed in a suitable job queue to which the metadata from the manual check is appended.

All of the measurement results and pass/fail data are transferred to the DAM service

Initiating Actor
A 'digitise from tape (ingest)' process, or a 'file import' process

Supporting Actors
QC operator (to do the manual QC), DAM service (to move the file)

Inputs
File-based programme assets delivered by a DAM service
Information flows:

- Asset type (Programme part, Promotion, Commercial etc)
- Programme asset metadata: filename, title (series episode, version etc), SOM, duration,
- Essence coding information, wrapper etc.

Control flows

- Request QC service
- Job priority
- 'Complete by' date/time
- Abort job

Outputs
No content is returned by this service.

Information

- Asset status at completion (available, not available)
- Measurement results and pass/fail data
- Remedial action required (with related metadata)

States

- Expected time before process starts
- Expected time to completion (of automatic check)
- Queued for manual check (and position in queue)
- Process complete

Pre-conditions

1. The asset is ready for use, it is available in a suitable form in a DAM for transfer to the service (N.B.: the file may not be complete and the QC service may operate in 'tail-mode').
2. The asset type is described by metadata that is rich enough for the service to accept or reject the request

Post-conditions

1. Process complete status is set to TRUE

Non-functional requirements

- The service can receive files using the mechanism provided by the DAM service
- The checking service performs to agreed benchmarks
- The service should specify whether the file and data should be encrypted during transfer

Default flow

1. A programme asset is added to a system using a 'digitise from tape' (ingest) or 'file import' process
2. A request is made to the QC service to check a specified asset
3. A business rule is used to select a quality check process that is relevant to this particular asset
4. A file is requested from the DAM service
5. The file is delivered to an appropriate location from where it can be processed
6. A number of measurements are made according to a profile determined by the asset type
7. The results of the measurements are compared to a set of 'pass/fail' parameters and a set of results created
8. The pass/fail data is used to update the status of the asset
9. If all tests are passed the asset status is set to 'available'
10. Assets which have failed are added to job list for a subsequent manual check by an operator
11. The operator selects the job from a queue and checks the asset using a manual toolset. Every 'fail' point in the asset identified by the automatic process is checked; data is manually appended to the asset record indicating whether remedial action is required.
12. If remedial action is required the asset status remains 'not available', otherwise the asset is approved and the status is set to 'available'.
13. If the asset requires a rectification process it is placed in a suitable job queue to which the metadata from the manual check is appended.
14. The measurement results and pass/fail data are transferred to the DAM service.

Exception Handling

Exceptions include:

- The metadata indicates that the file should be readable, but it is rejected for an unknown reason (the codec is not as specified, for example)

11.2.4 Use-Case – EDL Processing

Name

EDL Processing

Summary

After script editing, storyboarding or off-line editing in drama production, the user (script editor, director, video editor) wants his creative output to serve as input for the craft editing process whereby the script is converted into an EDL that the craft editing system is able to process automatically.

Description

This use cases envisages a complex interaction between 4 different systems, i.e.

- The script editing or a storyboarding application used by the director, by which he has de facto defined an editing sequence before any production of raw material is executed;
- An acquisition system, usually made up of one or more cameras and production servers, which have produced sequences of unedited raw material;
- The Media Asset Management system, which is a secure deposit of raw material as it has been shot. Moreover the MAM system can store and forward EDLs to and from craft editing systems;
- The Craft Editing system processes the information of the script and/or output of the off-line editing process, assuming the MAM system has converted it into a correct EDL. The craft editing system can then automatically retrieve or process the correct material fragments from the MAM system based on the identifiers that have been put in the EDL by the MAM system and it displays the sequences in the correct order.

Initially, the director refines the shooting script eventually defining the complete storyboard of a scene or items. He defines camera positions and the sequence of the intended shots.

Using file-based cameras it is feasible to trace which files (shots, takes) have originated from which camera and thus it is possible to link them to the information in the script. So by converting the information in the script into a correct EDL-format, it should be feasible to issue script to the craft editing system and for the craft editing system to retrieve the correct media assets.

The problem is that, while the camera and the craft editing systems will use UMIDs for identification purposes, the script editing system has to use implicit references before the material has been actually shot. As soon as the material fragments have been acquired and identified by the MAM system, a service or system must process the EDL with implicit references in order to replace them by explicit material identifiers. This service is considered a core feature of a MAM system.

Initiating Actor

While this use case is typically initiated by the director or the editor, it is possible that the script editor executes it as well in order to get a pre-view of the contents represented by the script.

Supporting Actors

The script editor or the script supervisor has approved the script for further processing;

Inputs

A list of material assets have been ingested by the MAM system. A detailed script has been marked up by the director of the production:

- The director has indicated per scene a list of camera point of views, he has defined for each shot the shot type and the subject, and he has indicated for each shot which camera device will be used.
- The marked-up script contains an indication of the order and the duration of each intended shot.

The director or the editor uses the MAM system to initiate the craft editing process.

Outputs

Based on the information in the script/EDL, the Craft Editing System has been able to retrieve the specific material items (fragment)s that had been selected for further processing and it has placed them in the correct order on a timeline.

Pre-conditions

A detailed and marked up script have been stored by the MAM system;

The MAM system was able to process the script and to create a logical unit of work (an 'item') per scene; The director has been able to select and ingest material as part of the production process. The MAM system respects and stores the UMID that has been generated during the production.

Post-conditions

(see Output)

Default flow

1. Given a selected production (programme, project) and a user authorised to manipulate the material assets associated with the production, the back-end of the MAM system received the request to construct an EDL based on the available information, i.e. the marked up script and a list the available media assets;
2. In a first pass, the MAM system will try to resolve which media assets (aka different takes of individual shots within a scene);
3. In a second pass, the MAM system will try to link each individual shot with to one or more media assets and if necessary longer sequences are split into individual fragments. In the latter case, new media objects and UMIDs are defined;
4. The MAM system generates an EDL per scene, using the original UMIDs. In case of overlap or double takes, the EDL represents multiple layers of video.
5. The MAM system partially retrieves all fragments of media material referred to by the EDL and issues these to an appropriate project folder of the craft editor;
6. The MAM system issues the EDL to the craft editing system;
7. The craft editing system processes the EDL and retrieves all associated media assets from the appropriate project folder. If necessary, the craft editing system processes and checks in each individual media fragment;
8. The craft editing system displays the EDL and the associated media objects on a timeline.

Annex 1: Requirements of Respondents

This Annex details areas where the FIMS Task Force is seeking specific responses from Respondents. Readers of this RFT may find invitations in other sections to provide additional background information, but it is in this section that we have enumerated the specific information which we are seeking.

Respondents are asked to describe the framework of their proposed technology according to a number of defined categories. Please respond using the associated spreadsheet [FIMS_Respondent_Spreadsheet.xls]. For each category, say whether you have a response to that category or not. Then, please answer each of the specific requests for information that start '*Respondents shall* ...'. Respondents should then provide a pointer to the section of their response which addresses the specific request.

Partial responses to a category are acceptable and encouraged. The Task Force can choose to ignore responses not provided on the spreadsheet (or pointed to from the spreadsheet) or outside of the defined categories.

To illustrate their frameworks, Respondents are then asked to characterize three service definitions – Interactive Capture, Transform Media, and Transfer Media in the context of their framework. A spreadsheet template has been provided, based on the framework categories and the response shall be according to this template. Additional supporting materials may be provided, such as a schema defining a service.

Finally, the Task Force invites you to provide your vision of a service-oriented media industry in the form of an industry-decomposition for media. This decomposition will inform the direction and prioritisation of the work of the Task Force going forward. No template is provided for this but please try to limit all responses to no more than 5 diagrams and 10 pages of description.

Discovering a Service

1.1 Respondents shall indicate whether their response addresses the discovery of services.

- Discovering a service should be considered from two different perspectives:
  - **infrastructure**: infrastructure mechanisms to discover a service.
  - **service**: the service itself has to provide appropriate information to make it discoverable.
- There might be variations for services local to the infrastructure and those on remote sites - for example services in a cloud deployment versus those running on company infrastructure.

1.2 Respondents shall describe in detail how the dynamic discovery of a service is realized in their proposed technology.

1.3 Respondents shall describe in detail how a dynamic allocation of services and service instances is achieved in their proposed technology.

1.4 Respondents shall describe their concept and mechanism for the dynamic and safe execution of multiple service instances at run time.

- For example, within a sand-boxed environment and/or through the use of resource pools.

1.5 Respondents shall describe in detail the information required to enable discovery,

identification, and execution of their service.

Respondents may wish to consider the following aspects in their responses:

1 Identification of available service registries and their hierarchy (e.g. master registries with regard to their slave counterparts);

2 Searching registries for services through the definition of service parameters and constraining criteria on these parameters;

3 defining and constructing new service registries;

4 Identification and description of available service patterns, i.e. groups of services intended to solve a specific piece of a standard business process (e.g. 'Ingest');

5 Definition and construction of new service patterns;

6 Access to SLA descriptions of discovered services/patterns;

7 Making cross-registry searches.

Service Lifecycle Management

2.1 Respondents shall indicate whether their response addresses Service Lifecycle Management.

2.2 Respondents shall describe how services are set up, registered, and made available or unavailable.

2.3 Respondents shall describe how a new service can replace an existing one following an upgrade or a change of service provider.

Service Instance Management

3.1 Respondents shall say if their response addresses service instance management.

3.2 Respondents shall describe how service instances can be paused, resumed, aborted, and reprioritised.

- The agent managing the service could be a rules engine, automated monitoring process or an operator.

3.3 Respondents shall describe generic mechanisms provided in their proposed technology to allow an operator to intervene in a running service.

- Long-running media services may require intervention by an operator to change their behaviour. For example, before a render operation completes, it may be necessary to change the resource allocated for processing that render without changing the result delivered by that service.

Resource-bound services

4.1 Respondents shall indicate whether their response covers resource-bound services.

- A resource-bound service depends on a particular resource such as a specialised processing device or a video input/output device.

4.2 Respondents shall describe how their proposal allows for reservation of resources for resource-bound services.

4.3 Respondents shall describe how deadlocks are avoided or resolved.

- The way in which a resource-bound service resolves deadlocks may be different to that in which those of a dynamic service are resolved.

4.4 Respondents shall describe how the use of hardware resources is optimised.

Dynamic services

5.1 Respondents shall indicate whether their response covers dynamic services.

- A dynamic service is not constrained to a particular resource and so it can be load-balanced or virtualised.

5.2 Respondents shall indicate how a dynamic service is allocated resources and instantiated within the proposed technology.

- For example, a common approach is to use a pool of resources, such as a thread pool, to efficiently allocate resource to a new service instance without significant overhead.

5.3 Respondents shall indicate how a dynamic service is finalised and resource is made available for other services within the proposed technology.

5.4 Respondents shall say how services may be load-balanced or virtualised.

Service Behaviour

6.1 Respondents shall indicate whether their response addresses service behaviour.

6.2 Respondents shall describe in detail how their framework supports the following:

6.2.1 Protocols and patterns for synchronous and asynchronous communication;

6.2.2 Status request and reporting;

6.2.3 Error and exception handling.

6.3 Respondents shall describe, for each of the categories listed above, provisions made in their proposed technology for working with large media files.

6.4 Respondents shall say how services are able to recover in the event of an unexpected failure in terms of both software errors/exceptions and system conditions/outages.

Service Level Agreement (SLA) Management

7.1 Respondents shall indicate whether their response addresses technical enablement of SLA management.

- Within a system, it may be convenient (e.g. for reasons of economy) to deploy different services that provide the same basic capability with different levels of performance. The expected performance characteristics should then be discovered in a SLA message.

7.2 Respondents shall indicate whether SLA provisioning and management tools are explicitly included in the proposed technology.

7.3 Respondents shall describe how performance characteristics shall be declared in SLA messages through a language or schema specification.

- It is expected that any such description should include:
  - A general dictionary, e.g. including general concepts like execution time, availability, storage capacity;
  - A media-specific dictionary, e.g. including transfer bandwidth, content-related error rates, managed A/V formats.
- A typical language should be expressive enough to describe functional levels of the service from the service's behaviour point of view (e.g., load curves varying with key functional parameters. For example, data read delay with respect to the average number of concurrent read requests for a storage service.

7.4 Respondents shall provide a definition of SLA metrics measured and reported by the proposed technology.

- Metrics information should be complemented with the indication of available sampling

frequencies.

- SLA specifications should be applicable to both atomic services and service patterns. When applicable, service pattern SLA should be derived by an automatic combination of SLAs of their atomic components

7.5 Respondents shall describe provision for the negotiation of SLA parameters between services to allow the selection of the most appropriate one to call.

- For example, a user or automated mediation service may want to call a service by balancing its distance in terms of network topology against its current performance.

7.6 Respondents shall describe SLA reporting tools and the way in which SLA infringements discovered by such tools are subsequently managed.

Large media files

8.1 Respondents shall indicate whether their response addresses large media files.

- Unlike many domains where SOA is used the professional media industry regularly works with large numbers of multi-gigabyte video & audio files.

8.2 Respondents shall indicate the architectural approach taken to handling large media files. Please indicate any maximum file sizes supported or other sizing constraints of your proposed technology.

- For example, a pure video codec may have a file size or duration limit. This may impact on associated audio files.

8.3 Respondents shall describe any aspects network, storage, or other resource allocation that is employed in order to accommodate very large file sizes.

Media Formats

9.1 Respondents shall indicate whether their response covers their approach to handling different media formats.

- The professional media industry uses a number of essence and wrapper formats, some of which are not so widely adopted in other sectors.

9.2 Respondents shall describe their approach to handling different media essence formats.

- For example, media essence formats include uncompressed, MPEG-2, H.264, MP3, AAC etc. This question is more about describing frameworks and strategies for managing different essence formats within the proposed technology than which essence frameworks are currently supported.

9.3 Respondents shall describe their approach to handling different media wrapper formats.

- For example, media wrapper formats include MXF and QuickTime™. This question is more about describing frameworks and strategies for making effective use of different essence wrappers within the proposed technology than which wrapper formats are currently supported.

9.4 Respondents shall describe their approach to adopting new formats.

- New cameras and/or video compression technologies are often introduced to the market. Please describe how your framework can be extended to adopt these formats.

9.5 Respondents shall say how their services can discover and adapt to media format of the content they are acting upon.

- A service should be able to provide the same business function for more than one type of applicable media format. For example, adding a graphic to every frame of video in a

content item should succeed whether the input video is uncompressed or compressed according to the MPEG-2 standards.

9.6 Respondents shall say if their proposed technology provides mechanisms for managing generational loss.

- Wherever possible, operations on media should be carried out on the original or an early generation of content.

Renditions and versions

10.1 Respondents shall indicate whether they are describing a means to manage more than one rendition and/or version of the same content.

10.2 Respondents shall describe how they identify and manage collections of renditions of the same content item.

- A rendition of a content item represents the same content, normally of the same length, but using a different resolution and bit rate of codec. It is common practice in media workflows to have more than one rendition available, for example.
  - A master format for editing;
  - A browse or proxy copy for desktop review;
  - A highly-compressed, reduced resolution version for streaming to a mobile device.

10.3 Respondents shall describe mechanisms within their proposed technology for selecting the most appropriate rendition of a content item for a client and/or in a given context.

10.4 Respondents shall describe how they identify and manage collections of versions of the same content item.

- A version of a content item is editorially different from another version in some way, often being of a different length and created for a specific delivery requirement. Different versions of a content item include:
  - A version with all swear words bleeped out;
  - A pre-watershed and post-watershed version;
  - A different language version with a dubbed audio track.

10.5 Respondents shall describe mechanisms within their proposed technology for selecting the most appropriate version of a content item for a client and/or in a given context.

10.6 Respondents shall describe any related mechanism that ensures that the wrong version is not displayed within a particular context.

- For example, a version of material with swearing cannot be played out in the context of a children's channel.

Intra-framework Media Transfer

11.1 Respondents shall indicate whether their response addresses intra-framework media transfer.

- In this RFT, *media transfer* means copying, moving, or accessing media files over local and/or wide area IP networks. Responses in this section should refer to mechanisms for media movement *provided within a service framework* to enable the federated services requiring access to media content items. This section does not ask for a description of services for media transfer in the context of business-to-business transactions or publication, although the technologies and approach may be similar.

11.2 Respondents shall describe their intra-framework architectural approach to transferring media files. The response should include details of the transfer

mechanisms that are supported, for example transfer protocols such as FTP/HTTP and shared file system such as CIFS.

11.3 Respondents shall describe how transfers can be verified.

11.4 Respondents shall describe their approach to transfer security, including information about authentication, encryption and transfer through firewalls as well as working between trust domains.

11.5 Respondents shall describe how their response adapts to different network speeds and network conditions (e.g. latency, congestion, disruption). Please describe any network acceleration techniques that are used.

11.6 Respondents shall describe how to signal transfer times, rates, network conditions, and in the case where SLA estimates were made and later, network conditions deteriorate, how this change is communicated within the system.

11.7 Respondents shall describe how slower than real-time, real-time, and faster than real-time file transfers are supported.

11.8 Respondents shall describe how transfers can be scheduled and prioritised.

11.9 Respondents shall describe how transfers are performed when the files are on removable media.

Partial Content Item Handling

12.1 Respondents shall indicate whether their response addresses partial content item handling.

- Many media services only need to be applied to part of a content item, for example to a single frame, a short sequence of frames or to only the audio tracks and not the video track. To apply such an operation to an entire content item represented by many gigabytes of data is inefficient. Services may need to be parameterized to operate on content items specified with an identifier and additional scoping details.

12.2 Respondents shall describe in detail how services can be directed to operate on only part of a given content item specified by in points and out points.

- For example, in points and out points may be specified by a pair of timecode values, timecode and duration, frame numbers or elapsed time.

12.3 Respondents shall describe how services can be directed to operate on a subset of tracks in a multi-track audio/video content item.

Metadata management

13.1 Respondents shall say if their response addresses metadata management.

- Media content items are made up from a combination of raw essence streams/files and data about that essence, known as metadata. Metadata can be of many different kinds, including technical, structural and descriptive.
- The relationship between essence and metadata will often be managed by a media asset management system. Responses to this section should not describe a media asset management system but rather the characteristics of the service interfaces to such a system in the context of metadata management.

13.2 Respondents shall describe in detail how content is uniquely and canonically identified.

13.3 Respondents shall describe how their proposed technology represents and handles structural metadata, including representations of the provenance of the content, with reference to any standards used.

- Structural metadata links the constituent parts of content together, such as linking and synchronising audio tracks to their associated video tracks. It should also describe what the content was made from and may describe how it was derived. For example, a video track may have been made using a video pull down operation from 24fps to 30fps.

13.4 Respondents shall describe how their proposed technology represents and handles technical metadata, with reference to any standards used.

- Technical metadata describes for format of the content, including parameters that describe whether a device can or should play the content. For example, this should include codec and resolution for video content and may include the outcome of technical review processes.

13.5 Respondents shall describe how their proposed technology represents and handles descriptive metadata, with reference to any standards used.

- Descriptive metadata represents information about the content in its business context and includes description of all or part of the content. For example, the title, episode and series of a programme that the content represents, and/or logging information describing what appears in a set of frames.

13.6 Respondents shall describe the lifecycle of metadata in relation to the essence that it describes, including the creation, update, reference, removal and filtering of that data.

- The lifecycle of essence and metadata is often related but not exactly the same. The essence itself is often immutable (unchanging) but its metadata is mutable. For example, changing metadata to fix spelling mistakes or to enhance its description.
- It may not be appropriate to include business sensitive information on business-to-business transactions, so filtering operations are often appropriate.

13.7 Respondents shall provide details of how essence can be found based on its metadata and how updates to the metadata can improve the ability to find the related essence.

- One of the primary reasons for adding metadata to essence is to make items of content that can be found and become revenue-generating business assets.

13.8 Respondents shall describe the relationship, if any, between metadata contained in the wrappers of audio/visual files and the metadata stored elsewhere in the system. The description should indicate difference in approaches between structural, technical and descriptive metadata.

- Common approaches in different system architectures may represent metadata wrapped in the same file as the related essence, for example, in a sidecar XML file related by filename in database tables. In a framework, the approach may have a significant impact on the design of interoperable services.

13.9 Respondents shall describe any difference of approach in the way in which the relationship between metadata and essence is handled within a system. The description should indicate difference in approaches between structural, technical and descriptive metadata.

- Of the common approaches (wrapped, sidecar, database), a different strategy may be used internally to that used for interoperability between services.

13.10 Respondents shall say if the operations of their services on content preserve existing metadata, including metadata that is dark to that service.

- It is important that loosely-coupled services do not have a side effect on the content they act upon. Loosely-coupled services should provide flexibility to reconfigure workflows without, for example, the need to add additional steps to repair metadata after the execution of a service.

13.11 Respondents shall say if the operations of their services on content preserves, extracts, and/or enhances genealogy/provenance metadata.

- Preservation of genealogy/provenance metadata enables efficient, assisted-by-automation enterprise reporting activities, such as music reporting and contributor identification.

Event Management and Event Handling

14.1 Respondents shall indicate if their response addresses event management and event handling.

- Events notify services, workflow engines, rules engines and people when a change of interest to them has occurred in a system. For media management, events are important because operations are often long-lived, especially in comparison to a banking or similar transaction, and users need to multi-task to make efficient use of their time.

14.2 Respondents shall describe the representation(s) their proposed technology uses for events, for example event messages. Please describe all the different kinds of events your technology supports and an overview of event structures of any event header and payload.

- Ideally, services within the same framework should share a common event framework.

14.3 Respondents shall say how their proposed technology routes events from the source to the interested targets.

14.4 Respondents shall indicate how an interested party registers to receive events and any contextual filtering mechanisms.

- Filtering mechanisms, such as message queue subscription, ensure an interested party receives mainly the events they are interested in and not too many more. Chatty network broadcast messages can significantly degrade network and processing performance.

14.5 Respondents shall describe any reliable event delivery mechanisms their proposed technology includes.

- It is desirable that some events get delivered, even if the subscribing system is not available when the event is generated.

14.6 Respondents shall indicate if their technology has a means to support real time do-it-now events, how this is achieved and any event prioritisation mechanisms available.

- For example, in live media operations, pressing a button in a gallery or production control room is expected to have an immediate effect on the next frame of video produced. All such events need to prioritised and delivered in a timely, frame-accurate way.

Federated Service Coordination

15.1 Respondents shall say if their response addresses federated service coordination.

- A potential benefit of a well structured SOA-framework is the ability to flexibly deploy, orchestrate and scale services strategically and tactically over distributed, possibly outsourced, infrastructures. Well designed services should work effectively from a standalone laptop right up to geographically-distributed data centres. This section requests information about the context in which services are coordinated within your proposed technology. This context may impact the design of those services.

15.2 Respondents shall describe how the messaging between services is coordinated in federated and distributed infrastructures, for example through the use of messaging

controlled by an enterprise service bus.

15.3 Respondents shall describe how service calls are routed, referencing any standards used.

15.4 Respondents shall describe how the integrity of transactions is maintained across distributed systems, referencing any standards used.

15.5 Respondents shall provide details of how the architecture scales, including any known limitations. Include details of how is application state managed between a server and a client (session IDs etc.) and what standards are used.

15.6 Respondents shall describe any distinction between service messaging for business data and media data.

- This is sometimes distinguished as an Enterprise Service Bus versus an Enterprise Media Bus.

15.7 Respondents shall describe how their proposed architecture supports loose-coupling of services, so that a service remains agnostic to the client that calls it.

15.8 Respondents shall indicate if their proposed technology provides support for message transformation and, if so, how this is achieved and what standards are used.

15.9 Respondents shall indicate if their proposed technology requires and/or provides mechanisms for the persistence of and reliable delivery of messages

Delayed Input (offline media)

16.1 Respondents shall indicate whether their response addresses delayed input.

- Delayed input of content items may be required when invocation of a service instance before part or all of its input essence or metadata are available, such that the service needs to wait for the missing input(s). For example, although some placeholder metadata about the media may be available, essence files may need to be retrieved from near-line storage.

16.2 Respondents shall describe service instance behaviour when part or all of its inputs are not available at invocation time.

16.3 Respondents shall indicate what information is provided to the client when input is delayed.

16.4 Respondents shall describe how a service instance can use information about when inputs will be available to reserve processing or other resources.

16.5 Respondents shall indicate whether a service instance can start working with the inputs that it does have, while it waits for ones that are missing.

Media Stream Management

17.1 Respondents shall indicate whether their response covers media stream management.

- The deadlines encountered in the professional media industry often require processing of media that is still being created or data that is being streamed between facilities.

17.2 Respondents shall describe how services operate on media files that are still being written, and whose final size is known.

17.3 Respondents shall describe how services operate on 'open-ended' media files whose final size is not known until it is closed.

17.4 Respondents shall describe how services operate on media streams, including information on how the stream is identified and connected to.

- Media inputs might take the form of streams typically provided over an IP network.

Although often provided at real time, streams may also be delivered at faster or slower than real time.

17.5 Respondents shall describe the approach used in their proposed technology to adapt to new types of streaming technology.

Real Time Interaction

18.1 Respondents shall indicate whether their response covers real time interaction.

- Live events represent an important aspect of media production, especially in broadcasting. An important consideration is the capability of services to react in a timely way, and support for real-time and synchronised operation.

18.2 Respondents shall describe their approach to enabling services to operate on media at the video frame frequency (typically between 24 and 60 fps, depending on the video standard and location).

18.3 Respondents shall describe their approach to minimising messaging latencies (or equivalent) in their proposal.

18.4 Respondents shall provide an indication of the maximum response time of a service.

18.5 Respondents shall describe their approach to scheduled triggering, i.e. where a service instance starts its operation at an absolute or relative particular time.

18.6 Respondents shall describe their approach to live triggering, i.e. where a service instance is started, but waits for a trigger (e.g. from an operator or an external system such as a VTR) to execute its operation.

Deadlines

19.1 Respondents shall indicate whether their response covers deadlines management.

- For some applications, the output(s) of a service is useful only if delivered before a given deadline.

19.2 Respondents shall describe how deadlines are specified and handled.

19.3 Respondents shall say what happens to a service instance and its resources when a deadline is likely to be missed.

Interaction with legacy hardware

20.1 Respondents shall indicate if their response covers interaction with legacy hardware.

- Legacy hardware means professional media equipment, typically optimised for real time performance and equipped with specific hardware interfaces (e.g. HD-SDI, AES/EBU), or proprietary IP-based interfaces.

20.2 Respondents shall describe their approach to integrating legacy equipment with hardware interfaces.

20.3 Respondents shall describe their approach to integrating legacy equipment with proprietary IP-based interfaces.

Ownership and access control

21.1 Respondents shall indicate whether their response addresses ownership and access control.

21.2 Respondents shall indicate whether their response supports content ownership based on:

21.2.1 Individuals;

21.2.2 Job roles (e.g. editor, producer, media manager);

21.2.3 Projects (e.g. that might correspond to a programme or a production);

21.2.4 Business units/organisation (e.g. a news department);

21.2.5 Groups of the above (e.g. a group of related productions might need access to the same shared content in a local library).

21.3 Respondents shall describe how the ownership model is managed, for example by defining new roles, associating individuals with roles and projects, and creating groups.

21.4 Respondents shall describe how users are authenticated.

21.5 Respondents shall describe how content owners can control read/write/delete access to that content (for example a project might decide to allow a particular organisation to have read-only access to some of its content).

21.6 Respondents shall describe how access control can be based on content's metadata, for example access might only be granted to content tagged as non-sensitive, or only on low resolution versions.

21.7 Respondents shall describe how access to manage media services is handled. For example, whether it is possible for a particular job role to stop or reprioritise service instances.

21.8 Respondents shall describe how their proposed technology deals with multiple simultaneous accesses to the same work product.

- It is often the case that multiple services, owned by either the same user or different users, need to act on the same essence at the same time.

Data encryption

22.1 Respondents shall indicate whether their response addresses data encryption.

- Data encryption is an important tool in securing information in shared infrastructures:
  - Many programmes have an embargo whereby the leak of a programme's content before a scheduled time would be either embarrassing or have legal implications
  - Archived content may be legally sensitive, such as an interview with a member of a terrorist organisation.
  - Content items require protection to prevent mass copying of master copies into file sharing networks to maximise returns through authorized resellers.

22.2 Respondents shall describe if they address data encryption in their communication infrastructure.

22.3 Respondents shall describe which data encryption mechanisms they provide securing data for the communication between services in the infrastructure.

22.4 Respondents shall describe their concepts and mechanisms to secure essence and their metadata using encryption technologies.

Network traffic control

23.1 Respondents shall say if their response addresses network traffic control.

- Management of bandwidth should take into account the quality-of-service requirements of clients, to support working towards urgent deadlines, etc.

23.2 Respondents shall say if their proposed technology provides a means to make service clients aware of the quality-of-service that they should expect based on current network topology and traffic.

23.3 Respondents shall indicate if they provide reservation of appropriate network bandwidth as part of a service and, if so, how this achieved.

- Network bandwidth reservation and management may be achieved through SLA services.

23.4 Respondents shall describe how their proposed technology monitors network traffic in the context of broader business service provision. This should include description of facilities to dynamically adjust the individual usage of services.

- Note that a description of a rules or policy engine is not required here. Rather, please describe service interfaces provided by the technology that enable such an engine, or a human operator, to make an adjustment in network traffic.

Content access times

24.1 Respondents shall say if their response addresses content access times.

- The time to access items of content will vary according to the storage and network infrastructure. The acceptable time to access items of content will vary according to the usage. For example, it may be acceptable to wait tens of minutes for high quality versions of archived content, while a recently ingested local content item will require near-immediate access by the local production team.

24.2 Respondents shall say how their proposed technology handles access times to content in terms of the service or user that is requesting and working with it.

- You may also wish to indicate if your proposed technology has a means to provide an alternative or less appropriate rendition until the requested resolution is available, allowing a user to continue to carry some aspects of their role.

24.3 Respondents shall say how the quality-of-service a user should expect given content access time information can be communicated to that user.

- For example, the framework should be able to work with different tiers of storage, including digital video tape, data tape and file server storage. Similarly, the content may be at a different site and require a network file transfer. The user of a service should be provided with information about the expected wait expected due to accessing content. See also the section on delayed input.
- Note that in contrast to network transfer, it will not often be possible to *speed up* storage-oriented aspects of the time taken to access content.

Monitoring

25.1 Respondents shall say if their response addresses monitoring.

- Monitoring here refers to both how a service operates within the context of business activity monitoring systems and the monitoring of resource usage.

25.2 Respondents shall say whether their proposed technology provides a means to produce timely reports in a dashboard or similar monitoring toolset.

- For example, appropriately authorised users should be able to view:
  - Amount of content that is being used in total, in terms of number of items, items size, by whom, in what projects, etc.
  - Services usage, in terms of number of invocations, resource consumption etc.
  - Current and historical information.

25.3 Respondents shall indicate how their proposed technology reports information about available capacity.

- It should be possible to view the capacity of individual processing resources, content

servers, etc.

25.4 Respondents shall say if and how monitoring information is provided in a technical format suitable for automatic parsing and log analysis.

- Monitoring information may be fed into rules engines or other software services to dynamically configure a system, or to provide an alarm to an operator that action is required. Ideally, a framework will provide a common data structure for such information independently from the service called or resource reported on.

Resource provisioning and management

26.1 Respondents shall say if their response addresses resource provisioning and management.

- Ideally, a framework should automatically reallocate service reservations to best use the current resources and provision resources appropriately quickly.

26.2 Respondents shall describe how appropriately authorised users can add and remove storage without significant effect on users or in-flight services.

26.3 Respondents shall describe how appropriately authorised users can alter network capacity without significant effect on users or in-flight services.

26.4 Respondents shall describe how appropriately authorised users can alter processing capacity without significant effect on users or in-flight services.

26.5 Respondents shall say how their proposed technology can ensure the availability and responsiveness of a service that is required to provide an interactive or immediate response.

Costs and quotas

27.1 Respondents shall indicate whether their response covers quotas and costs.

- Services may be provided by an out-sourced provider, in a multi-tenanted data centre and/or in a remote *cloud computing* environment. In this case, aspects of the invocation of a service may be required to be reported to an accounting system or system that manages fair use of a shared resource through quotas.

27.2 Respondents shall describe how services provide information about the cost of an invocation, for example based on bandwidth transferred or processor usage.

27.3 Respondents shall describe how they implement quotas on the use of services, and how such quotas are managed by authorised users.

Example Service Characterization

This section is intended to validate the framework-specific requirements in the context of three example services.

28.1 Respondents shall provide characterizations of each service in the context of their proposed framework. Responses shall be provided using the associated spreadsheet [FIMS_Respondent_Spreadsheet.xls] that contains the template described below.

Interactive Capture

This service creates one or more related content items from video and audio streams and metadata. This service often involves a human operator to start and stop operation.

- As a guide to the nature of this service, see the AV Content Capture Use Case.

Interactive Capture Template

The following template provides a set of questions that should be answered by the Respondent about the service in the context of their proposed technology. This template is provided as a worksheet in the associated spreadsheet [FIMS_Respondent_Spreadsheet.xls]. For more details about the topics, see the relevant subsection of the framework for media industry services section.

IC_1 Respondents shall describe the basic capability of the service.

IC_2 Respondents shall describe the core properties and parameters of the service.

- A core property would be an input or an output that is essential to a service's operation.

IC_3 Respondents shall describe the extended properties and parameters of the service.

- An extended property or parameter provides access to additional extended capabilities of a service, such as applying a colour filter.

IC_4 Respondents shall describe how the service itself is defined and described, through schema, data structure or other concept.

- As well as the description, the Respondent may attach an example, such as a WSDL file, XML schema or JSON data structure.

IC_5 Respondents shall describe any resources consumed by the service.

IC_6 Respondents shall describe how the service is discovered.

IC_7 Respondents shall describe how the service is set up, registered, made available and torn down.

IC_8 Respondents shall describe how a new service can replace an existing one.

IC_9 Respondents shall describe how a service instance is paused, resumed, aborted and reprioritized.

IC_10 Respondents shall describe how resources for the service are reserved, deadlocks avoided, and resource usage optimised.

IC_11 Respondents shall describe how the service may be load-balanced or virtualised.

IC_12 Respondents shall describe the protocols and patterns used by the service.

IC_13 Respondents shall describe how the service performs status request and reporting.

IC_14 Respondents shall describe how errors and exceptions are handled.

IC_15 Respondents shall describe how performance characteristics are specified in terms of SLAs.

IC_16 Respondents shall describe any limit on file sizes for the service.

IC_17 Respondents shall describe the approach to handling different formats in the context of this service.

IC_18 Respondents shall describe how the service works with multiple renditions of the same content item.

IC_19 Respondents shall describe how the service works with multiple versions of the same content item.

IC_20 Respondents shall describe how intra-framework media transfer into and out of the service is handled.

IC_21 Respondents shall describe how the service can operate on partial content items, between a specified set of in- and out-points.

IC_22 Respondents shall describe how the service can operate on a selection of tracks.

IC_23 Respondents shall describe how metadata is handled in the context of this service.

IC_24 Respondents shall describe how this service operates in a federated/distributed infrastructure.

IC_25 Respondents shall describe any events responded to, handled by or generated by this service.

IC_25 Respondents shall describe how the service behaves when its inputs are not available at invocation time.

IC_27 Respondents shall describe how the service supports operation on files that are still being written.

IC_28 Respondents shall describe how the service supports streaming.

IC_29 Respondents shall describe any real-time aspects of the service.

IC_30 Respondents shall describe how the service executes so as to meet deadlines.

IC_31 Respondents shall describe any interaction of the service with legacy hardware.

IC_32 Respondents shall describe how the service manages content ownership and service access control.

IC_33 Respondents shall describe any data encryption aspects of the service.

IC_34 Respondents shall describe how the service interacts with network traffic control.

IC_35 Respondents shall describe how the service provides context-dependent access times appropriate to its use.

IC_36 Respondents shall describe how the activity of the service is monitored.

IC_37 Respondents shall describe how resources are allocated to the service and may be altered in-flight.

IC_38 Respondents shall describe any cost or quota mechanisms associated with the service.

Transform Media

This service transforms one or more related content items into one or more new content items as part of a particular publishing task.

- As a guide to the nature of this service, see the Content Publishing Use Case.

Transform Media Template

The following template provides a set of questions that should be answered by the Respondent about the service in the context of their proposed technology. This template is provided as a

worksheet in the associated spreadsheet [FIMS_Respondent_Spreadsheet.xls]. For more details about the topics, see the relevant subsection of the framework for media industry services section.

TM_1 Respondents shall describe the basic capability of the service.

TM_2 Respondents shall describe the core properties and parameters of the service.

- A core property would be an input or an output that is essential to a service's operation.

TM_3 Respondents shall describe the extended properties and parameters of the service.

- An extended property or parameter provides access to additional extended capabilities of a service, such as applying a colour filter.

TM_4 Respondents shall describe how the service itself is defined and described, through schema, data structure or other concept.

- As well as the description, the Respondent may attach an example, such as a WSDL file, XML schema or JSON data structure.

TM_5 Respondents shall describe any resources consumed by the service.

TM_6 Respondents shall describe how the service is discovered.

TM_7 Respondents shall describe how the service is set up, registered, made available and torn down.

TM_8 Respondents shall describe how a new service can replace an existing one.

TM_9 Respondents shall describe how a service instance is paused, resumed, aborted and reprioritized.

TM_10 Respondents shall describe how resources for the service are reserved, deadlocks avoided, and resource usage optimised.

TM_11 Respondents shall describe how the service may be load-balanced or virtualised.

TM_12 Respondents shall describe the protocols and patterns used by the service.

TM_13 Respondents shall describe how the service performs status request and reporting.

TM_14 Respondents shall describe how errors and exceptions are handled.

TM_15 Respondents shall describe how performance characteristics are specified in terms of SLAs.

TM_16 Respondents shall describe any limit on file sizes for the service.

TM_17 Respondents shall describe the approach to handling different formats in the context of this service.

TM_18 Respondents shall describe how the service works with multiple renditions of the same content item.

TM_19 Respondents shall describe how the service works with multiple versions of the same content item.

TM_20 Respondents shall describe how intra-framework media transfer into and out of the service is handled.

TM_21 Respondents shall describe how the service can operate on partial content items, between a specified set of in- and out-points.

TM_22 Respondents shall describe how the service can operate on a selection of tracks.

TM_23 Respondents shall describe how metadata is handled in the context of this service.

TM_24 Respondents shall describe how this service operates in a federated/distributed infrastructure.

TM_25 Respondents shall describe any events responded to, handled by or generated by this service.

TM_25 Respondents shall describe how the service behaves when its inputs are not available at invocation time.

TM_27 Respondents shall describe how the service supports operation on files that are still being written.

TM_28 Respondents shall describe how the service supports streaming.

TM_29 Respondents shall describe any real-time aspects of the service.

TM_30 Respondents shall describe how the service executes so as to meet deadlines.

TM_31 Respondents shall describe any interaction of the service with legacy hardware.

TM_32 Respondents shall describe how the service manages content ownership and service access control.

TM_33 Respondents shall describe any data encryption aspects of the service.

TM_34 Respondents shall describe how the service interacts with network traffic control.

TM_35 Respondents shall describe how the service provides context-dependent access times appropriate to its use.

TM_36 Respondents shall describe how the activity of the service is monitored.

TM_37 Respondents shall describe how resources are allocated to the service and may be altered in-flight.

TM_38 Respondents shall describe any cost or quota mechanisms associated with the service.

Transfer Media

This service moves or copies one or more related content items from one location to another taking advantage of available network transport mechanisms.

Note: It is assumed that this service is used for inter-framework and/or business-to-business transfers rather than intra-framework internal media transfers.

- As a guide to the nature of this service, see the Content Publishing Use Case.

Transfer Media Template

The following template provides a set of questions that should be answered by the Respondent about the service in the context of their proposed technology. This template is provided as a worksheet in the associated spreadsheet [FIMS_Respondent_Spreadsheet.xls]. For more details about the topics, see the relevant subsection of the framework for media industry services section.

TRM_1 Respondents shall describe the basic capability of the service.

TRM_2 Respondents shall describe the core properties and parameters of the service.

- A core property would be an input or an output that is essential to a service's operation.

TRM_3 Respondents shall describe the extended properties and parameters of the service.

- An extended property or parameter provides access to additional extended capabilities of a service, such as applying a colour filter.

TRM_4 Respondents shall describe how the service itself is defined and described, through schema, data structure or other concept.

- As well as the description, the Respondent may attach an example, such as a WSDL file, XML schema or JSON data structure.

TRM_5 Respondents shall describe any resources consumed by the service.

TRM_6 Respondents shall describe how the service is discovered.

TRM_7 Respondents shall describe how the service is set up, registered, made available and torn down.

TRM_8 Respondents shall describe how a new service can replace an existing one.

TRM_9 Respondents shall describe how a service instance is paused, resumed, aborted and reprioritized.

TRM_10 Respondents shall describe how resources for the service are reserved, deadlocks avoided, and resource usage optimised.

TRM_11 Respondents shall describe how the service may be load-balanced or virtualised.

TRM_12 Respondents shall describe the protocols and patterns used by the service.

TRM_13 Respondents shall describe how the service performs status request and reporting.

TRM_14 Respondents shall describe how errors and exceptions are handled.

TRM_15 Respondents shall describe how performance characteristics are specified in terms of SLAs.

TRM_16 Respondents shall describe any limit on file sizes for the service.

TRM_17 Respondents shall describe the approach to handling different formats in the context of this service.

TRM_18 Respondents shall describe how the service works with multiple renditions of the same content item.

TRM_19 Respondents shall describe how the service works with multiple versions of the same content item.

TRM_20 Respondents shall describe how intra-framework media transfer into and out of the service is handled.

TRM_21 Respondents shall describe how the service can operate on partial content items, between a specified set of in- and out-points.

TRM_22 Respondents shall describe how the service can operate on a selection of tracks.

TRM_23 Respondents shall describe how metadata is handled in the context of this service.

TRM_24 Respondents shall describe how this service operates in a federated/distributed infrastructure.

TRM_25 Respondents shall describe any events responded to, handled by or generated by

this service.

TRM_25 Respondents shall describe how the service behaves when its inputs are not available at invocation time.

TRM_27 Respondents shall describe how the service supports operation on files that are still being written.

TRM_28 Respondents shall describe how the service supports streaming.

TRM_29 Respondents shall describe any real-time aspects of the service.

TRM_30 Respondents shall describe how the service executes so as to meet deadlines.

TRM_31 Respondents shall describe any interaction of the service with legacy hardware.

TRM_32 Respondents shall describe how the service manages content ownership and service access control.

TRM_33 Respondents shall describe any data encryption aspects of the service.

TRM_34 Respondents shall describe how the service interacts with network traffic control.

TRM_35 Respondents shall describe how the service provides context-dependent access times appropriate to its use.

TRM_36 Respondents shall describe how the activity of the service is monitored.

TRM_37 Respondents shall describe how resources are allocated to the service and may be altered in-flight.

TRM_38 Respondents shall describe any cost or quota mechanisms associated with the service.

Media Industry Decomposition

29.1 Respondents shall provide their own vision of a decomposition of sectors, domains and services for their particular media industry sector.

The purpose of this section is to ask the Respondent to provide guidelines and advice to the FIMS group as to how the Task Force should structure its activity going forward. To do this, the Task Force asks Respondents to provide their view of the organization and categorization of the entire media industry in terms of specific industry sectors, domains and services - a *media industry decomposition*.

For example, a decomposition of the entire media industry is illustrated in the following figure. Such decomposition may identify a broadcast production sector. This sector could be further sub-divided into domains and examples are provided below. These domains will contain some domain-specific services and some shared cross-domain services. Some example broadcast production services are provided below.

Each industry sector and content genre has specific requirements that impact on the organization of work and on the type of resources used. In general, a high percentage of human workflow is needed. This also implies that as part of the process of conceiving new technological platforms and processes, the culture of users has to be taken into consideration. Users involved are not only technical staff, but also creative people and freelancers external to the organization.

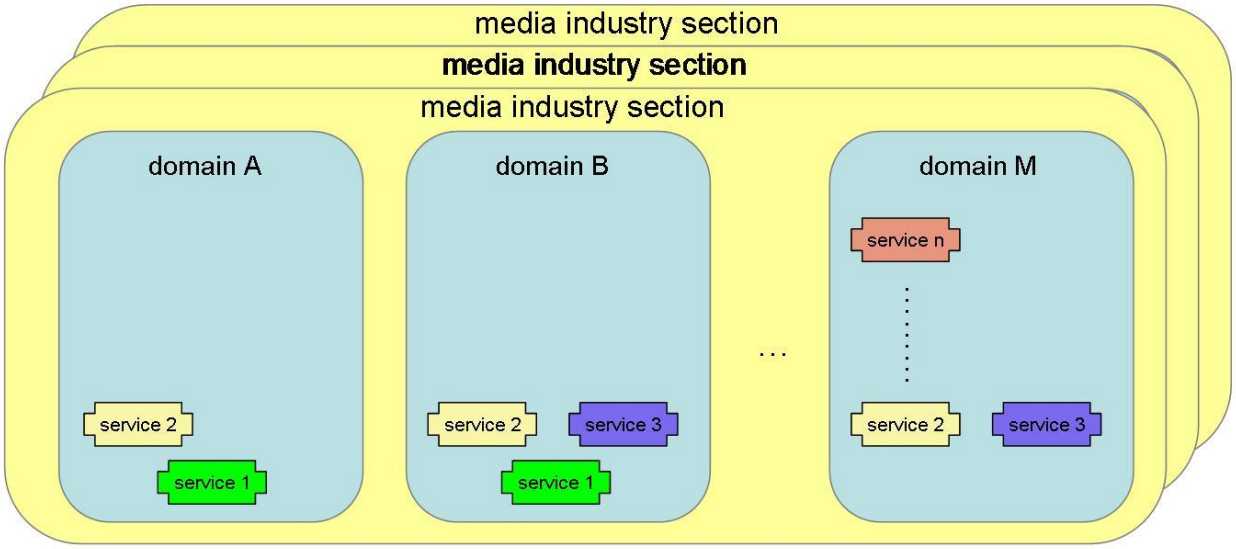

Decompositions will structure further Task Force activities in defining appropriate process models and identifying the basic service component of each domain. Single services may be part of multiple domains across the enterprise. A decomposition into domains with a breakdown of services will help us to decide which service definitions should be addressed in future RFTs.

Note that the terms *sector*, *domain* and *service* are considered the most appropriate by the Task Force at this time. However, you are not required to structure your response around these terms and an alternative approach to decomposition is acceptable and will be considered. For example, a decomposition could be based on *asset type*, *process* and *capability*.

Example Broadcast Production Domains

Some examples of domains in specific broadcast production environments are outlined in the following.

News is a case of production with a number of specific requirements that have a direct impact on the technology. High availability of the system is a top priority both because every day a number of news stories have to be produced and updated, and also because the rapidity with which breaking news is broadcast is a distinctive sign of a good news channel. Therefore, the workload of the system is typically bursty with high peaks close to broadcasting and, in the case of unexpected major events, the peaks of usage can greatly exceed average values.

User-perceived quality of service is very much related to the behaviour of the system in such situations. Workflows are generally stable, but the system must provide flexibility to rapidly overcome any sudden problem that risks undermining broadcasting capability. Media quality is generally less of an issue compared to production speed. Stories are assembled using purpose-shot content, content acquired from international news feeds or freelancers, retrieved from the archive and, occasionally, user-generated (consumer cameras, smart phones, downloaded from internet). Live contributions are sometimes inserted during broadcast. The abilities to edit stories in the field and to download content files to the broadcast station are increasingly requested features. Editing is often limited to splicing of video and voice-over insertion, while graphics and animations are being used more frequently, as well as virtual sets.

Sports production is another specific case that presents a number of similarities with news concerning the production of reports and highlights, but it diverges due to the predominance of live events. Signal quality is generally quite important and major events are nowadays covered in high definition only. System reliability during events is of course mandatory, but unlike news most of the activities are planned well in advance. Production consists mainly of cut edits, but special effects

(slow motion, 3D rendering, and animations) and graphics are used extensively. The ability to edit at the same time as acquiring is very important in order to be able to produce highlights whilst the event is still ongoing.

Serials production has totally different technical organization as each production tends to be a completely self-contained island, especially concerning the shooting phase. All the activities and resources are planned in advance and all the efforts are generally devoted to achieve the highest possible quality within the given budget. A small amount of system down time can be tolerable. Flexibility in the workflows is a necessity, as these productions are based largely around high quality of craftsmanship.

Example Broadcast Production Services

Some additional service examples related to broadcasting are given below, but these are for illustrative purposes only.

A reference broadcast production architecture usually includes four interconnected modules that are described briefly below:

Ingest: new media content is acquired from various sources including live feeds, cameras or media, and conformed to internal working formats. Identification and indexing metadata are also added.

Edit: ingested content, possibly complemented by archive clips, is edited into finished programmes or contributions to be used in live contexts. Audio and video are often processed on different platforms.

Play-out: programmes are organized into a rundown according to organised editorial and marketing concepts, and broadcast. Channel graphics and commercials are added.

Archive: Long-term preservation and access to content is provided to production and to external parties. Descriptive metadata are added.

Live broadcast can be assimilated to the same schema, but most of the production operations have to be carried out in real time, generally using dedicated hardware.

Aside from the main content production flow, many operations not directly involving media are performed; these include commissioning of programmes, planning of resources, run down scheduling and monitoring of processes and outcome.

In some situations, part of the work needs to be outsourced to specialized companies. In this case, security, monitoring and SLA management aspects assume a key role. With the availability of commercial SaaS offerings, it is foreseeable that this way of working will increase in future, especially concerning resource intensive automated processes such as rendering and transcoding.