Project Number:	**FP7-257123**
Project Title:	**CONVERGENCE**
Deliverable Type:	**Report**
Dissemination Level	**Public**
Deliverable Number:	**D4.2**
Contractual Date of Delivery to the CEC:	**30.09.2011**
Actual Date of Delivery to the CEC:	**12.11.2011**
Title of Deliverable:	**CONVERGENCE Rights Expression Language**
Workpackage contributing to the Deliverable:	**WP 4**
Nature of the Deliverable:	**Report**
Editor:	**Giuseppe Tropea, Richard Walker**
Authors:	**Nicola Blefari Melazzi, Giuseppe Tropea, Stefano Salsano, Giuseppe Bianchi (CNIT), Maria Teresa Andrade, Helder Castro (INESC), Alina Hang (LMU), Mihai Tanase (UTI), Daniel Sequieria (WIPRO), Andrea de Polo (ALINARI), Thomas Huebner (MORPHO), Francis Lemaitre (FMSH), Angelo Difino (CEDEO), Angelos-Christos Anadiotis, Aziz Mousas, Dimitra Kaklamani (ICCS), Panagiotis Gkonis (SIL)**
Abstract:	**This deliverable describes the MPEG-21 REL in the context of Convergence's VDI-based publish/subscribe approach to information distribution and discusses issues related to the implementation of the system's security features.**
Keyword List:	**MPEG-21, Licenses, Rights Expression, Publish/Subscribe, Certified Content.**

Executive Summary

In the digital media value chain, Rights Expression Languages (RELs) are used to enable controlled access to digital resources, addressing several different issues from the description of licenses to access and usage control, payments, etc. A REL is an essential component of any security infrastructure supporting differentiated controlled access to digital resources, and providing adequate protection of intellectual property rights.

There are relatively few existing RELs capable of supporting CONVERGENCE's needs. Among these the project has selected MPEG-21 part 5 [1]. This open standard, which can be implemented in XML, is one of the main current contenders for a general-purpose REL.

Using a REL in a given application or use scenario requires the identification of the REL elements necessary to describe the different types of licenses and access control mechanisms required by the application. In this report, therefore, we identify requirements on the REL from the four CONVERGENCE use scenario, in particular the need for custom REL verbs to describe rights to govern the system's publish/subscribe functionality. We go on to provide a full description of these new verbs.

The report describes CONVERGENCE's governance and licensing scheme, based on the MPEG-21 part 5 standard and on the specific content protection and rights management requirements, identified in the CONVERGENCE use scenarios. The scheme is designed in the light of CONVERGENCE's ability to distribute and manage any kind of digital resource in a large distributed environment.

Finally, the report explains how REL data is embedded into the CONVERGENCE data unit, the Versatile Digital Item (VDI) and introduces a basic set of security features, based on digital certificates, for the enforcement of the rights and conditions expressed in CONVERGENCE licenses.

INDEX

Glossary

Term	Definition
Access Rights	Criteria defining who can access a VDI or its components under what conditions.
Advertise	Procedure used by a CoNet user to make a resource accessible to other CoNet users.
Application	Software, designed for a specific purpose that exploits the capabilities of the CONVERGENCE System.
Business Scenario	A scenario describing a way in which the CONVERGENCE System may be used by specific users in a specific context or, more narrowly, a scenario describing the products and services bought and sold, the actors concerned and, possibly, the associated flows of revenue in such a context.
Clean-slate architecture	The CONVERGENCE implementation of the Network Level, totally replacing existing IP functionality. See “Integration Architecture” and ”“Overlay Architecture” and “Parallel Architecture”.
CoApp	The CONVERGENCE Application Level.
CoApp Provider	A user providing Applications running on the CONVERGENCE Middleware Level (CoMid).
CoMid	The CONVERGENCE Middleware Level.
CoMid Provider	A user providing access to a single or an aggregation of CoMid services.
CoMid Resource	A virtual or physical object or service referenced by a VDI, e.g. media, Real World Objects, persons, internet services. It has the same meaning of “Resource” and it is used only to better specify the term “Resource” when there is a risk of a misunderstanding with the term “CoNet Resource”.
Community Dictionary Service (CDS)	A CoMid Technology Engine that provides all the matching concepts in a user’s subscription, search request and publication.
CoNet Provider	A user providing access to CoNet services, i.e. the equivalent of an Internet Service Provider.
CoNet Resource	A resource of the CoNet that can be identified by means of a

	name; resources may be either Named-data or a Named service access point.
Content-based resource discovery	A user request for resources, either through a subscription or a search request to the CONVERGENCE system (from literature). See “subscription” and “search”.
Content-based Subscription	A subscription based on a specification of user’s preferences or interests, (rather than a specific event or topic). The subscription is based on the actual content, which is not classified according to some predefined external criterion (e.g., topic name), but according to the properties of the content itself. See “Subscription” and “Publish-subscribe model”.
Content-centric	A network paradigm in which the network directly provides users with content, and is aware of the content it transports, (unlike networks that limit themselves to providing communication channels between hosts).
CONVERGENCE Applications level (CoApp)	The level of the CONVERGENCE architecture that establishes the interaction with CONVERGENCE users. The Applications Level interacts with the other CONVERGENCE levels on behalf of the user.
CONVERGENCE Computing Platform level (CoComp)	The Computing Platform level provides content-centric networking (CoNet), secure handling (CoSec) of resources within CONVERGENCE and computing resources of peers and nodes.
CONVERGENCE Core Ontology (CCO)	A semantic representation of the CoReST taxonomy. See “CONVERGENCE Resource Semantic Type (CoReST)”
CONVERGENCE Device	A combination of hardware and software or a software instance that allows a user to access Convergence functionalities
CONVERGENCE Engine	A collection of technologies assembled to deliver specific functionality and made available to Applications and to other Engines via an API
CONVERGENCE Middleware level (CoMid)	The level of the CONVERGENCE architecture that provides the means to handle VDIs and their components.
CONVERGENCE Network (CoNet)	The Content Centric component of the CONVERGENCE Computing Platform level. The CoNet provides access to named-resources on a public or private network infrastructure.
CONVERGENCE node	A CONVERGENCE device that implements CoNet functionality and/or CoSec functionality.

CONVERGENCE peer	A CONVERGENCE device that implements CoApp, CoMid, and CoComp (CoNet and CoSec) functionality.
CONVERGENCE Resource Semantic Type (CoReST)	A list of concepts or terms that makes it possible to categorize a resource, establishing a connection with the resource's semantic metadata.
CONVERGENCE Security element (CoSec)	A component of the CONVERGENCE Computing Platform level implementing basic security functionality such as storage of private keys, basic cryptography, etc.
CONVERGENCE System	A system consisting of a set of interconnected devices - peers and nodes - connected to each other built by using the technologies specified or adopted by the CONVERGENCE specification. See "Node" and "Peer".
Digital forgetting	A CONVERGENCE system functionality ensuring that VDIs do not remain accessible for indefinite periods of time, when this is not the intention of the user.
Digital Item (DI)	A structured digital object with a standard representation, identification and metadata. A DI consists of resource, resource and context related metadata, and structure. The structure is given by a Digital Item Declaration (DID) that links resource and metadata.
Domain ontology	An ontology, dedicated to a specific domain of knowledge or application, e.g. the W3C Time Ontology and the GeoNames ontology.
Elementary Service (ES)	The most basic service functionality offered by the CoMid.
Entity	An object, e.g. VDIs, resources, devices, events, group, licenses/contracts, services and users, that an Elementary Service can act upon or with which it can interact.
Expiry date	The last date on which a VDI is accessible by a user of the CONVERGENCE System.
Fractal	A semantically defined virtual cluster of CONVERGENCE peers.
Identifier	A unique signifier assigned to a VDI or components of a VDI.
Integration Architecture	An implementation of CoNet designed to integrate CoNet functionality in the IP protocol by means of a novel IPv4 option or by means of an IPv6 extension header, making IP content-aware. See "Clean-state Architecture", "Overlay Architecture", "Parallel Architecture"
License	A machine-readable expression of Operations that may be

	executed by a Principal.
Local named resource	A named-resource made available to CONVERGENCE users through a local device, permanently connected to the network. Users have two options to make named-resources available to other users: 1) store the resource in a device, with a permanent connection to the network; 2) use a hosting service. In the event she chooses the former option, the resource is referred to as a local named-resource.
Metadata	Data describing a resource, including but not limited to provenance, classification, expiry date etc.
MPEG eXtensible Middleware (MXM)	A standard Middleware specifying a set of Application Programming Interfaces (APIs) so that MXM Applications executing on an MXM Device can access the standard multimedia technologies contained in the Middleware as MXM Engines.
MPEG-M	An emerging ISO/IEC standard that includes the previous MXM standard.
Multi-homing	In the context of IP networks, the configuration of multiple network interfaces or IP addresses on a single computer.
Named-data	A named-resource consisting of data.
Named resource	A CoNet resource that can be identified by means of a name. Named-resources may be either data (in the following referred to as “named-data”) or service-access-points (“named-service-access-points”).
Named service access point	A kind of named-resource, consisting of a service access point identified by a name. A named-service-access-point is a network endpoint identified by its name rather than by the Internet port numbering mechanism.
Network Identifier (NID)	An identifier identifying a named resource in the CONVERGENCE Network. If the named resource is a VDI or an indentified VDI component, its NID may be derived from the Identifier (see “Identifier”).
Overlay architecture	An implementation of CoNet as an overlay over IP. See “Clean-state Architecture” and “Integration Architecture” and “Parallel Architecture”
Parallel architecture	An implementation of CoNet as a new networking layer that can be used in parallel to IP. See “Clean-state Architecture” and “Integration Architecture” and

	”“Overlay Architecture”
Policy routing	In the context of IP networks, a collection of tools for forwarding and routing data packets based on policies defined by network administrators.
Principal (Rights Expression Language)	The User to whom Permissions are Granted in a License.
Principal (CoNet)	The user who is granted the right to use a *CoNet Principal Identifier* for naming its named resources. For example, the principal could be the provider of a service, the publisher or the author of a book, the controller of a traffic lights infrastructure, or, in general, the publisher of a VDI. A Principal may have several Principal Identifiers in the CoNet.
Principal Identifier (CoNet)	The Principal identifier is a string that is used in the Network Identifiers (NID) of a CoNet resource, when the NID has the form: NID = <namespace ID, hash (Principal Identifier), hash (Label)> In this approach, hash (Principal Identifier) must be unique in the namespace ID, and Label is a string chosen by the principal in such a way that hash(Label) is unique for in the context of the Principal Identifier.
Publish	The act of informing an identified subset of users of the CONVERGENCE System that a VDI is available.
Publisher	A user of CONVERGENCE who performs the act of publishing.
Publish-subscribe model	CONVERGENCE uses a content-based approach for the publish-subscribe model, in which notifications about VDIs are delivered to a subscriber only if the metadata / content of those VDIs match constraints defined by the subscriber in his Subscription VDI.
Real World Object	A physical object that may be referenced by a VDI.
Resource	A virtual or physical object or service referenced by a VDI, e.g. media, Real World Objects, persons, internet services.
Scope (in the context of routing)	In the context of advertising and routing, the geographical or administrative domain on which a network function operates (e.g. a well defined section of the network - a campus, a shopping mall, an airport -, or to a subset of nodes that receives advertisements from a service provider).
Search	The act through which a user requests a list of VDIs meeting a set of search criteria (e.g. specific key value pairs in the metadata, key words, free text etc.).

Service Level Agreement (SLA)	An agreement between a service provider and another user or another service provider of CONVERGENCE to provide the latter with a service whose quality matches parameters defined in the agreement.
Subscribe	The act whereby a user requests notification every time another user publishes or updates a VDI that satisfies the subscription criteria defined by the former user (key value pairs in the metadata, free text, key words etc.).
Subscriber	A user of CONVERGENCE who performs the act of subscribing.
Timestamp	A machine-readable representation of a date and time.
Tool	Software providing a specific functionality that can be re-used in several applications.
Trials	Organized tests of the CONVERGENCE System in specific business scenarios.
Un-named-data	A data resource with no NID.
User	Any person or legal entity in a Value-Chain connecting (and including) Creator and End-User possibly via other Users.
User (in OSI sense)	In a layered architecture, the term is used to identify an entity exploiting the service provided by a layer (e.g. CoNet user).
User ontology	An ontology created by CONVERGENCE users when publishing or subscribing to a VDI.
User Profile	A description of the attributes and credentials of a user of the CONVERGENCE System.
Versatile Digital Item (VDI)	A structured, hierarchically organized, digital object containing one or more resources and metadata, including a declaration of the parts that make up the VDI and the links between them.

1 Goals and structure of this document

Deliverable 4.2 (Rights Expression Language, REL) is the second deliverable from WP4: Definition of the Versatile Digital Item. The general goal of WP4 is to define **the Versatile Digital Item**, extending the scope of the MPEG-21 DI by including new classes of objects, including Real World Objects, services and people and supporting new classes of operation. The VDI is the basic unit for transaction used within CONVERGENCE.

CONVERGENCE will use a publish-subscribe paradigm to provide a secure environment for the distribution and transaction of digital resources represented as VDIs. This requires the definition of a security infrastructure to enforce digital rights management. An essential component of any such infrastructure is a Rights Expression Language (REL), allowing users to specify rights to digital resources.

Among the few RELs meeting CONVERGENCE needs, the project has selected MPEG-21 part 5. This open standard, which can be implemented in XML, is a leading current contender for a general-purpose REL, satisfying the requirements of a broad range of applications and scenarios. Chapter 2 provides further justification for this choice.

In this report, we present the subset of MPEG-21 part 5 required for the implementation of the four CONVERGENCE use scenarios, together with a CONVERGENCE governance and licensing scheme, applicable to all our use scenarios and flexible enough to be easily extended to new scenarios.

To achieve this, it was necessary to analyze not only the use scenarios and their requirements, but also the mode of operation of current and foreseeable technologies, understanding how rights can and should be expressed, how they will be dealt with in the system and their impact on system operations. Accordingly, this report includes a description of: 1) the content management and protection requirements for CONVERGENCE use cases; 2) the techniques used to embed REL data in VDIs; 3) the way REL is used by the system and the implications for publication, subscription and search processes.

The rest of the report is structured as follows. Chapter 2 motivates the choice of the MPEG-21 part 5 standard. Chapter 3 describes the requirements of the use scenarios in terms of content protection and rights management, providing details of the way REL elements will be embedded in VDIs. Chapter 4 presents the CONVERGENCE governance and licensing scheme, explaining the way the scheme uses elements in MPEG-21 part 5 and including examples of hypothetical licenses matching the requirements of CONVERGENCE scenarios. Chapter 5 describes a basic set of CONVERGENCE security mechanisms for the enforcement of the rights and permissions declared in the licenses and describes the impact of these techniques on publication and matching mechanisms. The security issue will be fully dealt with in next deliverables. Here we introduce a promising approach, which CONVERGENCE is exploring, based on so-called Attribute Based Encryption (ABE). Chapter 6 concludes.

2 Introduction

The explosive expansion of the Internet and spectacular associated developments in the computer industry have revolutionized the way people distribute and access content and information related services.

Until recent past, consumption of information was rigidly bound to physical objects containing the information (CDs, VHS cassettes, newspapers, etc.). The distribution infrastructure for these information-bearing objects (brick and mortar stores, newspaper shops, vending machines, etc.) was also physical, requiring heavy investment for start-up and maintenance. This investment came from dedicated economical entities (media retailers), which therefore became inescapable intermediaries between information creators and consumers.

The Internet and associated computer technologies changed all this, enabling the "dematerialization" of information-carrying objects and the related distribution infrastructure. At the same time, the capabilities provided by household PCs eliminated costs for the reproduction of information: with PCs, digital information objects can be reproduced without any consumption of physical resources. With the rise of the Internet, it became possible to distribute information goods as purely digital objects, drastically reducing distribution costs.

For the first time, information producers and consumers could **exchange information directly, without intermediaries** and at a very low cost.

These trends have fuelled the development of new technologies to facilitate, automate and manage content flow and service access over the Internet. The result has been an explosion in the exchange of legitimate (commercial and other legal ventures) and illegitimate (content piracy) media assets, which turned the Internet into the heart of the digital **economy**. However, harnessing its full potential requires new technologies for the management of intellectual property **rights**. If CONVERGENCE is to be a platform in the future Internet, this is an important issue. CONVERGENCE's technical structure and mode of operation has to safeguard the rights of all users, including final content consumers, original content producers and intermediating entities. In other words, CONVERGENCE has to manage and enforce access and use rights for the digital contents it distributes, independently of the terminal it connects (computers, mobile phones, other equipment connected through the Internet or via other telecommunications networks).

To satisfy these requirements within the CONVERGENCE environment, appropriate entities have to formally specify and authenticate relevant user rights, and preferences must be formally specified and authenticated by appropriate entities and store this information in information objects and services. These objects/services must then be securely stored/operated and made available to all authorized requiring entities. CONVERGENCE CoMid instances will use this information to enforce users' rights.

The formal specification of information pertaining to user rights and preferences, for the manipulation of digital items, is typically expressed in licenses. There already exist a number of tools for the specification of such licences. These include ODRL 6, ccREL [3], OMA DRM [4], and MPEG-21 REL [5].

CONVERGENCE, will use the MPEG-21 REL [1]. This choice is motivated the scope of the standard, the broad range of rights-related situations it can formally and precisely describe and the additional facilities offered by the rest of the MPEG-21 standard (RDD, IPMP, etc.).

3 CONVERGENCE Governance and Licensing Needs

CONVERGENCE offers a broad spectrum of possibilities for the exchange, manipulation and consumption of digital resources. As a result, the spectrum of needs for rights protection and content management is equally vast.

An exhaustive study of these needs would go beyond the scope of this document. In this chapter we state some of the most important requirements for rights management needs and licensing, arising from the operational characteristics and capabilities of the CONVERGENCE system. We begin with constraints arising from the design of the system architecture, and move on to discuss the specific requirements of the use scenarios.

We begin by clarifying the distinction between data and metadata, the structure of the VDI itself, and the way in which REL statements will be embedded in this structure. The description will build on architecture design patterns specified in previous WP3 and WP4 deliverables. The following sections will systematically analyze licensing requirements arising from the use cases developed in WP2 and their implementation in the first phase of the trials. Later phases, when more advanced functionalities will be deployed, may give rise to new requirements. The summary description should therefore be regarded as provisional.

It is nonetheless possible to derive a number of general conclusions. These are summarized in paragraph 3.4.5, which formulates the critical distinction between **description** of rights and **enforcement** of rights. Thus Chapter 4 describes design decisions and standardization proposals to extend the expressiveness and usefulness of the MPEG-21 REL in an ecosystem of digital items interconnected by a web of semantic links. Chapter 5 describes security technologies to enforce a subset of REL statements and to control publication and subscription mechanisms.

3.1 Structuring of VDIs in the Presence of REL Statements and Relationships

The MPEG-21 REL standard applies REL statements to **resources** contained within Digital Items. In the CONVERGENCE approach, on the other hand, data about resources (**metadata**, or descriptors of data) are as important as the resources themselves. To reconcile the differences between these approaches it is necessary to specify how CONVERGENCE will embed complex REL statements in VDIs.

Details of the structure of a VDI are reported in deliverable D4.1 [5]. Essentially a VDI is conceived as follows:

ITEM
  DESCRIPTOR
  RESOURCE
ITEM

This standard clearly represents the difference between data and data-about-data and allows very complex groupings and nesting of ITEMs. It is a flexible mechanism well-adapted to the human preference for two-level **distinctions**. However, it is also dangerous, since it leaves the door open to human interpretation of what to categorize as data and what as metadata.

When the concept of Digital Item was first introduced in MPEG, this danger was limited by the basic characteristics of the standard, in which:

- DESCRIPTORs were meant to hold simple information such as "The title of the song is Interstellar Overdrive";
- DESCRIPTORs where **not** designed to be crawled by search engines;
- relationships between ITEMs were mostly specified by the grouping and nesting of physical ITEMs.

As soon as the MPEG DI standard emerged, designers exploited the freedom provided by DESCRIPTOR tags, to enrich DI behaviour. In CONVERGENCE, too, we push the role of metadata DESCRIPTORs to their limits. In CONVERGENCE:

- DESCRIPTORs hold **valuable**, structured information;
- ITEMs are preferably shallow and not nested,
- DESCRIPTORs bear the load of linking distinct items into a structured web of relationships;
- said relationships are complex and semantically tagged;
- DESCRIPTORs are searchable;
- DESCRIPTORs are expressed in machine-readable schemas (ontologies).

The CONVERGENCE approach makes the distinction between data and metadata much more fuzzy. As an example, let us consider Samsam, a big hardware manufacturer, which is about to launch a new LED TV. The VDI for the TV needs to package the following pieces of information:

- name, brand, model
- dimensions, weight
- features
- warranty
- relationships with other similar products

How should the different pieces of information be packaged? Are they to be considered resources (data) or descriptors (metadata)?

The approach we favour in CONVERGENCE is to package valuable information as generic resources and then link resources by identifying which are data, and which contain information about that data. In the end, therefore, there is no encoded difference between resources and descriptors. All we have are digital objects linked to other digital objects.

CONVERGENCE has further decided to make a distinction between the packages used for publication and search (Publication and Subscription VDIs) and those used for self-standing resources (Resource VDIs).

Using this scheme allows us to cope with our example in the following way:

- warranty, hardware features, and the like, are packaged as RESOURCEs, in one or many R-VDIs.
- as they are resources, REL statements can be applied to them. They represent valuable information, which can be protected.
- when the resource is published in the semantic overlay (see deliverable D3.2 [6]), making it discoverable, the relevant information (snippets of the features, model, relationships it is involved in) is packaged as a **RESOURCE of the P-VDI**.
- a REL license can be applied to the resource of the P-VDI making it possible to control who can read it (e.g. who can search for the features of the TV).

This is an example of an important CONVERGENCE requirement, namely the need to express and manage rights to what may be considered as metadata for (digital) resources. By representing (meta) data itself as resources/items, linked to the "original" object, it becomes possible to apply the full power of the REL. When (meta) data is intended to be searchable, it can be embedded as a resource inside a P-VDI.

This makes it possible to mark the syntactic distinction between RESOURCE tags and DESCRIPTOR tags in more subtle ways than in the original standard. For instance:

- DESCRIPTORs can be treated as the "**public part**" of the package, holding simple descriptive sentences. The system retrieves a package, and automatically presents it to the user.
- DESCRIPTORS can be treated as the parts of the package to which REL **is not** applied. If metadata is valuable and people want to protect it, it can be packaged as additional RESOURCEs.
- DESCRIPTORS can be treated as the parts of the VDI that contain searchable hints and annotations, **guiding the** creation of publication VDIs out of resource VDIs
- DESCRIPTORs can be used to **carry system-level** information such as relationships, expiry dates, sequence identifiers etc.

3.2 Control of Resources in VDIs

In CONVERGENCE, resources whose manipulation needs to be controlled, are referenced by VDIs. The manipulating entities are either end-users/devices, which consume resources, or peers of the system responsible for search and match operations. These different types of resources have different requirements for control and licensing.

We therefore classify these requirements by type of VDI:

- Resource VDIs – R-VDIs represent/enclose actual consumable resources. They thus present the most pressing requirements for control and licensing.
    - The system has to control access to (consumption of) R-VDI resources, enforcing the corresponding licenses. Access should only be granted to authorized users.

- The system has to control publication of R-VDI resources by users (through P-VDIs) enforcing the corresponding licenses. A resource may only be published by a specific P-VDI, if the user issuing the P-VDI has the right to publish the resource in question. This means that a different user can (re)publish the same R-VDI, with different metadata, only if she has the right to do so. The original creator can publish the same resource multiple times exposing different aspects of the resource. None of the rights expressed in the original R-VDI should prevent users from creating independent R-VDIs commenting, criticizing, or recommending a resource.
- "Updating" of R-VDIs by new VDIs in the same sequence must be controlled by the system through the enforcement of the corresponding licenses. Only users authorized to manipulate the resource (typically the owner of the resource or the R-VDI creator) may update an R-VDI. This can be interpreted as the right to reuse the same sequence identifier for a newly created VDI.
- Revocation of R-VDIs must be controlled by the system through the enforcement of the corresponding licenses. Only authorized users (typically the resource owner or R-VDI creator), may revoke an R-VDI.

- Publication VDIs – P-VDIs declare the existence of resources in the space of searchable objects by publishing their metadata and a link to the corresponding R-VDIs. The main control/licensing needs pertaining to this type of VDI are the following:
  - The system shall control access to (inspection of the contents of) P-VDI content (metadata) and matching of subscriptions to P-VDI metadata, enforcing the corresponding licenses. The only system peer authorized to access its content are those belonging to the fractal where it was injected. Therefore, a search issued by a user will only find P-VDIs which the user is entitled to access and a peer will only perform a search operation for a P-VDI (more precisely, its fractal) if it has the right to do so.
  - The system shall control revocation of P-VDIs through the enforcement of the corresponding governing licenses. Only authorized users (typically the P-VDI creator), may update a P-VDI.
  - The system shall control reporting of subscription results to users, enforcing the corresponding licenses. Only authorized fractals (sets of peers) may issue such reports and these may be directed only to an authorized set of users.

- Subscription VDIs – an S-VDI declares user's subscription criteria. The main control/licensing requirements are the following:

- The system shall control access to (inspection of the contents of) S-VDIs through the enforcement of the corresponding licenses. Only authorized fractals may process the contents of these licenses and only at the service of authorized users. This right can be used to regulate terms for collecting statistics about user subscriptions to specific VDIs.
- The system shall control revocation of S-VDIs through the enforcement of the corresponding licenses. Only authorized users (typically the S-VDI issuer), may revoke an S-VDI.

3.3 REL Verbs

Below we list the twenty-one verbs whose semantics are defined in ISO/IEC 21000-5 and its three amendments. The first fourteen are defined in the ISO/IEC 21000-6 Rights Data Dictionary. This list was the starting point for our analysis of the requirements of the use scenarios.

ActType (Verb)	Definition
Adapt	To **ChangeTransiently** an existing **Resource** to **Derive** a new **Resource.**
Delete	To **Destroy** a **DigitalResource**.
Delist	The right to unlink or delist (the reference to) the related resource from a related playback control sequences description (i.e. play-list) for the optical disc when the play-list is newly created from an existing one
Diminish	To **Derive** a new **Resource** which is smaller than its **Source.**
Embed	To put a **Resource** into another **Resource.**
Enhance	To **Derive** a new **Resource** which is larger than its **Source.**
Enlarge	To **Modify** a **Resource** by adding to it.
Enlist	To link the related resource into a new playback control sequences description (i.e. play-list) for the optical disc.
Execute	To execute a **DigitalResource**.
Export	To export the associated broadcast program to another rendering or storage device
ExtendRights	To extend the rights which are the originally transmitted
GovernedAdapt	To adapt the resource and at the same time to result in certain rights being associated with the adapted resource.
GovernedCopy	To copy the resource and at the same time to result in certain rights being associated to the copied resource.
GovernedMove	To move the resource and at the same time to result in certain rights being associated to the moved resource.

Install	To follow the instructions provided by an **InstallingResource.**
Modify	To **Change** a **Resource,** preserving the alterations made.
Move	To relocate a **Resource** from one **Place** to another.
Play	To **Derive** a **Transient** and directly **Perceivable** representation of a **Resource.**
Print	To **Derive** a **Fixed** and directly **Perceivable** representation of a **Resource.**
Reduce	To **Modify** a **Resource** by taking away from it.
Uninstall	To follow the instructions provided by an **UninstallingResource**.

Table 1 — Standardized ActType supporting ISO/IEC 21000-5

3.4 REL in CONVERGENCE scenarios

In this section, we analyze the four CONVERGENCE use scenarios (see deliverable D2.2 [7]), in terms of their licensing requirements and their expression as basic REL elements (i.e. description of issuers, principals, rights, conditions, etc.).

In each subsection, we describe licenses for specific VDI types (R-VDI, S- VDI, P-VDI) and issuers. Readers are asked to note that the analysis is limited to the requirements of the first phase of the trials. Additional licensing requirements will be described in later deliverables (D.8.2 and D.8.3), and summarized in the final deliverable of WP4 (D4.3).

3.4.1 Photos in the cloud and analyses on the earth

3.4.1.1 R-VDI for photos

A photographer with a business relationship with Alinari (i.e. belonging to the group of Alinari photographers) uploads photos and other data to the Alinari server though the VDI Creation service run by Alinari.

Element	Subelements/instances	Examples
Resource	Photo	
Metadata		Decided by photographer
Issuer	Photographer	
Principal	Alinari	A group that includes Alinari personnel
Rights	**GovernedCopy**	Alinari may store photo and adaptations
	GovernedAdapt	Alinari may edit photo and make low-res
	Post	Alinari may make P-VDIs
	License	Right to sublicense (new right)

Conditions	Fee Condition	No condition (Low Res), price set (High Res): lump money, percentage

Freelancer or Alinari personnel create a Photo VDI.

Element	Subelements/instances	Examples
Resource	Low-res photo	
Metadata		According to Alinari metadata schema
Issuer	Photographer	Freelancer or Alinari personnel
Principal	Anybody	
Rights (Low-res)	**Play**	
	Print	
Rights (Hi-res)	**Play**	Licence is downloaded
	Print	Licence is downloaded
Conditions	Exercise Limit Condition	Number of times Valid until…
	Fee Condition	No condition (LR), price set (HR)
	Territory Condition	Country where rights may be exercised
ERR	When	Each time play and print rights are exercised
	ER recipient	Issuer

3.4.1.2 P-VDI for photos

Freelancer is primarily interested in selling her photos to a large company such as Alinari, creating a business relationship with the company.

Freelancer creates the P-VDI and sets the condition that, for a fixed number of days following publication (X in the table below), the P-VDI will be discoverable only by S-VDIs from certified Alinari personnel. In this case, she will sell the photo at Price1 = 1 euro. For the next Y days and until the P-VDI expiration date, the P-VDI will be discoverable by any S-VDI containing keywords in the metadata. In this case, she will sell the photo at a different price Price2 = 0.8 euros.

Element	Subelements/instances	Examples
Resources	Photo Metadata	
Metadata	No	
Issuer	Freelance photographer	
Principal	Fractal	
Conditions	Fulfiller Condition	Issuer of P-VDI should be Alinari personnel (for X days after injection)
	Validity interval condition	X days after the injection of the P-VDI
	Fulfiller Condition	Not applicable
	Validity interval condition	Y days after the X days until expiry date
Rights	Match	
	Notify	
Rights	When	Every time a Match is found
	ER Recipients	Issuer and Alinari personnel

Alinari advertises photo to the cloud.

Element	Subelements/instances	Examples
Resources	Photo Metadata	
Metadata	No	
Issuer	Alinari	
Principal	Fractal	
Rights	Match	
	Notify	
ERRs	When	Every time a Match is found
	ER Recipients	Issuer

3.4.2 Videos in the cloud and analyses on the earth

In the discussion below we will use the following acronyms:

VMO = Video Material Owner

VCO = Video Channel Owner

VCU = Video Channel User

VD = Video Distributor, a service provider streaming video content (in our scenario, FMSH/ESCoM)

CH = Channel VDI Holder, a service provider hosting video channels (in our scenario, FMSH/ESCoM)

INC = Peruvian National Institute of Culture.

- Videos are co-produced by VMO and the Peruvian National Institute of Culture. VMO still owns the video material, but has contractual obligations towards INC. INC has the exclusive right to broadcast the material on its own Video Channel for the first three months after publication. The INC Video Channel is YouTube-like, where videos don't need to be analyzed before being posted.

- After 3 months, VMOs make the videos available for:

    - Analysts and VCOs certified by FMSH

    - VCOs certified by the Peruvian government (who post videos without making analyses, like INC)

3.4.2.1 S-VDI for videos

The entities involved are the Analyst and the Fractal. Analyst subscribes to videos for purposes of analysis. The license in the subscription is used to express a validity interval for the query itself.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Analyst** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	videos about a subject of interest
REL Verbs	**Match**	
	Notify	
Conditions	Validity Interval Condition	Match, Notify: 1 year
ERRs	Recipient: issuer	For each Match

The parties involved are the Video Distributor and the Fractal. VD subscribes to videos produced by a group of VMOs, making them available for streaming. In this case, a typical subscription would contain a query such as "videos about a subject of interest & produced by a VMO belonging to group of FMSH-VMOs". No specific licensing of the S-VDI is needed, hence Issuer, Principal and Verbs fields are empty.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VD** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	Videos about a subject of interest & produced by a VMO belonging to group of

		FMSH-VMOs
REL Verbs	**Match**	
	Notify	
ERRs	Recipient: issuer	For each Match

The parties involved are INC Member and Fractal. INC subscribes to videos to post them on her YouTube-like video channel. This is similar to the previous case.

Element	**Examples of Sub-elements or different instances**	**Examples**
Issuer	**INC Member** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	Videos talking about Peruvian Cultural Heritage and produced by a VMO belonging to group of Peruvian Government-VMOs
REL Verbs	**Match**	
	Notify	
ERRs	Recipient: issuer	For each Match

3.4.2.2 R-VDI for videos

The parties involved are VMO and INC Members. As INC is producer of video, VMO ensures that she can watch, analyse and post the video.

Element	**Examples of Sub-elements or different instances**	**Examples**
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**INC Members** Cryptographic Key	Principal has to belong to group of INC-Members *(Video resource is encrypted: a key for principals shall be given)*
Digital Resources	**Video Resource**	
REL Verbs	**GovernedCopy**	*Downloading, storing video*
	Play	*Watching video*
	Post	*Posting video in a channel*
Conditions	Territory Condition	Principal localized in Peru
	Track report	Report to Issuer any publishing of an analysis of the video
ERRs	Recipient: issuer	For each GovernedCopy & Post For each Play

Involved parties are: FMSH and Analysts. VMO gives to FMSH-Analysts the right to analyze the video.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**Group of Analysts** Cryptographic Key	Principal has to belong to group of FMSH-Analysts *(Video resource is encrypted: a key for principals shall be given)*
Digital Resources	**Video Resource**	
REL Verbs	**GovernedCopy**	*Downloading video*
	Play	*Watching video (includes decryption)*
Conditions	Validity Interval Condition	GovernedCopy: 6 months
	Exercise Limit	GovernedCopy: once
	Track report	Report to Issuer any publishing of an analysis of the video
ERRs	Recipients: issuer, group of FMSH-VDs	For each GovernedCopy

The parties involved are VMO and Video Distributor. VMO gives to FMSH-VD the right to store and stream the video.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**VD** Cryptographic Key	Principal has to belong to group of FMSH-VDs *(Video resource is encrypted: a key for principals shall be given)*
Digital Resources	**Video Resource**	
REL Verbs	**GovernedCopy**	*Downloading & storing video*
	Post	*Streaming video*
Conditions	Validity Interval Condition	GovernedCopy: 4 years Post: 4 years
	Exercise Limit	GovernedCopy: once
	Territory Condition	Principal localized at FMSH premises in

		Paris, France
ERRs	Recipients: issuer, group of FMSH-VCOs	For each GovernedCopy

The parties involved are VMO and Video Channel Owners. VMO gives to FMSH-VCOs the right to post icons for the video on their channel.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**Group of VCOs**	Principal has to belong to group of FMSH-VCOs or group of Peruvian Government-VCOs
Digital Resources	**Video Resource**	
REL Verbs	**Post**	*Posting icon of video in a channel*
ERRs	Recipients: issuer, group of FMSH-VDs	For each Post

The parties involved are VMO and Anybody. VMO gives to anybody the right to watch the video.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**Any**	
Digital Resources	**Video Resource**	
REL Verbs	**Play**	*Watching a video*
ERRs	Recipients: issuer, group of FMSH-VDs	For each Play

3.4.2.3 P-VDI for videos

Parties: VMO – Fractal

VMO notifies the uploading of his video to INC & VCOs, certified by the Peruvian government

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license	

	creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Metadata**	Metadata of Video
REL Verbs	**Match**	
	Notify	
ERRs	Recipients: issuer, group of INC-Members, group of Peruvian Government-VCOs	For each Match

Parties: VMO – Fractal

VMO notifies the uploading of his video to FMSH-VCOs, 3 months later

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Metadata**	Metadata of Video
REL Verbs	**Match**	
	Notify	
Conditions	Validity Interval Condition	Match, Notify: 2 years
	Validity Start Time Condition	Match, Notify: after 3 months
ERRs	Recipients: issuer, group of FMSH-Analysts, group of FMSH-VDs	For each Match

3.4.2.4 S-VDI for analyses

Parties: Video Material Owner – Fractal

VMO subscribes to analyses of his videos

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VMO** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	analyses about his videos
REL Verbs	**Match**	

	Notify	
ERRs	Recipient: issuer	For each Match

Parties: Video Channel Owner – Fractal

VCO subscribes to analyses about a subject of interest

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VCO** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	analyses about a subject of interest and produced by Analysts belonging to group of FMSH-Analysts
REL Verbs	**Match**	
	Notify	
ERRs	Recipient: issuer	For each Match

3.4.2.5 R-VDI for analyses

Parties: Analyst - Video Channel Owners

Analyst gives to FMSH-VCOs the ability to store & post his analysis on their channel

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Analyst** Date & time of license creation, Issuer localization	
Principal	**Group of VCOs**	Principal has to belong to group of FMSH-VCOs
Digital Resources	**Analysis**	
REL Verbs	**GovernedCopy**	*Downloading analysis*
	Post	*Posting the analysis in a channel*
Conditions	Validity Interval Condition	GovernedCopy: 4 years Post: 4 years
	Exercise Limit	GovernedCopy: once Post: once
ERRs	Recipient: issuer	For each GovernedCopy
	Recipients: issuer, group of FMSH-CHs	For each Post

Parties: Analyst – Anybody

Analyst gives to anybody the ability to read the analysis

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Analyst** Date & time of license creation, Issuer localization	
Principal	**Any**	
Digital Resources	**Analysis**	
REL Verbs	**Play**	*Reading analysis*
ERRs	Recipient: issuer	For each Play

3.4.2.6 P-VDI for analyses

Parties: Analysts – Fractal

Analyst notifies the uploading of her analysis to FMSH-VCOs

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Analyst** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Metadata**	Analysis
REL Verbs	**Match**	
	Notify	
Conditions	Validity Interval Condition	Match, Notify: 1 year
ERRs	Recipients: issuer, group of FMSH-VCOs	For each Match

Parties: Analysts – Fractal

Analyst notifies the issuer of the analyzed video (VMO), as requested in the Video VDI

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Analyst** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Metadata**	Analysis
REL Verbs	**Match**	
	Notify	
Conditions	Recipient Conditions	Principal has to be the issuer of the Video

		VDI referenced in the Analysis VDI
ERRs	Recipient: group of FMSH-VMOs	For each Match (as requested by issuer of the Video VDI)

3.4.2.7 **S-VDI for channels**

Parties: Video Channel Users – Fractal

VCO subscribes to posts of analyses about a subject of interest and/or belonging to a specific channel. No specific licensing needs.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VCU** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Query**	Posts about a subject of interest and/or belonging to a specific channel
REL Verbs	**Match**	
	Notify	
ERRs	Recipient: issuer	For each Match

3.4.2.8 **R-VDI for channels**

Parties: VCO - Holder of Channel VDIs

VCO gives to CH the ability to store and post her channel metadata

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VCO** Date & time of license creation, Issuer localization	
Principal	**CH**	Principal has to belong to group of FMSH-CHs
Digital Resources	**Channel metadata**	
REL Verbs	**GovernedCopy**	*Storing channel metadata*
	Post	*Posting the channel metadata*
Conditions	Validity Interval Condition	GovernedCopy: 4 years Post: 4 years
	Exercise Limit	GovernedCopy: once Post: once
	Territory Condition	Principal localized at FMSH premises in Paris, France
ERRs	Recipient: issuer	For each GovernedCopy For each Post

Parties: Video Channel Owner – Anybody

VCO gives to anybody the ability to browse her channel

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VCO** Date & time of license creation, Issuer localization	
Principal	**Any**	
Digital Resources	**Channel metadata**	
REL Verbs	**Play**	*Browsing channel*
Conditions		
ERRs	Recipient: issuer	For each Play

3.4.2.9 **P-VDI for channels**

Parties: Video Channel Owner – Fractal

VCO notifies end-users of the posting of a new analysis on her channel

Element	Examples of Sub-elements or different instances	Examples
Issuer	**VCO** Date & time of license creation, Issuer localization	
Principal	**Fractal**	
Digital Resources	**Metadata**	Metadata of Analysis
REL Verbs	**Match**	
	Notify	
Conditions	Validity Interval Condition	Match, Notify: 1 year
ERRs	Recipient: issuer	For each Play
	Recipient: group of VCUs	For each Match

3.4.3 Augmented Lecture Podcast

3.4.3.1 **R-VDI for Annotations**

Issuer: Students

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Student**	Issuer: User of the ALP-Application]

		Issuer Details: Date and time of the license creation, revocation methods, etc.
Principal	**Student**	Other students using the ALP-Application, identified one by one, as well as a group.
REL Verbs	**Extend Rights** **Play**	Students may be given the right to: - Grant other users the right to play annotations - Reference annotations and re-share them with other groups
Digital Resources	Annotations	
Conditions	Validity Interval Condition	Students may set an expiry date for their annotations: - "End of term" - "Never" - "Date: DD/MM/YYYY"
	Fulfiller Condition	*not available*
	Exercise Limit	Unlimited
	Track Report	*not available*
	Fee Condition	*not available*
	Territory Condition	*not available*

3.4.3.2 **R-VDI for Podcasts**

Issuer: Lecturers

Element	Examples of Sub-elements or different instances	Examples
Issuer	Lecturers	Issuer: Lecturers who grant certain rights to students (e.g. download of podcast) Lecturers who grant certain rights to podcast service (e.g. post video) Issuer Details: Date and time of the license creation, revocation methods, etc.
Principal	Podcast Service Students	

REL Verbs	**Post**	The podcast service has the right to: - provide learning materials to students Students have the right to: - Download slides and video - Play / View slides and video
	Governed Copy	
	Play	
Digital Resources	Slides, Videos, Podcasts	
Conditions	Validity Interval Condition	Lecturers may set an expiry date for their learning material: - "End of term" - "Never" - "Date: DD/MM/YYYY"
	Fulfiller Condition	*not available*
	Exercise Limit	Unlimited
	Track Report	*not available*
	Fee Condition	*not available*
	Territory Condition	*Right to post: only LMU (in case of collaboration, this can be extended)* *Right to create GovernedCopy: all*

3.4.4 Smart Retailing

3.4.4.1 R-VDIs for Product Type (WIPRO Trial)

Parties: Manufacturer – Retailers

Manufacturer creates a Product Type VDI allowing Retailers to change information in it

Issuer: Manufacturer

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Manufacturer**	Issuer Details: Date and time of the license creation, revocation methods, etc.
Principal	**Retailer**	The Principal receives a cryptographic key from Manufacturer
REL verbs	**GovernedAdapt**	Add information to a Product Type VDI: consumer id, serial number and warranty details. This happens when the product is sold to a Consumer, and the Product Type VDI is used to spawn a new Product Instance VDI. It

		is the Retailer's clerk at the POS creates the Product Instance VDI.
Digital Resources		Barcodes (Metadata)
Conditions	Validity Interval Condition	The Manufacturer defines the time frame where the Retailer may exert the Right to "Modify" the product VDI. For example: one year from the time the product is released in the market
	Fulfiller Condition	The Manufacturer specifies who his certified Retailers are. If a Retailer is NOT certified (fulfiller condition) then he cannot sell the product ("Modify" the product VDI).
	Exercise Limit	Specifies the maximum number of times that the Retailer may exert the Right ("Modify"). In this case the number of products the Retailer purchased from the Manufacturer can be used as the number of times that he can exert the Right to "Modify" a product VDI and sell it to a Consumer.
	Track Report	Counts the number of times that the Retailer exerts the Right to "Modify" the product VDI. Number of sales.
	Fee Condition	Commission of the Manufacturer in each product the Retailer sells to a Consumer
	Territory Condition	The Right may be exerted only in geographical regions where the certified Retailers operate.

3.4.4.2 **Product Type R-VDI from Manufacturer/Supplier to Retailer (UTI Trial)**

Parties: Manufacturer/Supplier - Retailer

The manufacturer gives to the supplier the right to sell its products

Element	**Examples of Sub-elements or different instances**	**Examples**
Issuer	**Manufacturer/ Supplier**	Issuer Details Digital Signature of Manufacturer/Supplier
Principal	**Retailer**	Retailer In the case of Key Holder Principal, a certain cryptographic key is provided to the Principal by the Key Issuer
REL Verbs	**GovernedCopy**	Download Product VDI and publish own

	GovernedAdapt	VDIs
Digital Resources	Metadata	Product Metadata
Conditions	Validity Interval Condition	1 year
	Fulfiller Condition	Supplier must be certified by Manufacturer/Supplier
	Exercise Limit	N/A
	Track Report	Counts the number of times that a Principal exerts a GovernedAdapt Right
	Fee Condition	N/A
	Territory Condition	Romania
ERRs	Recipient: issuer	For each GovernedAdapt

3.4.4.3 **Retailer Promotion Product Type R-VDI to consumer (UTI Trial)**

Parties: Retailer - Consumer

The retailer gives to the supplier the right to view its promotions.

Element	**Examples of Sub-elements or different instances**	**Examples**
Issuer	**Retailer**	Issuer Details: Digital Signature of Retailer
Principal	**Anyone**	
REL Verbs	**GovernedCopy**	Download Product VDI Publish own VDIs
	GovernedAdapt	
Digital Resources	**Metadata**	Promotion Metadata
Conditions	Validity Interval Condition	2 weeks
	Fulfiller Condition	N/A
	Exercise Limit	N/A
	Track Report	Counts the number of times that a Principal exerts a GovernedAdapt Right
	Fee Condition	N/A
	Territory Condition	N/A
ERRs	Recipient: issuer	For each GovernedAdapt

3.4.4.4 **Retailer Promotion Product Type P-VDI (UTI Trial)**

Parties: Retailer - Fractal

The retailer publishes its promotions.

Element	**Examples of Sub-elements or different instances**	**Examples**
Issuer	**Retailer**	

		Issuer Details: Digital Signature of Retailer
Principal	**Fractal**	
REL Verbs	**Match**	
	Notify	
Digital Resources	Metadata	Promotion Metadata
Conditions	Validity Interval Condition	Match, Notify: 2 weeks
ERRs	Recipient: Retailer	For each match

3.4.4.5 Consumer Preferences Product Type S-VDI (UTI Trial)

Parties: Consumer - Fractal

The consumer subscribes to her shopping preferences.

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Consumer**	Issuer Details: Consumer Identifier (if exists)
Principal	**Fractal**	
REL Verbs	**Match**	
	Notify	
Digital Resources	Query	Shopping preferences
Conditions	Validity Interval Condition	Match, Notify: 2 months
ERRs	Recipient: Consumer	For each match

3.4.4.6 R-VDI for product instance (WIPRO Trial)

Parties: Retailer - Customer

Retailer creates a product instance VDI for the Customer when he buys a product

Issuer: Retailer

Element	Examples of Sub-elements or different instances	Examples
Issuer	**Retailer**	Issuer Details: Date and time of the license creation, revocation methods, etc.
Principal	Customer	
REL Verbs	**Governed Copy**	- Download Product Instance VDI - Create any other VDI linked to this one - Delete the Product Instance VDI
	Adapt	
	Delete	
Digital		N/A

Resources		
Conditions	Validity Interval Condition	Permanent
	Fulfiller Condition	Retailer must be certified by Manufacturer/Supplier
	Exercise Limit	N/A
	Track Report	N/A
	Fee Condition	N/A
	Territory Condition	N/A

3.4.5 The CONVERGENCE REL in the use scenarios

In what follows, we provide descriptive narrations of the use scenarios, showing how the **descriptive** power of REL can support complex scenarios, provided that rights and actors are appropriately grouped.

3.4.5.1 **A photographer publishes his work in the cloud**

A photographer wants to sell her pictures. Her market includes two classes of potential buyers:

- Class A: Large companies who print travel guides using photos from the cloud. This class of company may include companies like Alinari, which advertise and sell photos via their web - site.

- Class B: Other individual photographers.

In the P-VDI, the photographer defines sales conditions, based on the time when it was injected and the identity of the potential purchases (the creator of an S-VDI matching the P-VDI). More specifically:

- She defines the condition that for a certain period of X days after injection, the P-VDI will only be visible to S-VDIs issued by purchasers in Class A and that in this period the photo will be sold at a price Y euros. Companies who buy the photograph will also have to pay the photographer royalties on every copy of the travel guide sold and on every time a user views the photo on their the web-site. Alternately, the company can create a longer-term business relationship with the photographer who thus has a strong interest in selling to companies in this category.

- She defines a second condition that if the photograph is not sold in the first X days, it will also become visible to users in Class B, until the expiry date for VDI. Users in this category will be able to buy the photo at a price of Z euros.

To achieve this the photographer uses the REL. Note that in a classical IP network, the photographer would have to create her own web site. In this scenario, it would be difficult for her to maintain different conditions for different types of users.

3.4.5.2 Licensing Video Archive Material

In this scenario too, the REL makes it easy for FMSH Video Material Owners and Video Channel Owners to set rights over their digital resources. More specifically:

- A VMO will be able to define a "trusted analysts group" that has been certified by FMSH and has the right to download the video and analyse it.

- A VCO will be able to define the different classes of users that have the right to post an analysis on their web site.

3.4.5.3 Augmented Lecture Podcast application

Lecturers who provide learning materials to students often want the learning materials to be redistributed for educational purposes, under the terms of Creative Commons licenses (http://creativecommons.org/). The REL makes it possible to formalize the terms of these licenses in a machine-readable form and to bind them with actual resources. The type of license a lecturer chooses will depend on circumstances.

If the lecturer herself created the content of the learning material, she may include the following license:

- CC by attribution, share-a-like, non-commercial or
- CC by attribution, share-a-like, no derivation

If, on the other hand, she has included external content, she may want a license that honours the original copyright, for example by not allowing further redistribution. In both cases, the VDI binds the expression of these rights and the actual resource.

3.4.5.4 Licensing in Retail Scenario

Samsam is launching a new LED TV. Before the launch, the company creates a Product Type VDI for the new model. The VDI contain a full description of the TV: name, brand, model, dimensions, weight, features, resources, etc. and is valuable for the company's relationship with its customers. Therefore, protecting the information and ensuring it is always reliable is very important for the company. Here REL technology can play a valuable role, restricting the rights granted to resellers, and protecting the product VDI as it passes along the value chain from manufacturer to retailer to consumer. The REL allows Samsam to define different levels of rights, for different categories of user.

1. Certified resellers, selling the TV to consumers. From the point of view of these users the REL can prevent the use of the VDI on stolen, counterfeit or imitation products, and by unofficial retailers.

2. Retailers with limited rights to alter (create, edit, revoke) the Product Type VDI. The REL will prevent these users from freely altering information about the TV. For example if the information in the VDI specifies that the TV has 1920x1080 video resolution, the REL will prevent retailers for claiming otherwise or from removing the correct data certified by Samsam.

When a certified retailer of Samsam products is informed about the release of the new TV and decides to sell the TV, he obtains the Product Type VDI for the TV. The REL will give him the right to add information about promotions, special offers, related products, and to create a new Product Instance VDI, when he sells a TV to a consumer. This new VDI will contain additional info such as consumer ID, serial number, warranty details, etc.

4 CONVERGENCE Governance and Licensing Scheme

4.1 Overview

All the requirements outlined in section 3 reflect a basic underlying need for *access control*. Controlling and restricting the ability to consume a media resource or to process a metadata resource, basically means limiting users' rights to access the resource in a specific way. In other words, the control requirements described in section 3, can be handled through the concession or denial of specific rights to users or fractals.

In real world situations, such as those described in the CONVERGENCE use scenarios, the "goods" that need to be delivered and controlled can be manipulated in different ways by different categories of user. For instance, Alinari resources may be able to be consumed by a large number of authorized consumer users, and edited by a smaller set of editor users.

Issuing a specific license to every individual user for every good to which the user has some kind of right, would lead to a proliferation of licenses, and an increase in the burden of managing and enforcing them.

To avoid this, the CONVERGENCE licensing scheme will attribute licenses (and corresponding rights), not to individual users but to user groups. The only individual licenses will be those declaring that an individual belongs to a specific group. It follows that every scenario requires an entity (the Governance Entity), responsible for governing a specific class of content. This entity will then establish a set of "authority levels", each conferring a specific set of rights to a specific user group. For each level, (and for each individual Digital Item) an appropriately authenticated license will define the rights concerned and specify the user group to which they will be attributed.

Whenever a specific user acquires the right to belong to a specific user group, the Governance Entity will issue him a license, declaring that he is a member of the group. This will allow the

CONVERGENCE system to enforce the licenses, assessing if a specific user is allowed to perform a specified action on a particular resource.

Figure 1 illustrates the rights attributed to four user groups, in which each group has a specific set of rights (or "rights area"), represented by a blue rectangle. The greater the area of the rights set, the broader the rights that the corresponding user group possesses. For instance, the group of users who have adhered only to the free service (Non-Paying UserGroup), has the least rights, while the group of the users which pay the highest fees (High-Paying UserGroup), and the group of editor users has the most rights.

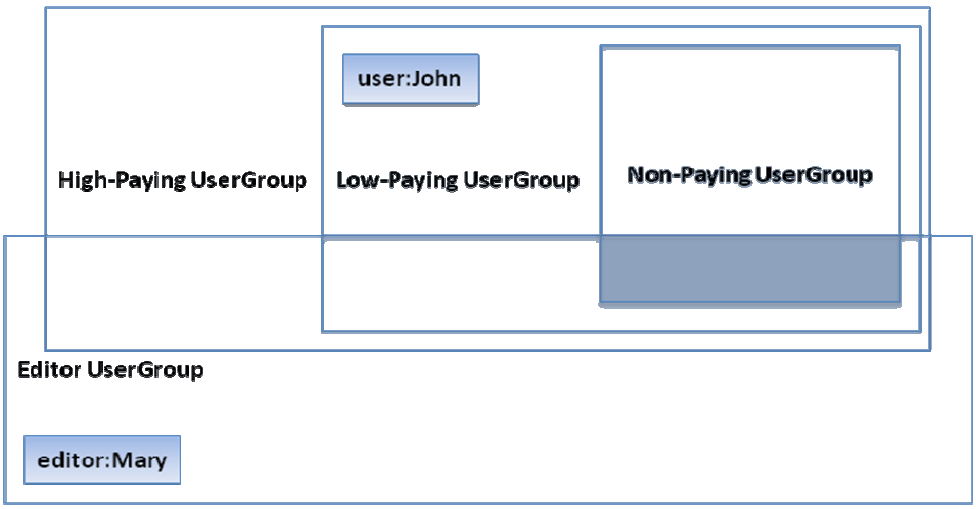

Figure 1 – Example of User Groups and related "Rights Sets"

In the situation described in Figure 1, there is one license defining user group rights for each groups, and for each resource to be controlled. Thus the Governance Entity would issue user John with an individual license declaring him to belong to user group "Low-Paying UserGroup" while editor Mary would receive an individual license declaring her to belong to user group "Editor UserGroup".

In this setting, Fractals are also groups, this time composed of multiple peers. CONVERGENCE's licensing scheme can thus attribute rights sets to Fractals, rather than to individual peers. Each individual peer would then be given a license, certifying that it belongs to a specific fractal. This situation is depicted in *Figure 2*.

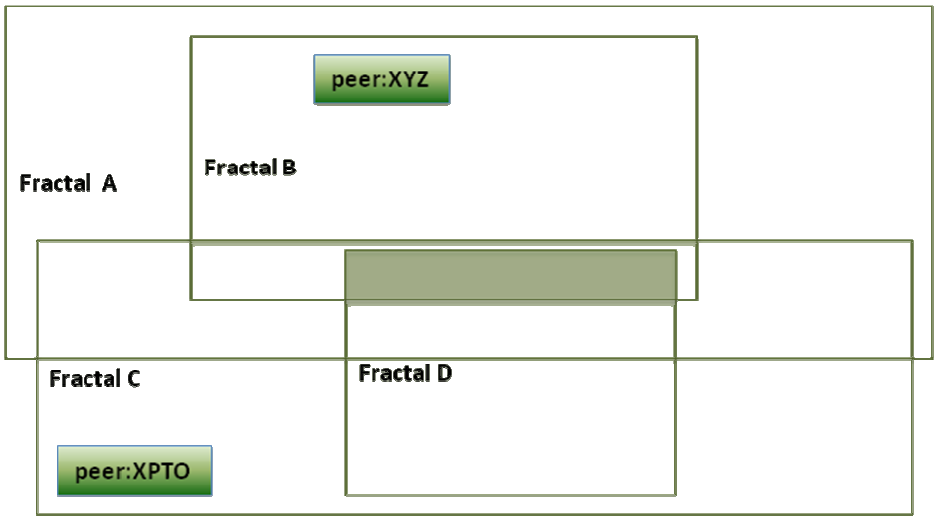

Figure 2 – Example of Fractal "Rights Sets"

4.2 Implementation with MPEG-21 REL

4.2.1 Base Mode

The governing and license scheme defined in 4.1 can be implemented in MPEG-21 REL through the two types of licenses presented in Figure 3.

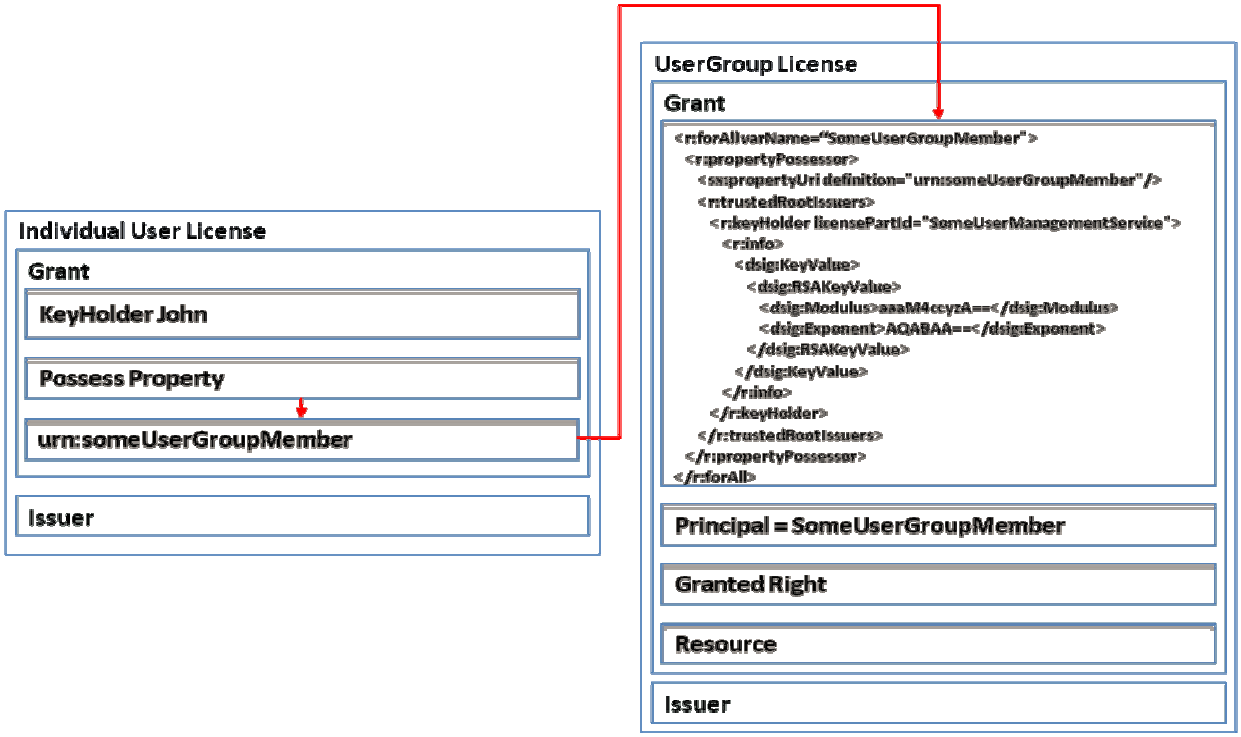

Figure 3 – CONVERGENCE Main Licenses

The right side of this image presents a User Group license. These licenses must:

- contain a *Grant* which, in its turn:
    - defines a variable (through the *forAll* element), which basically stands for all entities which possess some specific property. In the given graphical example it stands for the entities possessing the property of "*urn:someUserGroupMember*", that is, the property of being a member of user group "*someUserGroup*".
    - defines the *Principal* (the entity to which rights are being granted) as any entity of the type defined in the previous variable.
    - defines the *Right* which is being granted by the license. This should be a right that is performable over a digital item.
    - specifies the *Resource* over which the right is being granted (the digital item).
- identify the *Issuer* of the license. This should typically be an entity which operates a specific "content channel", owns the content that it distributes and thus sets the content accessing policy (the Governance Entity).

The left side of that image presents the individual user's license, (or User Membership License), which certifies his belonging to a specific user group. These licenses must:

- contain a *Grant* which, in its turn:
    - defines the *Principal* of the license. Employs the *keyHolder* element to uniquely identify a specific user.
    - defines the *Right* which is being granted by the license. This should be the possession of a specific property, which is defined (below) as the license's *Resource*.
    - specifies the *Resource* over which the right is being granted. This resource should be the property of belonging to some specific user group.
- identify the *Issuer* of the license. This should typically be an entity which operates a specific "content channel", owns the content that it distributes and thus sets the content accessing policy, that is, specifies the different user groups and their corresponding right and privileges and manages users' participation in said user groups (the Governance Entity).

In the example given in Figure 3, in the license on the left (the User Membership License), user (KeyHolder) John is granted a right which corresponds to the possession of the property "*urn:someUserGroupMember*", that is, the property of belonging to user group "*someUserGroup*".

In the license on the right (the UserGroup license), it is defined that all entities that possess the property "*urn:someUserGroupMember*" have a specific right over some specific resource.

Combining the two licenses, CONVERGENCE's content governance provisions are able to determine that John is entitled to exert the defined right over the specified resource.

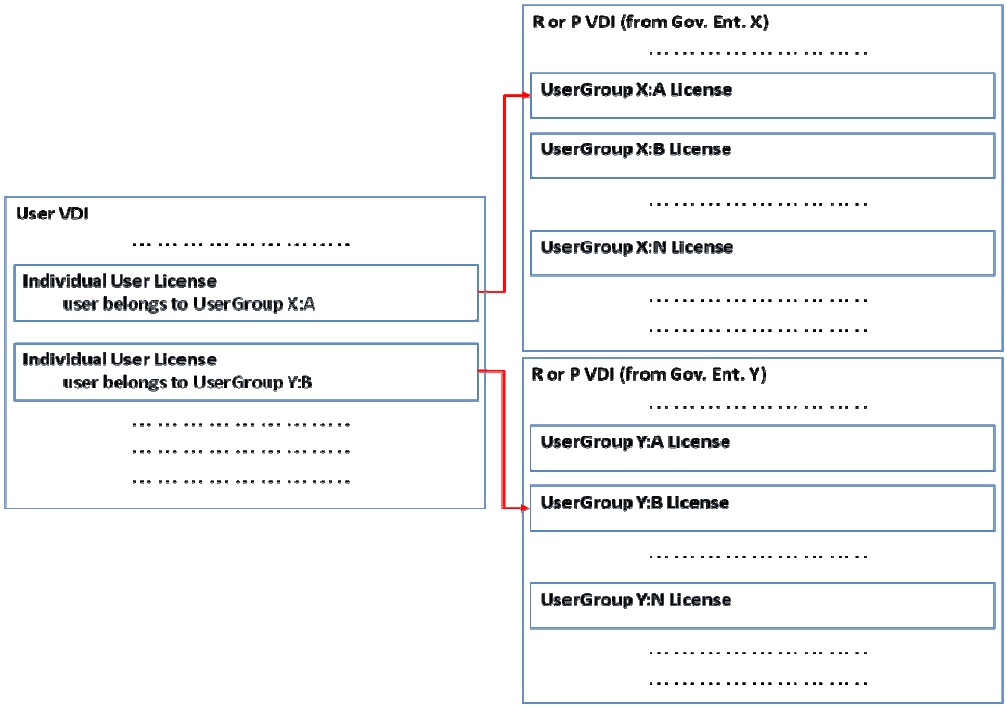

Figure 4 – License Placement in VDIs

In each CONVERGENCE scenario, the User Licenses (or User Membership Licenses) and UserGroup Licenses are issued by the appropriate Governance Entity (see section 4.1). As they refer to one specific resource, UserGroup Licenses are generated at VDI creation time, and need to be embedded within the appropriate VDIs (as presented in Figure 4). User Membership Licenses are generated whenever a specific user acquires the right to belong to a user group and are inserted in the user's VDI (as presented in Figure 4).

The situation for the rights of Fractals is very similar. Fractal rights are expressed in Fractal Licenses. The structure of the licenses and the production process are the same as for UserGroups Licenses. The only difference is that the Principals are Fractals (and not groups of users), and that the granted rights are different.

In this setting the equivalent of a User Membership License is a Peer Membership License. In this case, however, the structure of the license and the production process cannot be the same. Conceptually, individual Peers (Principals) are granted the right to belong to specific Fractals. However important distinctions break the symmetry. Fractals are highly dynamic in their nature, with Peers joining and leaving depending on the content they publish or subscribe. Moreover, it is impossible, given the size of the system, to maintain a centralized directory of the distribution of peers over fractals. Maintaining a decentralized, scalable system of "Fractals of Trust" is a major challenge. Such a system will need to guarantee the same kind of trust required for a Peer Membership License, but in a decentralized fashion. In the next chapter we present a preliminary design. One of the key tasks for future architecture and

implementation work-packages will be to finalize this design. The results will be reported in future deliverables.

4.2.2 License Optimization

In the scheme defined in section 4.2, every VDI in the CONVERGENCE system needs to embed not only Fractal Licenses but also UserGroup licenses for all relevant user groups.

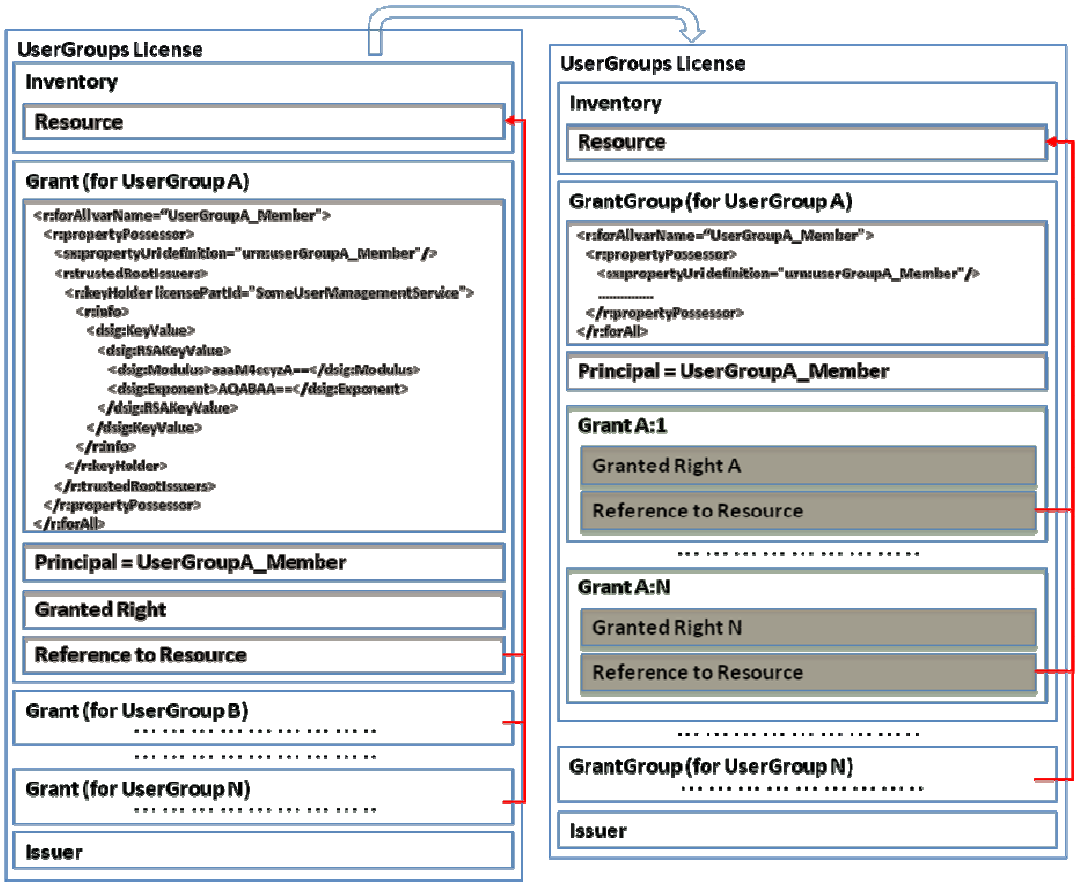

Figure 5 –UserGroup Licenses merged into a Single UserGroups License

To reduce duplication, it would be possible to merge all the grants, contained in different licenses, into a single license, as shown on the left side of Figure 5. The same procedure could also be applied to Fractal Licenses.

The license would thus contain a grant for each relevant user group (or fractal) and each grant would use its resource declaring element to reference the resource whose manipulation is controlled by the license. The resource would be declared only once, in the *Inventory* element, at the beginning of the license.

However, even in this scheme, the license would contain a potentially huge number of different grants (number of rights * number of UserGroups). There is thus room for further

optimization, providing an optimally ordered and minimally redundant internal structure for the license.

Each user group may hold more than one right over a particular digital resource. This means that *Grants* belonging to the same *Principal* (which may only specify one *Right* each) can be packed together, using the *grant Group* element.

Thus, if the UserGroups license contains:

- one *GrantGroup,* for each UserGroup, and the *GrantGroup* contains:
    - one declaration of the *Principal* defined in the form of a variable which represents all entities that possess a specific property (the property of belonging to a specific user group)
    - One *Grant* for each of the principal's rights over the controlled resource

the resulting license (presented on the right side of Figure 5) will define all the rights of all the user groups over the controlled resource.

Equivalently, a Fractals License can define the rights of all relevant Fractals over the controlled resource.

4.2.3 Expression of Some Specific Rights

In the light of the requirements described in section 3, it will be necessary to control at least the following actions:

- Regarding the R-VDI:
    - Playing or consuming the media resources of the R-VDI.
    - Deleting the R-VDI's resources (and consequently the R-VDI itself) from the CONVERGENCE system.
    - Revoking an old R-VDI and storing a new R-VDI.
- Regarding the P-VDI:
    - Performing matches on the metadata resources of the P-VDI.
    - Issuing event reports regarding the previous matches.
    - Collecting information from the metadata contained in the S-VDI, for statistical purposes.
    - Revoking the P-VDI from the CONVERGENCE system.
    - Revoking an old P-VDI and storing a new P-VDI.
- Regarding the S-VDI:
    - Collecting information from the query contained in the S-VDI, in order to aggregate it for statistical purposes.
    - Revoking the S-VDI from the CONVERGENCE system.

The focus of the following discussions will be on licenses for publications and subscriptions.

Not all of the actions just described have a corresponding expression in the MPEG-21 REL. The table below provides a summary description of actions not considered in the standard.

Action	Definition
Match	The right, (of Fractals), to store a set of metadata (P-VDI case) or a query formulation (S-VDI case), as the payload of a *did:Resource* element of the VDI, and to perform a matching of said resource against complementary resources, belonging to a specified set of users.
Notify	The right, (of Fractals), to report, (to a specific set of users), a match between a specific *Resource*, (specific metadata contents of a *did:Resource* of a VDI), and a user query or user subscription.
Aggregate	The right of peers (or crawlers external to the CONVERGENCE middleware), to report aggregated information about Rights that are executed on specific *Resources* (i.e. specific metadata contents of a *did:Resource* of a P-VDI, a specific user query or user subscription formulation of a S-VDI) to a specific set of users.

Thus, generally speaking, licenses in publications and subscriptions, control:

- what is to be collected or matched to what
- in which way such information is to be distributed to whom
- until when

All three actions listed in the table above require the specification of a list of users. Our preliminary solution is to include this list in the Condition fields of the license. However, this proposal needs to be validated by further investigations in WP3 and WP6.

At first, it might seem that Match and Notify are so tightly connected that only one of them is needed. However, in reality we have to treat and express both operations as separate rights. The reason stems from the situation depicted in *Figure 6*.

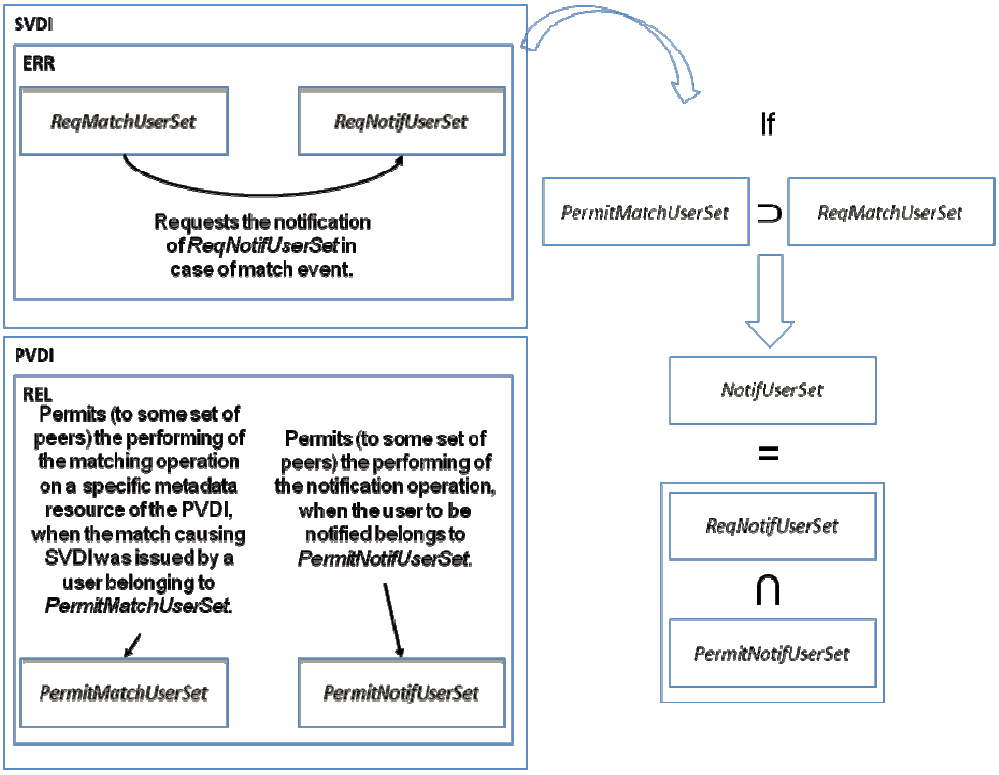

Figure 6 –Event Reporting and P-VDI Rights Scenario

The scenario shown in the figure shows an S-VDI. A user (*UserA*) uses MPEG-21 ERRs to specify the matching criteria for a subscription and the set of target users (*ReqNotifUserSet*) to be notified, with an ER, when matching P-VDIs are found.

UserA may be seen as a user set containing a single member (*UserA)*. We can call this set the *ReqMatchUserSet,* meaning the set of users who require a specific match.

When *ReqNotifUserSet* contains other users, besides *UserA*, *ReqNotifUserSet* ≠ *ReqMatchUserSet*. In this case the subscribing user is performing not only a traditional subscription, but also a *push* of notifications to third party users.

On the other hand, a content publishing user may wish to limit the set of users whose S-VDIs are matched against her P-VDIs (*PermitMatchUserSet*), or to limit the set of users who receive match notifications, for their P-VDI (*PermitNotifUserSet*).

If the condition *ReqNotifUserSet* = *ReqMatchUserSet* was mandatory for all situations, the only user who could receive the ER would be the original subscriber. Thus the Match and Notify operations would involve the same sets of users and, in terms of rights management, they could be treated in a joint manner.

However, CONVERGENCE will also support the condition when *ReqNotifUserSet* ≠ *ReqMatchUserSet*. Given that in this case the Match and Notify operations concern different sets of users, they must be treated separately in terms of rights management. In other words, *PermitMatchUserSet* ≠ *PermitNotifUserSet*.

Thus, the publishing user should be able to set Match and Notify rights separately and associate them with different sets of user (*PermitMatchUserSet, PermitNotifUserSet*).

The set of users that will actually receive an ER (*NotifUserSet*), from a specific Match event, is determined as follows:

- if *PermitMatchUserSet* ⊃ *ReqMatchUserSet:*
  - the intersection – *ReqNotifUserSet* ∩ *PermitNotifUserSet*
- else:
  - the empty set – ∅

To avoid pollution of the system by a myriad of notifications for every match, the system will assume *PermitNotifUserSet* = *PermitMatchUserSet*, unless a publisher specifies a different *PermitNotifUserSet*. In other words all users whose subscriptions match the P-VDI may receive a notification.

Similarly, the system will assume that *PermitMatchUserSet* = *ReqMatchUserSet* unless a publisher specifies a different *PermitMatchUserSet*. Thus any user's subscription can be matched against the P-VDI, but only the subscribing user has the right to be notified of the match.

Match and Notify rights will be expressed in UserGroups licenses, (or Fractal licenses) using the *rightUri* REL element, as illustrated below.

```
.....
<sx:rightUri definition="conv:right:match"/>
<sx:rightUri definition="conv:right:notify"/>
<sx:rightUri definition="conv:right:aggregate"/>
.....
```

An alternative would be to develop a CONVERGENCE-specific XML schema in which these rights would be expressed by extending the REL Right type, as illustrated below.

```
.....
<xsd:complexType name="Match">
  <xsd:complexContent>
    <xsd:extension base="r:Right"/>
  </xsd:complexContent>
</xsd:complexType>
```

4.3 Example of a License Scheme

Pictures Inc. is a company that has accumulated a vast photographic portfolio that it intends to exploit commercially, using the CONVERGENCE infrastructure.

In this scenario, there will be two types of users:

- Consumer Users –the basic users of the system - entitled to access photos on a pay-per-item basis.
- Editor Users – Pictures Inc.'s employees - responsible for the publication and continuous updating of company - owned content.

Consumer Users are only allowed to consume or subscribe to Pictures Inc. resources. Editor users are allowed to consume, subscribe, publish or delete Pictures Inc. resources.

As far as concerns Fractals (or, more specifically, peers belonging to Fractals), those where Pictures Inc.'s VDIs are published, will be granted the right to:

- Match user queries or user subscriptions to the metadata of a P-VDI
- Report matches to a specified set of users.

Implementing this scheme, will involve various types of licenses, in different locations, as described below.

4.3.1 UserGroups License

4.3.1.1 UserGroups License for R-VDIs

An R-VDI UserGroups License specifies all the rights of all Pictures Inc. users, over a specific R-VDI. This license will contain:

- An *Inventory* declaring the R-VDI's resource(s) (represented by *did:Resource* elements), whose manipulation is regulated by the license.
- Two *GrantGroups*:
    - A *GrantGroup* defining all the rights of the Consumer Users Group, over the governed resource(s). This will contain:
        - A variable that represents an entity that is characterized by the possession of the property of being a member of the Consumer UserGroup.
        - The *Principal*, to whom the rights are attributed, defined as an entity of the type defined in the variable above.

- A *Grant* specifying the right to play/consume the resource. This will contain:
  - A *Right* element of type *mx:play*.
  - A reference to the element (contained in the top *Inventory*), where the governed resource (over which the right is being granted) is specified.

- A *GrantGroup* defining all the rights of the Editor Users Group, over the governed resource(s). This must contain:
  - A variable that represents an entity that is characterized by the possession of the property of being a member of the Editor UserGroup.
  - The *Principal*, to whom the rights are attributed, defined as an entity of the type defined in the variable above.
  - Two *Grants*:
    - A *Grant* specifying the right to play/consume the resource. This will contain:
      - A *Right* element of type *mx:play*.
      - A reference to the element, contained in the top *Inventory*, where the governed resource (over which the right is being granted) is specified.
    - A *Grant* specifying the right to delete the resource. This will contain:
      - A *Right* element of type *mx:delete*.
  - A reference to the element, contained in the top *Inventory*, where the governed resource (over which the right is being granted) is specified.
- The *Issuer* of the license, typically the original owner of the R-VDI.

The text box below shows an example of this kind of license.

```
<?xml version="1.0" encoding="UTF-8"?>
<r:license
  xmlns:acme="urn:acme"
  sx:profileCompliance="acme:example"
  xmlns:r="urn:mpeg:mpeg21:2003:01-REL-R-NS"
  xmlns:sx="urn:mpeg:mpeg21:2003:01-REL-SX-NS"
  xmlns:mx="urn:mpeg:mpeg21:2003:01-REL-MX-NS"
  xmlns:dsig="http://www.w3.org/2000/09/xmldsig#"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:mpeg:mpeg21:2003:01-REL-R-NS rel-r-profile.xsd">
  <r:inventory>
    <!-- Specification of the Governed Resource -->
    <r:digitalResource licensePartId="governedMediaResource">
      <nonSecureIndirect URI="conv:vdi:VDI1:mediaResource"/>
    </r:digitalResource>
  </r:inventory>
```

```
<!-- GrantGroup for the definition of the rights of the Consumer Users UserGroup -->
<r:grantGroup>

  <!-- Variable which is employed to define the members of the Consumer Users UserGroup  -->
  <r:forAll varName="ConsumerUserGroupMember">
    <r:propertyPossessor>
      <sx:propertyUri definition="conv:consumerUserGroupMember"/>
      <r:trustedRootIssuers>
        <r:keyHolder licensePartId="ConvergenceUserManagementService">
          <r:info>
            <dsig:KeyValue>
              <dsig:RSAKeyValue>
                <dsig:Modulus>aaaM4ccyzA==</dsig:Modulus>
                <dsig:Exponent>AQABAA==</dsig:Exponent>
              </dsig:RSAKeyValue>
            </dsig:KeyValue>
          </r:info>
        </r:keyHolder>
      </r:trustedRootIssuers>
    </r:propertyPossessor>
  </r:forAll>

  <!-- Specification of the Principal who is entitled to the Rights over the Resource -->
  <r:principal varRef="ConsumerUserGroupMember"/>

  <!-- Definition of a Condition, the Validity Interval -->
  <r:validityInterval>
    <r:notBefore>2009-01-01T00:00:00</r:notBefore>
    <r:notAfter>2013-01-01T00:00:00</r:notAfter>
  </r:validityInterval>
  <r:grant>
    <!—
     Definition of the right to play the governed media Resource, to which the Principal is entitled
    -->
    <mx:play/>
    <r:igitalResource licensePartIdRef="governedMediaResource"/>
  </r:grant>
</r:grantGroup>

<!-- GrantGroup for the definition of the rights of the Editor Users UserGroup -->
<r:grantGroup>

  <!-- Variable which is employed to define the members of the Editor Users UserGroup  -->
  <r:forAll varName="EditorUserGroupMember">
    <r:propertyPossessor>
    <sx:propertyUri definition="conv:EditorUserGroupMember"/>
      <r:trustedRootIssuers>
        <r:keyHolder licensePartId="ConvergenceUserManagementService">
          <r:info>
            <dsig:KeyValue>
              <dsig:RSAKeyValue>
                <dsig:Modulus>aaaM4ccyzA==</dsig:Modulus>
                <dsig:Exponent>AQABAA==</dsig:Exponent>
              </dsig:RSAKeyValue>
            </dsig:KeyValue>
          </r:info>
        </r:keyHolder>
      </r:trustedRootIssuers>
    </r:propertyPossessor>
```

```
    </r:forAll>

    <!-- Specification of the Principal who is entitled to the Rights over the Resource -->
    <r:principal varRef="EditorUserGroupMember"/>

     <!-- Definition of a Condition, the Validity Interval -->
    <r:validityInterval>
      <r:notBefore>2009-01-01T00:00:00</r:notBefore>
      <r:notAfter>2013-01-01T00:00:00</r:notAfter>
    </r:validityInterval>
    <r:grant>
      <!—
        Definition of the right to play the governed media Resource, to which the Principal is entitled
      -->
      <mx:play/>
      <r:igitalResource licensePartIdRef="governedMediaResource"/>
    </r:grant>

    <r:grant>
      <!-- Definition of the right to delete the governed Resource, to which the Principal is entitled -->
      <mx:delete/>
      <r:igitalResource licensePartIdRef="governedMediaResource"/>
    </r:grant>
  </r:grantGroup>

  <!-- Specification of the Issuer of the license, which is the original VDI owner -->
  <r:issuer>
    <r:keyHolder licensePartId="conv:user:John">
      <r:info>
        <dsig:KeyValue>
          <dsig:RSAKeyValue>
            <dsig:Modulus>Fa7wo6NYfm==</dsig:Modulus>
            <dsig:Exponent>AQABAA==</dsig:Exponent>
          </dsig:RSAKeyValue>
        </dsig:KeyValue>
      </r:info>
    </r:keyHolder>
  </r:issuer>
</r:license>
```

An R-VDI UserGroups license should be contained within the R-VDI that it refers to (and possibly also within the corresponding P-VDI).

4.3.1.2 **UserGroups License for P-VDIs**

A P-VDI UserGroups License specifies all the rights of all Pictures Inc. user groups over a specific P-VDI.

The internal structure of this license is completely equivalent to the license for the R-VDI, described above. The main difference is that the governed resource(s) are now the P-VDI's resources rather than those of the R-VDI.

A P-VDI UserGroups License must be contained in its corresponding P-VDI.

4.3.1.3 UserGroups License for S-VDIs

An S-VDI UserGroups License specifies all the rights of all Pictures Inc user groups over a specific resource belonging to an S-VDI. This license will contain:

- An *Inventory* declaring the S-VDI resource(s) (represented by *did:Resource* elements), whose manipulation is regulated by the license.
- A *GrantGroup* defining all rights of the user who originally issued the S-VDI, over S-VDI resource(s). It will contain:
  - The *Principal*, to whom the rights are attributed (the user issuing the S-VDI).
  - Two *Grants*:
    - A *Grant* specifying the right to *play* the resource. This will contain:
      - A *Right* element of type *mx:play*.
      - A reference to the element (contained in the top *Inventory*), specifying the governed resource over which the right is being granted.
    - A *Grant* specifying the right to delete the resource. This will contain:
      - A *Right* element of type *mx:delete*.
      - A reference to the element, contained in the top *Inventory*, specifying the governed resource over which the right is being granted).
- The *Issuer* of the license, typically the issuer of the governed S-VDI.

The text box below shows an example of this kind of license.

```
<?xml version="1.0" encoding="UTF-8"?>
<r:license
  xmlns:acme="urn:acme"
  sx:profileCompliance="acme:example"
  xmlns:r="urn:mpeg:mpeg21:2003:01-REL-R-NS"
  xmlns:sx="urn:mpeg:mpeg21:2003:01-REL-SX-NS"
  xmlns:mx="urn:mpeg:mpeg21:2003:01-REL-MX-NS"
  xmlns:dsig="http://www.w3.org/2000/09/xmldsig#"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:mpeg:mpeg21:2003:01-REL-R-NS rel-r-profile.xsd">
  <r:inventory>
    <!-- Specification of the Governed Resource -->
    <r:digitalResource licensePartId="governedMetadataResource">
      <nonSecureIndirect URI="conv:vdi:VDI1:metadataResource"/>
    </r:digitalResource>
  </r:inventory>

  <!-- GrantGroup for the definition of the rights of the S-VDI Issuing User -->
  <r:grantGroup>

    <!-- Specification of the Principal who is entitled to the Rights over the Resource -->
    <r:keyHolder licensePartId="conv:user:John">
      <r:info>
        <dsig:KeyValue>
```

```
            <dsig:RSAKeyValue>
              <dsig:Modulus>Fa7wo6NYfm==</dsig:Modulus>
              <dsig:Exponent>AQABAA==</dsig:Exponent>
            </dsig:RSAKeyValue>
          </dsig:KeyValue>
        </r:info>
      </r:keyHolder>

      <!-- Definition of a Condition, the Validity Interval -->
      <r:validityInterval>
        <r:notBefore>2009-01-01T00:00:00</r:notBefore>
        <r:notAfter>2013-01-01T00:00:00</r:notAfter>
      </r:validityInterval>
      <r:grant>
        <!-- Definition of the right to play the governed Resource, to which the Principal is entitled -->
        <mx:play/>
        <r:igitalResource licensePartIdRef="governedMetadataResource"/>
      </r:grant>
      <r:grant>
        <!-- Definition of the right to delete the governed Resource, to which the Principal is entitled -->
        <mx:delete/>
        <r:igitalResource licensePartIdRef="governedMetadataResource"/>
      </r:grant>
    </r:grantGroup>

    <!-- Specification of the Issuer of the license, which is the original VDI owner -->
    <r:issuer>
      <r:keyHolder licensePartId="conv:user:John">
        <r:info>
          <dsig:KeyValue>
            <dsig:RSAKeyValue>
              <dsig:Modulus>Fa7wo6NYfm==</dsig:Modulus>
              <dsig:Exponent>AQABAA==</dsig:Exponent>
            </dsig:RSAKeyValue>
          </dsig:KeyValue>
        </r:info>
      </r:keyHolder>
    </r:issuer>
</r:license>
```

An S-VDI UserGroups License must be contained within its corresponding S-VDI.

4.3.2 User Membership License

4.3.2.1 User Membership License for Consumer Users

A User Membership License for Consumer Users certifies that a specific user belongs to the Consumer Users group.

The text box below shows an example.

```
<r:license
  xmlns:acme="urn:acme"
  sx:profileCompliance="acme:example"
  xmlns:r="urn:mpeg:mpeg21:2003:01-REL-R-NS"
```

```
xmlns:sx="urn:mpeg:mpeg21:2003:01-REL-SX-NS"
xmlns:mx="urn:mpeg:mpeg21:2003:01-REL-MX-NS"
xmlns:dsig="http://www.w3.org/2000/09/xmldsig#"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="urn:mpeg:mpeg21:2003:01-REL-R-NS rel-r-profile.xsd">
<r:grant>

  <!-- Specification of the Principal who exerts a Right (possession) over a Resource (property) -->
  <r:keyHolder licensePartId="conv:user:Mary">
    <r:info>
      <dsig:KeyValue>
        <dsig:RSAKeyValue>
          <dsig:Modulus>KtdToQQyzA==</dsig:Modulus>
          <dsig:Exponent>AQABAA==</dsig:Exponent>
        </dsig:RSAKeyValue>
      </dsig:KeyValue>
    </r:info>
  </r:keyHolder>

  <!-- Definition of the Right to which the Principal is entitled over the Resource -->
  <r:possessProperty/>

  <!-- The property which is possessed by the user -->
  <sx:propertyUri definition="conv:consumerUserGroupMember"/>

  <!-- License Revocation Freshness Period equal to one day-->
  <r:revocationFreshness>
    <r:priorToStart>P1D</r:priorToStart>
  </r:revocationFreshness>
</r:grant>

<!--The license Issuer (and signer), the ConvergenceUserManagementService.-->
<r:issuer>
  <dsig:Signature>
    <dsig:SignedInfo>
      <dsig:CanonicalizationMethod
        Algorithm="http://www.w3.org/TR/2001/REC-xml-c14n-20010315"/>
      <dsig:SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>
      <dsig:Reference>
        <dsig:Transforms>
          <dsig:Transform Algorithm="urn:mpeg:mpeg21:2003:01-REL-R-NS:licenseTransform"/>
        </dsig:Transforms>
        <dsig:DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
        <dsig:DigestValue>Jk9QbKOQCo941tTExbj1/Q==</dsig:DigestValue>
      </dsig:Reference>
    </dsig:SignedInfo>
    <dsig:SignatureValue>ABCqOhh5QQ==</dsig:SignatureValue>
    <dsig:KeyInfo>
      <dsig:KeyValue>
        <dsig:RSAKeyValue>
          <dsig:Modulus>aaaM4ccyzA==</dsig:Modulus>
          <dsig:Exponent>AQABAA==</dsig:Exponent>
        </dsig:RSAKeyValue>
      </dsig:KeyValue>
    </dsig:KeyInfo>
  </dsig:Signature>
  <r:details>
    <r:timeOfIssue>2012-01-01T00:00:00</r:timeOfIssue>
    <r:revocationMechanism>
```

```
        <r:revocationService>

          <!--This revocation service is represented with as a reference to a WSDL file -->
          <r:serviceReference>
            <sx:wsdlComplete xmlns:rs="convergence:user:management:service:revocation">
              <sx:wsdl>
                <r:nonSecureIndirect URI="http://www.conv.org/rswsdlfile.xml"/>
              </sx:wsdl>
              <sx:service>rs:aRevocationService</sx:service>
              <sx:portType>rs:aRevocationPortType</sx:portType>
            </sx:wsdlComplete>
          </r:serviceReference>
        </r:revocationService>
      </r:revocationMechanism>
    </r:details>
  </r:issuer>
</r:license>
```

All User Membership Licenses belonging to a specific user should be contained in her User VDI.

4.3.2.2 User Membership License for Editor Users

A User Membership License for Editor Users certifies that a specific user belongs to the Editor Users group.

The internal structure of this license is equivalent to that of the license described in the previous section. The only difference is that the *Principal* of the license will be declared to possess a property that places it in the Editors User Group, as shown in the text box below.

```
.......
    <!-- Definition of the Right to which the Principal is entitled over the Resource -->
    <r:possessProperty/>

    <!-- The property which is possessed by the user -->
    <sx:propertyUri definition="conv:editorUserGroupMember"/>
.......
```

4.3.3 Fractal License

Fractal Licenses are typically carried by P/S-VDIs only. Fractal Licenses specify which Fractals possess which rights over which parts of the VDI. Typically the Fractal(s) to which a VDI is published is given the right to:

- Match user queries or user subscriptions to specific portions of the VDI.
- Notify a specific set of users, when a VDI (a specific portion of the VDI) matches a user query or user subscription.

The resulting license is presented below.

```
<?xml version="1.0" encoding="UTF-8"?>
<r:license
  xmlns:acme="urn:acme"
```

```
sx:profileCompliance="acme:example"
xmlns:r="urn:mpeg:mpeg21:2003:01-REL-R-NS"
xmlns:sx="urn:mpeg:mpeg21:2003:01-REL-SX-NS"
xmlns:mx="urn:mpeg:mpeg21:2003:01-REL-MX-NS"
xmlns:dsig="http://www.w3.org/2000/09/xmldsig#"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="urn:mpeg:mpeg21:2003:01-REL-R-NS rel-r-profile.xsd">
<r:inventory>
  <!-- Specification of the Governed Resource(s) -->
  <r:digitalResource licensePartId="metadataResource1">
    <nonSecureIndirect URI="conv:vdi:VDI1:metadataRes1"/>
  </r:digitalResource>
  <r:digitalResource licensePartId="metadataResource2">
    <nonSecureIndirect URI="conv:vdi:VDI1:metadataRes2"/>
  </r:digitalResource>
  <r:digitalResource licensePartId="metadataResource3">
    <nonSecureIndirect URI="conv:vdi:VDI1:metadataRes3"/>
  </r:digitalResource>
</r:inventory>

<!-- GrantGroup for the definition of the rights of the VDI holding Fractal -->
<r:grantGroup>

  <!-- Variable which is employed to define the members peers of the VDI holding Fractal  -->
  <r:forAll varName="FractalXMember">
    <r:propertyPossessor>
      <sx:propertyUri definition="conv:FractalXMember"/>
      <r:trustedRootIssuers>
        <r:keyHolder licensePartId="ConvergencePeerMembershipManagementService">
          <r:info>
            <dsig:KeyValue>
              <dsig:RSAKeyValue>
                <dsig:Modulus>aaaM4ccyzB==</dsig:Modulus>
                <dsig:Exponent>AQABAA==</dsig:Exponent>
              </dsig:RSAKeyValue>
            </dsig:KeyValue>
          </r:info>
        </r:keyHolder>
      </r:trustedRootIssuers>
    </r:propertyPossessor>
  </r:forAll>

  <!-- Specification of the Principal who is entitled to the Rights over the Resource(s) -->
  <r:principal varRef="FractalXMembers"/>

  <!-- Definition of a Condition, the Validity Interval -->
  <r:validityInterval>
    <r:notBefore>2009-01-01T00:00:00</r:notBefore>
    <r:notAfter>2013-01-01T00:00:00</r:notAfter>
  </r:validityInterval>
  <r:grant>
    <!-- Definition of the right to perform a matching on the metadata contained in Resource 3,
         to which the Principal is entitled
    -->
    <sx:rightUri definition="conv:right:match"/>
    <r:igitalResource licensePartIdRef="metadataResource3"/>
  </r:grant>
  <r:grant>
    <!-- Definition of the right to notify users, of a match over the metadata of Resource 3,
```

```
            to which the Principal is entitled
        -->
        <sx:rightUri definition="conv:right:notify"/>
        <r:igitalResource licensePartIdRef="metadataResource3"/>
      </r:grant>
    </r:grantGroup>

    <!-- Specification of the Issuer of the license, which is the original VDI owner -->
    <r:issuer>
      <r:keyHolder licensePartId="conv:user:John">
        <r:info>
          <dsig:KeyValue>
            <dsig:RSAKeyValue>
              <dsig:Modulus>Fa7wo6NYfm==</dsig:Modulus>
              <dsig:Exponent>AQABAA==</dsig:Exponent>
            </dsig:RSAKeyValue>
          </dsig:KeyValue>
        </r:info>
      </r:keyHolder>
    </r:issuer>
</r:license>
```

In the license presented above all peers belonging to *FractalX* are given the right to:

- Perform matching operations on the metadata of VDI Resource 3
- Report matches to users

4.3.4 Fractal Membership License

A Fractal Membership License for an individual Peer certifies that the Peer belongs to a specific Fractal. Please notice that, according to considerations above, these Fractal Membership Licenses cannot be self-attributed to peers by themselves, whenever they join a Fractal (only applicable if security concerns were not relevant). All the same, they cannot be attributed by a possible fixed "watchdog" peer of the fractal, devoted to "secure" the fractal, because watchdog peers cannot in their turn be centralized and certified entities, given the highly dynamical process of fractal creation. This situation will be discussed in greater detail in section 5.5.

4.4 Summary of CONVERGENCE Approach to REL

CONVERGENCE's licensing scheme is based on the use of licenses to attribute rights over sets of resources to "collective" Principals.

These "collective" licenses are public knowledge and are delivered with the relevant VDIs.

Individual entities, such as users or peers, have their rights attributed to them, not by directly and individually expressing their possession of said rights, but by expressing the fact that they belong to specific collective Principals.

This approach makes it possible to reduce the number of licenses that need to be issued and managed, in the system, and simplifies the bulk attribution of rights to large groups of users.

4.5 Summary of new verbs required in the CONVERGENCE REL

In this section we briefly extend the analysis of the real world scenarios in chapter 3, examining more complex business scenarios that might require new verbs. The results of this investigation should be regarded as preliminary and may well change after further studies in work packages 8 and 9.

We start from the following Table, which summarizes the new verbs suggested by the discussion in the rest of this document. The second column provides a tentative definition of their semantics.

Verb	Definition
Match	To **compare** two different metadata sets
Notify	To **inform** a User or a Peer that an Event (i.e. the execution of another verb) has occurred
Aggregate	To **collect and aggregate** information from the resources of VDIs, complying with the requirements of existing legislation, obfuscating references to personal data and link-backs to specific individuals. In the case of publications or subscriptions this maps to the right to collect statistics about what is being published by users or most wanted search content.
License	To **exploit a resource for business purposes**, given to the principal by the issuer for specific rights over the resource.
Post	To **expose a resource** on an electronic mass media communication system (a web page, a mobile portal, a TV program guide on a broadcast channel) **for a third party** (everybody, subscribers to the electronic mass media communication system, etc.) **to play**

Table 2 — New REL verbs proposed by CONVERGENCE

Please notice we have added a row for the *License* Right, in addition to what has been discussed so far. Our study suggests it may be necessary to grant a *right to license*, i.e. the right given to a principal to make money from (or, in generic terms, to exploit) a given resource in question. In this case the issuer does not want to give the principal rights such as GovernedCopy or Play. However, she does want to give him the right to **sell** the Play right, provided he gives back a 50% share of his revenues.

A preliminary analysis of the MPEG-21 REL suggests it offers no easy way of implementing this scenario. We therefore propose to add a new *License* verb to the standard.

This analysis needs to be validated by means of on-going work of WP8 and WP9. For example, specific attention should be paid to sub-licensing scenarios, i.e. situations where the right that is sold to the principal by the issuer of the License, might in turn be License itself, resulting in the right for the principal to sub-license.

5 Techniques to implement REL in CONVERGENCE

5.1 Controlling the Matching Process

An issue that has been debated at length during the design of inter-VDI relationships is the governance of these relationships. Some of the use-cases raised concerns such as those illustrated by the example below

A manufacturer releases a VDI for a product to a supplier. The supplier signs a deal with a retailer to sell the product. The retailer will now create a VDI, advertising his sales conditions, linking back to the original product VDI through the kind of semantic relationship described in deliverable D4.1 [5]. How can the original manufacturer control such a relationship, when there is no direct business connection between the manufacturer and the retailer? In other words, how can the system forbid a black market reseller from linking to the manufacturer's resource VDI claiming to be the official reseller or setting the recommended price for the product in a given city?

At first sight it seems that it would be useful to provide the manufacturer with this sort of control. Such a mechanism, if technically feasible, would provide:

- Better control over the flow of information through the network;
- A new incentive for users to adopt the platform;
- An obstacle to false claims.

Further analysis suggests that users should be free to link their opinions (comments, annotations) to a VDI but that we might want to prevent them from claiming an official/business relationship with the VDI owner. In the case just described, a manufacturer would not be interested in a technology that allowed a third party to make false claims about his product – and provided no way to check the authenticity of the information provided. At the same time, however, strong controls could pose a danger for freedom of expression.

We therefore propose a solution that does not encourage censorship but which maintains the advantages of control. This solutions consists in giving to users the **option** of formulating subscription requests that filter out resources that do not come from certified sources.

In other words, we would provide a ranking and filtering complement to the Match TE (see [6]) making it possible to distinguish genuine items from items that people might have published in bad faith, matching only VDIs that are signed by a specific person/entity/company.

This solution would limit misbehaviour and limit some denial of service attacks, removing the incentive for dishonest users to flood the network with fake VDIs.

This kind of mechanism could be used to ensure that children will always download genuine cartoons and do not run the risk of reaching inappropriate content, and could be built into software as in parental control systems. This would not remove users' technical freedom to attach improper (or undesirable, or, conversely, critical) content to the resources of reputable brands, but it would make it possible to filter out this content when needed.

The following section describes technical aspects of this approach, and the specific issues they raise.

5.1.1 Searching for certified content

Consolidated tools such as Public Key Certificates (PKC), Attributes Certificates (AC) and digital signatures already make it possible to filter out uncertified content. For example, AC technology [8] allows a company to certify that a person/business is an authorized reseller for an area, and that it has the right to set the price of its products.
In the following example Samsam, the producer of the novel BrightStar tv, has certified Novelties Inc. as a *Certified_Partner* for all its products. In this setting *Certified_Partner* is an attribute associated to Novelties Inc. in an AC issued by Samsam.

Samsam creates and publishes a Product VDI for the TV, packaging the specifications and maybe a detailed brochure.

Novelties Inc. has subscribed to all Samsam products. Once it receives the Product VDI for the new line of tv sets, it creates a Reseller VDI that packages special offers, a map to find the Novelties Inc shop and a link to the website. A semantic relationship links the Reseller VDI to the Product VDI. Once this is done, Novelties Inc. creates and publishes a P-VDI letting users find the Reseller VDI, and advertising a recommended retail price. In other words, the P-VDI basically sets a price for the tv set and publishes it in the semantic overlay of CONVERGENCE.

The content of the P-VDI, is expressed as RDF statements (see deliverable D3.2), as shown below:
{
NOVELTIES_R-VDI *isResellerOf* BRIGHTSTAR_R-VDI,
BRIGHTSTAR_R-VDI *hasPrice* 300
}

In our solution, Novelties associates the first RDF statement with the AC they have received from Samsam. Thus the P-VDI is packaged as follows:

- RDF statements block
- Attributes Certificate for the *Certified_Partner* attribute
- Public Key Certificate (for instance issued by VeriSign) for Novelties' digital identity

This package, signed by Novelties. is sent to the peers of the semantic overlay.

User Nicola creates a subscription that asks: "Things that are *isResellerOf* BrightStar". CONVERGENCE matches **all** R-VDIs declaring the owner to be a resellers of the product.
But if Nicola formulates the query as follows: "Things that are *isResellerOf* BrightStar AND property *isResellerOf* certified by Samsam", and packages the S-VDI as follows:

- SPARQL equivalent of the above query
- PKC of Samsam trusted by Nicola (for instance issued by VeriSign)

This allows CONVERGENCE to filter out bogus resellers and match Novelties Inc. only.

In practice, Novelties has backed up its statement that it is a reseller (the *isResellerOf* property) with a certificate, from Samsam, for a somewhat different and specific attribute (i.e. *Certified_Partner*. Thus Samsam has **not** directly certified either the *isResellerOf* property, or the "NOVELTIES_R-VDI *isResellerOf* BRIGHTSTAR_R-VDI" statement. In other words Samsam is endorsing Novelties as a certified partner, not as a reseller (nor generic nor specific for reselling BrighStar products). It is Novelties' initiative to endorse that statement with the *Certified_Partner* AC. In other words, it is not Samsam who is establishing a **synonymy** between *Certified_Partner* and *isResellerOf*, but Novelties. It is the end-user who decides whether or not to accept the association.

Nicola will receive notifications that contain an *isResellerOf* property, which has been coupled, by the creator of the P-VDI, with a legitimate AC certificate issued by Samsam. He will thus know that:

- The P-VDI has been created by Novelties Inc.
- That it has not been tampered with
- That Novelties Inc. possess "Certified_Partner" certification from Samsam
- That it is it indeed Samsam Corp. that has issued the certificate
- That Novelties Inc. wishes to endorse its status as a reseller of BrightStar tv sets with said certificate

It is up to Nicola whether to accept the notification and to use the associated Reseller VDI (that then links back to the Product VDI). For this reason, the system flags the answer with a warning that the certificate is for attribute *Certified_Partner* while the declared associated property is *isResellerOf*. A properly designed Application should allow Nicola to check with the certified attributes with his own eyes, and check them against the claims that are being made.
Of course a simplified approach is possible, as explained later, when the user just asks for content certified by Samsam, without specifying a precise attribute to look for in the query. In this case the user does not need to confirm the matching of strings, because he has declared Samsam as trusted by him once and for all.

Reasoning along the same lines, Novelties Inc. could associate the same *Certified_Partner*, to the "BRIGHTSTAR_R-VDI *hasPrice* 300" statement. This would be perfectly legitimate.

In this scenario, if Nicola requests "BrightStar that *hasPrice* 300 AND property *hasPrice* certified by Samsam", he will get the Novelties Inc. resource again. In this case, however, the system shows that his request for a certified *hasPrice* is backed up by Novelties only by means of a vague "certified partner" attribute from Samsam, he might decide to discard Novelties as non-trustworthy. Or he might accept the Novelties claim. It is worth remembering that a match can occur in several peers of the same overlay fractal. There is thus a risk that the same subscription could yield **multiple potentially useless** results.

If we move a step further in the design process, we see that, it might be the interest of Samsam to create a service that can check what properties/relationships Samsam considers as equivalent to the *Certified_Partner* attribute. Nicola or an application working on his behalf might could then contact a service and find out that, for Samsam, *Certified_Partner* only implies *hasRetailPrice*.
In other words Samsam could deploy a private CDS-based service that expands a generic *hasRetailPrice* relationship into a *Certified_Partner* one and Nicola could add the service to the list of default CDS providers he uses in his private matching process. His private CDS can

then use Samsam's standards to expand his queries.

However, a public distributed matching system such as CONVERGENCE's semantic overlay has two goals:

- To avoid a private company like Samsam from taking control of the "public" CDS by dictating which one of its certified relationships match English words or ontology properties users might use in generic queries.
- To scale the matching process to a large number of peers, avoiding the need for a user to contact an external 3rd party CDS to validate attributes contained in certificates, when comparing publications and subscriptions.

In theory, the distributed nature of CONVERGENCE's middleware would make it possible to deploy a broad range of solutions. On the one hand filtering could be performed locally by the end-user device, which would have to handle a flood of answers from the network and filter them afterwards. On the other hand core peers could perform the filtering before results are notified to the end-user, using **public and common** dictionaries. Many intermediate solutions are also possible.

Filtering by peers in the overlay would obviously be desirable. In this way, however, the network would not be able to exploit the user's own trusted synonyms.
A conceptually straightforward solution to the problem would be to:

- Clearly distinguish public, common CDS repositories with respect to private CDS repositories
- Only use common CDS when expanding terms during distributed match
- Allow for a subset of private expansion rules be packaged together with the S-VDI
- Use those suggestions from the user, extracted from her private CDS repository, to complete the matching process.

Let us continue with the BrightStar example.

Alina has installed a private CDS repository on her device, sourcing expansion rules from Samsam's knowledge base. Now Alina formulates the same query Nicola formulated previously: "BrightStar that *hasPrice* 300 AND property *hasPrice* certified by Samsam". Her device packages the S-VDI as follows:

- SPARQL equivalent of the above query
- Samsam PKC trusted by Alina (for instance issued by VeriSign)
- snippet of expansion rule that declares: *Certified_Partner* only implies *hasRetailPrice*

Peers that receive this subscription can now filter out publications from Novelties Inc. declaring that *hasPrice* is backed up by *Certified_Partner* certification.

Another very important case, covered by this design, is the case when the user has no knowledge of the specific ACs Samsam has issued, and formulates a broad query such as: "Content that is endorsed by Samsam". This subscription does not match P-VDIs that are signed by Samsam, but P-VDIs signed by individuals who are holders of ACs issued by Samsam, regardless of the attribute string.

In summary, the solution described in this section guarantees that semantic expansion of

query terms, a crucial part of CONVERGENCE match technology, is secure against infection by private information which could distort the results provided to the general public, but still allows private expansions.

The crucial issue here is governance of CDS repositories. In other words, who decides when and **under what conditions a CDS repository is authoritative enough to be considered of public and common interest**.

As a concluding remark, we observe that the solution we propose confirms our decision that VDIs should be self-contained in terms of assertion of trust. S-VDIs and P-VDIs carry with them all certificates necessary for verification algorithms to run offline and in parallel in the core of the semantic network of CONVERGENCE. Distributed networks often have trade-offs between CPU time and the bandwidth required by the network overlay. The self-contained chain of trust within VDIs makes it possible to explore the terms of this trade-off, dropping any dependency on external services when certifying content.

This solution does not prevent anybody from linking to “official resources unless some conditions are met”, but lets subscribers verify that content is endorsed by sources they trust.

5.2 A Secure Environment for CONVERGENCE Technologies

In previous chapters we have demonstrated the power and flexibility of the REL. Obviously, however, the critical issue is the enforcement of the rights, which REL describes. The only way to achieve this is to run the whole Security TE, and possibly other middleware technologies, within a trusted sandbox on the devices acting as CONVERGENCE peers.

To better understand the issues, it is sufficient to imagine that an end-user downloads a photo VDI containing a licence to Play (view) an encrypted photo. The CONVERGENCE peer will decrypt and present the photo but, given that the middleware runs in an unprotected environment, it could also be tricked to Store, Adapt etc. the photo even though these rights have not been granted in the original licence.
Hence, a complete enforcement of licenses shall assume a trusted "sandbox" device. In order to enforce a license of a content provider it is **the provider** that has to trust the sandbox to honour licenses. The owner of the device, who wishes to "consume" the content, might, in fact, be glad if restrictive licenses were not honoured.

CONVERGENCE deployments that can count on peers equipped with fully-fledged TPM (Trusted Platform Module)-like secure hardware, will be able to check that the middleware behaves validly. Since management and presentation of VDIs is delegated to media engines at the middleware level, and mediated via the Security TE, it is possible to limit the user’s actions on resources she has no right to manipulate further.

On the other hand, in an open environment, the only small and limited - but very trustworthy - "sandbox" can be a user’s smartcard. In such a deployment, the smart card performs the crucial step in content decryption - namely the "unwrapping" of the content-key. It is of great interest for the project to explore how to support REL in this kind of scenario.

The security issue will be fully dealt with in next deliverables. Here we introduce a promising approach, which CONVERGENCE is exploring, based on so-called Attribute Based

Encryption (ABE).

5.3 State of the art of the ABE technology

ABE (Attribute Based Encryption) is an elegant form of key escrow allowing an issuer of content (or in general, the sender of a message) to encode specific access rules into his "ciphertext" (i.e. the encrypted content). ABE allows for an **attribute-based access structure** with logical gates like AND- and OR-gates (even threshold-gates), so that an issuer can encode complex **decryption policies** using these gates.
ABE can thus be described as a type of encrypted access control, where access control policies are either embedded in the user private keys or in the ciphertexts. An example (taken from [9]) will help clarify the approach.

A head FBI agent may want to encrypt a sensitive memo so that only personnel that have certain credentials or attributes can access it. For instance, the head agent may specify the following access structure for accessing this information: ((“Public Corruption Office” AND (“Knoxville” OR “San Francisco”)) OR (management-level > 5) OR “Name: Charlie Eppes”). By this, the head agent could mean that the memo should only be seen by agents who work at the public corruption offices at Knoxville or San Francisco, FBI officials very high up in the management chain, and a consultant named Charlie Eppes.

In this example, ABE techniques can guarantee that agents not satisfying the encoded attributes are not able to decrypt the memo.

The general "high-level" setup is quite easy to describe: To begin with, (like with IBE, Identity Based Encryption) there is the need for an **absolutely trustworthy master instance** (a CA, "Central Authority") which generates system parameters, holds one **master secret key** it never reveals to anyone, and issues **one common public key** for the ABE-instance at hand. (Note the difference with IBE, where each user is associated with a public key, or even several public keys, of his own). For each member of the ABE-instance, the CA generates an individual private key consisting of components with a list of attributes "woven" into them.
A "message" is encrypted by encoding a decryption policy described through an "access structure" built up from logical combinations of attributes. A user can decrypt the corresponding "ciphertext" if and only if the access tree used for its encryption matches the sets of attributes woven into his or her private key.

ABE also tackles a fundamental problem immediately rising from such techniques: namely **collusion**. Two or more users, whose individual private keys may not satisfy the access structure of a given ciphertext, may still collude ("combine their private key material") if the **union** of attributes allocated to them allows decryption.
A simple example: user1 is > 18 years old, but not a member of clubxy; user2 is a member, but < 18 years. Take a message encrypted under an access structure requiring any receiver to be member in the clubxy **and** > 18 years old. Neither user1 nor user2 satisfy both attributes, but joining together they do.

In reality, ABE is collusion resistant: the attributes "woven" into the private key components are "bound" to each other by encoding them with user-specific random numbers. Thereby attribute related components from **different** keys are incompatible. During decryption the embedded random numbers will only "cancel out" if compatible key components (with

identical embedded random numbers) are used.

In plain words, a private key can only decrypt a ciphertext if its underlying access structure is satisfied by the attributes associated to components within **one and the same** private key. As a consequence, disjoint colluding users cannot simply "merge" their private keys to obtain a "combined" private key capable of decrypting a ciphertext whose access structure would only be satisfied by their joint set of attributes, while none of the colluding users could alone decrypt it. This means that in a particular ABE-instance, there is no need for all recipients to get their individual private keys from the CA.

5.3.1 Existing ABE Schemes

The CP-ABE scheme by Bethencourt, Sahai and Waters, cited above, describes a scheme requiring **one** trusted CA for each ABE instance holding a master secret.
It is this CA that is responsible for governance and in particular for the derivation of decryption keys. For each such decryption key generation, the CA takes a set of attributes and outputs a key that identifies with that set. Given that the master key is required, only the CA can generate such keys.
To avoid offering a single point of attack, the CA can be split into different CAs using secret sharing techniques. Some schemes even lend themselves to homorphic encryption, making it possible to derive individual keys using secret shares and yet prevent recovery of the master secret.

Another very interesting ABE-variant uses a "decentralized" approach. "Decentralizing Attribute Based Encryption" ([10]), proposes a scheme which, unlike the schemes discussed so far, has no **single master secret.**

This decentralized approach allows much more flexibility: various trusted authorities can share attribute administration, with each authority retaining responsibility for "its" attributes. Global parameters exist, but they are all public, and there is no need for one master secret, or one single authority to be absolutely trustworthy. Users can encrypt designated messages to attribute-based access structures each involving various attributes from various authorities. Including users' trustworthy global IDs in key components makes the scheme collusion resistant.

Simply speaking, each user "collects" his key components from various authorities according to the attributes assigned to him. These key components are "masked" with the user's global ID. The mask cancels out during decryption of a message only if the encoded access structure is satisfied by attributes for which the receiving user possesses all necessary components (i.e. all masked with "his" ID, though possibly arising from different authorities). By contrast, colluding users can only combine components masked with different IDs, which will not cancel out.

5.4 Challenges in ABE Support to REL

In ABE, attributes are usually thought of as describing properties of recipients rather than attributes of the content itself. However the strong point of ABE is that basically any (binary) string can become an attribute. This is achieved by invoking a suitable hash-function mapping arbitrary strings to elements of a specific group.

The construction of the group involves mathematics; usually it is a prime order subgroup of an elliptic curve. Access trees have attributes assigned to their leaf-nodes, while each interior node works as a threshold gate determining the minimum number of its child nodes, which need to be satisfied.
Although, generically speaking, keys can be generated according to designated sets of attributes, further research from the project is needed to have ABE-based techniques integrated with the architecture we have so far.

5.4.1 Mapping complex REL statements to ABE

Mapping complex REL statements to ABE will require a specific method for **transferring REL descriptions to ABE**. A promising approach, that CONVERGENCE is exploring, is to allow users to access portions of VDIs only if they possess a certain set of credentials or attributes.
Encrypting specific parts of content - portions of VDIs – would make it possible to administer complex rights. For instance, we might freely access low-resolution photos, but need additional decryption to decode high-resolution components.

Below we show how we could use ABE to code the right to download a video in the FMSH application scenario:

User	**Conditions on user characteristics**
FMSH Analysts	(FMSH certified) AND (university level >= M1)
FMSH broadcaster	(FMSH certified) AND (localized in Paris, France)
FMSH VCO	(FMSH certified) AND (paid a 1 € fee)
INC	(INC certified)
Peruvian government	(Peruvian government certified) OR (localized in Peru)
SUMMARY	
	((FMSH certified) AND ((university level >= M1) OR (localized in Paris, France) OR (paid a 1 € fee))) OR (INC certified) OR (Peruvian government certified) OR (localized in Peru)

Success in this task demands a clean and simple interface that lets us deal with two fundamental aspects of ABE techniques in a technology-agnostic way:

- encrypt/decrypt operations.
- procedures for releasing attributes and for interaction with one (or more) authorities.

The interface should provide functionalities for temporal management of attributes, for instance to check if attributes are still valid and to issue replacements. Ideally it should also deal with attributes' revocation, but this is technically very much challenging since most ABE solutions do not support such concept.

The interface would encapsulate the specific details of the ABE solution chosen, hiding them from CONVERGENCE technologies.

5.4.2 Coping with a highly decentralized system

We have seen that there are two approaches for distributed ABE solutions: shared-secret among multiple authorities, and truly decentralized techniques.

The problem with shared-secret approaches is relates primarily to the management of the system, rather than to any intrinsic limitation of the shared-secret concept. There exist distributed key generation algorithms (see Pedersen [11]) that make it possible to easily build a shared-secret that is not known by the participants, and that allows verification in a way that participants cannot fake. The main problem lies in the steps authorities have to take to make the system work. This problem can only be alleviated by dynamic incremental setup techniques.

The decentralized approach from Lewko does not involve any shared-secret, hence avoiding setup procedures, and this is a strong point. We note, however, that any setup procedures required by a specific ABE scheme will be hidden from CONVERGENCE.

Furthermore, Lewko's decentralized technique suffers from three criticalities:

- It is extremely difficult to integrate revocation methods. To the best of our knowledge there are no known solutions for this. This is problem common to all decentralized techniques. The only schemes that offer limited opportunities for revocation are single authority schemes.

- There is no formal proof that Prime groups are secure. This can be proved for composite groups. However their huge size makes them impractical. Of course, this does not mean the scheme is insecure, only that a formal proof is missing.

- Performance and overhead. This technique requires the addition of a point to the elliptic group, **plus** three more points per attribute. Thus, with 10 attributes in a policy, this sums up to 31. The technique also requires a high number of elliptic exponentiations for each attribute. We are currently carrying on detailed investigations on this issue. However, our implementation is still not optimal and are our results are not yet final.

These three points, while not extremely critical, are a challenge for the distributed technique. This implies that, if well-thought and implemented shared-secret schemes compensate such criticalities, it could be worth paying the price of the required setup phases.

Another example shows the difficulty of choosing the right ABE scheme for CONVERGENCE. To the best of our knowledge there are currently no **short ciphertext**

techniques (i.e. schemes in which size is independent of the number of attributes) for multiple authority schemes. To date this has been achieved only for a single authority. It is almost sure that if and when such techniques are developed, extensions will apply to multiple authorities schemes and not to decentralized ones.

5.4.3 Fitting the validation to the smart-card

We are currently conducting detailed experiments to evaluate the performance of ABE schemes in traditional computational environments before transferring them to smart-cards. In our experience so far, contrary to the claim that the performance problem is ultimately a problem of pairing, exponentiation on elliptic curves is more costly than pairing.

Since, ultimately, the goal is to bind REL statements to the validation of license-conditions, these will have to be checked on-card prior to unwrapping. At least some license rights will need to be encoded in a card-verifiable certificate. Even with advanced schemes like IBE or ABE the smart card will still need to act as a "sandbox" attached to the consumer's device. The smart card operates as a secure repository for the user, allowing him to safely decrypt ABE-messages on machines with open software, or even on untrustworthy machines (e.g. in internet cafes) without having to worry about his private key being compromised.

5.4.4 Dealing with post-decryption rights

How should we deal with rights, such as "read only" / "read + modify" / "store", etc., which apply to the resource even after it has been decrypted?
Within this model, we cannot really prevent a photo from being distributed, manipulated, etc after decryption. Likewise, a video or soundtrack, once streamed, can easily be recorded and processed, even though such rights may not have been granted by a license. There is the chance of taking "a posteriori" actions (like tracing illegal distribution through embedding robust watermarks), but only after the damage has been done. We may also require a signature from everyone who uploads (possibly modified) content, but this would simply transfer the issue of enforcement to yet another level: how do we / can we prevent the upload of non-signed content?
Our conclusion is that only a trusted middleware running on trusted hardware platforms such as TPM(Trusted Platform Module)-based devices allows fully-fledged enforcements of rights, capable of capturing all the possibilities the REL is capable of describing.

5.5 Fractals of Trust

The core idea lying behind the CONVERGENCE publish/subscribe model is the distribution of peers to fractals, according to the content they are contributing to the system. So, when a peer publishes some content within a fractal, the content is propagated to a number of other peers in the fractal. The same holds for subscription. In other words a fractal is a set of peers sharing some common interests.
In this section we propose our approach to create a fractal based on criteria other than the users' interests (as expressed through the semantic descriptions of the resources), but also based on level of trust of a fractal. This is on-going work. The scheme described below is thus provisional. Future deliverables in WP3 and WP6 will provide more specifics.

5.5.1 Beyond Basics

Users' specification of the fractal of interest may be expressed in terms of a semantic keyword, a flat keyword or an operation on these keywords. For example, we can have the fractal MOVIE∩ACTION for action movies. However, the fractals are not necessarily bounded by the content circulated inside them; they can also be based on the properties of other peers, such as trust.

This is crucial not only for giving owners control over content, but also for balancing the need for performance against the need to protect content. If we assume that the owner encrypts the VDI using some license and then propagates it blindly over some fractals, selected only by content, some receivers will not be able to handle the VDI. Moreover, there may be peers that are allowed to have access to that content, but will never receive it (two reasons: (i) probabilistic protocols & (ii) limiting propagation depth). The technique we propose provides another means of limiting the number of target peers, contributing to the scalability of the protocol.

5.5.2 Fractals of Trust

The question that arises at this point is how we can define such a fractal. In our current approach, we consider the creation of a fractal to be dynamic and totally up to the user. In parallel we are working on automatic extraction based on the descriptive metadata in the P-VDI. For example, if a user says that the content he is about to publish is MOVIEZ, then, the MOVIEZ fractal is created. We should always keep in mind that fractals are just a virtual organization of peers and, hence, a peer may belong to more than one fractal.

Apart from the increased matching performance, another important property of a fractal of trust is that it allows a set of peers to create closed organizations, accessed by peers that correspond to certain criteria, keeping the symmetric nature of the system. This is achieved during the registration process, by protecting the fractal registry. At this point, we should observe that, when a peer requests to enter the fractal, it communicates with another peer and asks for a part of the fractal registry. If this registry is protected, then the peer will be able to use it only if it has the required properties.

Another issue is how to manage/enforce rights on the fractal registry and, thus, how to support fractals of trust. Obviously, a solution based on one certificate per fractal does not scale. This is then a case in which Attribute Based Encryption can be valuable.

Consider that each peer has a set of attributes, perhaps stored in a smart card. Thus, the creator of the fractal who initially owns the registry, encrypts it using the access policy for this particular fractal of trust. From this point on, only peers that have attributes satisfying this access policy will be able to read the registry entries they need to discover their neighbours, so only they will be able to access the fractal.

6 Concluding Remarks

The CONVERGENCE system aims to become a platform for the future Internet, enabling the secure distribution of resources and the realization of diversified and advanced business models in which resources transit through complex value chains. Digital resources are often covered by a license or contract, specifying a business relationship, including access and usage rights and conditions. Accordingly, the CONVERGENCE system has to offer a complete security infrastructure that enables the formulation and enforcement of complex and diversified access and usage rights and conditions.

Rights Expression Languages provide a means to declare those rights and conditions in a machine-readable manner. Most existing RELs have been developed for a specific application domain or do not aim at exerting subsequent enforcement of digital rights. Many lack formality in the definition of their elements and are thus unsuitable for automated operation.

The MPEG-21 part 5 standard can be seen as a highly flexible, general-purpose rights expression language, well suited to producing machine-readable assertions of rights. The ultimate aim of the standard is to kick-start technologies that correctly enforce access and usage rights over different types of digital resources and across a broad range of application domains.

Having already developed detailed application scenarios and considering the diversity of digital resources that it will be possible to distribute and transact within the CONVERGENCE system, the project has defined a CONVERGENCE governance and licensing scheme based on its own patterns and a sub-set of MPEG-21 part 5 REL elements. This scheme, whilst supporting all requirements already identified in terms of content protection and digital rights management, is sufficiently flexible to be extended to accommodate further needs. The way in which MPEG-21 REL elements will be used has been thoroughly explored. It should thus be possible to make future extensions to the scheme at low cost.

The project has formulated detailed proposals for an extension of the (already rich) set of REL verbs, identifying complex business cases that the current REL is unable to accommodate easily.

We are also engaged in an on-going effort to design novel integration patterns of advanced security technologies, making it possible to enforce at least the fundamental aspects of the CONVERGENCE REL scheme, taking into account the challenging trade-offs that arise when we abandon (maybe-impractical) requirements for fully-trusted hardware and software in favour of trust chains and lightweight smart-cards.

7 References and relevant literature

[1] Information technology — Multimedia Framework — Part 5: Rights Expression Language, ISO/IEC FDIS 21000-5:2003(E), July 2003.

[2] ODRL Website http://odrl.net/

[3] ccREL Website, http://creativecommons.org/ns

[4] OMA DRM Release 2.2 Website, http://www.openmobilealliance.org/Technical/release_program/drm_v2_2.aspx

[5] CONVERGENCE project. Deliverable D4.1. Preliminary Definition of the Versatile Digital Item. http://www.ict-convergence.eu/deliverables

[6] CONVERGENCE project. Deliverable D3.2. System architecture. http://www.ict-convergence.eu/deliverables

[7] CONVERGENCE project. Deliverable D2.2. Use cases and requirements for CONVERGENCE development work. http://www.ict-convergence.eu/deliverables

[8] IETF Network Working Group. An Internet Attribute Certificate Profile for Authorization. Request for Comments: 3281. April 2002. http://www.ietf.org/rfc/rfc3281.txt

[9] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-Policy Attribute-Based Encryption. In Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP '07). IEEE Computer Society, Washington, DC, USA, 321-334. DOI=10.1109/SP.2007.11 http://dx.doi.org/10.1109/SP.2007.11

[10] Allison Lewko and Brent Waters. 2011. Decentralizing attribute-based encryption. In Proceedings of the 30th Annual international conference on Theory and applications of cryptographic techniques: advances in cryptology (EUROCRYPT'11), Kenneth G. Paterson (Ed.). Springer-Verlag, Berlin, Heidelberg, 568-588.

[11] Torben P. Pedersen. 1991. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In Proceedings of the 11th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO '91), Joan Feigenbaum (Ed.). Springer-Verlag, London, UK, 129-140.