Project Number:	**FP7-257123**
Project Title:	**CONVERGENCE**
Deliverable Type:	**Report**
Dissemination Level	**Public**
Deliverable Number:	**D 8.1**
Contractual Date of Delivery to the CEC:	**31.08.2011 (date agreed with the PO; the original delivery date was 31/5/2011)**
Actual Date of Delivery to the CEC:	**31.08.2011**
Title of Deliverable:	**Plan for trials and experimentations**
Workpackage contributing to the Deliverable:	**WP 8**
Nature of the Deliverable:	**Report**
Editor:	**Richard Walker**
Author(s):	**Richard Walker, Andrea Pede, Barbara Benincasa (XIW), Andrea de Polo, Sam Minelli (Alinari), Leonardo Chiariglione (CEDEO), Andrea Detti, Nicola Blefari Melazzi, Stefano Salsano, Giuseppe Tropea (CNIT), Peter Stockinger, Valérie Legrand, Francis LeMaitre (FMSH), Iakovos S. Venieris, Charalampos Z. Patrikakis, Dimitra I. Kaklamani, Georgios Lioudakis, Vasso Giotopoulou (ICCS), Maria Teresa Andrade, Helder Castro (INESC), Heinrich Hussman, Alina Hang (LMU), Angelos-Christos Anadiotis (NTUA), Pantelopoulos Stelios (SIL), Panagiotis Gkonis (SIL), Mihael Tanase (UTI), José Ribas, Daniel Filipe Pereira Sequeira (WIPRO)**
Abstract:	**This deliverable describes plans for the evaluation of CONVERGENCE in four real-life business trials and in network experiments and simulations. An appendix provides specifications for the VDIs that will be used in the trials.**
Keyword List:	**Trials, evaluation, performance, reliability, focus groups, interviews, use cases**

Executive Summary

The goal of this document is to outline detailed plans for testing the VDI concept, the CONVERGENCE publish-subscribe model, the CONVERGENCE middleware, CONVERGENCE networking concepts, and CONVERGENCE applications. The tests will take place in two "tracks".

The first track will consist of four real world trials**,** providing input for the design of CONVERGENCE applications and the CONVERGENCE framework:

- Photos in the cloud and down to earth (Alinari)
- Videos in the cloud and analyses on earth (FMSH)
- Augmented Lecture Podcast (LMU)
- Smart Retailing (WIPRO/UTI)

These trials will be organized into three phases, each of which will implement a broader and richer set of functionality than the preceding phase. Design teams in WP7 will use feedback from earlier phases to refine the applications and workflows used in later ones. The final phase of the trials will consist of a "summative evaluation" assessing the acceptability and usefulness of CONVERGENCE for end-users, its potential for commercial exploitation and its longer-term social and economic impact. These trials will be based on a single server and a plain IP network.

The second track, organized in two phases, will consist of a series of **network experiments** and simulations. In the network experiments, VDIs will be distributed and transported across nodes on the CONVERGENCE content-centric network (CONET). Initial small-scale experiments, in the first phase, will use small-scale laboratory networks and simulations; larger ones, in the second phase, will use the OFELIA Wide Area Network, and simulations. Additional work will test the CONVERGENCE publish-subscribe infrastructure when implemented on top of the CONET.

Reviewers will observe that the description above differs from CONVERGENCE's original plan to deploy a content-centric network layer on top of IP (the *overlay* approach) and compare its performance against that of a content-centric network-layer (the *clean-slate* approach). This change is motivated by the introduction of a novel *integration* approach [15] that extends the IP layer with a new header option that makes IP content-aware. Given its advantages (see [15] and D3.2), we will focus our studies on this approach. We will test the integration approach: i) in a configuration in which content-centric functionality is deployed on end-nodes and the core network is formed of IP routers; ii) in a configuration in which all nodes support content-centric functionality; this is equivalent to completely changing the IP layer.

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

This document outlines detailed plans for testing the VDI concept, the CONVERGENCE publish-subscribe model, the CONVERGENCE middleware, CONVERGENCE networking concepts, and CONVERGENCE applications.

As described in the Description of Work, the tests will take place in two "tracks" (see Figure 1). The first track will consist of four **real world trials** providing input for the design of CONVERGENCE applications and the CONVERGENCE framework. The trials will evaluate the acceptability and usefulness of CONVERGENCE for end-users, its potential for commercial exploitation and its longer-term social and economic impact.

The second track will consist of a series of **network experiments** and simulations. In the network experiments, VDIs will be distributed and transported across nodes on the CONVERGENCE content-centric network (CONET). Initial small-scale experiments will use small-scale laboratory networks and simulations; larger ones will use the OFELIA Wide Area Network, and simulations. Additional work will test the CONVERGENCE publish-subscribe infrastructure when implemented on top of the CONET [1].

Each track will be divided into phases. In track 1, the first phase will consist of a "formative evaluation" in which real users test CONVERGENCE applications under laboratory conditions. The results of the trial will be used to refine the design of the applications ready for the second phase in which "friendly users" will test the applications in "real-life" conditions. Feedback from this phase will provide further feedback for application design. The acceptability and usefulness of CONVERGENCE will be tested in the third phase, which will consist of a preliminary summative evaluation, involving users from outside the CONVERGENCE consortium.

In track 2, the two phases of the study will investigate the CONVERGENCE APIs, CONVERGENCE network functionality, and CONVERGENCE performance, first using a small scale laboratory network and simulations (phase 1) and then using a much larger

---

[1] Reviewers will observe that the description above differs from CONVERGENCE's original plan to deploy a content-centric network layer on top of IP (the overlay approach) and compare its performance against that of a content-centric network-layer (the clean-slate approach). This change is motivated by the introduction of a novel integration approach [15] that extends the IP layer with a new header option that makes IP content-aware. Given its advantages (see [15] and D3.2), we will focus our studies on this approach. We will test the integration approach: i) in a configuration in which content-centric functionality is deployed on end-nodes and the core network is formed of IP routers; ii) in a configuration in which all nodes support content-centric functionality; this is equivalent to completely changing the IP layer. It is to be noted that the integration approach allows analyzing both the overlay and the clean-slate deployment, by properly configuring the network framework.

network involving multiple OFELIA sites and simulations (phase 2). An additional study will investigate the performance of CONVERGENCE publish/subscribe mechanisms.

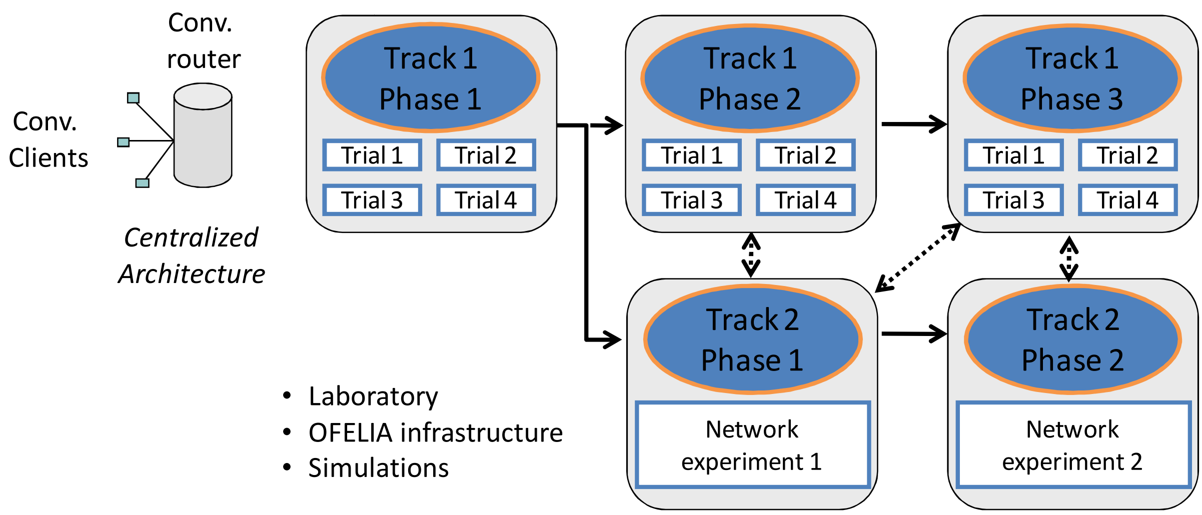

Figure 1 - Overview of trials and network experiments

The structure of this report reflects this underlying logic. The first half of the report (Chapters 2, 3 and 4) describes the trials planned in track 1. Chapter 2 describes the features to be implemented and tested in each phase of the trials. Chapter 3 describes the implementation of the trials (participants, procurement and deployment of hardware, installation of software, training and technical support etc.). Chapter 4 describes the methodology that will be used to evaluate trial results. The second half of the report (Chapter 5) describes the networking experiments planned in track 2 and presents a number of preliminary simulation results.

An Appendix contains technical specifications for the VDIs used in track 1 applications. These descriptions should be seen as an extension and revision of the corresponding, but less detailed, specifications provided in D7.1.

2 Trial Description and planning (track 1)

2.1 The scenarios

Track 1 will test CONVERGENCE concepts in four real-life business scenarios:

- Photos in the cloud and down to earth (Alinari) (trial 1)
- Videos in the cloud and analyses on earth (FMSH) (trial 2)
- Augmented lecture podcast (LMU) (trial 3)
- Smart retailing (WIPRO/UTI)[2] (trial 4)

A detailed description of the scenarios can be found in D2.2, which describes them from an end-user perspective, and in D7.1, which provides a technical specification of required tools, applications and Technology Engines.

2.2 Organization by phases

The track 1 trials will be organized into three phases, each of which will implement a broader and richer set of functionality than the preceding phase. Design teams in WP7 will use feedback from earlier phases to refine the applications and workflows used in later ones.

2.2.1 Small-scale pilot trial in laboratory conditions (phase 1)

Goal: collect user feedback to improve the design of workflows and applications.

End-Users: 3 separate groups of users, each containing at least one person for each role described in the scenario walkthrough.

Developers: all technical staff involved in the development of applications for the trial and/or interfaces to legacy software.

Expert review: the trial will be preceded by an expert review of the trial software by XiWrite (XS). XS will produce a list of recommended changes. The software development will prioritize the changes and implement those that it considers to be feasible and useful. XS will check that the changes have been carried out correctly.

The trials: each group will be tested separately for approximately one afternoon. The tests will take place under laboratory conditions. Each user will fill in a brief questionnaire before starting the trial. A representative of the partner organizing the trial will introduce the trial and its goals and the tasks users are supposed to perform. The users will then perform their tasks in the order foreseen in the walkthrough. A second representative of the partner will observe

[2] This trial will test a common set of applications on two separate sites managed by WIPRO and UTI respectively.

the users as they go through the walkthrough and take notes of any difficulties they encounter. After the trial the users will be debriefed in a focus group.

Evaluation of developers: at the end of the trial, the evaluation team will assess the experience of developers, while developing the trial applications. The assessment will be based on a focus group and/or in depth phone interviews.

2.2.2 Pilot trial with friendly users in real-life conditions (phase 2)

Goal: collect input for improving the design of workflows and applications.

Users: a single group of "friendly" users (users associated with the partner) large enough to represent a real life scenario (between 10 and 20 users). If necessary for purposes of realism the group may be divided into 2 or more subgroups.

Developers: as in phase 1.

Expert review: as in phase 1,

The trials: each trial will last long enough to test the usefulness and usability of the application under real life conditions (at least one month). At the end of the trial, participants in the trial will fill in a short questionnaire (designed by XS) and participate in a focus group.

2.2.3 Pilot trial with normal users (some from outside the Consortium) in real-life conditions (phase 3)

Goal: evaluate the usefulness, usability and marketability of the application under real-life conditions.

Users: a representative sample of the target population (12 users in the Alinari trial, >20 users in the other trials) including as many participants from outside the consortium as possible. If necessary for purposes of realism the group may be divided into 2 or more subgroups.

Developers: as in phase 1.

Expert review: as in previous phases.

The trials: each trial will last long enough to test the usefulness and usability of the application under real life conditions (at least one month). At the end of the trial, participants in the trial will fill in a short questionnaire (designed by XS) and participate in a focus group.

2.3 Schedule

The schedule for the trials is given in Table 1. Reviewers are asked to note that the first phase of the trials is scheduled to begin two months later than originally planned (see DOW). This delay is necessary to accommodate delays in WPs 3-6. The second phase of the trials is scheduled to begin with one month's delay. The third phase is scheduled to respect the original deadlines.

Phase	Activity	Period	
Planning	D 8.1	July – August 2011	M14-M15
Phase 1	Preparation	July –October 2011	M14-17
	Trial	November 2011	M18
	Partners Report	December 2011	M19
	Deliverable 8.2	December 2011	M19
Phase 2	Preparation	January 2012	M20
	Trial	February - March 2012	M21-M22
	Partners Report	April 2012	M23
	Deliverable 8.3	May-June 2012	M24-M25
Phase 3	Preparation	July 2012	M26
	Trial	August - November 2012	M27-M30
	Partners Report	December 2012	M31
	Deliverable 8.4 and 8.5	January-February 2013	M32-M33

Table 1 - Schedule for CONERGENCE real-life trials

2.4 Common tools

All the trials will use the set of common tools listed in Table 2. The rationale for the choice of these tools is described in D7.1. Implementation of the tools will be completed before the beginning of Phase 1.

Name	Developer
User Registration	SAFRAN - MORPHO
Content Registration	SAFRAN – MORPHO
Publish VDI	ICCS/SIL
Subscription to VDI	ICCS/SIL
Revoke VDI	INESC
Browse Event Report	UTI
Create Annotation	LMU

Table 2 - Common tools used in all trials

2.5 Applications

Each trial will test a set of trial-specific applications, specified in D7.1. Some of these applications will be available from the first phase of the trial. Others will become available only in later phases. Below we specify the applications to be tested in each phase.

2.5.1 Photos in the Cloud and down to Earth (Alinari)

Name	Phase 1	Phase 2	Phase 3	Developer
Photographer publishes Photo VDI	X	X	X	SIL
Rights holder unpublishes Photo VDI	X	X	X	SIL
Alinari subscribes and downloads photo	X	X	X	SIL
Photographer uploads Photo VDI	X	X	X	SIL
Alinari publishes new Photo VDI		X	X	SIL
Customer visits Alinari Museum		X	X	SIL
Customer purchases interesting photos		X	X	SIL

Table 3 - Trial-specific applications to be tested in each phase, trial 1

2.5.2 Videos in the Cloud and Analysis on Earth (FMSH)

Name	Phase 1	Phase 2	Phase 3	Developer
Video Material Owner creates and uploads a Video VDI	X	X	X	SIL/ICCS/MORPHO
Video Material Owner creates and injects a Publication VDI	X	X	X	SIL/ICCS
Analyst creates and injects a Subscription VDI	X	X	X	SIL
Video Material Owner revokes a Video VDI	X	X	X	SIL/ICCS
Video Material Owner revokes a Publication VDI	X	X	X	SIL/ICCS
Analyst and FMSH/ESCoM download Video	X	X	X	SIL/MORPHO
Analyst creates an Analysis and an Analysis VDI		X	X	SIL/ICCS/FMSH
Analyst creates and injects a Publication VDI		X	X	SIL/ICCS/FMSH
Video Channel Owner creates and injects a Subscription VDI		X	X	SIL
Analyst revokes an Analysis		X	X	SIL

VDI				
Analyst revokes a Publication VDI		X	X	SIL
Video Channel Owner creates and stores a Channel and a Channel VDI		X	X	SIL/ICCS/FMSH
Video Channel Owner posts an Analysis and an Analysis VDI		X	X	SIL/ICCS/FMSH
Video Channel User creates and injects a Subscription VDI		X	X	SIL
Video Channel Owner creates and injects a Publication VDI		X	X	SIL
Video Channel Owner unposts an Analysis VDI		X	X	SIL
Video Channel Owner revokes a Publication VDI		X	X	SIL
Video Channel User browses an analysis on a channel		X	X	SIL /FMSH

Table 4 - Trial-specific applications to be tested in each phase, trial 2

2.5.3 Augmented Lecture Podcast (LMU)

Name	Phase 1	Phase 2	Phase 3	Developer
Lecturers create slide/video/lecture podcast VDI	X[3]	X[4]	X[3]	LMU
Lecturers publish slide/video/ lecture podcast VDI	X	X	X	LMU/ICCS/SIL
Lecturers update slide/video/lecture podcast VDI	X	X	X	LMU/ICCS/SIL
Lecturers revoke slide/video/lecture podcast VDI		X	X	LMU/INESC
Students search for slide/video/lecture podcast VDI		X	X	LMU
Students retrieve slide/video/lecture podcast VDI by their VDI/Sequence identifier		X	X	LMU
Students download		X	X	LMU

[3] With manually generated metadata

[4] With CDS-support generated metadata

slide/video/lecture podcast VDI				
Students subscribe to slide/video/lecture podcast VDI		X	X	LMU/ICCS/SIL
Students register to the service		X[5]	X[6]	LMU
Students search for slide/video/lecture podcast VDI	X[7]	X[8]	X[7]	LMU
ALP application dowloads video and slides from CoNet and displays them in sync.	X[9]	X[8]	X[10]	LMU
Students subscribe to slide/video/lecture podcast VDI		X	X	LMU/ICCS/SIL
Students create and publish Annotation VDI	X[11]	X[12]	X[11]	LMU
Students revoke Annotations		X	X	LMU/INESC
Students set expiry date to Annotations	X	X	X	LMU/INESC
Lecturers browse event report			X	LMU/UTI

Table 5 - Trial-specific applications to be tested in each phase, trial 3

5 Username/password based

6 Add-on: smartcard-based solution

7 Preliminary semantic search

8 Fully-fledged semantic search

9 Download

10 Tentative streaming

11 No privacy options

12 With privacy options

2.5.4 Smart Retailing (WIPRO/UTI)

Name	Phase 1	Phase 2	Phase 3	Developer
Creation and storage of VDI	X	X	X	UTI/WIPRO
Creation and injection of Publication VDI		X	X	UTI/WIPRO
Unpublish VDI			X	UTI/WIPRO
Creation and injection of Subscription VDI		X	X	UTI/WIPRO
Browse VDI	X	X	X	UTI/WIPRO
Browse Event Reports			X	UTI
Generate statistics			X	UTI

Table 6 - Trial-specific applications to be tested in each phase, trial 4

3 Preparing the trials

3.1 Overview

Executing the trials requires a number of preparatory actions:

- Creation of development teams
- Creation of management, technical support and evaluation teams
- Procurement and/or deployment of hardware
- Creation of interfaces to legacy software (only in some trials)
- Installation of software and services
- Recruitment and training of participants

In what follows, we describe the specific steps planned for each trial.

3.2 Creation of development teams

3.2.1 Overview

The partners have completed the definition of the developer teams required for the trials. The role of the developers will be to develop the applications for the trials and (where necessary) to implement and test interfaces to legacy software. The tables below define the teams for each trial.

3.2.2 Photos in the Cloud and down to Earth (Alinari)

Name	Partner	Role
Angelo Difino	CED	Adaptation of the generic Convergence server environment to the trial's needs. Support to integration of middleware in peers used in the trial
Leonardo Lemmi	ALI	Data porting and integration (from Alinari toward Convergence middleware).
Panagiotis Gkonis	SIL	Coordination of SIL's development team for the deployment of applications described in Table 3

3.2.3 Videos in the Cloud and Analysis on Earth (FMSH)

Francis Lemaitre	FMSH	Development activities by ESCoM-FMSH will focus on: - Developing 2 technical engines for converting ontologies and analyses into OWL format thus allowing interaction with the CDS
Richard Guerinet		

		- Implementing those technical engines in FMSH's current (non-CONVERGENCE) Video Analysis Application - Updating the existing FMSH Video Channel Application to allow posting and browsing analyses on video channels - Deploying video channels
Panagiotis Gkonis	SIL	Coordination of SIL's development team for the deployment of applications described in Table 4

3.2.4 Augmented Lecture Podcast (LMU)

Alina Hang	LMU	Development activities by LMU will focus on: - Implementation of a podcast creation application using CONVERGENCE engines - Implementation of the augmented lecture podcast application using CONVERGENCE engines - Implementation of a general tool to create annotations within CONVERGENCE

3.2.5 Smart Retailing (WIPRO/UTI)

Daniel Sequeira	WIPRO	Development activities by WIPRO will focus on: - Customization of a Point-of-Service (POS) application, equipping it with CONVERGENCE engines - Customization of a Retail Merchandise System product's database according to CONVERGENCE needs
Lucian Corlan	UTI	Development activities by UTI will focus on: - Analysis, design and development of the application - Coordination with the retailer for the deployment of the system at his premises

3.3 The management team

3.3.1 Overview

The partners have completed the definition of the management, technical support and evaluation teams for the trials. The tables below define the teams for each trial.

3.3.2 Photos in the cloud and down on earth (Alinari)

Name of person responsible for trial	Andrea de Polo
Names of staff providing technical support and training	Leonardo Lemmi, Angelo Difino

Names of staff carrying out evaluation protocols (focus groups etc.)	Francesca Tavanti, Andrea de Polo, Sam Minelli, Rita Scartoni and other staff, depending on the number of focus groups

3.3.3 Videos in the Cloud and Analysis on Earth (FMSH)

Name of person responsible for trial	Peter Stockinger
Names of staff providing technical support and training	Francis Lemaitre, Richard Guerinet
Names of staff carrying out evaluation protocols (focus groups etc.)	Francis Lemaitre, Peter Stockinger, Elisabeth de Pablo

3.3.4 Augmented Lecture Podcast (LMU)

Name of persons responsible for trial	Prof. Dr. Heinrich Hußmann Alina Hang
Names of staff providing technical support and training	Alina Hang
Names of staff carrying out evaluation protocols (focus groups etc.)	Alina Hang
Note	Alina Hang will provide first level support to users for technical support. Where necessary she will use second level support from other partners

3.3.5 Smart Retailing (WIPRO)

Name of person responsible for trial	José Ribas, Daniel Sequeira
Names of staff providing technical support and training	José Ribas, Daniel Sequeira
Names of staff carrying out evaluation protocols (focus groups etc.)	José Ribas, Daniel Sequeira

3.3.6 Smart Retailing (UTI)

Name of person responsible for trial	Mihai Tanase
Names of staff providing technical support and training	Lucian Corlan

Names of staff carrying out evaluation protocols (focus groups etc.)	Lucian Corlan

3.4 Procurement and/or deployment of hardware

3.4.1 Hardware used in all trials

The hardware listed below will be provided by CNIT and will be used in all the trials.

Phase	**Backend hardware: servers etc.**	**End user hardware (PCs, laptops etc.)**	**Plans for installing hardware**
1	CONET Server	None	The backend hardware is already installed at CNIT site.
2	CONET Server	None	The backend hardware is already installed at the CNIT site.
3	CONET Server	None	The backend hardware is already installed at the CNIT site.

3.4.2 Photos in the Cloud and down to Earth (Alinari)

Phase	**Backend hardware: servers etc.**	**End user hardware (PCs, laptops etc.)**	**Plans for installing hardware**
1	Database server	Laptops (OS: Windows 7 Professional) PC desktops (OS: Windows 7 Professional)	HW is already installed at ALI location
2	Database server	laptops/PC desktops	HW is already installed at ALI location
3	Database server	Laptops/PC desktops	HW is already installed at ALI location

3.4.3 Videos in the Cloud and Analysis on Earth (FMSH)

Phase	Backend hardware: servers etc.	End user hardware (PCs, laptops etc.)	Plans for installing hardware
1	Streaming server Database server Web server Application server	PC, laptop (Windows)	ESCoM-FMSH servers and PC are already installed. Application server will be installed by ICCS/SIL in Athens
2	Streaming server Database server Web server Application server	PC, laptop (Windows)	Application server will be installed by ICCS/SIL/FMSH in Paris
3	Application server	Smart cards and card readers	All smart cards and readers will be supplied by Morpho

3.4.4 Augmented Lecture Podcast (LMU)

Phase	Backend hardware: servers etc.	End user hardware (PCs, laptops etc.)	Plans for installing hardware
1	*Not needed*	1 PC (windows) 1 Laptop (OS X)	The end-user hardware is already available at LMU site.
2	Application Server	1 PC (windows)	The backend hardware will be installed at LMU site. The end-user hardware is already available at the LMU site.
3	Application Server Streaming Server	1 Laptop (windows) 1 Smartcard Readers + Smartcards	The backend hardware is already installed at the LMU site. The end-user hardware is already available at LMU site.
Note	In the third phase users will also be able to use their own hardware (PC, laptop)		

3.4.5 Smart Retailing (WIPRO)

Phase	Backend hardware: servers etc.	End user hardware (PCs, laptops etc.)	Plans for installing hardware
1	RMS server (linux) POS server (linux) Oracle Application Server (linux)	1 IBM SurePOS (retailer/windows) 1 Laptop (retailer/windows)	All already installed at WIPRO location
2	RMS server (linux) POS server (linux) Oracle Application Server (linux)	1 IBM SurePOS (retailer/windows) 1 Laptop (retailer/windows) 1 Smartphone	All already installed at WIPRO location (or external data center provider)
3	RMS server (linux) POS server (linux) Oracle Application Server (linux)	1 IBM SurePOS (retailer/windows) 1 Laptop (retailer/windows) Smart card, Smart card reader	Only the smart card and the smart card reader are not installed

3.4.6 Smart Retailing (UTI)

Phase	Backend hardware: servers etc.	End user hardware (PCs, laptops etc.)	Plans for installing hardware
1	JBoss Application Server Oracle Database Server	3 PC/Laptop (Windows or Linux)	The backend hardware is already installed at the UTI site. The end-user hardware is already available at the UTI site.
2	JBoss Application Server Oracle DataBase Server	10 PC/Laptop (Windows or Linux)	The backend hardware is already installed at the UTI site. The end-user hardware is already available at the UTI site.
3	JBoss Application Server Oracle DataBase Server	10 PC/Laptop (windows) 1 Tablet PC (Linux) 1 Smartphone (Android) RFID tags and RFID Readers	The backend hardware is already installed at the UTI site. The end-user hardware is already available at the UTI site.

3.5 Legacy software

3.5.1 Photos in the Cloud and down to Earth (Alinari)

Name of software	AIM (Alinari Image Management) photo image retrieval Application
Function	Image management, photo annotation and advanced search and retrieval
Role in trial	Browsing and downloading photographs to the end user computer
Note	Web access and browsing through most of browsers (is. Safari, Firefox, Explorer) on Mac and PC

Name of software	VENUS (Virtual musEum Navigation picture accesS) content access, sharing and retrieval Application
Function	Allow museum visitor to capture, share, integrate, print, enrich, publish content
Role in trial	Browsing, enriching and downloading photographs to the end user device
Note	Web access interface

3.5.2 Videos in the Cloud and Analysis on Earth (FMSH)

Name of software	ESCoM-FMSH Video Analysis Application
Function	Virtual segmentation and analysis of a video
Role in trial	Interacting with the CDS to define metadata for inclusion in an Analysis VDI
Note	Standalone Windows app.

Name of software	ESCoM-FMSH Video Channel Application
Function	Managing and viewing a video channel
Role in trial	Posting an Analysis on a Channel Browsing an Analysis on a Channel
Note	Web app., running on any Web Browser / any OS

3.5.3 Augmented Lecture Podcast (LMU)

Name of software	Internet Browser
Function	Search/display/retrieve/download VDIs
Role in trial	Download lecture podcasts and components to the local device of the user
Note	Standalone Windows program + plugin for CONVERGENCE

3.5.4 Smart Retailing (WIPRO/UTI)[13]

Name of software	Oracle Retail Point-of-Service 13.3 Client
Function	Register products' sales to customers
Role in trial	Create new products VDIs: - Associate a customer to a product - Associate a serial number to a product - Generate warranties
Note	Java app

Name of software	Retailer's system
Function	Manage and control the products information
Role in trial	Used as a data source for importing data into CONVERGENCE
Note	-

3.6 Software and services to be installed

3.6.1 Software used in all trials

The software described below will be installed and managed by CNIT.

[13] In addition to the systems described in this section, the WIPRO/UTI team is also considering the use of other legacy systems. These will be specified in a later deliverable.

Phase	Backend software/services	End-user software	Plans for installing software
1	CONET Server	CONET TE	CONET Server is already installed at Roma Tor Vergata, CONET TE is ready to be installed on end user devices at the start of the phase
2	CONET Server	CONET TE	CONET Server is already installed at Roma Tor Vergata, CONET TE is ready to be installed on end user device at the start of the phase
3	CONET Server	CONET TE	CONET Server is already installed at Roma Tor Vergata, CONET TE ready to be installed on end user devices at the start of the phase

3.6.2 Photos in the Cloud and down to Earth (Alinari)

Phase	Backend software/services	End-user software	Plans for installing software
1	Oracle Java SDK Apache Tomcat Microsoft SQL server/Postgres	Internet Browser	All server software will be installed and updated remotely by CED via ssh.
2	Oracle Java SKD Apache Tomcat Microsoft SQL server/Postgres	Internet Browser	All server software will be installed and updated remotely by CED via ssh.
3	Oracle Java SKD Apache Tomcat Microsoft SQL server/Postgres	Internet Browser	All server software will be installed and updated remotely via ssh by CED.

3.6.3 Videos in the Cloud and Analysis on Earth (FMSH)

Phase	Backend software/services	End-user software	Plans for installing software
1	Apache Tomcat/JAVA SDK/Hibernate Microsoft SQL Server	Videos Manager Video Analysis Application Java Runtime	Back-End software will be installed by SIL in Athens. End-user software is already installed

		Environment Microsoft .Net Framework 4	
2	Apache Tomcat/JAVA SDK/Hibernate Microsoft SQL Server Microsoft IIS 7	Videos Manager Analyses Manager Channel Content Manager Channel Browser Video Analysis Application Video Channel Application Java Runtime Environment Microsoft .Net Framework 4	Back-End software will be installed by SIL in Paris. End-user software is already installed
3	Apache Tomcat/JAVA SDK/Hibernate Microsoft SQL Server Microsoft IIS 7	Videos Manager Analyses Manager Channel Content Manager Channel Browser Video Analysis Application Video Channel Application Java Runtime Environment Microsoft .Net Framework 4	Back-End software will be installed by SIL in Paris. End-user software is already installed
Note	Videos Manager, Analyses Manager, Channel Content Manager and Channel Browser are web applications running with Java Runtime Environment. Video Analysis Application and Video Channel Application are legacy software developed by FMSH.		

3.6.4 Augmented Lecture Podcast (LMU)

Phase	Backend software/services	End-user software	Plans for installing software
1	Tomcat for JSP MySQL Database JAVA	Podcast Application Podcast Creator Internet browser	Back-End and end-user software will be installed by LMU.

2	Tomcat for JSP MySQL Database JAVA	Podcast Application Podcast Creator Internet browser	Back-End and end-user software will be installed by LMU.
3	Tomcat for JSP MySQL Database JAVA	Podcast Application Podcast Creator Internet browser	Back-End software will be installed by LMU. End-user software will be installed on Students / Lecturers own machines with support by LMU.

3.6.5 Smart Retailing (WIPRO)

Phase	Backend software/services	End-user software	Plans for installing software
1	POS server RMS server Oracle Application Server Oracle Database	Oracle Retail Point-of-Service 13.3 Client Oracle Retail Merchandise System 13.3	Back-End and end-user software will be installed by WIPRO.
2	POS server RMS server Oracle Application Server Oracle Database	Oracle Retail Point-of-Service 13.3 Client Oracle Retail Merchandise System 13.3	Back-End and end-user software will be installed by WIPRO.
3	POS server RMS server Oracle Application Server Oracle Database	Oracle Retail Point-of-Service 13.3 Client Oracle Retail Merchandise System 13.3	Back-End and end-user software will be installed by WIPRO.

3.6.6 Smart Retailing (UTI)

Phase	Backend software/services	End-user software	Plans for installing software
1	Java EE JDK JBoss Application Server Oracle Database Server	Internet Browser	Backend software is already installed. End-user software is legacy software embedded in the OS.

2	Java EE JDK JBoss Application Server Oracle Database	Internet Browser	Backend software is already installed. End-user software is legacy software embedded in the OS.
3	Java EE JDK JBoss Application Server Oracle Database	Internet Browser	Backend software is already installed. End-user software is legacy software embedded in the OS.

3.7 Recruitment of participants

3.7.1 *Photos in the Cloud and down to Earth (Alinari)*

Phase	Category of user	#	Description of how users will be recruited
1	Photographers	2	Alinari will contact users who have already participated in previous trials of Alinari services
	People from cultural heritage sector	6	
	Students	2	
	Managers	2	
2	Photographers	2	
	People from cultural heritage sector	6	
	Students	2	
	Managers	2	
3	Photographers	2	
	People from cultural heritage sector	6	
	Students	2	
	Managers	2	
Note	The Alinari trial will include only 12 participants, as described above. All participants (except the students) will be highly skilled. Alinari expects that the number of participants will be sufficient to provide valuable feedback for designers and to evaluate the overall usefulness and value of the application		

3.7.2 *Videos in the Cloud and Analysis on Earth (FMSH)*

Phase	Category of user	#	Description of how users will be recruited
1	Video Material Owners	3	The FMSH user population will consist of: - 2 research engineers from ESCoM-FMSH staff, playing the role of Video Material Owners and Video Channel Users - 1 university lecturer from ESCoM-FMSH staff, playing the role of Video Channel Owner
2	Video Material Owners	3	
	Analysts	10	
	Video Channel Owner	2	
	Video Channel Users	12	

3	Video Material Owners	3	- 1 doctoral researcher from ESCoM-FMSH staff, located in Peru, playing the role of Video Channel Owner - 1 Peruvian anthropologist, outside of consortium, playing the role of Video Material Owner - 10-15 students in intercultural communication, from outside the consortium, playing the role of Analysts and Video Channel Users. In addition to these participants, focus groups will also involve partners from outside the Consortium: - Researchers and institutions owning video material (BnF, INA, Museum of Quay Branly, EHESS, etc.) - Researchers and teachers interested in building video channels (Azerbaijani Cultural Heritage, History of Mathematics, Archaeological Heritage in Seine-Saint-Denis, etc.) All planned users in the trials are already in contact with FMSH staff. Therefore, there will be no need for a special recruitment effort.
	Analysts	10	
	Video Channel Owner	2	
	Video Channel Users	12	
Note	Participants in the trial will test the application in their normal day-to-day activities. Thus the ESCoM-FMSH staff who will test the application are already engaged in the management of audiovisual materials; the students in the trial already have to produce video analyses for their lessons; the Peruvian Anthropologists already own videos.		

3.7.3 Augmented Lecture Podcast (LMU)

Phase	Category of user	#	Description of how users will be recruited
1	Students	3	Since our main target groups are students, we will reach them through the lectures and courses they attend. We will also contact them via email and through different communication platforms provided by the university. Lecturers will be contacted by email or personally.
	Lecturers	3	
	Developers	1	
2	Students	12	
	Lecturers	8	
	Developers	1	
3	Students	12	
	Lecturers	8	
	Developers	1	
Note	In our trial, the role of “lecturer” will mainly be played by teaching assistants responsible for organizational matters (e.g. providing learning materials to students).		

3.7.4 Smart Retailing (WIPRO)

Phase	Category of user	#	Description of how users will be recruited
1	Retailer	1	This trial will be executed in a controlled environment. The users will be WIPRO staff, not involved in the development of the application, who will simulate the interaction between retailers and customers.
	Customers	3	
2	Retailer	1	
	Customers	3	
3	Retailer	1	
	Customers	15	

3.7.5 Smart Retailing (UTI)

Phase	Category of user	#	Description of how users will be recruited
1	Retailer (simulated)	1	Participants in this phase of the trial will be members of UTI staff not involved in CONVERGENCE.
	Customers on the CONVERGENCE web	3	
2	Retailer	1	In this phase the participants in the trial will be staff members from a retailer already using UTI systems. Retailer staff will play the role of the retailer and simulate the role of customers
	Customers on the CONVERGENCE web	10	
3	Retailer	1	In this phase, retailer staff will play the role of the retailer. Customers will be real customers selected from the retailer customer database and invited to participate.
	Customers on the CONVERGENCE web	10	
	Customers in store	15	
Note	The first phase of the trial will take place on UTI premises. Integration with the retailer system will be simulated. In the second phase the smart retailing application will be integrated with the actual retailer system. The trial will focus on web functionality. The third phase of the trial will take place on retailer premises and will test the full range of web and in-store functionality described in D2.2 and D7.1.		

4 Evaluation Methodology

4.1 The goal of the evaluation

In phases 1 and 2, the main goal of the evaluation will be to collect user and developer feedback to refine the design of applications and workflows. This is what is known as *formative evaluation.* Phase 3 will consist of a preliminary *summative evaluation* – designed to measure the usefulness and usability of CONVERGENCE-based applications for potential users. The specific aims of the evaluation may be summarized as follows

4.1.1 Formative evaluation (phases 1 and 2)

- Identify technical, performance and usability issues that prevent users from making effective use of CONVERGENCE applications
- Identify user attitude/expectations that facilitate/hinder take-up of CONVERGENCE applications
- Identify new applications functionality/modifications that would increase the value of applications to users
- Identify critical issues in the deployment of CONVERGENCE hardware/software and in integration with legacy software
- Identify ways in which CONVERGENCE applications could improve current workflows
- Identify issues in the design of the CONVERGENCE middleware (CoMid) that hinder the development of applications
- Identify new middleware functionality that would facilitate the development of applications

4.1.2 Summative evaluation (phase 3)

- Measure overall end-user satisfaction
- Measure end-user satisfaction with specific aspects of CONVERGENCE applications (workflow logic, implementation of specific functions, usability, graphics interface etc.).
- Characterize the impact of CONVERGENCE applications on the user's work
- Characterize the business impact of CONVERENCE applications
- Compare user opinions of CONVERGENCE and of alternative solutions
- Elicit user opinions on broader societal issues related to CONVERGENCE (access to media, privacy, other issues identified by users)
- Measure overall developer satisfaction
- Measure developer satisfaction with specific aspects of CONVERGENCE middleware (design and functionality of APIs, documentation, tools)
- Characterize patterns of use (user characteristics, frequency of use, variations in use over time, drop up rates).

- Measure software stability on the front-end and the back-end (numbers of reported malfunctions).
- Measure software performance on the front-end and the back-end (speed of response to user commands, speed of search, quality of search results)
- Measure the time required to develop novel application functionality.

4.2 Selection criteria for trial participants

As far as possible, end-users in the trial will be representative of the target population for specific applications. Every effort will be made to ensure the gender, age and skill mix appropriate for specific applications. All trial participants will be aged at least 20.

4.3 Evaluation of the CONVERGENCE trials

The trials will test four different scenarios, at three different points in time (see phases plan in the previous chapter). Evaluation data will be collected and summarized in three reports, one at the conclusion of each experimental phase. Data from the reports in phases 1 and 2 will feedback into the development process.

In each phase the evaluation will be organized as follows:

- **At the beginning of applications development**
  Developers will be asked to fill in a brief questionnaire describing their personal characteristics, capabilities and expectations.
- **At the beginning of the trial**
  The partner responsible for the organization of the trial will introduce participants to the aims of the project, the technology involved, and the specific aims of the trial. Participants will be asked to fill in a brief questionnaire describing their personal characteristics, capabilities and expectations.
- **At the end of the trial**
  Each end-user in the trial will be asked to fill in a questionnaire providing information on:
    - His/her overall satisfaction with the applications tested in the trial
    - His/her satisfaction with specific application functionality
    - His/her satisfaction with the overall usability and performance of the applications
    - The impact of the applications on his/her work
    - His/her opinion on the potential business impact
    - Technical, usability and performance issues encountered during the trial
    - Suggestions for new functionality/modifications in functionality
    - Other suggestions for improvement

  After analysis of questionnaires, end-users will be invited to take part in a focus group designed to elicit specific suggestions for improvement.

Developers will also be asked to fill in a questionnaire. The questionnaire will elicit information on:

- The ease of developing applications using the CONVERGENCE middleware (CoMid)
- Technical/performance issues related to the middleware
- Technical/performance issues related to the applications
- Suggestions for improvement

Use of the system and dropout rates will be computed from data contained in system log files.

4.4 Tools for CONVERGENCE evaluation

The tools used in the evaluation will be designed by XiWrite (XS) and revised at the end of each phase. The table below summarizes the tools that will be used during the evaluation and the points at which they will be deployed.

	Tool	Administered	Method	Mode of administration
End users	Initial questionnaire	On the first day of the trial	Questionnaire	In person
	Exit questionnaire	At the end of the last day of the trial	Questionnaire	In person
	Focus group	In the first week following the end of the trial	Semi-structured discussion	In person
Developers	Initial questionnaire	As soon as possible after the start of development work	Questionnaire	In person
	Final questionnaire	At the end of the trial	Questionnaire	In person
Other	Analysis of log files	At the end of the trial	Statistical analysis	

4.5 Data analysis

The analysis of the data from the trials will be carried out by the partners responsible for the trials, using a template provided by XS. The analysis will cover:

- Results from end-user questionnaires
- Results from end-user focus groups
- Results from developer questionnaires
- Results from log analysis.

The results will be included in the reports due at the end of each phase (D 8.2, D8.3, and D8.4).

5 Network experiments (track 2)

5.1 Testing CONVERGENCE networking

5.1.1 Overview

As reported in Chapter 1, track 2 in the CONVERGENCE trials will consist of a series of experiments designed to test CONVERGENCE content-centric networking concepts in two different networking scenarios. The first scenario is the *Laboratory Network* depicted in Figure 2. This consists of two end-nodes and two serving-nodes spread across three CONET Sub Systems (CSSs), interconnected by two border-nodes (BNs). CSSs n.1 and n.3 are formed by a point-to-point IPv4/v6 link; CCS n.2 is an IPv4 network with plain IPv4 routers (R) and an internal-node. A common Name System Node supports the routing-by-name (lookup-and-cache) operation described in D5.1 and D3.2.

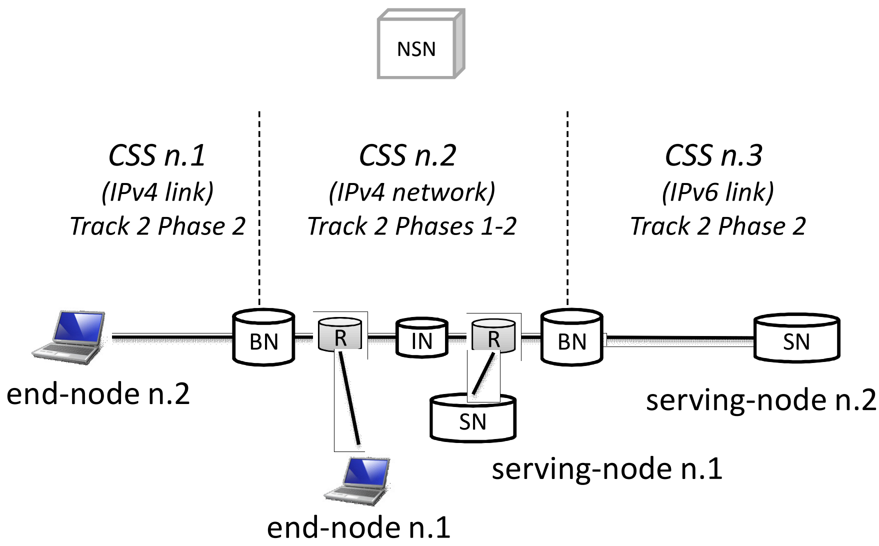

Figure 2 - Laboratory scenario for CONET tests

The second scenario uses the OFELIA platform [2] to provide a *Wide Area Network* that is qualitatively similar to the network shown in Figure 2, but with more devices and CSSs. The OFELIA network include so-called islands in Rome, Catania, Berlin, Gent, Zurich, Barcelona, Essex. Some OFELIA islands will be coupled by IP to form a single wide-area CSS. Others will form autonomous CONET Sub Systems. We envisage three main categories of test: API tests, networking tests and performance tests. Table 7 summarizes the CONET test plan. Results from preliminary tests are reported later in this chapter.

Test	Phase	Instrument
API test	1	Laboratory network depicted in Figure 2 but formed only by a the IPv4 CSS n.2
API test	1	Laboratory network depicted in Figure 2
Networking test	1	Laboratory network depicted in Figure 2 but formed only by the IPv4 CSS n.2
Networking test	2	Laboratory network depicted in Figure 2
Performance test	1	Laboratory network depicted in Figure 2 but formed only by the IPv4 CSS n.2
Performance test	2	Laboratory network depicted in Figure 2
Performance test	2	OFELIA infrastructure

Table 7 - Test Plan for CONVERGENCE networking concepts

5.1.2 API test

5.1.2.1 Description of the test

The API test aims at verifying the correct operation of the functions offered by the CONET API, as specified in D3.2. The evaluation will be based on the examination of the outputs provided by the API with respect to specific inputs. The test will not attempt to measure performance (e.g. execution time, network throughput, etc.).

5.1.2.2 Implementation framework

The trials in phase 1 provide an *implicit* test of the API: if they succeed this is evidence that the API is functioning correctly. In what follows, we describe an *explicit* test, based on the *laboratory* scenario depicted in Figure 2. In this scenario, two end-nodes perform GET, SendToName and SendToLocation operations. Two serving-nodes are the endpoints for the GET, SendToName and SendToLocation operations performed by end-nodes. Serving-nodes also perform Advertize, Update and Revoke operations. End-nodes and serving-nodes are spread across three CONET Sub Systems (CSSs), interconnected by two border-nodes (BNs). CSSs n.1 and n.3 are formed by a point-to-point IPv4/v6 link; CCS n.2 is an IPv4 network with plain IPv4 routers (R) and an internal-node. A common Name System Node supports the routing-by-name (lookup-and-cache) operation described in D5.1 and D3.2.

5.1.2.3 Implementation Plan and Evaluation Methodology

The CONET API test will be performed during track 2 phases 1 and 2. In phase 1, we will deploy a single IPv4 CSS, i.e. the implementation will be limited to CSS n.2 in Figure 2. This test will verify the API for the case of an end-point that communicates with a serving-node in the same CSS. In phase 2, we will extend the network to include CSS n.1 and CSS n.3. In this case, we will test the API when endpoints are on different CSSs, e.g. when end-node n.2 communicates with serving-node n.2.

5.1.3 Networking test

5.1.3.1 Description of the test

The aim of the Networking test is to verify the correct operation of the CONVERGENCE network functionality supporting API services, as described in D5.1 section 5. The evaluation will test the correct execution of specific operations without attempting to measure performance. The test will focus on the following operations:

- Packetization of named-data
- Interactions with Name System Nodes to insert and lookup name-based routing information
- In network caching
- Lookup-and-cache
- Route-by-name
- Route-by-path-info

5.1.3.2 Implementation Plan and Evaluation Methodology

The CONET networking test will be performed during track 2 phases 1 and 2. In phase 1 we will deploy a single IPv4 CSS (CSS n.2 in Figure 2). We will test:

- The packetization of named-data items
- The interaction between a serving-node and a name system node during insertion and revocation of routing information for a named-resource
- In-network caching functionality on an internal-node
- Lookup-and-cache and routing-by-name on an end-node.

In phase 2 we will extend the network to include CSS n.1 and CSS n.3. In addition to the tests performed in phase 1 we will additionally test:

- Routing by path info
- In-network caching by border-nodes
- Lookup-and-cache and routing-by-name on border-nodes.

5.1.4 Performance test

5.1.4.1 Description of the test

The aim of the CONET performance test is to assess the effectiveness of the CONVERGENCE content-centric network with respect to off-the-shelf IP technology and to alternative proposals contained in the literature [1]. The performance measures used in the test will be of interest both to Internet Service providers and to users.

5.1.4.2 Implementation framework

The tests will use both the laboratory network described in Figure 2 and the OFELIA wide area network. Furthermore, specific measurements will be carried out also using simulation, to test algorithms prior to deployment on the real network and for analyses requiring very high numbers of nodes or named-resources.

5.1.4.3 Implementation Plan and Evaluation Methodology

5.1.4.3.1 Laboratory and OFELIA tests

In phase 1 we will deploy a laboratory network formed by a single IPv4 CSS (CSS n.2 in Figure 2).

In phase 2 we will extend the laboratory network to include CSS n.1 and CSS n.3. Moreover, we will deploy a single CSS distributed across different OFELIA sites and perform intra-CSS tests in a wide area network. Then, we will increase the scope of the wide area network including other OFELIA islands, forming different CSSs. This will enable us to perform inter-CSS tests in a wide area network. In what follows, we provide a preliminary description of planned measurements that we will carry out both in laboratory and with the OFELIA network.

Throughput measurements

In these tests, we will compare the average bit-rate (*throughput*) when downloading a file using the CONET, with the average bit-rate when downloading the same file via a plain HTTP connection. The test will not exploit in-network caching; hence, we expect that the two throughputs will be quite similar. The goals of the tests will be to show that: i) current technology can support the complexity of CONET forwarding functionality with no penalty in terms of bit-rate, and ii) CONET traffic can coexist with TCP/IP services.

Routing-by-name measurements

These measurements will test the effectiveness of lookup-and-cache, the CONVERGENCE approach to routing-by-name (see D5.1 section 5). The advantage of lookup-and-cache is that it limits the size of the name-based routing tables that contain the NIDs of named-resources. The disadvantage is that border-nodes sometimes have to query name-system-nodes, introducing additional delay in the communication path. The planned measurements will compare delay and throughput during the lookup-and-cache operation against the impractical case in which all possible names of named-resources are included in CONET routing tables. The measurements will be based on the inter-CSS configuration planned for track 2 phase 2.

<u>*In-network caching measurements*</u>

These tests will compare the time needed to retrieve a named-data item (*retrieval time)*, using CONET in-network caching, against the performance of a plain TCP/IP network. To this end, we will position end-nodes (HTTP-clients) and serving-nodes (HTTP-servers) in available OFELIA islands. Serving-nodes (HTTP-servers) will provide different content. End-nodes (HTTP-clients) will randomly choose contents to download, according to content popularity statistics. CONET nodes will provide in-network caching.

5.1.4.3.2 Simulation

During the development process and in cases where we need to verify performance with very large numbers of nodes and/or named resources, tests will use NS3 simulation and (if required) ad-hoc tools.

At the time of writing we are already using a NS3 simulation to assess:

- The performance of in-network caching mechanisms
- The effectiveness of a receiver-driven implementation of TCP transport protocol
- The scalability of CONET routing, i.e. the CONVERGENCE lookup-and-cache algorithm.

5.1.5 Results from preliminary studies

5.1.5.1 Overview

In this section, we report results from preliminary studies designed to test three key CONET concepts: the CONVERGENCE lookup-and-cache scheme, the CONET IP option (see D5.1 section 5) and the in-network caching functionality.

5.1.5.2 Lookup-and-cache

We recall that routing-by-name is restricted to interest-CIUs; data-CIUs are routed back to the end-node by means of source-routing (see D5.1 section 5). The CONET nodes involved in routing-by-name are either end-nodes or border-nodes. In the case of end-nodes, the lookup-and-cache approach resembles the interaction between an Internet host and a DNS server, where the host implements a local DNS cache service. This is evidence that lookup-and-cache is definitely feasible on end-nodes. In what follows, we will focus on its feasibility in border-nodes.

We assume that a standard TCP session between a client and a WEB server is replaced with a CONET session (exchange of CONET CIUs) between an end-node and a serving-node, or an intermediate cache. Specifically:

- An URL <http://IP address:80 (or domain-name)/path> is replaced by the network-identifier: namespace=”www”, principal=”IP address:80”, label=”path”;

- TCP segments are replaced by carrier-packets that convey segments of named-data CIUs;
- TCP ACKs are replaced by carrier-packets that convey interest CIUs.

With these assumptions, we can map a real Internet trace, formed by TCP segments and ACKs, to a "CONET trace", formed by carrier-packets. In our study, we applied this re-mapping to two Internet traces: the first captured from a 10 Gbit/s interface on a tier-1 router [4], the second captured from a 10 Mbit/s interface on a tier-3 router [5].

For the purposes of our study, we fed the two re-mapped traces to a software emulation of a CONET border-node. We considered a PLHB (Principal/Label Hash Based) NID namespace and following the approach suggested in [3]. We further assumed that routing-by-name is performed only on the base of the principal identifier. This means that name-based routing entries have the form <namespace, hash (principal),*> and that all named-data for a given principal are stored in a serving-node (and its replicas, if any). Finally we assumed that the route cache adopted a Least Recently Used (LRU) caching policy, discarding the least recently used item first.

Figure 3 plots the number of name-lookups per second issued by the border-node to the name-system, against the size of the route cache. The results show that the number of look-ups decreases in a log-like fashion with increases in the size of the cache. In the case of the tier-1 node, with a route cache of 8k entries we observed 10 lookups per second and a cache-miss probability of about 10^{-4}. In the case of the tier-3 node, with a route cache of 2k entries, we observed about 2 lookups per second and a cache-miss probability of about 10^{-3}. Considering that current BGP routers handle about 350k entries and that 2 or 10 lookups per seconds are reasonable values, we conclude that the CONVERGENCE lookup-and-cache scheme is probably feasible with current technology.

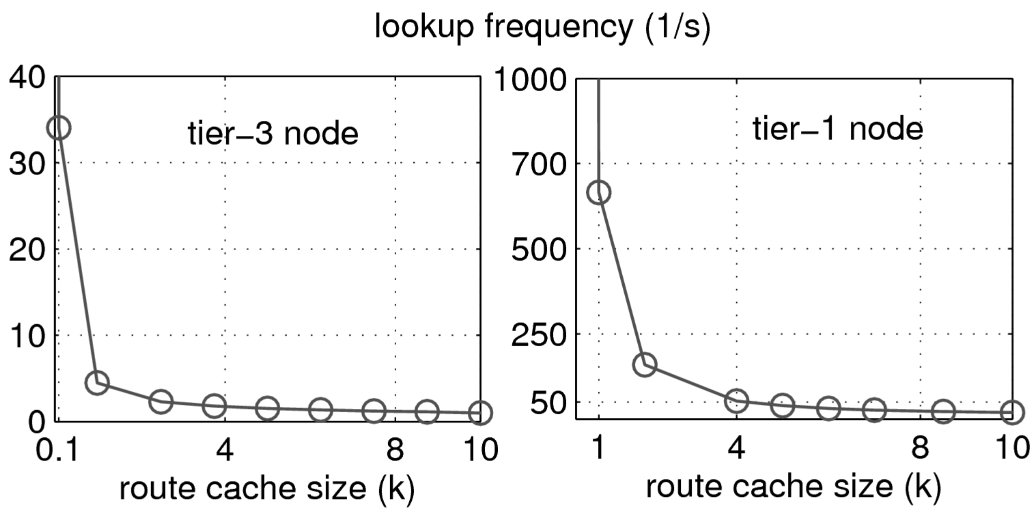

Figure 3 - Lookup frequency of a tier-3 and a tier-1 border-node

5.1.5.3 CONET-IP integration

In this study, we tested the feasibility of conveying the header of carrier-packets in an IPv4 option, as described in D5.1. The rationale for the test was that IP routers tend to process packets with IP options in the slow forwarding path. Given that in a real deployment scenario, CONET nodes would have to coexist with plain IP routers, it is important to test how far this slows down throughput.

To check the behaviour of IP routers, we transmitted IP packets with and without our CONET option from a PlanetLab node at our premises in Rome, Italy to eleven other PlanetLab nodes, in Asia, Europe, North America, and Australia. In each case we measured the difference in terms of throughput and round-trip-delay (i.e. the available capacity between a sender and a receiver). All measurements were repeated ten times. The mean results are shown in Figure 4.

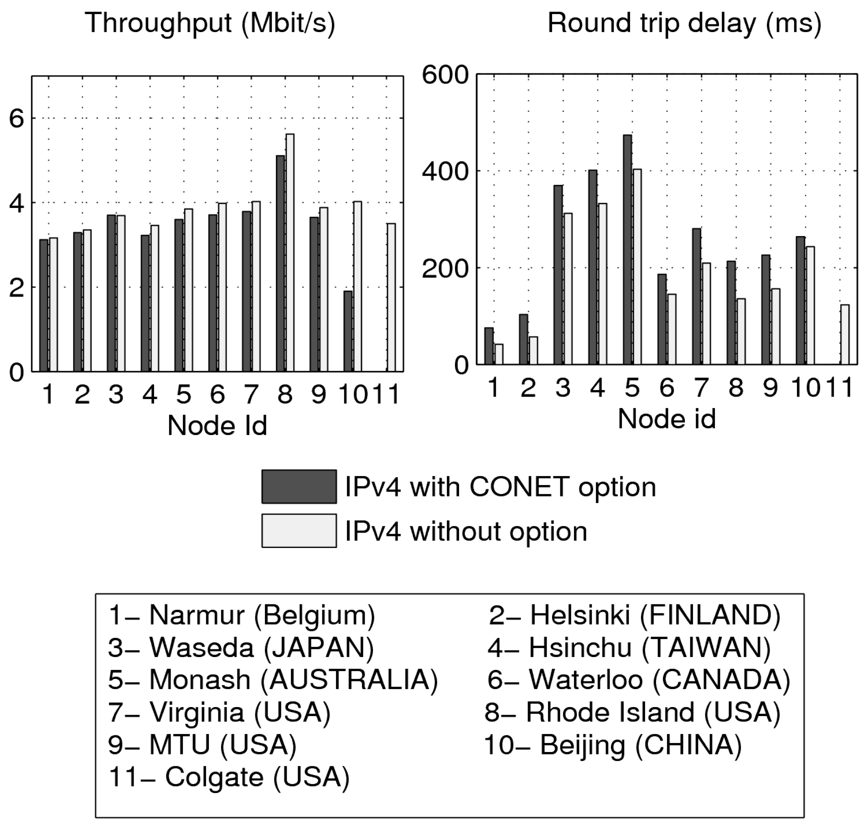

Figure 4 - Throughput and round-trip-delay of IP packets with and without CONET options on different Internet paths

As far as concerns throughput, the performance of the first nine PlanetLab end-nodes was almost constant regardless of whether we were using the CONET option or not. However, two end-nodes showed significant differences in performance. Further analysis revealed that: i) the Beijing-Rome path includes a router that statistically drops half of all packets with IP options; ii) the Colgate-Rome path includes a router (in Australia) that drops all packets with IP options. These problems affect only a minority of the routers we examined and depend on the configuration of the router software. Policies that prevent the use of IP options are probably designed to prevent DoS attacks [6] and could easily be modified to accept CONET carrier packets without restrictions. As far as concerns round-trip-delay, we observed a small increase in latency for packets with the CONET option. Overall, our measurements show that properly configured IP routers would not create a critical performance bottleneck for the CONET option (see [7] for a similar analysis).

5.1.5.4 In-network Caching

5.1.5.4.1 Background

In this study, we made a preliminary analysis of the CONET in-network caching functionality. We recall that in-network caching, aka *en-route* caching [9], provides network nodes with the capability to cache named-data CIUs (chunks of named-data items, see D5.1) they forward. When a node receives an Interest CIU (i.e., a request for a named-data CIU) that refers to a cached named-data CIU, the node returns the cached data without further propagating the Interest CIU. Caching operations are performed directly by nodes between the end-node and the serving node, and do not require the re-routing of traffic to off-path caching/proxy servers. Given that some requests for named-data CIUs are satisfied by nodes closer to the end-node than the serving-node, this reduces network traffic.

Since a named-data CIU is only a small part of a whole named-data item, an end-node has to transmit several Interest CIUs to download a single named-data item. Nodes then have to execute a lookup for every Interest CIU they receive (see D5.1). This requires fast cache lookup. Lookups are usually based on an *index table* that keeps track of content in storage. Index tables are stored using fast memory technology, e.g. TCAM (4 ns) or DRAM (55 ns). These memory technologies are costly and have limited storage capacity. At the time of writing, TCAM offers a maximum capability of about 20 Mbits at a cost of $US200/Mbit. The maximum capacity of DRAM is 4 Gbyte at a cost of 0.016$/Mbyte [11]. Index tables are usually implemented as hash tables. Assuming 40 bit hash items this implies that a 4 Gbyte DRAM could hold at most 800 million items. This is quite a small number compared to the number of content items available on the Internet.

A fundamental element of a caching mechanism is the *caching policy* (aka caching strategy). This i) determines whether or not a new named-data CIU should be stored in cache and ii) selects a cached named-data CIU to kick out to free memory for the new data. When there are multiple caches in a network, caching policies may be either *coordinated* or *uncoordinated*. In coordinated caching policies, the caches share information about cached content in order to optimize the use of the overall caching space; for instance, a cache policy could avoid storing content in a cache that is already stored in a neighbouring cache. In uncoordinated policies, caches take caching decisions autonomously. Coordinated policies improve caching performance. However, it can be very difficult to maintain a consistent vision of a network of caches in which each cache hold thousands of frequently changing items.

At the time of writing, the CONVERGENCE project is investigating the possibility of using uncoordinated caching policies in CONET nodes. In our study, therefore, we carried out a preliminary performance analysis using a well-known caching policy, namely the Least Recent Used (LRU) policy. In LRU, each named-data CIU received by a node is stored in the cache where it replaces the least recently used item. LRU is widely implemented in proxy software (e.g. Squid) and is relatively simple to implement.

There is already an extensive literature on LRU performance with a single cache [10]. This literature shows that, in these conditions, LRU tends to maintain the most frequently used named-data CIU's in cache thereby maximizing cache hit probability (the probability of satisfying an Interest CIU with a cached named-data CIU) without attempting to monitor the frequency of all possible CIUs. To date, however, there have been few studies on the collective performance of "networks of caches" such as CONET.

To investigate this issue, we designed a simulation tool based on Network Simulator 3 [8]. This tool enables us not only to analyze the effectiveness of LRU but also to test novel caching policies that could be implemented in the future. In the simulation model, an end-node sends out an Interest CIU, which is forwarded toward the serving-node. If an en-route node has cached the related named-data CIU, it returns the data to the end-node; if no en route node has cached the data, the data is returned directly by the serving-node.

The simulator does not model possible segmentation and reassembly operations in which named-data CIU are split into underlying carrier-packets. However, these have no impact on caching performance. As the simulator does not model actual packet exchange, it should be considered as a session-level, rather than a packet level simulator.

5.1.5.4.2 Analysis of LRU with an Internet like topology

Our first simulation study assessed the effectiveness of an LRU caching policy in an Internet like network topology. The network topology used in our simulation contained i) a set of core and edge nodes forming a *backbone* network; and (ii) a second set of end-nodes and serving-nodes forming an *access* network.

The backbone network topology consisted of 500 nodes and was obtained by rescaling the Autonomous System (AS) Internet topology, published by The Cooperative Association for Internet Data Analysis (CAIDA) [12]. For rescaling, we used the ORBIS tool [13]. This made it possible to scale down the network topology while maintaining the graph structure of the current ASs. Nodes in the backbone network were classified either as *edge* nodes or as *core* nodes. Edge nodes were defined as nodes with a single link to a core node. Core nodes were defined as nodes with more than one link to other nodes. The final network contained 250 edge nodes and 250 core nodes.

The access network consisted of 250 end-nodes and 250 serving nodes. Each edge node in the core network was connected to one end-node and one serving node in the access network.

Network links were assumed to have unlimited transfer capacity. This assumption means that the simulation did not have to take account of lost data units – a problem that would have complicated the analysis of caching performance. In the simulation, each serving-node made available 4000 named-data CIUs. The network thus contained 1 million named-data CIUs.

For the purposes of the study, we considered a traffic workload in which each end-node issued Interest CIUs according to a Poisson process. The mean time between two consecutive Interest CIUs was 1 sec. Named-data CIUs were assumed to have varying popularity. The selection of named-data CIUs requested by end-nodes was modelled as a random process based on a Zipf distribution with parameter α=1. This distribution provides a good representation of the popularity of current Internet Web content [14]. Each simulation modelled requests for data over a period of 10000 seconds.

The performance of the caching algorithm was measured as the ratio between the number of named-data CIUs transferred across links in the backbone network in the presence and in the absence of caches. Thus a ratio of 0.5 implies that in-network caching halves backbone traffic with respect to the case in which there is no caching. All caches in a given simulation contained the same number of storable named-data CIUs. In what follows, we represent cache size in terms of the proportion of all named-data CIUs (i.e. 1 million) stored in a single cache.

Figure 5 shows the reduction in backbone traffic with different sizes of cache, comparing results when caches are deployed only in edge nodes and when they are deployed in core nodes as well. The results show that in-network caching provides a valuable reduction in backbone traffic. With a cache size of 0.1%, for instance, backbone traffic is roughly halved. We note, however, that adding caches to core nodes provides only a small gain in performance. When an end-node downloads a named-data CIU, all nodes between the end-node and the node providing the named-data CIU store the CIU. As a result cache content on edge and core nodes becomes closely correlated reducing the usefulness of the latter.

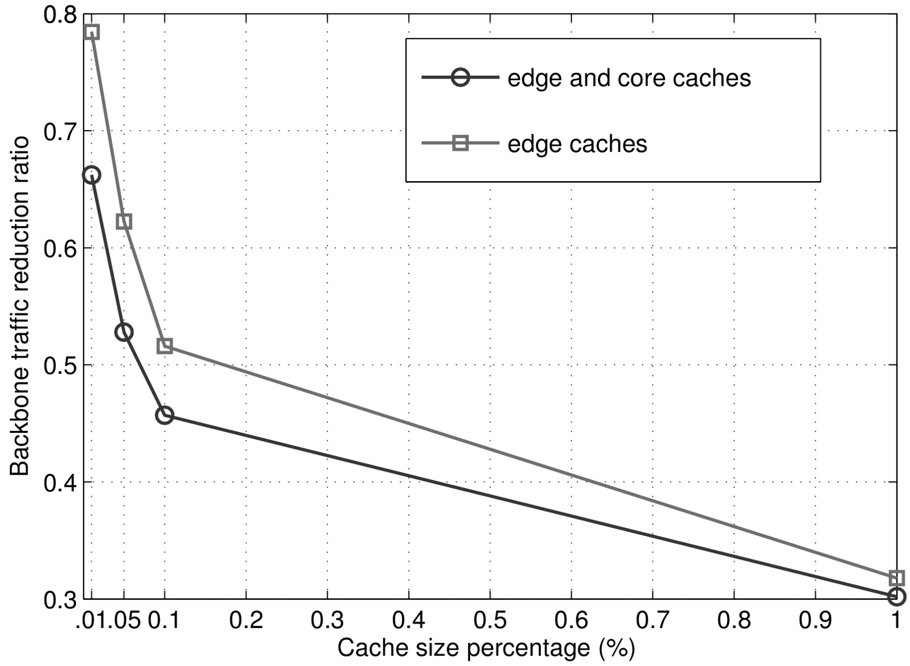

Figure 5 – Reduction of Backbone traffic with an LRU caching policy and an Internet like network topology

5.1.5.4.3 Analysis of LRU in a tree topology

To better understand the potential of an LRU caching strategy in a network of caches, we studied the performance of the strategy with the tree topology shown in Figure 6. In this topology, the tree has a depth of 3. Each node has 3 descendent nodes. A single serving-node (SNs) located at the root of the tree (depth=0) hosts all named-data CIUs (1 million). In the figure, each node is labelled with its depth in the tree. In the simulation, each node with depth 3 is connected to an end-node. Only one of these nodes is shown in the figure. Caches are deployed exclusively at nodes at depths 1, 2 or 3.

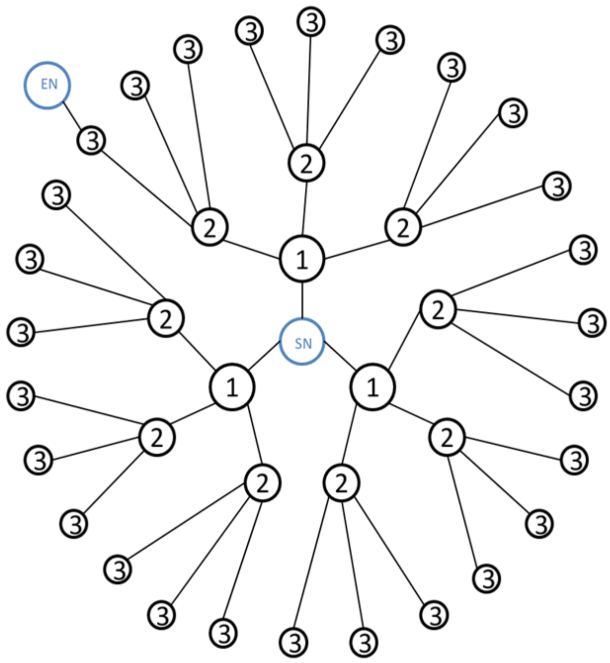

Figure 6 - Tree Topology

It can be shown that, in this topology, performance is optimal when, for cache size N, each node with depth=3 contains the 1st to the Nth most popular named-data CIUs, when each node with depth=2 contains the N+1-th to the 2N-th most popular CIUs and when each node with depth=1 contains the 2N+1-th to the 3N-th most popular CIUs.

Performance with this configuration can be compared against performance with an LRU caching strategy, thereby providing an estimate of the distance between this performance and the theoretically optimal strategy.

Figure 7 shows the reduction in traffic obtained with LRU and compares it against the optimal value. Traffic reduction is measured as the ratio between the numbers of named-data CIUs transferred over network links (excluding EN-3 links) in the presence and in the absence of caches. The results show a valuable reduction of traffic. However the reduction achieved is smaller than the theoretical optimum value.

Figure 8 shows the correlation among content cached by network nodes with a cache size of 0.5%, measuring correlation as the number of named-data CIUs present in two "sequential" caches, as a percentage of cache size. Separate curves show the number of CIUs shared between connected nodes with depth=3 and with depth =2, and the number shared between connected nodes with depth=2 and depth=1. At the start of the simulation, the caches are filled-in optimally. There is thus no correlation between the content of connected nodes.

However, as the simulation procedes, cache contents become more and more similar, converging towards a state in which about 35% of content is shared between sequentially connected caches. We conclude that in a network of caches such as CONET, LRU can provide good perfomance. Given, however, that performance is less than optimal, it may still be possible to devise a better caching strategy.

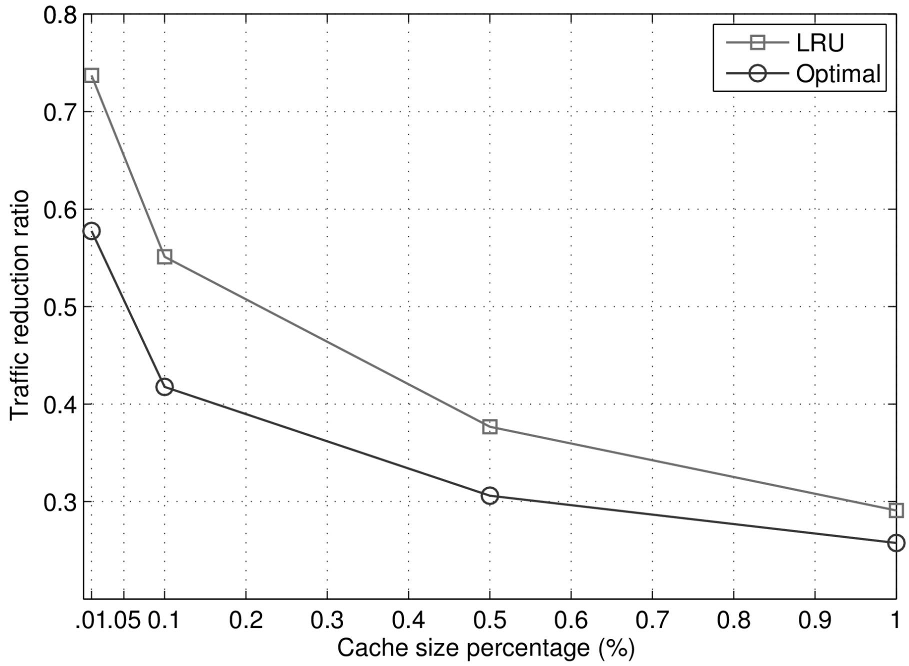

Figure 7 - Traffic reduction ratio in the case of a tree network

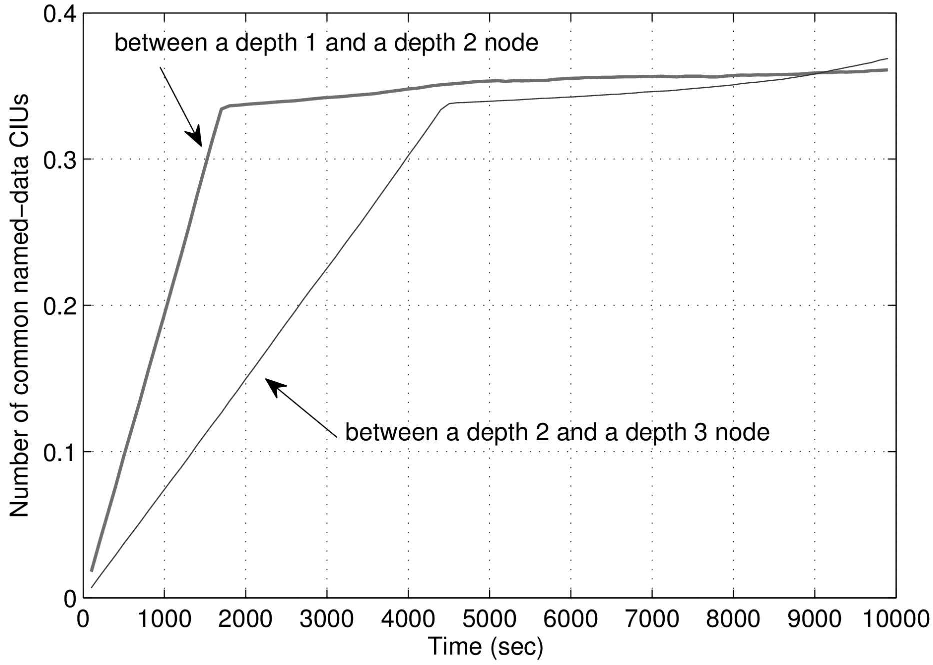

Figure 8 - Correlation of caches with LRU caching policy

5.2 Evaluation of the CONVERGENCE publish/subscribe infrastructure

5.2.1 Overview

The CONVERGENCE publish/subscribe infrastructure is based on a middleware engine – the *overlay technology engine (Overlay TE).* The engine maintains a registry of semantic fractals and implements the operations necessary to send and receive messages to/from these fractals. The CONVERGENCE publish/subscribe mechanism uses the engine to propagate publication and subscription messages to a specific set of peers, which perform the matching between publications and subscriptions. More details on the overlay engine and on the publish/subscribe functionality can be found in D3.2.

In the following paragraphs we will describe plans to test the API used to access the overlay engine, its functionality and its performance. In the final paragraph of this section, we provide results from preliminary simulations.

5.2.2 API Test

5.2.2.1 Description of the test

The goal of this test is to verify that the output produced by calls to the Overlay TE corresponds to expectations.

5.2.2.2 Implementation framework

All peers used in tests will use a LINUX image loaded with VirtualBox. Tests will focus on the following operations:

- Peer registration to fractal.
- Message buffering.
- Propagation targets selection.
- Message gossip to peers.

5.2.2.3 Implementation plan and evaluation methodology

The first phase of testing will use only a small number of peers. This will make it easy to monitor calls to the API. We expect that the majority of problems with the API will be detected and eliminated at this stage. The second phase of the work will test the API in a network configuration in which publication and subscription requests are issued over the CONET and the registry resides on a CONET peer. The goal of the tests will be to identify changes required to allow the proper operation of the API over the CONET.

5.2.3 Functional test of the publish/subscribe

5.2.3.1 Description of the test

Tests of the API will be based on a pre-defined set of test cases. This second set of tests, will test its functionality under more realistic assumptions and will complement the feedback coming from the real-life trials.

5.2.3.2 Implementation plan and evaluation methodology

The planned tests planned in track 2 will systematically test the following specific functionality:

- Construction of a multi-dimensional overlay.
- The effect of peers entering and leaving the semantic overlay (or fractals of the semantic overlay) on local/remote registries.
- Dynamic construction of a set of target peers (multiple dimensions – multiple fractals per dimension).
- Effective peer fan-out (efficient selection of peers receiving messages avoiding reselection of peers that have already received them).

5.2.4 Performance test

5.2.4.1 Description of the test

Performance tests will focus on the efficiency with which the system propagates and matches publication and subscription messages as measured by subscription hit rates and hit delays (i.e. the time between the start of content propagation and the first match). Simulation will incorporate several real world assumptions including peers going online and offline at arbitrary points in time and limited buffer sizes.

5.2.4.2 Implementation framework

Some of the tests will use the OFELIA platform; others will be based on simulation.

5.2.4.3 Implementation plan and evaluation methodology

5.2.4.3.1 OFELIA test

Experimentation on the OFELIA network will allow us to monitor the behaviour of the publish/subscribe infrastructure in a real environment, which incorporates an actual CONET implementation in the network layer. This will make it possible to test the functionality of the registration module (which exploits the CONET content centric approach) and to compare the performance of the gossip protocol currently used for CONVERGENCE publish-subscribe operations with alternative protocols (e.g. DHT-based protocols).

5.2.4.3.2 Simulation

Simulation-based tests will make it possible to test the performance of the publish/subscribe mechanism on networks with very large numbers of nodes handling very large volumes of data. Planned tests will use simulation to investigate hit rates and hit delays in dynamic networks and in networks with restricted buffer sizes. Additional simulations will investigate the optimal trade-off between the number of messages transmitted to other nodes, hit rates and hit delays.

5.2.5 Results from Preliminary Studies

5.2.5.1 Overview

In what follows, we report results from preliminary studies in which we measured hit rates and hit delays for a gossiping protocol in a fractal architecture. The main characteristics of our simulations can be summarized as follows.

- A taxonomy was used to partition peers into fractals with depth 3, following the scheme described in D3.2. In this scheme, the first level of a taxonomy consisted of a single node (Resource), the second level consisted of four nodes (Digital Resource, Person, Real World Object, Service). Lower levels contained between 5 and 15 nodes randomly linked to parent nodes at the next highest level.

- Each taxonomy contained between 1 and 3 dimensions, all with depth 3.

- Each peer had 1 to 3 interests per dimension (i.e. it participated in between 1 and 3 fractals for each of its dimensions). The exact number of interests within this range was chosen randomly from a uniform distribution.

- Peers were distributed to fractals according to a Zipf distribution.

- The time was slotted.

- In different simulations, peers changed their status (online to offline and vice versa) with probabilities 0, 0.05, 0.1, 0.3 and 0.5.

- Different simulations considered overlays of size 10000 to 40000 peers with step 5000.

- Each simulation included 6000 publications and 4000 subscriptions. The simulation procedure began by choosing the semantic type of a publication and proceeded to choose a peer in the corresponding fractal where it would be published. Selection probabilities were uniform for all semantic types and all peers.

- Publication and subscription matching were based on an identifier contained in each message.

5.2.5.2 Hit Rate

The following tables and figures contain hit ratios for the subscriptions in semantic overlays with one, two and three dimensions. As we observe, in Table 8 and in Figure 9 hit rates are consistently high, even when peers changes their status with 50% probability. Hit rates are higher for larger networks. This is a clear indication that the publish/subscribe operation scales well. It is also a sign that the operation is resilient in dynamic networks.

dim = 1	10000	15000	20000	25000	30000	35000	40000
0	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
0.05	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
0.1	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
0.3	98.33%	99.13%	98.85%	99.18%	99.40%	99.33%	99.53%
0.5	94.50%	94.60%	95.75%	96.50%	97.33%	96.90%	98.03%

Table 8 - Hit Rate for 1 dimension

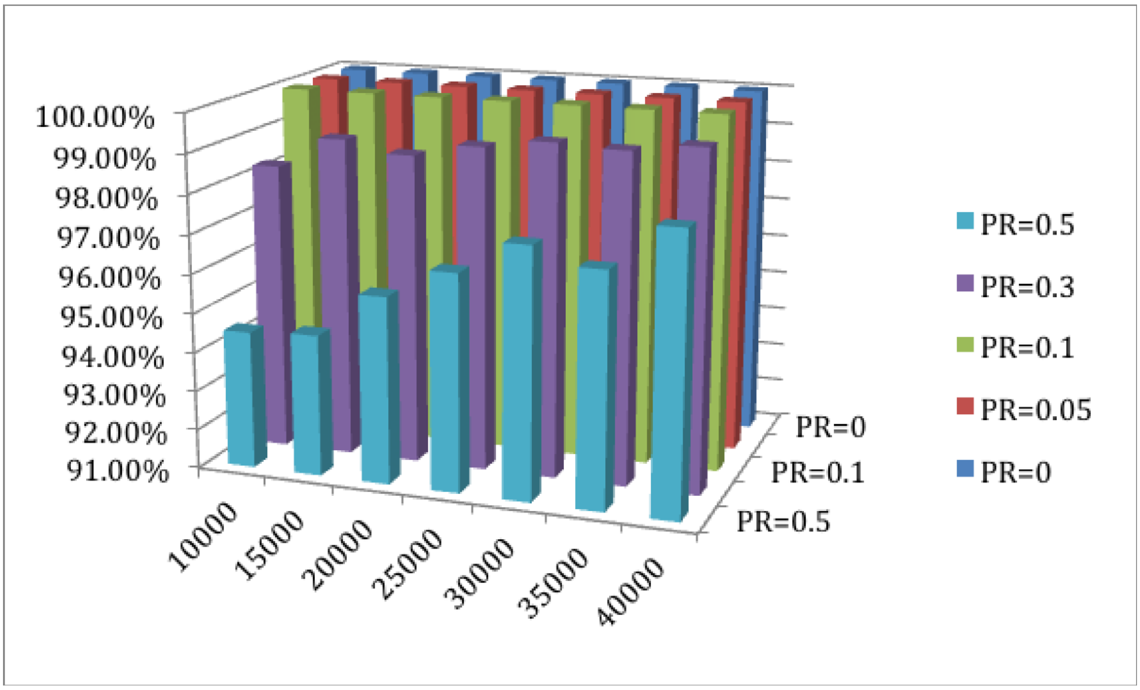

Figure 9 - Hit rate for 1 dimension

Table 9 and Figure 10 show hit rates for a two-dimensional overlay. Although the hit rates are still high and grow with network size, they are lower than those observed in the previous case. This is because this simulation uses smaller fractals than the previous example (intersections of higher level fractals) and the gossip protocol used to propagate message is designed for large scale systems. We are currently studying schemes that allow statistical adjustment of the protocol parameters to the size of the fractal.

dim = 2	10000	15000	20000	25000	30000	35000	40000
0	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
0.05	99.68%	99.78%	99.90%	99.98%	99.98%	99.93%	99.98%
0.1	99.08%	99.43%	99.50%	99.83%	99.80%	99.83%	99.88%
0.3	93.73%	94.48%	97.08%	96.58%	96.80%	96.95%	97.20%
0.5	88.55%	89.63%	91.85%	92.50%	91.83%	93.15%	94.80%

Table 9 - Hit rate for 2 dimensions

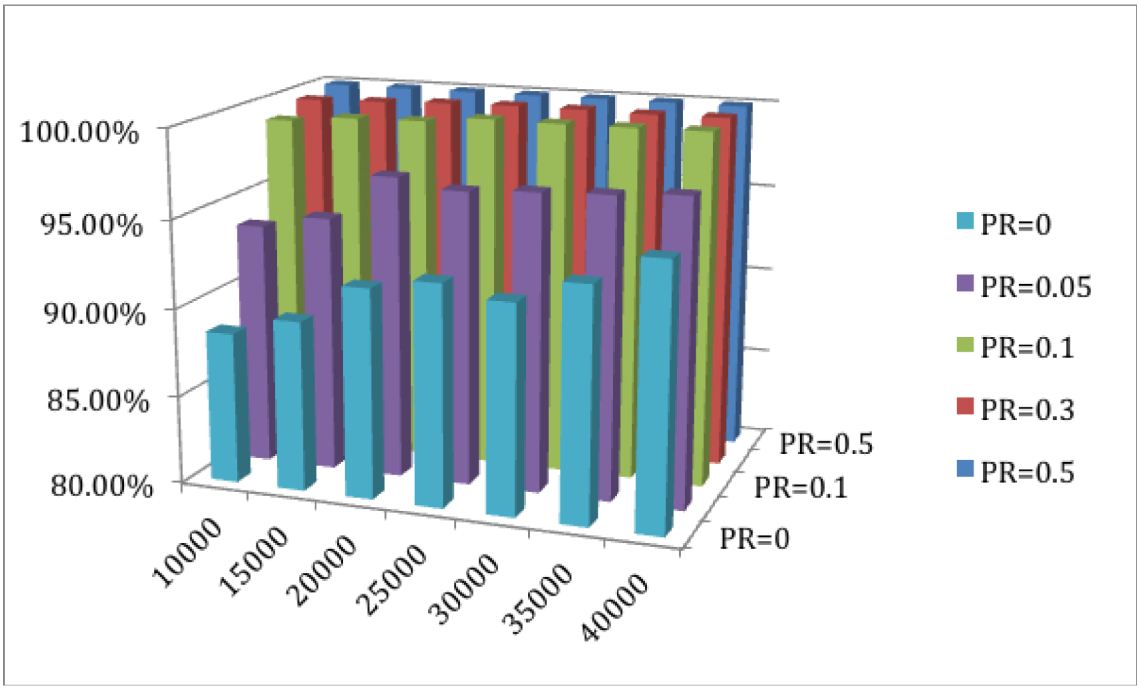

Figure 10 - Hit rate for 2 dimensions

Finally, Table 10 and Figure 11 show the hit rate for a three-dimensional semantic overlay. Again we observe that the hit rate grows with the network size and that it is relatively high even with very dynamic networks. This case also confirms our previous observation about the behaviour of the protocol with small network sizes and strengthens the case for statistical correction of protocol parameters.

dim = 3	10000	15000	20000	25000	30000	35000	40000
0	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
0.05	98.53%	99.25%	99.23%	99.10%	99.18%	99.35%	99.43%
0.1	96.60%	97.63%	97.95%	97.95%	98.43%	98.23%	98.68%
0.3	88.63%	89.43%	90.83%	90.35%	92.23%	91.98%	93.38%
0.5	82.63%	83.40%	84.43%	84.95%	87.53%	86.88%	86.38%

Table 10 - Hit rate for 3 dimensions

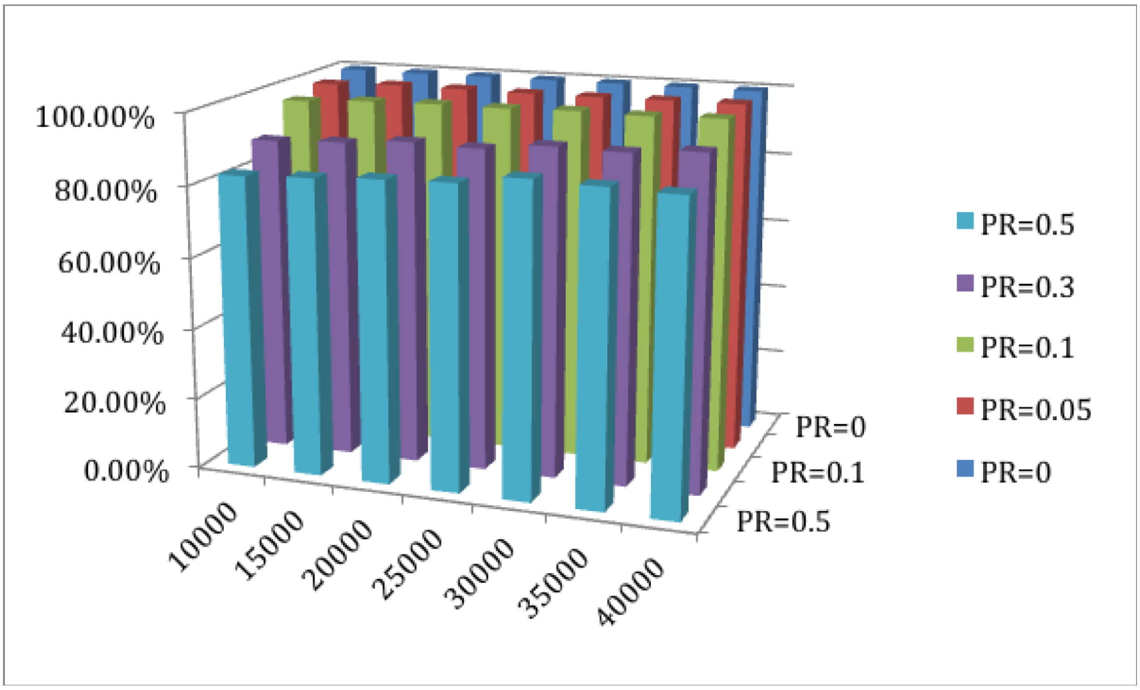

Figure 11 - Hit Rate for 3 dimensions

5.2.5.3 Hit Delay

As we can observe from Table 11 and Figure 12, the hit delay for one dimension is very small, and is always less than 1. This means that in most cases in which we register a hit, the hit occurs on the same peer that issued the subscription. As the network grows, the delay varies. Higher probabilities of status change are associated with longer delays. This result is expected, given that messages have to travel further to find a match.

dim = 1	10000	15000	20000	25000	30000	35000	40000
0	0.025	0.023	0.028	0.033	0.031	0.036	0.033
0.05	0.049	0.055	0.057	0.062	0.061	0.063	0.064
0.1	0.092	0.106	0.091	0.107	0.095	0.079	0.068
0.3	0.414	0.374	0.407	0.427	0.393	0.424	0.404
0.5	0.588	0.626	0.657	0.656	0.654	0.649	0.640

Table 11 - Hit delay for 1 dimension

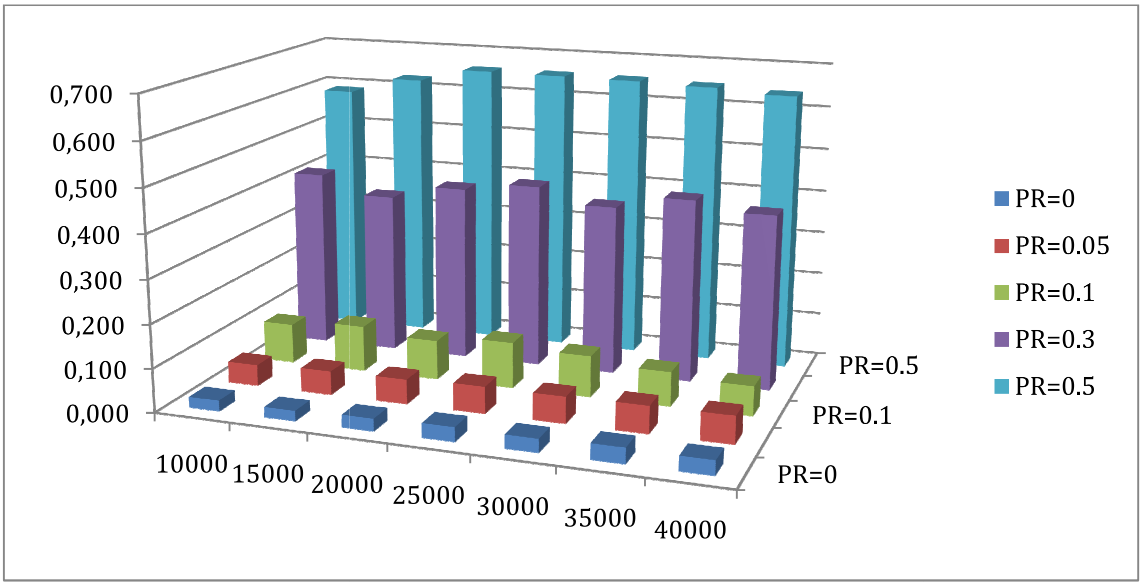

Figure 12 - Hit delay for 1 dimension

Table 12 and Figure 13 show the two-dimensional case. Here, since the network is smaller and matches are found more rapidly, the hit delay is lower than in the previous case. Though there is some variation in the delays, overall delay grows with the size of the network, which is natural.

dim = 2	10000	15000	20000	25000	30000	35000	40000
0	0,011	0,008	0,016	0,017	0,017	0,015	0,020
0,05	0,025	0,031	0,035	0,042	0,041	0,042	0,039
0,1	0,069	0,069	0,073	0,077	0,076	0,069	0,073
0,3	0,307	0,298	0,323	0,330	0,318	0,299	0,334
0,5	0,462	0,477	0,497	0,498	0,509	0,550	0,531

Table 12 - Hit delay for 2 dimensions

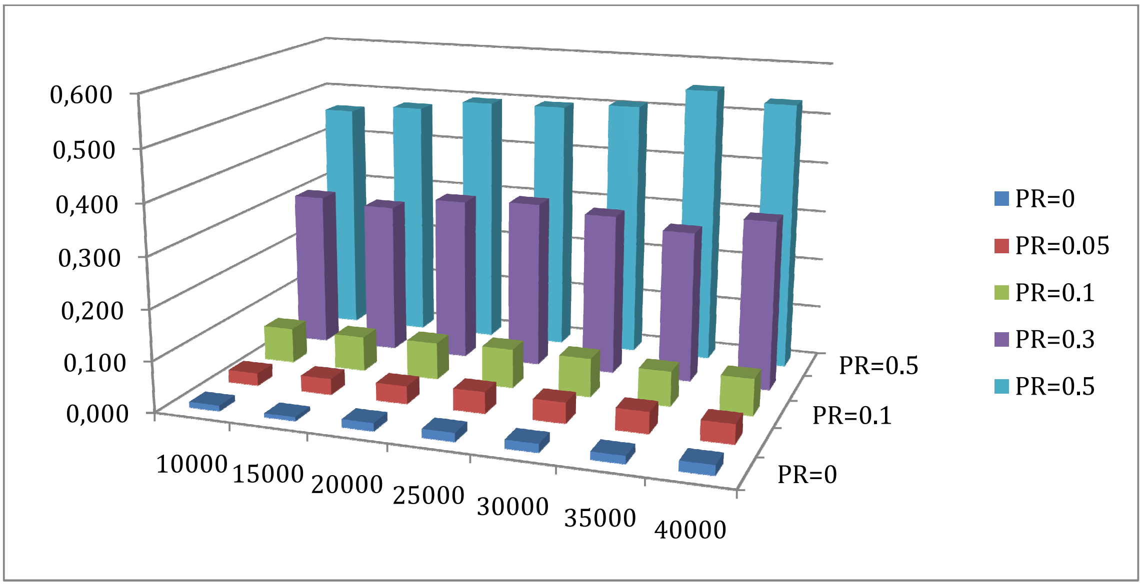

Figure 13 - Hit delay for 2 dimensions

Finally Table 13 and Figure 14 show the three dimensional case, in which the fractal size is smaller and hit delays lower than in the two previous cases. Here again increased network size is associated with increased delay.

dim = 3	10000	15000	20000	25000	30000	35000	40000
0	0,011	0,011	0,016	0,012	0,017	0,009	0,019
0,05	0,028	0,030	0,037	0,035	0,035	0,037	0,035
0,1	0,052	0,066	0,063	0,070	0,068	0,067	0,062
0,3	0,247	0,281	0,285	0,287	0,284	0,302	0,284
0,5	0,374	0,398	0,438	0,449	0,443	0,424	0,477

Table 13 - Hit delay for 3 dimensions

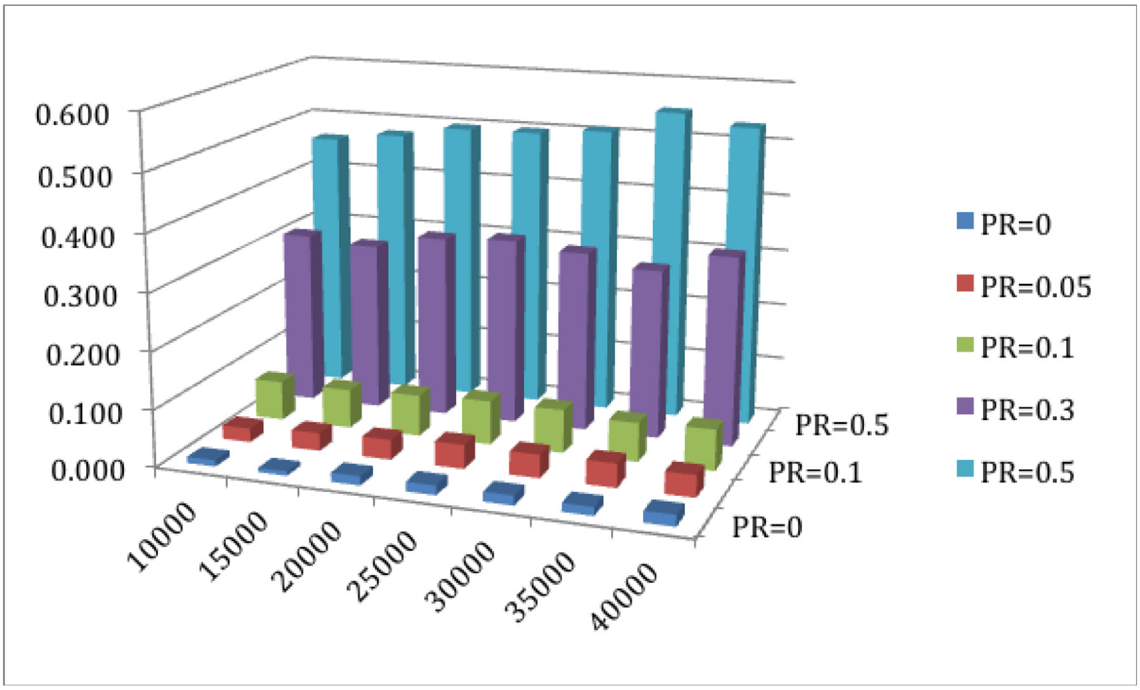

Figure 14 - Hit delay for 3 dimensions

5.2.5.4 Overhead

In this Section, we examine the overhead that the publish/subscribe protocol imposes on the CONVERGENCE system. Since the goal of the protocol is to propagate messages to fractals, we plot the number of messages sent against the size of fractals (or the union/intersection of fractals).

Table 14 and Figure 15 demonstrate that when peers are not changing their status, the overhead is large, with a number of messages 68.99% higher than the size of the target fractal. However, as the size of the network grows, the overhead is reduced, more evidence that the behaviour of the network improves with increased network size. Furthermore, Table 15 and Table 16 and their corresponding figures, Figure 16 and Figure 17, show that there is little difference in performance between overlays of different dimensions. We observe that although the protocol does not cover the whole fractal, hit rates and hit delays are very good. This result supports the findings reported in the previous sections.

dim = 1	10000	15000	20000	25000	30000	35000	40000
0	68.990%	66.873%	64.600%	58.133%	51.841%	47.381%	44.406%
0.05	16.462%	10.860%	9.506%	8.434%	3.811%	-0.172%	-3.591%
0.1	-20.355%	-20.160%	-23.978%	-25.641%	-27.115%	-30.540%	-28.566%
0.3	-75.764%	-76.076%	-79.001%	-78.863%	-78.126%	-80.466%	-81.431%
0.5	-85.154%	-86.504%	-87.512%	-87.958%	-87.870%	-88.706%	-88.877%

Table 14 - Overhead for 1 dimension

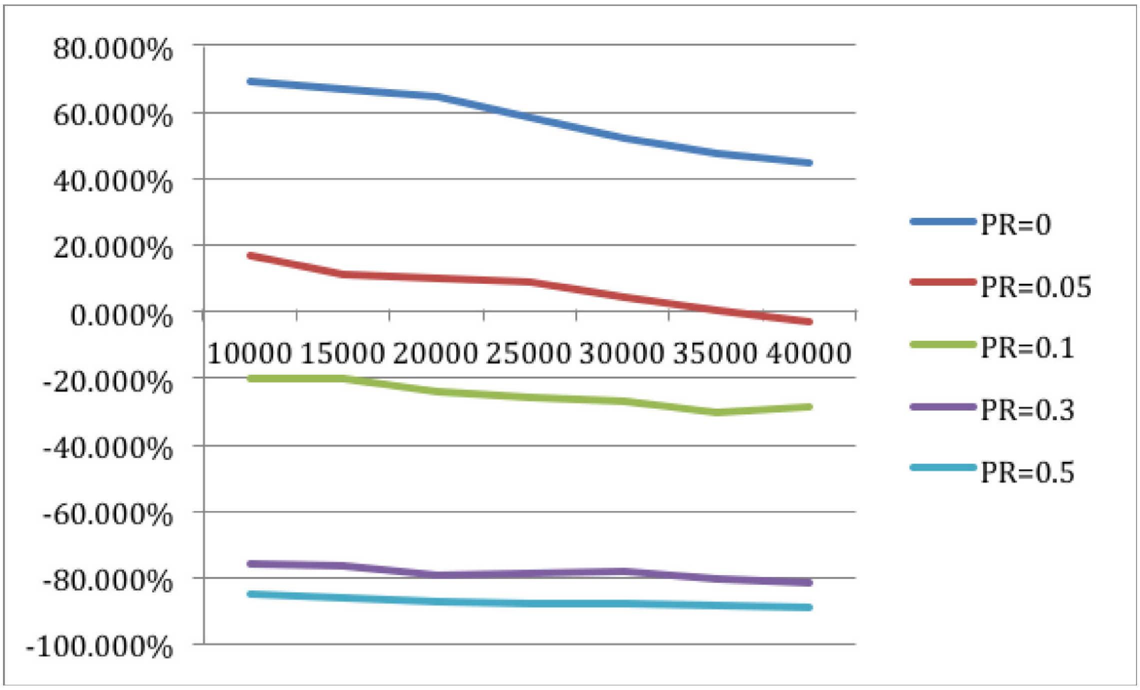

Figure 15 - Overhead for 1 dimension

dim = 2	10000	15000	20000	25000	30000	35000	40000
0	72,025%	68,602%	58,962%	55,599%	50,263%	47,226%	42,171%
0,05	15,651%	13,004%	11,790%	8,275%	5,155%	0,633%	-0,503%
0,1	-20,889%	-21,845%	-21,538%	-26,776%	-27,663%	-25,766%	-30,003%
0,3	-76,323%	-76,030%	-77,844%	-78,614%	-80,220%	-79,824%	-80,601%
0,5	-84,537%	-86,219%	-86,239%	-87,824%	-87,941%	-88,525%	-88,639%

Table 15 - Overhead for 2 dimensions

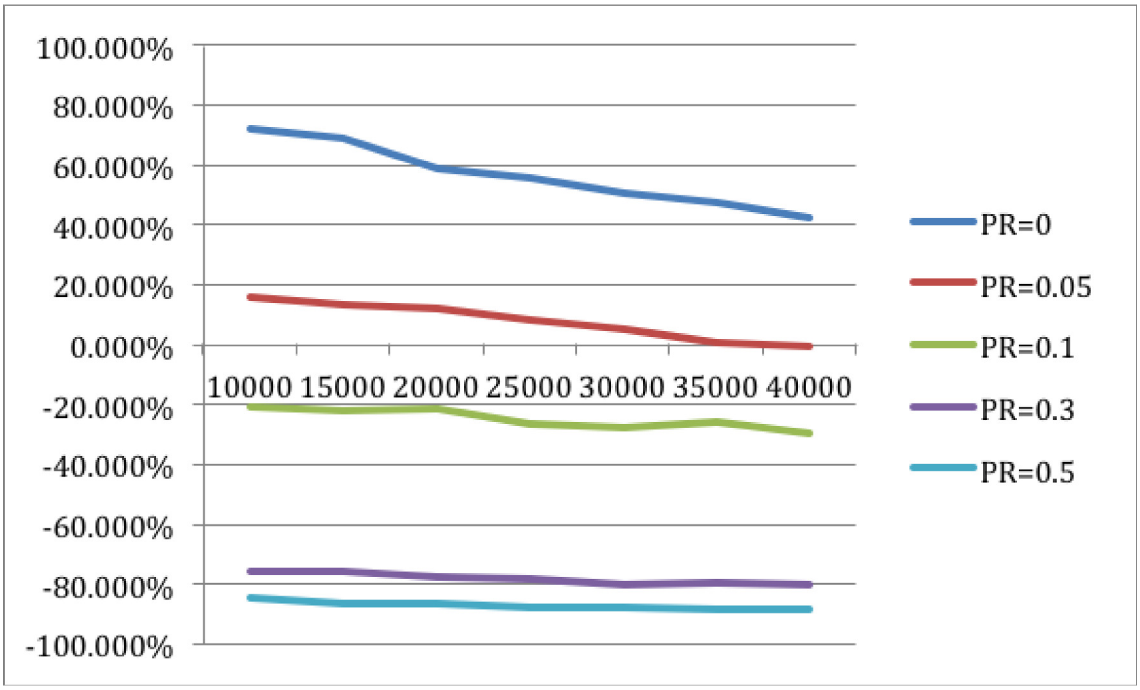

Figure 16 - Overhead for 2 dimensions

dim = 3	10000	15000	20000	25000	30000	35000	40000
0	69,665%	64,391%	58,186%	55,298%	51,727%	46,149%	45,966%
0,05	14,736%	15,116%	8,126%	5,954%	3,058%	-1,178%	-0,785%
0,1	-19,296%	-22,347%	-24,421%	-25,915%	-28,887%	-29,704%	-32,377%
0,3	-73,964%	-77,841%	-77,929%	-78,583%	-80,433%	-81,784%	-80,749%
0,5	-84,594%	-86,004%	-86,194%	-87,029%	-87,579%	-88,663%	-88,848%

Table 16 - Overhead for 3 dimensions

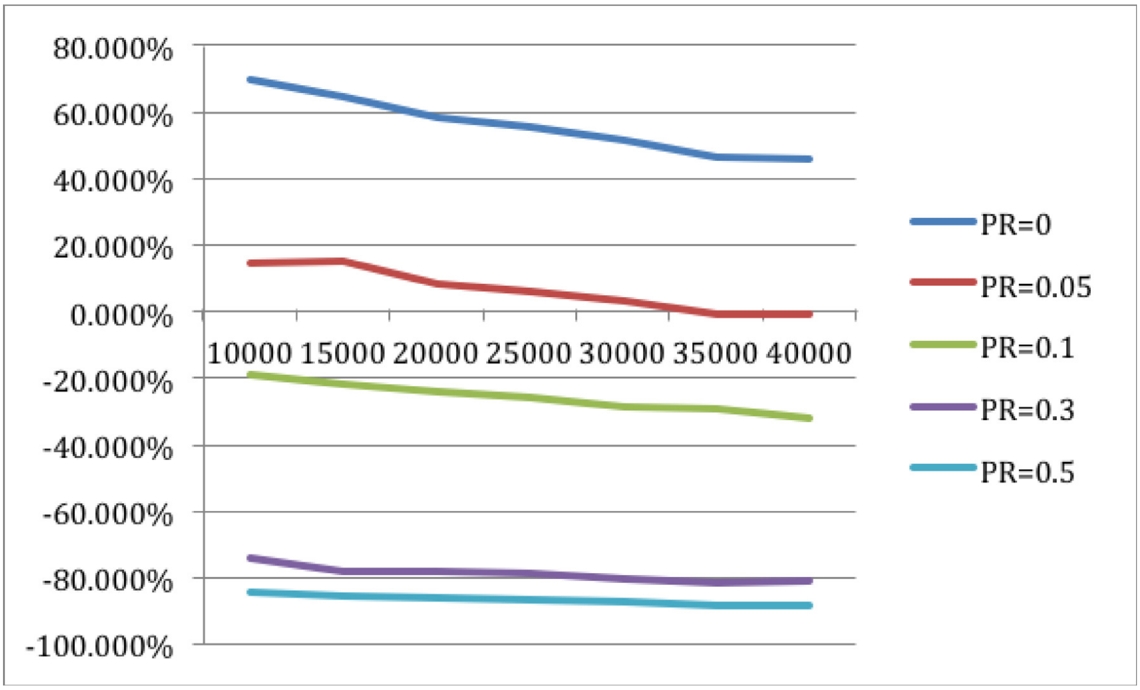

Figure 17 - Overhead for 3 dimensions

5.2.5.5 Conclusions

The results presented in the previous Sections indicate that the CONVERGENCE publish/subscribe protocol is well-adapted to the needs of the large scale dynamic networks, typical of modern settings. The simulations also show that the protocol has a number of drawbacks when used with small fractals. We are currently investigating a solution allowing, statistical adjustment of the protocol to the size of the fractals.

REFERENCES

[1] V. Jacobson, D. K. Smetters, J. D. Thornton et al., ”Networking named content”, ACM CoNEXT 2009.

[2] OFELIA project website: www.fp7-ofelia.eu/.

[3] T. Koponen, M. Chawla, B.G. Chun, et al.: “A data-oriented (and beyond) network architecture”, ACM SIGCOMM 2007.

[4] CAIDA Internet Trace Storage, https://data.caida.org/datasets/passive-2010/equinix-sanjose/20101118/.

[5] Waikato Internet Trace Storage, http://www.wand.net.nz/wits/waikato/1/20050815-000000-0.php.

[6] F. Gont, S. Fouant, “IP Options Filtering Recommendations”, Internet Draft, draft-gont-opsec-ip-options-filtering-00.txt.

[7] R. Fonseca, G. Porter, R. Katz, S. Shenker, and I. Stoica, “IP options are not an option”, Technical report, EECS Department, University of California, Berkeley, 2005.

[8] Network Simulator 3, http://www.nsnam.org.

[9] X. Tang, S. T. Chanson, “Coordinated En-Route Web Caching”, IEEE Transaction on computers, vol. 51, no. 6, June 2002.

[10] S. Podlipning, L. Boszormenyi, “A survey of Web Cache Replacement Strategies”, ACM Computing Surveys, vol. 35, no. 4, December 2003.

[11] D. Perino and M. Varvello, “A Reality Check for Content Centric Networking”, in proc. of ACM SIGCOMM’11, ACM SIGCOMM Workshop on Information-Centric Networking (ICN 2011), August 19, 2011, TORONTO, ON, CANADA.

[12] http://www.caida.org/research/topology/.

[13] P. Mahadevan, C. Hubble, D. Krioukov, B. Huffaker, A. Vahdat, " Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies" in proc. of ACM SIGCOMM ‘07

[14] L. Breslau, et al. “Web caching and Zipf-like distributions: Evidence and implication”, in proc. of IEEE INFOCOM’99..

[15] A. Detti, N. Blefari-Melazzi, S. Salsano, M. Pomposini: ”CONET: A Content Centric Inter-Networking Architecture”, ACM SIGCOMM Workshop on Information-Centric Networking (ICN 2011), August 19, 2011, TORONTO, ON, CANADA.

Appendix: VDIs involved in trial scenarios

1 Photos in the cloud and down to earth

1.1 VDIs used

1.1.1 Photo VDI stored on CONET

Semantic relationship	**Element**	**Sub element**	**Comments**
Is identified by	Photo identifier		
Is described by	Photo metadata		Everybody uses the Alinari photo metadata schema
Includes	License		
		Issuer	
		Principal	Alinari
		Rights	Create derivative work Resell derivative work
		Conditions	
		Encrypted key	
Includes	ER-R		
		Verbs	Play
		Destination addresses	Photographer
Is identified by	R-VDI identifier		
Is signed by	R-VDI signature		

1.1.2 Photo VDI stored on Alinari server

Semantic relationship	**Element**	**Sub element**	**Comments**
Is identified by	Photo identifier		
Is described by	Photo metadata		
Includes	License		
		Issuer	

		Principal	Alinari
		Rights	Create derivative work Resell derivative work
		Conditions	
		Encrypted key	
Includes	ER-R		
		Verbs	
		Destination addresses	
Is identified by	R-VDI identifier		
Is signed by	R-VDI signature		

1.1.3 Publication VDI for Photo VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Photo identifier		
Is described by	Photo metadata		
Includes	License		
		Issuer	
		Principal	Anybody
		Rights	Play P-VDI
		Conditions	
Includes	ER-R		
		Verb	Play P-VDI
		Destination addresses	Photographer
Is identified by	P-VDI identifier		
Is signed by	P-VDI signature		

1.1.4 Subscription VDI for Photo VDI

Semantic relationship	Element	Sub element	Comments
	SPARQL query		
Includes	License		
		Issuer	

		Principal	Anybody
		Rights	Play S-VDI
		Conditions	
Includes	ER-R		
		Verbs	Play S-VDI
		Destination addresses	Photographer
Is identified by	S-VDI identifier		
Is signed by	S-VDI signature		

1.1.5 Refined Photo VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Photo identifier		
Is described by	Photo metadata		
Includes	License		
		Issuer	
		Principal	
		Rights	
		Conditions	
		Encrypted key	
Includes	ER-R		
		Verbs	
		Destination addresses	
Is identified by	R-VDI identifier		
Is signed by	R-VDI signature		

1.1.6 Customer Photo VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Photo identifier		
Is described by	Photo metadata		
Includes	License		
		Issuer	

		Principal	
		Rights	
		Conditions	
		Encrypted key	
Includes	ER-R		
		Verbs	
		Destination addresses	
Is identified by	R-VDI identifier		
Is signed by	R-VDI signature		

1.2 Relationships between VDIs

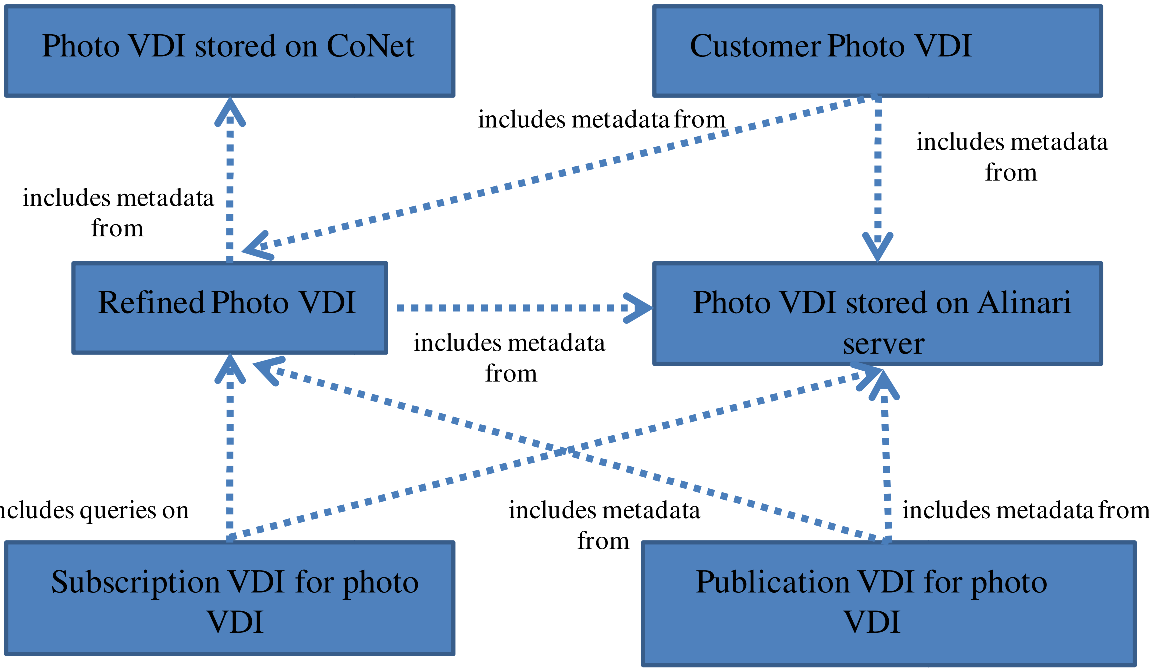

Figure 18 - Relationships between VDIs, trial 1

2 Videos in the cloud and Analyses on earth

2.1 VDIs used

2.1.1 Video VDI

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Video Identifier		
Includes	Video Resource		Encrypted video resource
Is Described By	Video Metadata		Free text
		Title	
		Subtitle	
		Video Type	
		Authors	
		Producers	
		Date	
		Location	
		Spoken Languages	
		Short Description	
		Media Format	
		Media Duration	
Includes	Licenses		
		Issuer	Video Material Owner
		Principal	1/ FMSH/ESCoM 2/ Analyst 3/ Video Channel Owner 4/ Video Channel User
		Rights	1/ store and stream 2/ decrypt, download, analyze 3/ post 4/ watch
		Conditions	defined by Video Material Owner
		Encrypted Key	
Includes	Event Report Requests		
		Verbs	1/ video downloaded and/or decrypted 2/ video analyzed 3/ video posted

			4/ analysis of the video revoked or un-posted
		Destination Addresses	address of the Video Material Owner
Is Identified By	R-VDI identifier		
Is Signed By	R-VDI signature		Video Material Owner signature

2.1.2 Analysis VDI

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Analysis Identifier		
Is Described By	Analysis Metadata		OWL metadata interacting with CDS
Includes	Licenses		
		Issuer	Analyst
		Principal	1/ Video Channel Owner 2/ Video Channel User
		Rights	1/ post 2/ read
		Conditions	defined by Analyst
Includes	Event Report Requests		
		Verbs	1/ analysis posted 2/ analysis un-posted
		Destination Addresses	address of the Analyst
Is Identified By	R-VDI identifier		
Is Signed By	R-VDI signature		Analyst signature

2.1.3 Channel VDI

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Channel Identifier		
References	Analyses		
		Analysis Identifier	identifier of each analysis posted on the channel
Is Described By	Channel Metadata		
		Title	
		Alias	
		Url	
		Short Description	
Includes	Licenses		
		Issuer	Video Channel Owner
		Principal	1/ ESCoM 2/ Video Channel User
		Rights	1/ store 2/ browse
		Conditions	defined by Video Channel Owner
Includes	Event Report Requests		
		Verbs	1/ analysis posted 2/ analysis un-posted 3/ subscription of a Video Channel User
		Destination Addresses	address of the Video Channel Owner
Is Identified By	R-VDI identifier		
Is Signed By	R-VDI signature		Video Channel Owner signature

2.1.4 Publication VDI created and injected by Video Material Owner

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Video Identifier		
Is Described By	Video Metadata		
Includes	Licenses		
		Issuer	Video Material Owner
		Principal	1/ Analyst
		Rights	1/ read
		Conditions	defined by Video Material Owner
Includes	Event Report Requests		
		Verbs	1/ Publication VDI read
		Destination Addresses	address of the Video Material Owner
Is Identified By	P-VDI identifier		
Is Signed By	P-VDI signature		Video Material Owner signature

2.1.5 Publication VDI created and injected by Analyst

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Analysis Identifier		
Is Described By	Analysis Metadata		
Includes	Licenses		
		Issuer	Analyst
		Principal	1/ Video Channel Owner
		Rights	1/ read
		Conditions	defined by Analyst
Includes	Event Report Requests		
		Verbs	1/ Publication VDI read
		Destination Addresses	address of the Analyst

Is Identified By	P-VDI identifier		
Is Signed By	P-VDI signature		Analyst signature

2.1.6 Publication VDI created and injected by Video Channel Owner

Semantic relationship	Element	Sub-element	Comments
Is Identified By	Analysis Identifier		
Is Described By	Analysis Metadata		
Is Described By	Channel Metadata		
Includes	Licenses		
		Issuer	Video Channel Owner
		Principal	1/ Video Channel User
		Rights	1/ read
		Conditions	defined by Video Channel Owner
Includes	Event Report Requests		
		Verbs	1/ Publication VDI read
		Destination Addresses	address of the Video Channel Owner
Is Identified By	P-VDI identifier		
Is Signed By	P-VDI signature		Video Channel Owner signature

2.1.7 Subscription VDI created and injected by Analyst

Semantic relationship	Element	Sub-element	Comments
Has Conditions	SPARQL query		conditions on metadata and/or the issuer of the license of a video
Includes	Licenses		
		Issuer	Analyst
Includes	Event Report Requests		
		Verbs	1/ new Publication VDI matching subscription conditions
		Destination Addresses	address of the Analyst
Is Identified By	S-VDI identifier		
Is Signed By	S-VDI signature		Analyst signature

2.1.8 Subscription VDI created and injected by Video Channel Owner

Semantic relationship	Element	Sub-element	Comments
Has Conditions	SPARQL query		conditions on metadata and/or the issuer of the license of an analysis
Includes	Licenses		
		Issuer	Video Channel Owner
Includes	Event Report Requests		
		Verbs	1/ new Publication VDI matching subscription conditions
		Destination Addresses	address of the Video Channel Owner
Is Identified By	S-VDI identifier		
Is Signed By	S-VDI signature		Video Channel Owner signature

2.1.9 *Subscription VDI created and injected by Video Channel User*

Semantic relationship	Element	Sub-element	Comments
Has Conditions	SPARQL query		conditions on metadata of a posted analysis and/or a specific channel
Includes	Licenses		
		Issuer	Video Channel User
Includes	Event Report Requests		
		Verbs	1/ new Publication VDI matching subscription conditions
		Destination Addresses	address of the Video Channel User
Is Identified By	S-VDI identifier		
Is Signed By	S-VDI signature		Video Channel User signature

2.2 Relationships between VDIs

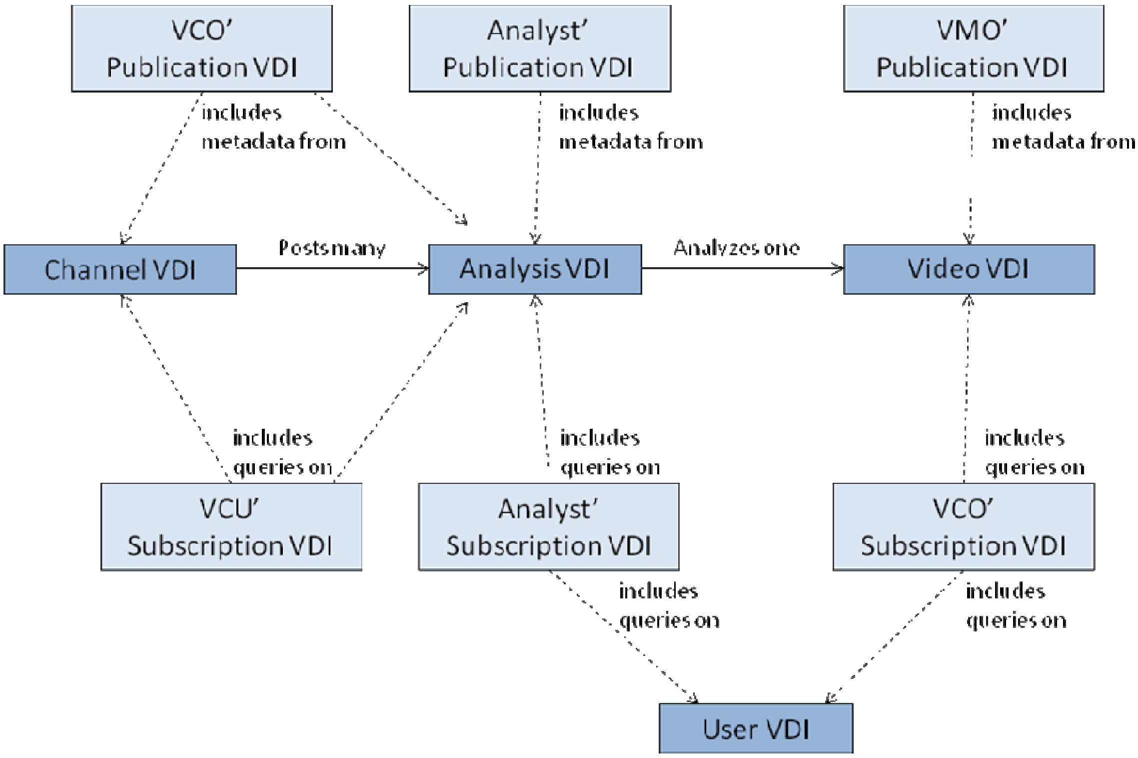

Figure 19 - Relationships between VDIs, trial 2

3 Augmented Lecture Podcast

3.1 VDIs used

3.1.1 Slide VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Slide identifier		
Contains	Sequence identifier		
Is described by	Slide metadata		e.g.: author, keywords
References	Slide resource		as PDF
Includes	Event report		
		Destination address	

3.1.2 Slides VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Slides identifier		
Contains	Slides sequence identifier		
Is described by	Slides metadata		e.g.: author, keywords, title, term
References	Slide sequence identifier (from Slide1 to SlideN; n=number of slides)		References VDI in 4.3.1
Includes	Event report		
		Destination address	

3.1.3 Video VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Video identifier		
Contains	Video sequence identifier		
Is described by	Video metadata		e.g.: file format, available

			resolutions, author, title, duration
References	Video resources		May reference multiple resources (i.e. for different resolutions)
Includes	Event report		
		Destination address	

3.1.4 Podcast Episode VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Podcast identifier		
Contains	Podcast sequence identifier		
Is described by	Podcast metadata		
		Synchronization Information	
		Other metadata	e.g.: author, title, date, duration
References	Video sequence identifier		
	Slides sequence identifier		
Includes	Event report		
		Destination address	
Includes	License		
		Issuer	
		Principal	
		Rights	e.g.: private use only, distribution rights only granted to service provider and specified organizations.

3.1.5 Annotation VDI

Semantic relationship	Element	Sub element	Comments

Is identified by	Annotation identifier		
Contains	Annotation Sequence identifier		
Is described by	Podcast metadata		
		Timestamp to Video	
		Annotation type	e.g.: private, semi-private, public
		Other metadata	e.g.: author, creation time, expiration date
Includes	Text resource		
References	Slide / Video sequence identifier and Slide/Video VDI identifier		
Includes	License		
		Issuer	
		Principal	
		Rights	Who is allowed to see the annotation (personal, shared, public annotation)
Includes	Event report		
		destination address	
Is signed by	VDI signature		Can be signed by specific user or can be signed anonymously

3.2 Relationships between VDIs

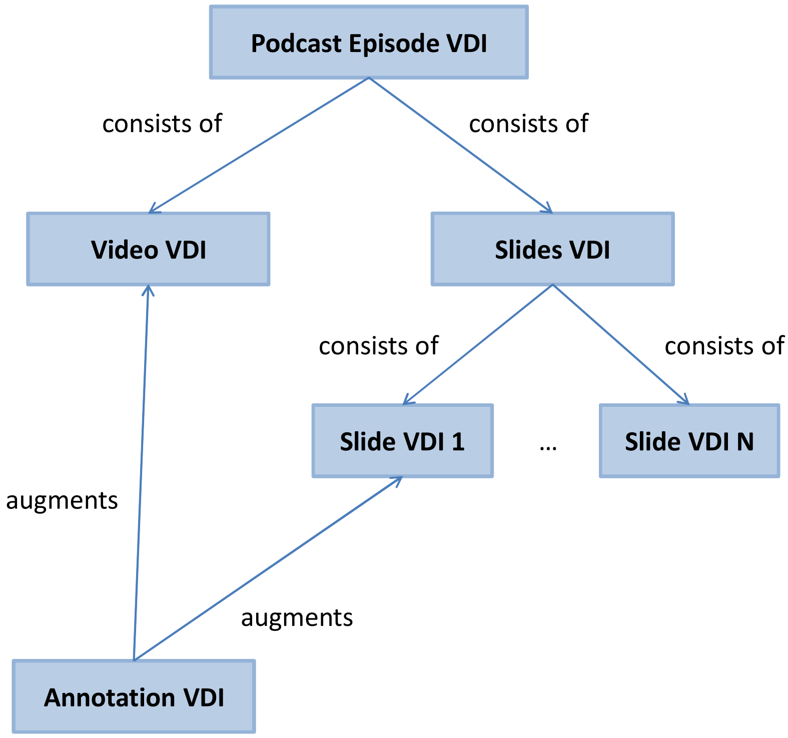

Figure 20 - Relationships between VDIs, trial 3

4 Smart Retailing

4.1 VDIs used

4.1.1 Product Type VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Product identifier		
Includes	Product resource		e.g.: images, user manual, other documentation
Is described by	Product metadata		
		Product information	e.g.: descriptions, dimensions, weight, features, etc.
Includes	License		
		Issuer	
		Principal	
		Rights	
		Conditions	
Includes	Event Report		
		Verb	
		Destination addresses	
Is identified by	VDI identifier		
Is signed by	VDI signature		

4.1.2 Product Instance VDI

Semantic relationship	Element	Sub element	Comments
Is identified by	Product identifier		
Includes	Product resource		e.g.: images, user manual, other documentation

	Other information		
		Barcode	
		Serial number	
		Warranty details	Date of purchase
Is described by	Product metadata		
		Product information	e.g.: descriptions, dimensions, weight, features, etc.
References	Product type identifier		
	Owner (Buyer)		Consumer information
Includes	License		
		Issuer	
		Principal	
		Rights	
		Conditions	
Includes	Event Report		
		Verb	
		Destination addresses	
Is identified by	VDI identifier		
Is signed by	VDI signature		

4.2 Relationships between VDIs

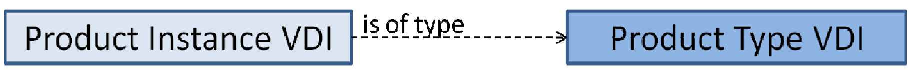

Figure 21 - Relationships between VDIs, trial 4