Publish/Subscribe over Information Centric Networks: a Standardized Approach in CONVERGENCE

Nicola BLEFARI MELAZZI[1], Stefano SALSANO[1], Andrea DETTI[1], Giuseppe TROPEA[1], Leonardo CHIARIGLIONE[2], Angelo DIFINO[2], Angelos-Christos G. ANADIOTIS[3], Aziz S. MOUSAS[3], Iakovos S. VENIERIS[3], Charalampos Z. PATRIKAKIS[3]
[1]*CNIT research unit of Rome Tor Vergata, Via del Politecnico, 1, Rome, 00133, Italy*
Tel: +39 06 7259.7501, Mail: blefari@uniroma2.it
[2]*CEDEO.net, Via Borgionera, 103, Villar Dora, Turin, 10040, Italy*
Tel: +39 011 935.04.61, Mail: info@cedeo.net
[3]*National Technical University of Athens, 9 Iroon Polytechniou str., Athens, 15780, Greece*
Tel: +30 210 7722424, Fax: +30 210 7724324, Email: aca@icbnet.ntua.gr

Abstract: Originally conceived as a "network of hosts", the Internet is evolving into an Internet of services, an Internet of media, an Internet of people and an Internet of "things". This implies a strategic shift from "host-centric" to "content-centric" and "data-centric" networking. CONVERGENCE proposes to enhance the Internet with a novel, information-centric, publish-subscribe service model, based on the Versatile Digital Item (VDI): a common container for all kinds of digital content, derived from the MPEG-21 standard. Results in terms of standardization activities and software implementation are presented.

Keywords: Future Internet, Information Centric Networks, Digital Items.

1 Introduction

The goal of the CONVERGENCE project is to enhance the Internet with an information-centric [1], publish-subscribe service model [2], based on a common container for any kind of digital data, including representations of people and Real World Objects (RWOs). We call this container the Versatile Digital Item (VDI), the basic unit of distribution and transaction in the CONVERGENCE network. VDIs can incorporate every possible kind of information, including signalling and control, and therefore minimize the need to store external information and states outside the data unit. The definition of VDIs is derived from the MPEG-21 Digital Item Declaration Standard [3].

The introduction of VDIs corresponds to a shift from "host-centric" to "information-centric" networking, in which the network layer provides users with content, instead of providing communication channels between hosts, and is aware of information about this content, at least in the sense it knows its name. This shift is analogous to the switch from circuit to packet switching: in circuit switching a PCM slot contains only user data; in packet switching an IP datagram also contains destination addresses. Similarly, in "information switching", the VDI contains a complete package of user data and meta-data describing content and how to handle it.

The VDI is a container to encapsulate any kind of digital information: not only classical media files, but also data about services, people and RWOs (e.g. items of merchandise identified with an RFID). VDIs bind *meta-information* (describing the content and structure of the item) and *resources* (other VDIs, audio, images, video, text, descriptors of RWOs, descriptors of people etc.). The meta-data describing the VDI may include: structural information; cryptographic keys allowing robust authentication and protection of information included in the VDI; rights information defining rights to use the item; an expiry date, support for "digital forgetting". VDIs are identified by a unique identifier, which is translated (or simply equal) to a network-level name used to route the VDI.

A key feature of CONVERGENCE is the support for a publish/subscribe service model: *publishers* advertise resources (data and service-access-points) on the system and *subscribers* express their interest in specific resources, and are asynchronously notified of events generated by publishers. Publish/subscribe effectively decouples the application end-points in space, time and synchronization. This allows for greater scalability, a more dynamic network topology and a much enlarged and flexible typology of services.

Every resource in CONVERGENCE is associated with a VDI, and VDIs are used to transport and represent both publications and subscriptions. Subscriptions express criteria that can be verified by inspecting VDI meta-data. Therefore, CONVERGENCE system supports *content-based* subscriptions as defined in [2].

Considering the goal to define an open system, the project is pursuing the standardization of this approach under the MPEG umbrella as discussed in section 3. A reference open-source implementation of CONVERGENCE is on-going; current results are described in section 4. Section 2 here below, describes the architecture.

2 Architecture and key features

A CONVERGENCE system consists of a set of interconnected peers, based on a 3-level architecture. From the top down (see Figure 1):

1. Application level. CONVERGENCE-compliant Applications package domain-specific resources and metadata into VDIs, and consume VDIs and their components, by extending basic middleware functional blocks and using its services to create descriptors. Hence both resources and VDIs are exchanged at the interface between Applications and Middleware. Re-usable Application elements are called Tools.
2. Middleware level. The CONVERGENCE Middleware (CoMid) is the level responsible for manipulating and processing VDIs. CoMid allows users to publish VDIs and search for them using semantic subscriptions. It builds on and extends the MPEG-M standard [4], which provides a distributed eco-system of Protocol Engines (PE), Technology Engines (TE) and Aggregated Services.
3. Computing Platform level. The Computing Platform level provides novel information-centric networking (CoNet) and secure handling (CoSec) of data. The computing platform also provides interfaces to access local hardware of the peer.

2.1 – The VDI

VDIs are XML structures containing: (a) *identifiers*; (b) (links to) *resources*; (c) (links to) semantically-rich *metadata* describing resources; (d) (links to) *licenses* expressing, by means of Rights Expression Language statements [5], what rights are given to act on resources; (d) *event report requests* (ERR) instructing a peer to issue an event report (ER) to a specific user/peer in the event certain actions (e.g. play, store or match) are performed on a resource.

2.2 – Descriptors and Dictionaries

CONVERGENCE supports semantic descriptions of resources. Both well-known (hence worldwide accepted for a specific domain), as well as user-designed ontologies, are typically used to describe resources. The former come from manufacturers, research institutions or service providers, while individual users create the latter, so that describing resources with metadata extracted from custom semantic taxonomies is a common case. The Community Dictionary Service middleware component (CDS TE), maintains dictionaries that help translate concepts and properties from one ontology model to another. The CDS is exploited when users describe resources and when descriptions of what is being published do not match the terms used in subscriptions. In these cases, the CDS attempts to translate between concepts and resolve the match.

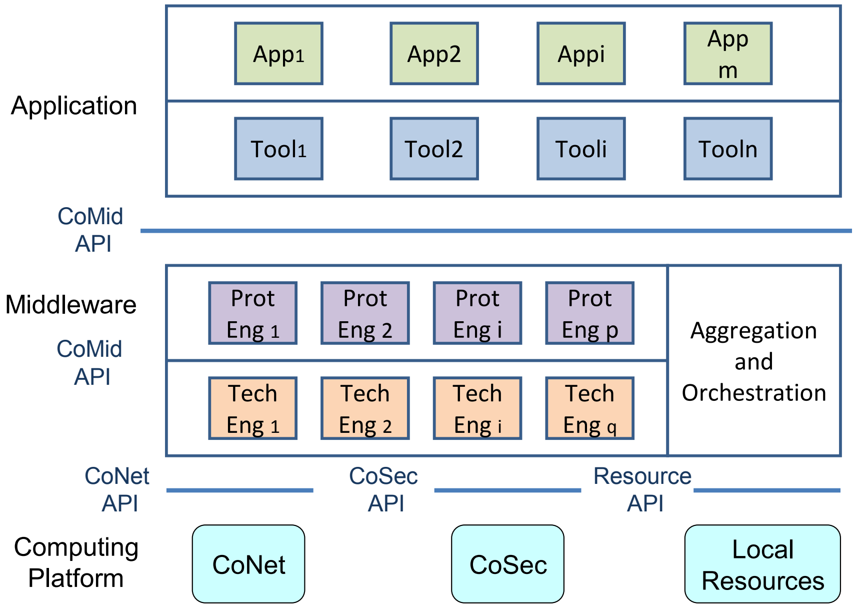

Figure 1. Architecture of a CONVERGENCE peer

2.3 – Publish/Subscribe

The key interaction pattern for CONVERGENCE users is based on a publish/subscribe paradigm. Users publish information about resources into a *semantic overlay* automatically maintained by the peers, the topology of which is based on the same semantic taxonomies employed for descriptors; peers join or leave the overlay regions, based on what users are currently publishing and subscribing to.
Users search for resources in specified regions of the semantic overlay and content can be delivered even when the match is not exact. An extended MPEG-M Search Content PE executes a semantic query, representing the subscription. The query is successful when published content matches the semantic subscription criteria. The Event Report TE takes care of notifying subscribers. Both publications and subscriptions are packaged and transported in VDIs, just like any other resource in CONVERGENCE, so that licenses are applicable to them, in order to restrict to certain sets of users the scope of publications, or limit their validity temporal span.

2.4 – Information Centric Network Infrastructure

Retrieval of resources and communication between peers is handled through the CoNet information-centric network. Peers simply refer to remote named-resources (as opposed to remote hosts as in current Internet), which can be (a) *named-data*, i.e. a sequence of bits, like a VDI or the resource the VDI refers to; (b) *named-service-access-points* (named-sap), i.e. a network endpoint from which a Protocol Engine receives CoMid messages.
In both cases, the named-resource is identified by a *network-identifier*, i.e. a *name* like “foo:VDI1”. In our current implementation, network identifiers coincide with the identifiers assigned to VDIs by CoMid services, but to ensure scalability and stability, middleware and

network remain conceptually decoupled: a VDI, which for the middleware represents a "CoMid data-unit", is simply a sequence of bits addressed by its network-identifier for CoNet, which has its own data-unit. CoNet is aware of the network location of named-resources and uses *routing-by-name* to route the request to the copy of the named-resource held by the network node that is closest to the requesting user. CoNet supports *built-in caching/replication* functions; to improve access to popular resources, the same named-resource can be replicated in different network nodes. CoNet provides users with access to the most convenient replica. Replica nodes could be pre-provisioned, as in Content Delivery Networks, or opportunistically selected by *in-network caching* mechanisms. In-network caching prevents denial of service events called "flash crowds", i.e. situations when a very large number of users simultaneously access a popular resource. Unlike state-full and off-the-shelf transparent proxy technologies, CoNet performs stateless caching. This speeds up caching and reduces the cost of implementation.
CoNet provides *information-based quality of service*. Network nodes can differentiate performance in terms of bandwidth and storage (caching) on the basis of the name of the resource they are serving. For instance, the named-data "foo:VDI-high-priority" might have a higher transmission priority and a higher probability of being cached locally than the named-data "foo:VDI-low-priority". Unlike current IP technology, information-based QoS mechanisms do not require complex and slow deep packet inspection (DPI).
CoNet also handles *digital forgetting.* Owners of named-resources may request CoNet to remove them from all serving-nodes and caches. This can be achieved either by specifying an expiration time for the request or by making an explicit request for removal.

2.5 – Security

Security is an essential feature of CONVERGENCE. The system's main security features are: i) assurance of VDI integrity (and authenticity); ii) governance of VDI access restrictions (confidentiality); iii) user identification and authentication; iv) issuing and enforcement of licenses; v) protection of user privacy; and vi) network security.
Most of these features are provided at the middleware level, some at the computing platform level. CoMid security features are provided by the Security TE, which in turn may exploit CoSec functionalities. Security features of the CoNet are independent from the middleware ones, and may in turn exploit CoSec, if needed.
Security TE can: i) create new credentials and manage certificates; ii) generate keys and encrypt/decrypt data or keys; iii) store confidential information (e.g. licenses and keys) in the secure repository; iv) certify the integrity of engines.
Other Engines rely on the Security TE to perform the following operations: i) signing of VDIs (VDI TE); ii) symmetric encryption/decryption of resources (Media Framework TE); iii) asymmetric encryption/decryption of a key (REL TE); iv) user Identification (Identify User TE); v) user authentication (Authenticate User TE).
CoSec is responsible for handling the majority of cryptographic protocols and security related tasks. Although the architecture diagram shows it as a single monolithic block, it has a distributed architecture encompassing several independent (and possibly distant) components, each of which includes software as well as hardware.
Most components of CoSec are located on client computers (e.g. end-user laptops), smart cards, application servers and network peers. The majority of protocols processed within CoSec involve several of these entities.
CoNet supports security and privacy mechanisms aimed at preserving the integrity of the networking service and, where required, the anonymity of owners and consumers of named-resources. A distinguishing aspect of CoNet security is the use of data-centric security: security information is embedded in CoNet data-units. Data-centric security makes it

possible for user and network nodes to verify the validity of named-resources, avoiding the caching and dissemination of fake versions. Protecting information at the source (i.e. protecting the data-unit) is more flexible and robust than delegating this function to applications, or securing only the communications channels.

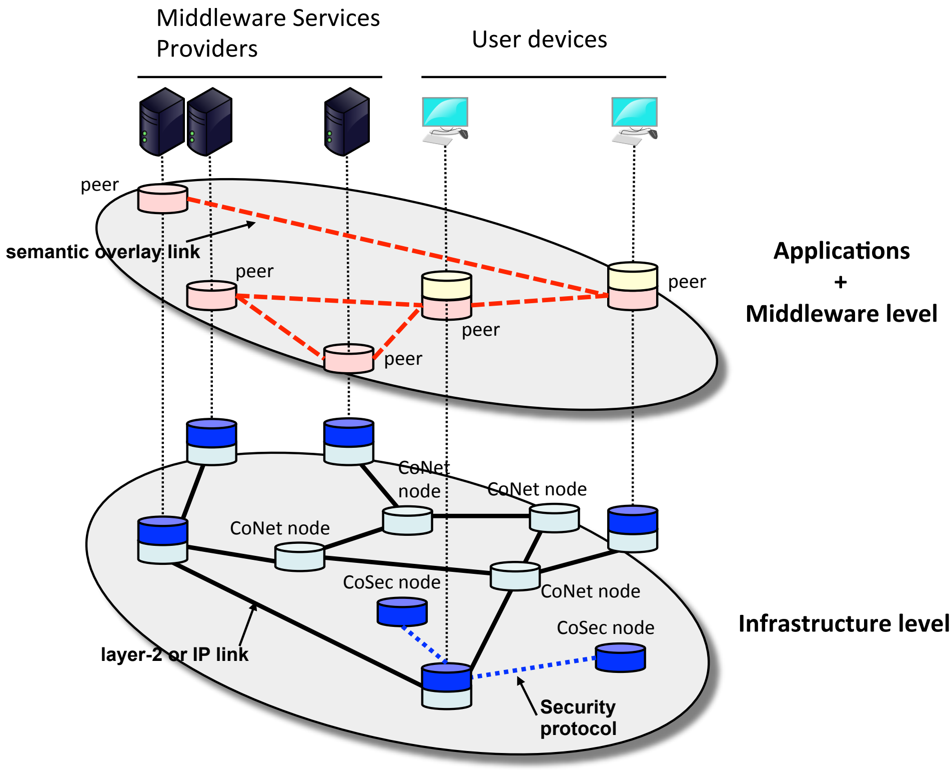

Figure 2. Distributed view of the CONVERGENCE architecture

Figure 2 above depicts the architecture from a distributed point-of-view. The functionalities at the computing platform level of each device are collectively distributed across an *Infrastructure level.* Special devices that only run CoNet (core network routers) or only run CoSec (specialized security servers) are referred to as *nodes.* Devices that additionally run the CoMid are referred to as *peers.* The middleware is **distributed across all peers**: however some of them can be more specialized in giving certain services, some of them (typically end-user devices) will additionally run CONVERGENCE-compliant applications. CoMid peers exchange messages with peer entities of the semantic overlay. CoNet nodes are connected via physical or traditional IP links and communicate over the information-centric network. CoSec nodes talk using distributed, specialized security protocols.

3 Standardization Results

CONVERGENCE has achieved some important results, concerning standardization activities based on its key requirements for (a) Digital Items to support semantic relationships and (b) for MPEG-M middleware to support publish/subscribe operations.
The standardization proposal for semantic relationships stems from our work on integration of DI and RDF/OWL (see [6]). It consists of a new element in the official DI schema definition, plus a simple ontology for Digital Items. It has been laid-out in the form of an Amendment to MPEG-21 part 3 standard for Digital Item Identification to introduce a new semantic Relationships element in the specification. This proposal for amendment was presented at the 97th MPEG meeting in Turin and 98th MPEG meeting in Geneva, and after some rounds of plenary discussion it has been moved to the Committee Draft status, with approval of Italy, Korea, Spain, US and Switzerland MPEG National Bodies. Most important, this proposal has already been referenced in other new initiatives of MPEG,

which deal with designing a standard for document preservation (Multimedia Preservation Description Information Requirements Document draft), because it allows for sequencing and versioning of Digital Items, key enablers of the preservation mechanisms.

The second major standardization proposal is for a new Post Content protocol and Overlay technology in the MPEG-M part 4 and part 2 specifications, respectively. It proposes to introduce them in the MPEG-M architecture, to consolidate the information-based resource discovery mechanisms that naturally fit in a large ecosystem of distributed MPEG-M devices. The current MPEG-M approach, even though it considers a distributed environment of multiple actors, it is designed in a server-client way, and does not propose any efficient distributed content discovery.

CONVERGENCE is suggesting to MPEG that MPEG-M should follow the evolution of cloud computing and define a standard way of:

- Creating and maintaining a cloud of MPEG-M devices
- Injecting (post) content into the cloud
- Accessing content that exists or *will exist* into the cloud

Standardization of discovery mechanisms of resources is crucial, in large distributed networks. The present approach of MPEG-M architecture neglects the dissemination of information about how to efficiently locate who in the network possesses specific and pertinent descriptions of some resource. Support of semantic searches over the DI space is thus practically not possible, unless a centralized and hardcoded, directory-based deploy is assumed. The proposed discovery functionality is based on a Post protocol and an Overlay technology to manage and insert content into the overlay. The key features this design has achieved are:

- Support for creation of a clustered overlay for efficient content distribution
- Support for dynamic systems (actors going in and out)
- Publish-Subscribe approach for content advertisement and discovery
- Integration of multiple technologies coming from literature and research projects

A key point is its ability to adapt both to fully symmetric approaches as well as to more centralized ones. For instance, on the one hand, this MPEG-M overlay technology can be based on a peer-to-peer, gossip based protocol, to support a semantic network of peers that publish or subscribe to content (this is exactly the current implementation in CONVERGENCE), which well represents the symmetric approach. On the other hand, a different implementation, but with **same interfaces**, could just as easily be used by a Content Delivery Network (CDN) provider as well as any network provider or even an ISP, to deploy reliable computing and content distribution clouds. This proposal is currently officially under discussion within the MPEG-M sub-group, and subject of National Bodies comments, prior to incorporation in parts 2 and 4 of the specification.

Further results are the incorporation of a CONVERGENCE-based use case of the MPEG-M middleware specification, to become Annex 1 of MPEG-M part 1, and the production of and IETF draft specifications for the CoNet information-centric technologies [7][8].

4 Implementation Results

Implementation of CoMid is integral with the on-going MPEG-M part 3 (Reference Software) development. CONVERGENCE has started from old MXM[1] implementations, tested, extended and integrated them to conform to the current status of the specification.

[1] MXM stands for MPEG eXtensible Middleware and it is the implementation branch of MPEG-M. Part 3 specifies the MXM project hosting the source code of the reference software.

Additionally, CONVERGENCE has specified and implemented new engines, outside the MPEG-M scope, exploiting the extensible nature of MXM.
CoMid implementation, following the MXM project setup, is split into four modules:

1. *Core*: this module defines the APIs (interfaces and abstract classes) of the engines. These APIs are either standard, extracted from MPEG-M part 2, or CONVERGENCE specific. Developers can rely on the core to use the middleware, independently of their implementation.
2. *Engines*: this module contains the implementation of all Protocol and Technology Engines. In essence, this is the implementation of the APIs defined in the core module.
3. *Dataobject*: this module contains the bindings of the XML schemas used for the protocols and the data structures.
4. *Webapp*: this module contains the server-side of the Elementary Services (see [4]).

A key aspect of the middleware architecture is that it can combine services or engines to compose aggregations and orchestrations respectively. Figure 3 presents an abstract example of a chain of Protocol and Technology Engines; MPEG-M does not yet define any standard way to implement the orchestrator and aggregator, although CONVERGENCE considers a BPMN [10] approach using executable processes. In detail, starting from MPEG-M parts 4 and 5, which are specifying the BPMN workflows for the elementary and the aggregated services, we are extending these workflows to executable ones, to be fed into a unified workflow environment that can run BPMN executable processes.

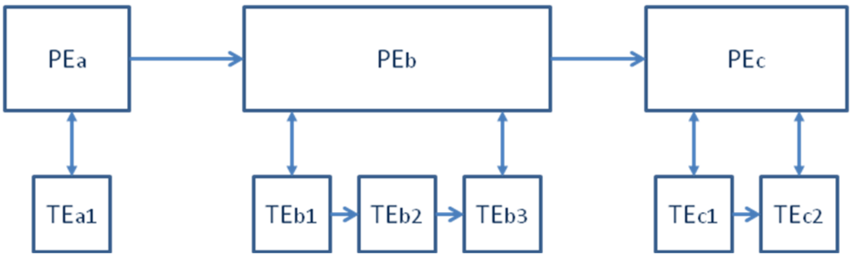

Figure 3. Chain of Protocol and Technology Engines

By further extending this approach, which is only about services, we are also introducing the concept of workflows of technology engines. This way, we can define complex workflows that describe a sequence of protocols and technologies to be used, independently of their implementation. Then, we can pass this workflow to the Aggregator/Orchestrator engine (CONVERGENCE is also taking care of the implementation of this standard component of the MPEG-M architecture) and execute it using the device's middleware.
As a design choice, middleware engines cannot have any interdependencies, so that a device can contain only a particular subset of them without forcing any dependencies. This is the reason why we have introduced the concept of the Dataobject: since there are interdependencies between the schemas used by the engines, if the schemas were defined in the context of the corresponding engine, we would have had dependencies between engines. Using the Dataobject approach, we are providing flexibility to the developers, since they can use the schemas locally and make calls to engines running remotely.
Finally, the middleware provides another level of flexibility concerning the decoupling of the engines specification from their implementation. The developer can rely solely on the core module (API specification) without caring about the actual implementation. This is accomplished through the use of the so-called MXMConfiguration file, which points to the implementations of the engines. This way, developers of CONVERGENCE-compliant Application can use particular parts of different middleware implementations, depending on what best suits their needs, by just switching pointers to the implementations they want, in the MXMConfiguration file.

Space constraints do not allow for detailed reporting of results, both in terms of technical solutions and performance, of the CoNet. The reader is referred to [9].
A prototypal application for publishing and subscribing of photographs has been shown at the 98th MPEG meeting, and is built upon a first packaging of the MPEG-M libraries and middleware, plus CONVERGENCE-specific extensions, which is called the Convergence Peer Kit. It is available for download at the project's website (see [11]).

5 Conclusions and Future Work

The most important activities of the project are now focused on evaluating the scalability of both the semantic overlay and the information-centric infrastructure up to a large number of devices and resources.
Additionally, a novel approach to enforcement of rights and licenses expressed within VDIs is under study, which is based on Attributes Based Encryption techniques.

6 Acknowledgements

This research was partially supported by the European Commission under contract FP7-ICT-2009-5 CONVERGENCE, but it does not necessarily represent the official position of the project itself and of its partners. Authors are solely responsible for the views, results and conclusions contained in this work. The authors would like to thank all people of the CONVERGENCE project.

7 References

[1] *Networking Named Content*. V. Jacobson, D. K. Smetters, et al., Fifth ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), 2009.
[2] *The Many Faces of Publish-Subscribe*. P.T. Eugster, P.A. Felber, R. Guerraoui, A. Kermarrec, in ACM Computing Surveys (CSUR), Volume 35, Issue 2, 2003.
[3] ISO/IEC 21000-2 – Information technology – Multimedia framework (MPEG-21) – *Part 2: Digital Item Declaration*.
[4] ISO/IEC 23006 – Information Technology – *Multimedia Service Platform Technologies* (MPEG-M).
[5] ISO/IEC 21000-5 – Information technology – Multimedia Framework (MPEG-21) – *Part 5: Rights Expression Language*, July 2003.
[6] *Exploring Semantic Relationships Across Internet Resources.* H. Castro, M. T. Andrade, F. Almeida, G. Tropea, N. Blefari Melazzi, L. Chiariglione, A. S. Mousas, D. I. Kaklamani. 7th International Conference on Next Generation Web Services Practices (NWeSP 2011), October 19-21, 2011, Salamanca, Spain.
[7] *IPv4 and IPv6 Options to support Information Centric Networking*. A. Detti, S. Salsano, N. Blefari-Melazzi. IETF, draft-detti-conet-ip-option, http://tools.ietf.org/html/draft-detti-conet
[8] Receiver driven trasport protocol for CONET ICN. S. Salsano, M. Cancellieri, A. Detti. IETF, draft-salsano-rditp, http://tools.ietf.org/html/draft-salsano-rditp
[9] CoNet technical report. http://netgroup.uniroma2.it/twiki/pub/Netgroup/CoNet/CONET_tech_report-v1ss.pdf
[10] Object Management Group/Business Process Management Initiative. http://www.bpmn.org/
[11] The Convergence Project, *deliverable D3.2 "System architecture"*. From the project's web site http://www.ict-convergence.eu/