CONVERGENCE: extending the media concept to include representations of Real World Objects

Nicola Blefari Melazzi

University of Rome, Tor Vergata. blefari@uniroma2.it

Abstract The Information Society relies on an increasingly diverse mass of digital information, a trend accentuated by the Internet of Things in which Real World Objects (RWOs) have their own Resource Identifiers and related data. In this setting, the goals of the CONVERGENCE framework are to enhance current media handling with new functionality and extend the traditional concept of media to include digital representations of RWOs. The CONVERGENCE framework is based on the concept of Versatile Digital Item (VDI), a structured package of digital content and meta-information, inspired by the MPEG-21 standard, but designed to address a broader range of application domains, including the management of RWOs in the Internet of Things. The VDI is supported by an innovative middleware and by tools and applications. The CONVERGENCE framework incorporates six innovations: 1) VDIs provide uniform mechanisms to handle different classes of information, including Real World Objects; 2) VDIs are intrinsically dynamic, allowing both providers and consumers to update content; 3) VDIs support "digital forgetting" (automatic "un-publishing", automated garbage collection of VDIs after a user-defined expiry date); 4) VDIs meet security and privacy needs of both professional and non-professional consumers and providers; 5) VDIs support new modes of semantic search; 6) VDIs allow easy sharing of information across multiple, heterogeneous devices. In this position paper we present the CONVERGENCE framework, its main concepts and ideas, some proposals to make it real, some use cases and we show how our proposed concept can be a fundamental enabler for the Internet of Things.

1 Introduction

One of the key enablers of today's media revolution has been the emergence of broadly accepted multimedia standards – in particular the standards produced by the MPEG community. A large part of the current MPEG standards are centred on "classical" needs of the media industry. Today, however, the distinction between

media products and other digital resources is increasingly blurred. In both cases, users need information resources to be constantly available, up to date and sharable. In both cases, they need to synchronize information over a complex mesh of devices, to guarantee its integrity and to exert control over the way it is accessed and used. In both cases, they need the ability to "revoke" information that is no longer true or valid or which they no longer wish to make available to others. In fact, these needs extend to digital information describing Real World Objects (RWOs) - products, companies, people, and locations that users wish to buy, meet or visit.

In brief, there is a convergence between the requirements of media users, and those of other consumers and producers of digital information.

Several of these needs are addressed by existing MPEG standards. MPEG-21 already defines standard ways of providing meta-information and standard ways of describing the content and structure of complex "Digital Items" [1]. The CONVERGENCE framework aims at extending the ability to manage and trade digital objects to a broader range of digital objects, including descriptors for Real World Objects (RWOs) and to extend the possible actions that we can perform on digital objects. We call these new, more adaptable, digital objects Versatile Digital Items (VDIs).

To achieve these goals, we will start from the Digital Item Declaration standard (ISO/IEC 21000-2) from MPEG, extend it to cope with requirements derived from new application scenarios, and develop a complete environment to handle VDIs. More specifically, the CONVERGENCE framework aims at:

1. handling new needs associated with the emergence of an "Internet of Things". In the Internet of Things, RWOs are enhanced with machine-readable digital identifiers such as barcodes and RFID tags which link to additional information describing their properties, dynamic state and context. The CONVERGENCE framework will define common mechanisms for handling descriptors for different classes of RWO. Such mechanisms will lay the basis to enhance existing information services that reference RWOs, by enabling a huge potential for integrating services on the basis of common data items. Example services include e-auction sites such as E-Bay, location based services, such as "Friend Finder" and even multiplayer games.
2. being intrinsically dynamic. In today's networked world, the content of the information exchanged between providers and consumers is increasingly volatile. Catalogs, CVs, technical specifications, descriptions of locations or metadata for physical items like books refer to a world that is continuously changing. The producers of information need the ability to update the information they have released and consumers need mechanisms to check if a digital resource is up to date, to request an update, and to select between several versions of the same item. Automatic updating of distributed information will thus be one feature of the CONVERGENCE framework. The update mechanism will include mechanisms allowing producers to "push" updates to consumers and enabling "consumers" to pull their own updates.

3. supporting “digital forgetting”, i.e. the guarantee that content generated at one period of a user’s life does not come back to haunt her. To meet this need, the CONVERGENCE framework will provide mechanisms allowing users to “un-publish” VDIs. A second mechanism will allow users to define expiry dates for whole VDIs or for specific items of information. Together these mechanisms will allow sites and services to perform automatic garbage collection deleting expired information. In this way, the CONVERGENCE framework will act as an enabling technology for site owners and regulators wishing to provide or enforce protection for user privacy.
4. providing novel security and privacy mechanisms supporting the needs of all users of VDIs (information providers and information consumers). The design will take account of the differing needs of professional information providers, whose products are designed for use by the general public, often for a fee, and non-professional providers, who will be enabled to express privacy requirements concerning the use of specific kinds of information and exert expiry dates for the validity of the information. An example is the protection of the privacy of users making queries about VDIs or requesting updates of VDIs. Similarly to a DNS server, which, by knowing the queries, knows the sites visited by users, it is technically possible that the creator or the maintainer of a VDI would know all the actions and therefore the interests of people requesting or updating her VDIs, a feature which may be helpful in some scenarios or has to be avoided in other scenarios. Additional mechanism will protect the privacy of information consumers, preventing unwanted profiling of their searches, requests for updates etc. Protecting the user information at the source (i.e. the VDI) is better than delegating this function to applications and makes the big difference between trustworthiness and trust.
5. incorporating metadata technologies from multimedia standards and Semantic Web technology, providing a homogeneous way of searching and handling digital items of all kinds. The framework will be integrated closely with RDF (Resource Description Framework) technology and will support ontologies for describing metadata. To be compatible with other popular solutions, it will also include a standard mechanism for folksonomy tagging, allowing users to “tag” (or search for) items of digital information across multiple web sites, services and applications.
6. facilitating users to contribute to the production of media. Many digital consumers have become producers of digital media, which they increasingly share with other users in Web 2.0 applications. The CONVERGENCE framework will provide users who generate their own content with a globally unique identity for their work together with standard mechanisms for bundling it with meta-data. In addition, the CONVERGENCE framework will provide users with a way to easily share information between heterogeneous devices and with other users.

2 The CONVERGENCE Framework

The framework consists of four main elements:

1. Versatile Digital Items (VDI): the basic unit of transaction and exchange within the framework.
2. Middleware: software allowing CONVERGENCE and third party applications to define and perform operations on VDIs. This includes network entities implementing the distributed CONVERGENCE logical architecture.
3. Tools and Applications: software tools and applications, based on the middleware and allowing end-users to manipulate VDIs.
4. A Community Dictionary Service (CDS): an RDF-based semantic framework for ontology sharing.

Fig. 1 shows the relationships among these elements and the way in which the CONVERGENCE approach extends the existing MPEG-21 framework. Below we describe the first three elements of our proposed framework.

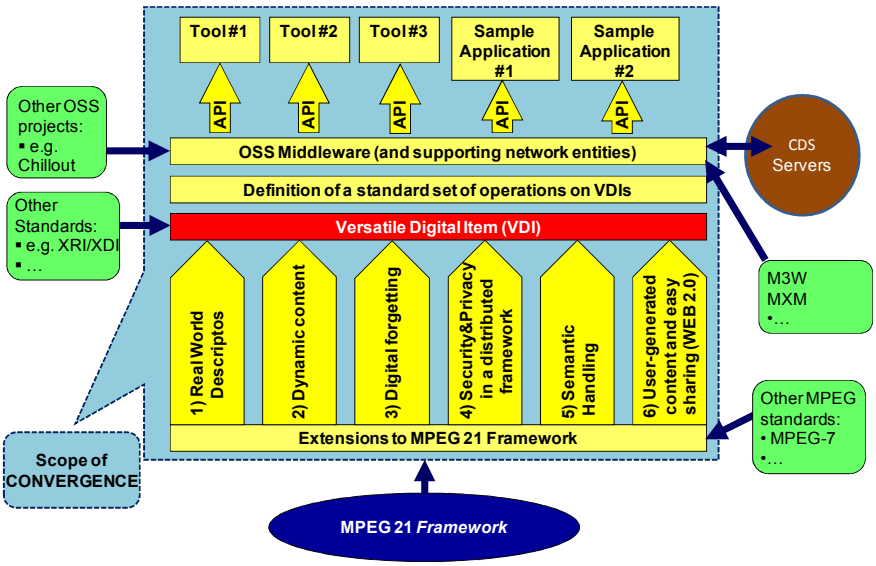

Fig. 1: CONVERGENCE framework.

2.1 Versatile Digital Items

Within the framework, the fundamental unit of distribution and transaction will be the Versatile Digital Item (VDI). The VDI is technically an extension of the MPEG-21 concept of a Digital Item. It is meant as a very general data format

holding information related to any virtual or physical item. Like MPEG-21 DIs, VDIs will bind together meta-information (describing the content and structure of the item) and resources (other VDIs, audio, images, video, text, complex descriptors of RWOs etc.). Like MPEG-21 DIs, they will include identifiers (e.g. Digital Item Identifiers as specified by ISO/IEC 21000-3) and statements expressed in a Rights Expression Language (REL) allowing users to define the ways in which they can be used. Unlike DIs, they will be deliberately designed to take account of the dynamic, infinitely variable nature of the information world beyond classical multimedia.

VDIs will be designed to provide broad support to references to RWOs such as products, locations and people, making it possible to store, synchronize and certify RWO descriptors in the same way as conventional media objects. In addition, CONVERGENCE will define the metadata included in MPEG-21 DIs to support the new functionality provided by the framework, for instance synchronization and deletion of VDIs that have already been released, and semantically enhanced search and interpretation of VDI metadata.

2.2 The CONVERGENCE Middleware

Like MPEG-21 DIs, VDIs are self-contained: all the information necessary to handle a VDI is contained in the VDI itself. This is analogous to the situation in connectionless packet-switching protocols where all the information necessary to transport packets is contained in the packets. The CONVERGENCE partners believe that, just as the state-less nature of IP packets have contributed to the success of the IP protocol, so, self-contained VDIs can contribute to the success of the CONVERGENCE framework.

But a data format on its own is merely a facilitator for operations on DIs. These operations have to be provided by a separate layer. The MPEG community has understood this need, and has already published the Committee Draft of MPEG eXtensible Middleware (MXM, ISO/IEC 23006) [2]. MXM's goal is to specify an architecture and a set of standardized APIs for handling a variety of MPEG technologies, including DIs. The goal of the CONVERGENCE middleware is to play a similar role with respect to VDIs.

The CONVERGENCE middleware will provide APIs to dynamically define and encapsulate new classes of content and related meta-information, to create VDIs packaging different classes of information resource, to guarantee their security and privacy and integrity, to name them, to support semantic interpretation of metadata and tags using Community Dictionary Services, to search for them, filter them, read and write their attributes and content, adapt them for use on different machines, copy them, test their validity and efficiently synchronize them across multiple machines.

Compared to MPEG-21 and MXM, the CONVERGENCE middleware will significantly extend opportunities for interactions among different actors along the

value chain.

Fig. 2 describes a classical value chain for an MPEG-21 DI, in this case a DVD. Actor 1 creates the DI, defining the title of the video, Rights and Permissions etc. (meta-information) and inserting video tracks (resources). Actor 2, a dubbing company, adds sound tracks (additional resources). Actor 3, the censorship board, adds information on the video's certification (new meta-information). The chain continues until DI reaches the final Actor (i.e. the consumer). The only way to edit or read a DI is locally, when it is physically available on a given actor's premises. Thus, a downstream actor has no way of modifying choices that have been made upstream. And once the DI has been delivered to the next actor down the chain there is no way the upstream actor can update the information he or she has provided.

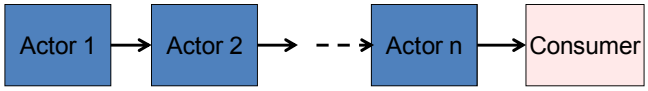

Fig. 2: Sequence of Actors involved in a classical value chain.

In CONVERGENCE, by contrast (see Fig. 3) VDIs are living objects, that can always be modified by any authorized actor of the value chain. In other words, CONVERGENCE extends the life of the value chain beyond the actions performed in the first delivery phase (i.e. when the VDI physically reach the consumer for the first time); as a consequence, CONVERGENCE enables new forms of business models. What is more, actors can operate on VDIs remotely, even when they are physically located on the premises of other actors. In brief, we can consider the CONVERGENCE value chain as a "bus" allowing any station attached to the bus (any actors) to interact with any other station.

The CONVERGENCE "bus" enables a broad range of operations that are not straightforward with conventional DIs. For example:

- Distributors (with appropriate rights) can remove certain portions of a movie they do not wish to distribute to their market;
- Producers and government agencies can inform customers who have bought a product (and its associated VDI) that the product has been found to be dangerous and should not be used;
- The author of a CV (distributed as a VDI) can update the copy she has sent to perspective employers;
- The buyer of a movie in English can sell it to an Italian user, acquiring and inserting the Italian audio track;
- A team of doctors can handle a patient's medical record (a VDI) collaboratively with each doctor updating the record locally and propagating the updates to the other copies.

As regards the technical approach for middleware, below we present some key features of our proposal.

Middleware model

Dynamism of information is easily managed when the involved resources are always available on the network. Nevertheless, we cannot assume that the devices storing VDIs will always be connected. To remove the "always-on" assumption, CONVERGENCE will develop an Asynchronous Middleware that "remembers" when an operation has been requested and ensures the request is executed when networking facilities allow it, that is when the middleware detects a specified event. To support this functionality, the middleware will adopt an event-based approach, implemented via Publish/Subscribe technologies.

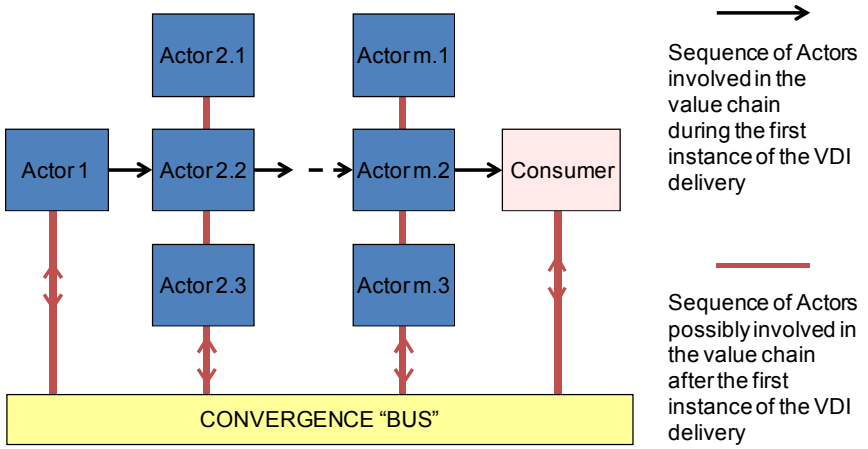

Fig. 3: Sequence of Actors in the CONVERGENCE value chain.

In the publish/subscribe paradigm, "subscribers" are interested in particular events generated by "publishers". Depending on the way subscribers express their interests, we can distinguish two main classes of publish/subscribe systems, namely topic-based and content-based systems. In topic-based systems, subscribers join a group interested in a specific topic of interest, and publications on the topic are broadcast to all group members. In other words, publishers and subscribers have to explicitly specify the group they wish to join. In content-based publish/subscribe systems, on the other hand, matching of subscriptions and publications is based on content and no prior knowledge is required. These systems also have the advantage that they allow users to set up permanently active queries referring to large numbers of different "terms". For a system such as CONVERGENCE, which provides descriptors for highly heterogeneous RWOs this is an extremely important feature. A second distinction is between centrally managed and distributed systems. The major disadvantage of centralized systems is their lack of scalability and fault-tolerance. Distributed systems overcome these limitations. Given the enormous scale of the "Internet of Things" that CONVERGENCE seeks to support, the CONVERGENCE middleware will adopt a distributed content-based strategy.

Networking Approach

The CONVERGENCE publish/subscribe mechanism will exploit distributed P2P overlay infrastructures. More specifically the CONVERGENCE middleware will implement a content-based publish/subscribe system based on a Distributed Hash Table (DHT). DHT approaches to network look-up are described in [3][4][5] and have been shown to be very effective in performing scalable, fault tolerant resource lookup on large peer-to-peer networks. Within the DHT framework, CONVERGENCE will implement content-based publish/subscribe mechanisms specifically designed to operate on VDIs. The approach chosen will allow CONVERGENCE to design effective subscription-peer mapping schemas to store subscriptions into the network, to define predicate-based expressive query semantics to efficiently match and deliver events to subscribers, and to ensure a uniform distribution of load among peers.

To improve scalability, P2P connectivity in CONVERGENCE will be organized as a hierarchically-structured federation of organizations and communities of interest (e.g., a virtual marketplace). In this way a limited set of nodes will be responsible for managing the federation for performing DHT-based publication/subscriptions and searches in a hierarchical way.

The DHT approach will enable the CONVERGENCE middleware to perform semantically enabled search for VDIs and to react to specific VDI related events. However, the middleware also needs functionality for remote handling of the VDIs that the DHT lookup has found. CONVERGENCE will thus exploit invocation protocols to communicate between the two middleware entities concerned. Candidate solutions for implementing this functionality include the XML Remote Procedure Call, JSON-RPC, SOAP as well as REST protocols. Fig. 4 summarizes the functionality to be provided by the CONVERGENCE Middleware.

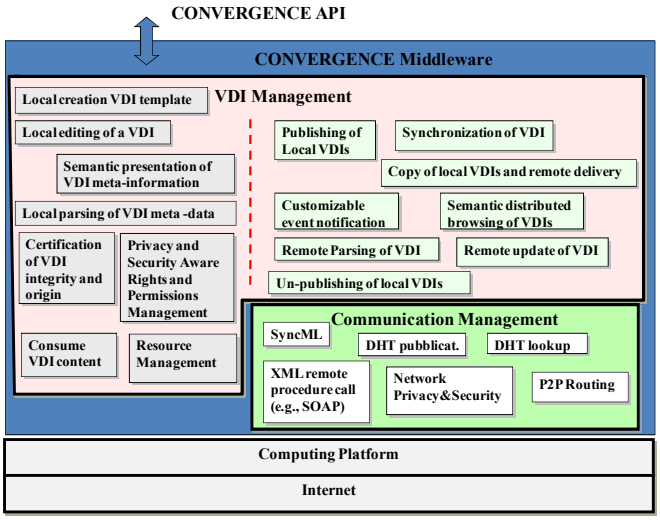

Fig. 4: Middleware functionality.

The middleware will offer two classes of functions: i) VDI management, enabling operations on local and remote VDIs; ii) Communication management, supporting distributed communication among middleware modules. The dashed line in Fig. 4 separates VDI Management modules that require the Communication Function (on the right) from those that do not require such functionality (on the left). A subset of these functions will be exported to the upper layer by means of standardized CONVERGENCE APIs.

2.3 Tools and applications

To demonstrate the usefulness and usability of the CONVERGENCE framework for developers and end-users, CONVERGENCE will design three general purpose tools that exploit the CONVERGENCE middleware to operate on VDIs.

- VDI Creator: a tool which takes a collection of digital information and wraps it into the VDI format;
- VDI Manager: a tool to operate on VDIs, where the operations allowed depend on the invoker and the rights defined in the VDI;
- VDI Search: a tool to find VDIs in a worldwide distributed storage infrastructure.

A key goal will be to offer an effective implementation of user security, rights and privacy, defined inside the VDI. The operations allowed in VDI Manager will depend on the rights of the invoker defined in the VDI. For example a typical publisher of VDIs would have the right to push updates to subscribers, change the metadata etc; a typical subscriber would have the right to view VDIs, create links to them and pull updates from the publisher. There are many possible variants to this model. It is quite possible, for instance, that an online artist might authorize other (unknown) artists to modify and redistribute her work.

The CONVERGENCE framework will provide many different ways of creating new applications. The simplest and most flexible way will be to use CONVERGENCE tools in combination with other standard software such as a browser. Alternatively, application software can use scripts or make direct calls to the APIs provided by the middleware. These will take the form of Web services, providing loose coupling with operating systems, programming languages and other related technologies, and facilitating the creation of Mashups.

3 Use cases

The CONVERGENCE framework and the VDI concept offer novel features and functionality of clear benefit to end-users. Below we describe use cases that illustrate these features.

A Dynamic Logbook

This use case illustrates the use of VDIs as a "dynamic logbook" for a technological device. In this example, the device is a handset but it could also be a washing machine, a dishwasher, a digital camera, a car. The functionality provided would be useful in a broad range of business to consumer and business to business applications.

Every handset that Big Telco sells to its customers comes with a pre-installed "digital logbook" – a VDI associated with the serial number for the phone. The logbook stores details of the date of purchase, the warranty, the firmware, pre-installed software, a user guide, a maintenance record; When the user purchases a SIM from an operator the logbook can be updated to include the SIM code and the owner's tariff plan (another VDI). Optionally, if the user agrees, the VDI may also contain personal identity information, which she can make available to some users and not to others.

Copies of the logbook are stored at several different locations. The manufacturer has a copy, Big Telco has a copy, the distributor has a copy, the operator has a copy, and a copy is pre-installed on the phone. When a customer buys and takes the phone home, she waves it over a reader and transfers a copy of the logbook to her local PC.

From this point on, each of the actors along the value chain can change the parts of the logbook for which it is responsible. The manufacturer can update the user manual. The phone operator can update the information on the user's tariff plan. If the user sells the phone to a friend she can change the ownership data. The different actors can use CONVERGENCE search facilities to identify phones with certain characteristics, and send messages which will appear to their owners (even if the identity of the owner is not known). Thus, a manufacturer who discovers a dangerous battery fault, can send a recall message to owners of phones with the defective part; an operator who has cut charges for certain calls can target owners on a certain tariff plan; a distributor can warn a customer when her warranty is about to wear out.

When a customer experiences problems with a phone, she sends a copy of the logbook to the call-center operator. The operator feeds the logbook to special simulation and diagnostic tools, which rapidly reproduce the problem and propose a solution.

It is important to stress that the functionality exemplified in this use case (i.e. a communication link between the producer of an object and the user(s) of that object), could be very useful in numerous situations.

Digital forgetting –automatic garbage collection

In the past when we wrote a diary or a letter to a friend, when we stayed up till the small hours talking politics over wine or beer, what we wrote and said remained in private drawers and the memory of friends. While today, when we write a blog or an email, or "chat" with a friend, we cannot know who will read it twenty years later. To fully enjoy the freedom offered by modern communications, to use it to the full, we need a right to "digital forgetting". In this use case,

we illustrate how CONVERGENCE can contribute to this goal.

Jaap is 20, very bright, cynical, and a natural rebel. When his father comes home with his friends from the board room, he listens in to their conversations. He hears how easily they talk about billions of dollars, how indifferent they are to the effects of what they are doing. But what strikes him is the way they talk, the tone, the language. Jaap has a good ear and a quick pen. His blog, "The Bank", – is a huge success – not least among his father's junior staff who think it is very true and very funny...

But now many years have passed and Jaap needs a job. He is no less bright, no less cynical and no less of a rebel than before, but it would be better if future employers did not know. "The Bank" is only one of his blogs. Going through everything he has ever posted takes two days, and for some of the postings there is no remedy. He can't remember the web sites address, or he's lost the password, or the service no longer exists. Fortunately Safeblog – his blogging service – is CONVERGENCE-based. VDI Search gives him his old postings. Thankfully he clicks on the "unpublish" button. A second later, the postings are flagged for deletion, not just on Safeblog, but on every other site that has a copy of the postings: friends' home pages, search engines that have cached copies, even the WayBack Machine. What happens next will depend on the policy of the site, but most respectable sites respect users' wishes, and in the Netherlands it is now a legal offence for companies to use "unpublished" digital information on job candidates.

Jaap is relieved he has managed to unpublish "the Bank" but he also knows he's run a huge risk. For his next very personal postings, he sets an expiry date just three months after publication. CONVERGENCE makes him feel safe.

VDIs in a store

Goods entering the store will be stamped with a RFID tag and associated with a VDI. The VDI will store information about the item's type, model/ version/ features/ firmware. The VDI will also store additional information such as stock entry date, expiration date (if applicable), inventory number or location in the storage area. This will make it easy to find the item via simple search processes.

Once the item is on the shelf, employees will be able to scan the RFID tag to obtain the associated VDI and use it to check/update the shelf stock or consult the inventory. Customers will be able to interact with the VDI on the spot, either by using their own CONVERGENCE-enabled smart-phones or by using an interactive display mounted on the shelf. They will be able to check the item's features, query for additional information and search for other compatible/complementary items, thus enabling cross-selling.

Once the item is checked-out, the information contained in the associated VDI will be updated in several ways. The store internal information (ex. inventory number) will be deleted or encrypted (in case of item returns), the receipt number stored and the ownership of the VDI changed. Other possible operations include notification of sales to the producer/supplier, and automated generation of warranties, stored directly on the VDI. Such warranties avoid loss of the warranty certificate by the customer.

Once home, customers will be able to search the web for further products details, manuals, new versions, firmware updates, product-related forums or users that bought similar items, related to the VDI. If the user brings the item to a CONVERGENCE-enabled repair shop, the technicians will use the associated VDI to check for warranty issues or search for technical information (schemes, parts details, repair methods, etc). The digitally signed details of the intervention will be stored in the VDI, thus keeping a complete history of the item lifecycle.

4 Conclusions

We believe that the CONVERGENCE concept can be a fundamental enabler for the Internet of Things as it provides both a common and flexible format for representing information related to Real World Objects and sophisticated functionality to handle such representations.

In addition, we remark that all the information necessary to handle a VDI is contained in the VDI itself. This makes it possible to handle networking functions at the "VDI layer", i.e., at the application layer, following the trend of overlay communications. In this setting, networked VDIs can support innovative scenarios and applications, helping boosting Internet of Things exploitation.

Acknowledgments **This position paper is based on a project proposal. The author thanks all the partners that contributed to the proposal and acknowledges specific contributions and suggestions to this paper from the following individuals: Maria Teresa Andrade, Leonardo Chiariglione, Andrea de Polo, Heinrich Hussmann, Dimitra Kaklamani, Stelios Pantelopoulos, José Ribas, Carsten Rust, Peter Stockinger, Mihai Tanase, Richard Walker, Andrea Detti.**

References

[1] ISO/IEC 21000 (MPEG-21) series of International Standards, ISO/IEC 21000-2:2005, http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=41112
[2] http://www.chiariglione.org/mpeg/working_documents/explorations/mxm/MXM-Reqs.zip
[3] D. Tam, R. Azimi, and H.-A. Jacobsen, Building Content-Based Publish/Subscribe systems with distributed hash tables, in Proceedings of the 1st Intl. Workshop on Databases, Information Systems, and P2P Computing (DBISP2P), Berlin, Germany, Sept. 2003
[4] Terpstra, W.W., Behnel, S., Fiege, L., Zeidler, A., Buchmann, A.P., A peer-to-peer approach to content-based publish/subscribe. In: Proc. of Workshop on DEBS, San Diego, CA (2003)
[5] X. Yang and Y. Hu, A DHT-based infrastructure for content-based publish/subscribe services, in Proceedings of the 7th IEEE International Conference on Peer-to-Peer Computing (P2P'07), Galway, Ireland, Sep 2-5, 2007.