Signals and Communication Technology

Fernando Almeida · Maria Teresa Andrade
Nicola Blefari Melazzi · Richard Walker
Heinrich Hussmann · Iakovos S. Venieris
Editors

Enhancing the Internet with the CONVERGENCE System

An Information-centric Network Coupled with a Standard Middleware

Springer

Signals and Communication Technology

Fernando Almeida · Maria Teresa Andrade
Nicola Blefari Melazzi · Richard Walker
Heinrich Hussmann · Iakovos S. Venieris
Editors

Enhancing the Internet with the CONVERGENCE System

An Information-centric Network Coupled with a Standard Middleware

Springer

Signals and Communication Technology

For further volumes:
http://www.springer.com/series/4748

Fernando Almeida · Maria Teresa Andrade
Nicola Blefari Melazzi · Richard Walker
Heinrich Hussmann · Iakovos S. Venieris
Editors

Enhancing the Internet with the CONVERGENCE System

An Information-centric Network Coupled with a Standard Middleware

FEUP

Springer

Editors
Fernando Almeida
Maria Teresa Andrade
Telecom and Multimedia Division
INESC Porto Campus da FEUP
Porto
Portugal

Nicola Blefari Melazzi
Dipartimento di Ingegneria Elettronica
Università degli Studi di Roma
Tor Vergata
Rome
Italy

Richard Walker
Blue Brain Project
Ecole Polytechnique Federale de Lausanne
Lausanne
Switzerland

Heinrich Hussmann
Medieninformatik
Ludwig-Maximilians-Universität
Munich
Germany

Iakovos S. Venieris
School of Electrical and Computer Engineering
National Technical University of Athens
Athens
Greece

ISSN 1860-4862 ISSN 1860-4870 (electronic)
ISBN 978-1-4471-5372-6 ISBN 978-1-4471-5373-3 (eBook)
DOI 10.1007/978-1-4471-5373-3
Springer London Heidelberg New York Dordrecht

Library of Congress Control Number: 2013953644

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

This book describes the main findings and achievements of the CONVERGENCE project, co-funded by the European Union, within the seventh Framework Programme. The project was conceived and submitted to the European Union services in October 2009. It started in June 2010, and ended in May 2013.

The rationale of the project lies in the evolving nature of the Internet. The Internet was originally conceived as an "Internet of Hosts", whose underlying protocols were designed to support exchange of simple unstructured information between well-identified nodes. Today, by contrast, it is becoming an Internet of Things (devices and appliances associated with their own IP address), an Internet of Services (in which users in different localities access different functionalities on different hosts), an Internet of Media (shared and managed across different networks) and an Internet of People (boosted by the explosion of social networking and the emergence of the Web 2.0 paradigm). In this evolving scenario, the key elements are no longer "hosts" but data and services. In other words, what we are observing is a shift from a "host-centric" Internet to a "data-centric" or "content-centric" or "information-centric" Internet.

A new Information-Centric Internet should support not only traditional service models such as client–server and conversational services but also peer-to-peer and publish–subscribe service models. In client–server models, the server provides a function or service to one or many clients, which initiate requests for such services. In peer-to-peer models, each node plays both the role of client and server for other nodes in the network allowing shared access to various resources such as files, peripherals and sensors without the need for a central server. In publish-subscribe models, subscribers register their interest in an event, or a class or a pattern of events, and are subsequently asynchronously notified of events generated by publishers. In practical terms, publishers (i.e. information providers but also simple users) make available in the network content that they produce; subscribers declare their interest for a given content or a pattern or a class of content. Publications and subscriptions are fully decoupled in space but also in time. Normally, subscriptions precede publications, and in this case the subscriber will receive the required content only if and when a publisher will make that content available in the network. In addition, we envisage that subscriptions can also follow publications, and in this case the subscriber will promptly receive the required content as in a traditional request/response or pull model. Subscribers can express interest both for known specific content (e.g. the repair manual for a piece of equipment) or content satisfying a given search criteria (e.g. all songs by a certain musician).

A new Information-Centric Internet should be able to handle any kind of digital data: classical media contents, data about services, but also digital representation of real-world objects (e.g. items of merchandise identified with an RFID) and people (e.g. profiles). Thus, a new Information-Centric Internet should provide structured and flexible data units allowing to contain and transport different kinds of information, each with its own characteristics and needs, together with meta-information, describing the information itself, its characteristics and needs, and possible actions that the network must exert on it.

More specifically, we believe that a new Information-Centric Internet should allow to:

- create a data unit, in a modular and flexible way, defining licenses and rights that apply to the information that it contains;
- sign and/or encrypt a data unit;
- support client/server, conversational, peer-to-peer and publish–subscribe service models; the latter implies the ability to publish a data unit and to subscribe to a data unit, meeting specified criteria;
- search and retrieve a data unit, with the meta-information contained in the data unit easing semantic searches and operation of search engines;
- verify the authenticity of a data unit;
- allow the publisher or owner of a data unit to monitor its use in the network;
- allow to communicate with owners or publishers of a data unit;
- versioning a data unit and linking it to other data units;
- update a published data unit;
- delete a data unit, allowing implementing digital forgetting and garbage collection, deleting 'expired'/obsolete contents from the whole network;
- do all of this in an efficient way, simplifying operations for all the involved players, including providers, business users and final users.

Now, to reach these objectives and provide the functionality and service models listed above, we were faced with two alternatives: (i) design and deploy different applications, each with its own data unit, operating on top of the current Internet in an overlay fashion and adding more and more patches to the basic operation of the network layer; (ii) define a common, standard unit of distribution and transaction; define a middleware implementing complex functionality at the service of all interested applications; transform the current network layer in a new information-centric network.

The project CONVERGENCE adopts the second alternative, providing new results in the following four areas.

CONVERGENCE defines a unit of distribution and transaction, called Versatile Digital Item (VDI), which contains all necessary supporting information, including signalling, control and security/privacy information.

CONVERGENCE defines a middleware supporting sophisticated functionalities, which we think are too complex to be implemented at the network layer in all routers, namely publish/subscribe services, searching functions, security functions; the data unit of our middleware is the VDI. It is important to observe however that not all CONVERGENCE communications must necessarily use a publish/subscribe paradigm. The CONVERGENCE middleware also accepts direct requests to immediately provide specific requested data, with a traditional request–response service model.

CONVERGENCE defines an information-centric network (ICN) layer. Information Centric Networking (ICN) is a new paradigm in which the network layer provides users with content, instead of providing communication channels between hosts, and is aware of the name (or identifiers) of the contents. The basic functions of an ICN infrastructure are to: (i) address content by adopting an addressing framework based on names, without a reference to the current content location (i.e., location-independent names); (ii) route a user request, based only on the content-name, towards the "closest" location containing the required content; potential locations include not only the origin server of that content but also network caches or even devices of other users that downloaded the same content beforehand; (iii) deliver the content back to the requesting host.

CONVERGENCE defines specific tools and of applications showing the CONVERGENCE potentiality. Tools are reusable Application elements, which facilitate reuse of code in Applications; an Application can make use of several tools. Our tools and applications exploit the VDI concept and make use of our middleware and network functionality, so offering to end-users the advantages brought about by our system. The project designed and implemented four main applications to show the usefulness of CONVERGENCE in four real-life scenarios. The four scenarios are: (i) management of audio–visual material; (ii) management of a large photo archive; (iii) customer relationship management and logistics for the retailing sector, exploiting information about Real World Objects; (iv) augmented lecture podcast service, enabling a collaborative learning environment. Two other applications have been built later on by integrating the four main applications; the first integrated application merges the first (video) and fourth (podcasts) original applications; the second integrated application merges the second (pictures) and third (retail) original applications. The aim of the integrated applications is to show that our system is flexible enough to combine different applications in one and to exploit common VDIs.

The end result of the project is the definition and evaluation of a complete ICT system able to provide the following advantages.

Advantages Deriving from Using an Information-Centric Network Layer

1. Efficient content-routing. Even though today's Content Delivery Networks (CDNs) offer efficient mechanisms to route contents, they cannot use network resources in an optimal way, because they operate over-the-top, i.e. without knowledge of the underlying network topology. ICN would let ISPs perform native content routing with improved reliability and scalability of content access. This would be a built-in facility of the network, unlike today's CDNs;
2. In-network caching. Caching enabled today by off-the-shelf HTTP transparent proxies requires performing stateful operations. The burden of a stateful processing makes it very expensive to deploy caches in nodes that handle a large number of user sessions. ICN would significantly improve efficiency, reliability and scalability of caching, especially for video;
3. Simplified support for peer-to-peer like communications, without the need of overlay dedicated systems. Users could obtain desired contents from other users (or from caching nodes) thanks to content-routing and forward-by-name functionality, as it is done today with specialized applications, which, once again, do not have a full knowledge of the network and involve only a subset of possible users;
4. Simplified handling of mobile and multicast communications. As regards handovers, when a user changes point of attachment to the network, she will simply ask the next chunk of the content she is interested in, without the need for storing states; the next chunk could be provided by a different node than the one that would have been used before the handover. Similar considerations apply for multicasting. Several users can request the same content and the network will provide the service, without the need for overlay mechanisms, multicast is an inherent capability in ICN;
5. Content-oriented security model. Securing the content itself, instead of securing the communications channels, allows for a stronger, more flexible and customizable protection of content and of user privacy. In today's network, contents are protected by securing the channel (connection-based security) or the applications (application-based security). ICN would protect information at the source, in a more flexible and robust way than delegating this function to the channel or the applications. In addition, this is a necessary requirement for an ICN: in-network caching requires to embed security information in the content data-unit, because content may arrive from any network or user node and we cannot trust all nodes; thus, end-users must be able to verify the validity of the received data; caching nodes must make the same check, to avoid caching fake contents;
6. Content-oriented quality of service differentiation (and possibly pricing); provision of different performance in terms of both transmission and caching. Network operators (especially mobile ones) are already trying to differentiate quality and priority of content, but they are forced to use deep packet inspection

technologies. ICN would let operators differentiate the quality perceived by different services without complex, high-layer procedures and off-load their networks via caching, a very handy functionality, particularly for mobile operators who can differentiate quality and priority of content transferred over the precious radio real estate;

7. Content-oriented access control, providing access to specific information items as a function of time, place (e.g. country) or profile of user requesting the item. This functionality also allows implementing: (i) digital forgetting, to ensure that content generated at one period in a user's life does not come back to haunt the user later on, and (ii) garbage collection, deleting from the network expired information;
8. Possibility to create, deliver and consume contents in a modular and personalized way; ICN provides opportunities for better customization of the interests of users and the content that is published by providers. This will enable more efficient consumption of content because of better "granularity" in how content is described and identified;
9. Network awareness of transferred content, allowing network operators to better control information and related revenues flows, favouring competition between operators in the inter-domain market and better balancing the equilibrium of power towards over the top players;
10. Support for time/space-decoupled model of communications, simplifying implementations of publish/subscribe service models and allowing "pieces" of network, or sets of devices to operate even when disconnected from the main Internet (e.g. sensors networks, ad hoc networks, vehicle networks, social gatherings, mobile networks on board vehicles, trains, planes). This last point is a very important one, especially to stimulate early take up of ICN in selected (and possibly isolated) environments.

A final overall advantage of ICN, which in a way comprehends the specific advantages listed above, is a simplification of network design, operation and management. Currently, content and service providers have to "patch" shortcomings and deficiencies of IP data delivery by using several "extra-IP" functionalities, such as HTTP proxies, CDNs, multi-homing and intra-domain multicast delivery, to name a few. This implies the involvement of several parties, the use of several specific protocols, the deployment of ad hoc devices and the interplay of different functionalities, often offered and managed by different companies and businesses. Apart from technical complexity, such operations also add management and administrative complexity. In an ICN environment, such diverse functions can be integrated in the network in a smooth and seamless way, e.g. by supporting inherently data replication, caching, multi-homing and multicast delivery.

Advantages for End-Users

CONVERGENCE allows end-users to:

1. Publish VDIs to the CONVERGENCE networks, using the same publish/subscribe services for different categories of information;
2. Manage and protect the information contained in VDIs, using the same security and privacy standards for all transactions;
3. Define and enforce licenses;
4. Monitor the use of published materials;
5. Communicate with the users of VDIs;
6. Operate on VDIs at all stages in the data life cycle, examine different versions of the same VDI;
7. Update VDIs that have already been published;
8. Delete VDIs that have already been published (digital forgetting);
9. Search for VDIs using semantic search mechanism;
10. Subscribe to VDIs;
11. Define how much identity information they wish to disclose when publishing or searching for information (multiple options available from full anonymity to full disclosure);
12. Enjoy better performance thanks to ICN;
13. Exchange data with nearby connected devices also when disconnected from the main Internet.

Advantages for Developers

1. Exploit the functionality of the CONVERGENCE middleware avoiding the need to use proprietary solutions;
2. Develop basic functionality using the CONVERGENCE tools and focus on adding value to the CONVERGENCE framework.

Advantages for Network Operators

1. Distribute content over the network more efficiently, reducing costs;
2. Differentiate Quality of Service for different categories of content;
3. Exert better control over information transfer and related revenues flows;
4. Compete more effectively with "over the top" players.

This book is organized into ten chapters. The introductory chapter contextualizes the work developed by the CONVERGENCE consortium, analysing the main limitations of the current Internet architecture to support the increasingly stronger user- and content-centricity trends. It provides an overview of the emergent Information Centric Networking (ICN) paradigm, addressing design principles, evolution scenarios and briefly describing the main research initiatives conducted in the last years.

Consequently, the next eight chapters strive to describe to the reader how CONVERGENCE has embraced the quest for answering the above-referred challenges and to contribute with innovative solutions for the future Internet. These chapters provide high-level descriptions of the concepts and design principles behind CONVERGENCE as well as details on how the different components operate individually and how they simultaneously co-operate among them, to deliver the desired functionality whilst facilitating the deployment of advanced business models. The conclusions chapter summarizes the main achievements and contributions of CONVERGENCE towards the future Internet, drawing paths for using and extending the CONVERGENCE outcomes.

Chapter 2 describes the overall architecture of the CONVERGENCE system. The objective of this chapter is to provide a concise, yet complete, understanding of the concepts and high-level functionality of the CONVERGENCE architecture and how its components interact in concrete deployment scenarios.

Chapter 3 covers the design principles, functionality and implementation details of the CONVERGENCE network level, the CoNet. It shows how this Information-Centric Network extends the content-centric paradigm in several aspects, including routing scalability and security handling, among others, whilst providing a graceful incremental solution, backwards compatible with the current Internet.

The Content Level of CONVERGENCE, conceived and implemented as the CONVERGENCE Middleware (CoMid), is described in Chap. 4. This chapter starts by presenting MPEG-M, which is the standard that provides the foundations for CoMid. It then describes the key components of the Content level, which comprise a diversified set of middleware engines to manipulate Versatile Digital Items (VDIs) and extending MPEG-M, as well as semantic tools, namely the Community Dictionary Service (CDS). The chapter describes how the semantic mechanisms are used together with the standardized technology to support the semantic publish–subscribe paradigm. It explains how CONVERGENCE manipulates rich metadata within a content-aware semantic overlay, running on top of the information-centric networking platform, namely on top of the CoNet.

Chapter 5 provides the definition of the Versatile Digital Item (VDI), the basic unit for data distribution used within the CONVERGENCE system. It explains how the VDI extends the scope of the MPEG-21 Digital Item specification to provide a self-contained data package that can be used to encapsulate any kind of digital information in an information centric, publish–subscribe framework.

Security aspects are analysed in Chap. 6. This chapter presents the distributed architecture of the CONVERGENCE security core component (CoSec), describing the use of smart cards as a secure token to provide sensitive security functions on a tamper-resistant device. The chapter starts by introducing the concepts and the architecture of the security infrastructure, then presenting the high-level security functionality, based on a description of the CONVERGENCE basic cryptographic primitives, as well as advanced cryptographic schemes.

Chapter 7 describes the CONVERGENCE's licensing scheme and its governance, using and extending the MPEG-21 part 5 standard to support specific content protection and rights management requirements of the future Internet. It explains how Rights Expression Language data is embedded into the CONVERGENCE data unit, the Versatile Digital Item (VDI) and how digital certificates can be used to enforce the rights and conditions expressed in CONVERGENCE licenses.

The overall CONVERGENCE functionality is explored and benefits are demonstrated, through the use of applications that were specifically designed and developed to implement a set of four use cases, which are likely to become commonplace in the future Internet. Accordingly, Chap. 8 provides a description of those four realistic use cases and the corresponding applications: *Photos in the Cloud and Analyses on the Earth* (under the responsibility of partner Alinari); *Videos in the Cloud and Analyses on the Earth* (under the responsibility of partner FMSH); *Augmented Lecture Podcast* (under the responsibility of partner LMU); and *Smart Retailing* (under the responsibility of partners WIPRO and UTI). This chapter also discusses the impact of the CONVERGENCE technology to the end-users, according to the user feedback obtained from the conducted field trials, as well as further exploitation scenarios of the developed applications. Such additional scenarios were exemplified through an extra use case that was built by integrating three of the CONVERGENCE applications. The integrated scenario proved that CONVERGENCE is flexible enough to combine different applications in one and to exploit common VDIs.

Finally, Chap. 9 presents the envisaged CONVERGENCE business Models for the commercial and non-commercial exploitation of CONVERGENCE applications and technology. This chapter discusses the feasibility and implications of alternative exploitation strategies, identifying competing products and services, as well as market risks and threats.

CONVERGENCE Researchers

CNIT		ALINARI		CEDEO.net	
Nicola	*Blefari Melazzi*	*Andrea*	*De Polo*	*Leonardo*	*Chiariglione*
Giuseppe	*Bianchi*	*Sam*	*Habibi Minelli*	*Riccardo*	*Chiariglione*
Marco	*Bonola*			*Angelo*	*Difino*
Lorenzo	*Bracciale*				
Stefano	*Brogi*				
Matteo	*Cancellieri*	ICCS		FMSH	
Alberto	*Caponi*	*Angelos*	*Anadiotis*	*Elisabeth*	*De Pablo*
Andrea	*Detti*	*Vasso*	*Giotopoulou*	*Valérie*	*Legrand-Galarza*
Andrea	*Fratini*	*Georgios*	*Lioudakis*	*Francis*	*Lemaitre*
Andrea	*Giancarli*	*Dimitra*	*Kaklamani*	*Manuela*	*Papino*
Sam	*Habibi Minelli*	*Sofia*	*Kapellaki*	*Peter*	*Stockinger*
Leonardo	*Linguaglossa*	*Aziz*	*Mousas*		
Francesco	*Lombardo*	*Charalampos*	*Patrikakis*		
Giacomo	*Morabito*	*Iakovos*	*Venieris*		
Claudio	*Pisa*				
Matteo	*Pomposini*				
Bruno	*Ricci*				
Stefano	*Salsano*	INESC Porto		LMU	
Salvatore	*Signorello*	*Fernando*	*Almeida*	*Alina*	*Hang*
Dimitri	*Tassetto*	*Maria Teresa*	*Andrade*	*Heinrich*	*Hussmann*
Giuseppe	*Tropea*	*Helder*	*Castro*	*Evgeniva*	*Ivanova*
MORPHO CARDS		Singular Logic		XIWRITE	
Silke	*Geisen*	*Panagiotis*	*Gkonis*	*Barbara*	*Benincasa*
Thomas	*Huebner*	*Stelios*	*Pantelopoulos*	*Raffaele*	*Di Fuccio*
Andreas	*Kohlos*			*Andrea*	*Pede*

(continued)

(continued)

Carsten	*Rust*			*Richard*	*Walker*
Amit	*Shrestha*				
UTI		WIPRO			
Bogdan	*Ardeleanu*	*José*	*Horta*		
Ileana-Catinca	*Bobric*	*João*	*Pereira*		
Danut	*Clapa*	*José*	*Ribas*		
Lucian	*Corlan*	*Inês*	*Santos*		
Costin	*Ochescu*	*Daniel*	*Sequeira*		
Mihai	*Tanase*				