Computer Networks 57 (2013) 3116–3127

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

From content delivery today to information centric networking

G. Carofiglio [a], G. Morabito [b,*], L. Muscariello [c], I. Solis [d], M. Varvello [e]

[a] Bell Labs, Alcatel-Lucent, France
[b] CNIT RU at University of Catania, Italy
[c] Orange Labs, France Telecom, Italy
[d] Palo Alto Research Center, Italy
[e] Bell Labs, Alcatel-Lucent, USA

ARTICLE INFO

Article history:
Available online 12 July 2013

Keywords:
ICN
Information centric networks
CDN
Content delivery networks

ABSTRACT

Today, content delivery is a heterogeneous ecosystem composed by various independent infrastructures. The ever increasing growth of Internet traffic has encouraged the proliferation of different architectures to serve content provider needs and user demand. Despite the differences among the technology, their low level implementation can be characterized in a few fundamental building blocks: *network storage*, *request routing*, and *data transfer*. Existing solutions are inefficient because they try to build an information centric service model over a network infrastructure which was designed to support host-to-host communications. The Information-Centric Networking (ICN) paradigm has been proposed as a possible solution to this mismatch. ICN integrates content delivery as a native network feature. The rationale is to architect a network that automatically interprets, processes, and delivers content (information) independently of its location. This paper makes the following contributions: (1) it identifies a set of building blocks for content delivery, (2) it surveys the most popular approaches to realize the above building blocks, (3) it compares content delivery solutions relying on the current Internet infrastructure with novel ICN approaches.

1. Introduction

The Internet was originally conceived to enable communication between remote hosts. Then, broadband access penetration and mobile terminal evolution have transformed the Internet into a medium to connect people with *content*: users search for information over Google, watch videos on YouTube, and share files via BitTorrent. Content dissemination has caused an exponential growth of Internet traffic volume; this trend is expected to continue in the near future and is currently driven by video-centric services.

Today "content delivery", i.e., getting content from a producer to a consumer, involves many technologies (e.g., content delivery networks (CDNs) and peer-to-peer networks), and players (e.g., CDN providers, ISPs and content providers). The evolution of content delivery has mainly been driven by market needs rather than a coherent architectural plan. Consequently, today content delivery is highly heterogeneous. In addition, solutions proposed in these domains are characterized by inefficiencies due to the fact that they try to match the content centric problem with an infrastructure which is built on top of a host-to-host communication model.

To overcome the friction caused by the difference between the content delivery problem and the host-to-host communication model, the research community has recently focused on *Information-Centric Networking* (ICN) [25,23], a novel networking paradigm that integrates content delivery as a native network feature. The rationale is to architect a network that automatically interprets, processes, and delivers content (information) independently of its location. The communication shift is realized by replacing host addresses with content names. Named data

* Corresponding author. Tel.: +39 095 7382355; fax: +39 095 7382397.
E-mail addresses: giacomo.morabito@dieei.unict.it, morabito.giacomo@gmail.com (G. Morabito).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2013.07.002

Table 1
Projects on information-centric networking.

Project	Website
AsiaFI	http://www.asiafi.net/
COMET	http://www.comet-project.org/
ANR Connect	http://www.anr-connect.org/
Convergence	http://www.ict-convergence.eu/
GreenICN	http://www.greenicn.org/
NDN	http://www.named-data.net/
NetInf	http://www.netinf.org/
PSIRP	http://www.psirp.org/
PURSUIT	http://www.fp7-pursuit.eu/
4WARD	http://www.4ward-project.eu/
SAIL	http://www.sail-project.eu/

is exchanged upon user demand, and can be cached by any network component (including routers) equipped with storage.

Since its proposal, the research community has payed an increasing interest in the ICN approach as demonstrated by the large number of research projects [2] (we report a list of relevant projects along with the links to their websites in Table 1) and and by the recent establishment of the *Information-Centric Networking Research Group* (ICNRG, see http://irtf.org/icnrg) within the Internet Research Task Force (IRTF). In this paper, we look at the basic building blocks required for content delivery and how they map to the general ICN concepts. Using this as a base we then finish off by comparing the ICN model with the current content delivery model.

The remainder of the paper is organized as follows. In Section 2 we describe existing solutions for content delivery that rely on the current Internet. More specifically, we characterize such solutions referring to the three major building blocks which can be identified in content delivery:*request routing*, *network storage*, and *data transfer*. In Section 3 we provide an overview of the approaches that ICN solutions propose for the above building blocks. A comparison of today's content delivery model versus the one proposed by ICN is the subject of Section 4. Finally, in Section 5 we will draw our concluding remarks.

2. Content delivery

In this section, we overview the low level mechanism used by state-of-the-art content delivery mechanisms. Despite the differences among the technology, their low level implementation can be characterized in few fundamental building blocks: *network storage*, *request routing* and *data transfer*, which will be addressed in Sections 2.1, 2.2, and 2.3, respectively.

2.1. Network storage

Content-delivery solutions can serve a piece of content from a set of network locations on behalf of the original content provider. We call this function *network storage*. Network storage allows improving end-user response time while reducing network traffic and server load; also, it allows increased service availability. Fig. 1 graphically shows how network storage solutions have been deployed over the years, from simple Web caching to Content Delivery Networks and ISP transparent caching. We will go over each of these solutions while analyzing their strengths and weaknesses.

Fig. 1. Network storage evolution over the years.

Web caching [39] is the first widely adopted network storage solution. Web caching leverages a set of Web proxies augmented with caching capability. Requests arrive from the clients browsers which are configured (statically or dynamically) to use the proxy. Web caches are normally organized in a hierarchical structure so that cache misses can be resolved upstream if possible. The major limitations of Web caching are the following: (i) *scalability*: traffic overload may still arise at the top of the hierarchy in the presence of flash crowds [51]; (ii) *content consistency*: there is no coordination between Web caches and content providers; (iii) *lack of transparency*: users may need to manually configure their browsers with the proxy address.

Content Delivery Networks (CDNs) [33] were introduced to overcome the limitations of Web caching. These limitations have been highlighted by the widespread of Web services like dynamic content, video on demand and live streaming. CDNs are complex networks of servers distributed across the Internet that cooperate for content delivery. CDNs differ from Web caches as they are completely transparent to the users, and they proactively replicate content at network edge. Also, CDNs maintain a direct relationship with content providers. Furthermore, by leveraging on specific markup languages (such as the *Edge Side Include*[1] (ESI)) CDNs support the dynamic assembly and delivery of content, e.g., *dynamic content*. Recent CDNs operate on content *chunks*, i.e., fractions of a file, allowing fine-grained load balancing even for large-file content retrievals [32].

Today's CDN model is based on at least three main players: the content providers, the CDN providers and the network operators. The relations between the three are mainly based on explicit contracts with clear customer/provider roles. There are no standard mechanisms to implement content publication and delivery down to the final users. Most of the time the interconnection is based on manual configuration and based on the very specific content and CDN technology being used. As a result, the three worlds are usually distinct and evolve as silos. CDN providers are IT (Information Technology) companies and focus on optimizing their software, while network operators on managing reliable interconnection and quality of service. In some cases network operators, as Level 3, are able to offer global CDN services over their international backbone network, implementing an IT/Network convergent service. In other cases nation-wide ISPs deploy their own CDN service installing servers in regional and national PoPs to reduce traffic costs in cross exchange points. IT/

[1] http://www.esi.org.