Journal of Network and Computer Applications 47 (2015) 85–98

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

Scalable data replication in content-centric networking based on alias names

Jaime Garcia-Reinoso*, Norberto Fernández, Ivan Vidal, Jesús Arias Fisteus

Universidad Carlos III de Madrid. Avda. de la Universidad 30, 28911 Leganés, Madrid, Spain

ARTICLE INFO

Article history:
Received 11 March 2014
Received in revised form
11 August 2014
Accepted 1 October 2014
Available online 13 October 2014

Keywords:
Content-Centric Networking
Data replication
Alias name

ABSTRACT

Content-Centric Networking (CCN) is a clean-slate proposal to redesign the current Internet by focusing on the content itself, instead of the classical computer-to-computer communication. In this paper we address scalability issues of the Forwarding Information Base (FIB) in CCN. Our solution proposes both the use of hierarchical names assigned by access providers and a novel *alias name* architecture. With the former, we allow the aggregation of entries at the routing tables of CCN content routers, while the latter reduces the processing load at those routers when replicas exist in different parts of the network. With some minor changes to the original proposal, we provide a scalable solution for data replication in CCN, which inherently supports content mobility at the same time. We validate our scheme by (1) comparing the scalability of CCN against our proposal and by (2) implementing and testing a proof-of-concept software based on CCNx, to prove the viability of this approach.

1. Introduction

Nowadays, the Internet users have a plethora of applications to download content, like web browsers, file transport applications, peer-to-peer (P2P), etc. After the users provide the identifier of the content, how the application downloads it is transparent to them: the content may be stored in a single server located "far away" from the user; it may be replicated in several servers if the service provider is using a CDN (Content Delivery Network); or, in the case of P2P, different parts of a file could be downloaded from different peers. Independently of how the application downloads content, there is an end-to-end communication where intermediate routers are just used to forward packets. In other words, the Internet (as opposite to users) cares about end-to-end communications, not content.

This model focused on computers instead of content has several drawbacks. First of all, it is necessary to secure the content exchange to guarantee confidentiality, integrity, availability, authenticity and non-repudiation of the content. Second, as communications are usually point-to-point (multicast is only available in certain networks) in the current Internet, it is not possible to use replicas of the objects in case they exist, and overlay networks (CDNs or P2P for example) have to be built on top of it to use these distributed replicas. Third, the Internet protocols (both IPv4 and IPv6) use one single locator/identifier value to route and identify computers. This is the main problem we have to face for mobility communications, as when a device changes its point of attachment to the network it must change its locator address. In IP, when a mobile node changes its IP address it is changing its locator as well as its identifier, which is not the desired behavior.

In recent years, several projects and organizations have proposed minor and major changes to Internet protocols to minimize or eliminate the aforementioned problems. IPsec (IP security) (Kent and Seo, 2005), TLS (Transport Layer Security) (Dierks and Rescorla, 2008) and MIP (Mobile IP) (Perkins, 2010) are examples of standards proposed to overcome some of those problems. Other initiatives are still under discussion or with little penetration like LISP (Locator/ID Separation Protocol) (Farinacci et al., 2013), RELOAD (REsource LOcation And Discovery) (Jennings et al., 2014) and PPSP (P2P Streaming Protocol) (Bakker, 2011) for example.

All the previous initiatives are focused on modifying or enhancing existing protocols, but other works try to go beyond that by proposing clean-slate approaches. Among them, we will focus on Information-Centric Networks (ICN), where everything is built around the content itself, independently of where it is stored, and based on *Publish/Subscribe* messages. Although there are different proposals around the ICN concept (Lagutin et al., 2010; Koponen et al., 2007; Gritter and Cheriton, 2001; Jacobson et al., 2012), almost all of them have to solve four main problems (Choi et al., 2011): (1) the naming structure of the content, (2) the mechanism to find the content, (3) how to deliver the content to the requester and (4) caching the content inside the network. The authors in Xylomenos et al. (2013) present a survey describing the

* Corresponding author.
E-mail addresses: jgr@it.uc3m.es (J. Garcia-Reinoso), berto@it.uc3m.es (N. Fernández), ividal@it.uc3m.es (I. Vidal), jaf@it.uc3m.es (J. Arias Fisteus).

http://dx.doi.org/10.1016/j.jnca.2014.10.003
1084-8045/© 2014 Elsevier Ltd. All rights reserved.

most important ICN proposals, including a comparison of the key functionalities described before. For example, and very related with our paper, the survey presents a comparison between hierarchical and flat naming, where the authors conclude that the former allows scalability when aggregation is possible while the latter avoids the location-identity binding. In other words, both approaches have their advantages and disadvantages, so other alternatives are necessary.

One significant initiative in the ICN paradigm is Content-Centric Networking (CCN) (Jacobson et al., 2012), which uses a hierarchical name scheme similar to Uniform Resource Identifiers (URIs), such as /es/uc3m/it/joe/documents/paper.tex. In addition, every individual Content Router (CR), which is a router with caching capabilities, has to know how to forward Subscribe messages (or Interest in CCN terminology) using its own Forwarding Information Base (FIB) table. With potentially billions of objects, the scalability of the FIB is a clear issue, and further study is necessary in this particular point. This can be even more problematic as content can be replicated, and replicas with the same content name can be distributed among different parts of the network domains (Xylomenos et al., 2013). Replication has also impact in CCN routing, which has to use a *Strategy Layer* in order to retrieve the content from the best source (with the lowest delay or the highest bandwidth, for example). In the case of core CRs receiving millions of packets, it would be advisable to do the source selection at the network end points, reducing the complexity of those intermediate nodes.

To solve the issues related to the scalability of the FIB and the processing overhead at the CRs, in this paper we propose a scalable data replication scheme for CCN. Firstly, our solution uses hierarchical provider-assigned names to facilitate aggregation, as it has been suggested in Zhang et al. (2010). This aggregation comes at the cost of requiring different CCN names for the same content replicated in different provider networks. Secondly, a novel *alias name* architecture is introduced, so that replicas in different parts of the network with different names can be identified as objects with the same content. We extend the Interest packet format, including a new field to transport an alias name as well as the content name. The advantage of our proposal is that CCN routers can check if the requested content is stored in its cache or not, as the content name is carried in the Interest. This way, consumers can select the alias name they want to use to retrieve each individual piece of data, or even try different alias in parallel, maintaining the benefits of using caching in the routers. Altogether, the proposed mechanisms achieve a scalable data replication, as FIB tables do not have to include entries for replicas.

To accomplish these goals, we introduce a new functional entity in the network called the *Alias Name Manager* (ANM), which could be placed by the access service provider inside its own domain network (but it can be anywhere in the network). Apart from the ANM, we extend the basic CCN proposal by introducing the notion of *Alias Routing Name* (ARN), which can be included as an optional field in CCN Interest packets, as stated before.

Besides the scalability advantages of our data replication scheme, it inherently supports mobility of content. When a content is moved to a different network it is assigned an alias name, which in turn has to be registered in the ANM. This entity allows accessing the content by using its original name.

The rest of this paper is structured as follows. In Section 2 we present a survey of CCN to provide some background for the rest of the paper. In Section 3 we describe our proposal explaining in detail all the modifications introduced to CCN. Section 4 compares regular CCN against our scheme in terms of scalability, by means of simulations. Section 5 presents a proof-of-concept software, implemented to show the feasibility of our proposal, and to evaluate its performance in terms of delay, when multiple copies of an object are replicated in several domains. Finally, Section 6 closes the paper with the main conclusions and the future work.

2. Background on Content-Centric Networking

Content-Centric Networking (CCN) (Jacobson et al., 2012) is a novel clean-slate design of the Internet, based on the concept of named content. Like other ICN initiatives, it focuses on retrieving the content by its name, instead of locating and establishing a communication with the end host that holds the content (as in the current Internet). In CCN, names are hierarchically structured into a set of components, and applications can choose any naming convention for an appropriate operation. As an example, the CCN name /es/uc3m/it/research/papers/paper.pdf/_v2/_s1, could be used by an application to retrieve the first segment of version 2 of this paper. In this example, the convention followed by the applications dictates to use the marker _*v* to indicate the version number and the marker _*s* to identify the file segment. CCN names with subsequent segment numbers would allow the application to retrieve the whole paper for display. On the other hand, from the perspective of the CCN transport, names are opaque (i.e. the transport does not need to understand name semantics) and are composed by a set of binary encoded components.

CCN defines two types of packets, *Interest* and *Data*. When a receiver decides to retrieve a given content, it generates and sends an Interest packet that includes the CCN name of the desired content in a field that, for convenience, we will name Content Name Identifier (CNI) from now on. The Interest packet is routed by CCN routers towards a source of the specified content. If this packet reaches a node (i.e. a source or an intermediate CCN router) that holds some content that matches the Interest, this node can directly answer back with a Data packet including the desired content. A content matches an Interest packet if the content name includes the CCN name indicated in the Interest.

Figure 1 illustrates the forwarding model of a CCN node. Whenever an Interest packet is received on a face (network or logical interface), if the CCN node cannot satisfy the Interest, it stores the CCN name included in the Interest and its incoming face in a Pending Interest Table (*PIT*). Then, the Interest is forwarded to the next hop towards a source of the content, according to the information stored in a Forwarding Information Base (*FIB*). In case several faces to the content exist, a router has to select the proper one by running the algorithm implemented by its *Strategy Layer*, which selects the optimal next hop to use. This face selection implies extra processing at the content router.

The FIB can be built at each node by the execution of a routing protocol that, similar to current IP networks, would be used to propagate content name prefixes between CCN routers (routing protocols such as OSPF or BGP could be adapted to this end Jacobson et al., 2012). Apart from content name prefixes, it is also possible to register content in a domain different from the original provider (i.e., when the prefix of the home provider does not match the names published by the visiting entity). In those cases, when the routing protocol is executed by the content routers, some of those CCN routers will not be able to aggregate names at their FIB. This occurs when the home domain of the content, and the visited domain where the content is located, are reachable from different faces of a given CCN router. This may lead to scalability issues in the FIB of the CCN routers as will be shown in Section 4.

When an Interest reaches a node that maintains a matching content, that Data packet is transmitted in response. This Data packet is routed back to the receiver via the reverse path followed by the Interest, based on the information stored in the PITs of the