

NetInf: An Information-Centric Design for the Future Internet

Christian Dannewitz, *University of Paderborn, Germany*, christian.dannewitz@upb.de

Abstract—The current Internet architecture has been very successful and has scaled well beyond the original aspirations. Nevertheless, the architecture suffers from multiple well documented problems, many of which result from the fact that the usage of the Internet has changed drastically. For example, the original Internet architecture is designed to provide access to specific nodes in the network whereas today’s usage is dominated by information dissemination where the information requester does not care about the source location. To adjust to those changes, we develop a new, information-centric network architecture called *Network of Information (NetInf)* in the context of the FP7 EU-funded 4WARD project. This architecture can significantly improve large scale information distribution. Furthermore, it supports future mobile networks in situations with intermittent and heterogeneous connectivity and connects the digital with the physical world to enable better user experience.

NetInf is built upon an extended identifier/locator split which is based on a simple but powerful information model. Main aspects of NetInf include the persistent naming of information, a world-wide scalable name resolution mechanism for flat identifiers, and improved information availability and dissemination via caching and storage integrated into the network architecture.

I. INTRODUCTION

An ongoing discussion in the scientific community (e.g. Jacobson *et al.* [1]) argues for a transition to the third generation of networking. The first generation dealt with connecting wires and laying down infrastructure. The second one placed end nodes, instead of the interconnecting points, at the forefront, leading to the emergence of the WWW and widespread Internet adoption. The third generation, which we call *Network of Information (NetInf)* in the context of the EU-funded FP7 project 4WARD, will refocus the point of attention to what humans care about the most: information. By taking information *per se* as the starting point of an *information-centric* network architecture, it will be possible to design a communication infrastructure which is much better adapted to the task of distributing and exchanging information compared to today’s host-centric approach. This paper gives an overview of an information-centric architecture. It is a condensed and extended version of previous work [2]–[4].

An information-centric architecture has major advantages in large-scale information dissemination scenarios like the distribution of Web pages and multimedia content. Furthermore, it offers support for intermittent connectivity, information access in (partially) disconnected scenarios, and mobility scenarios. The NetInf architecture also accommodates non-dissemination

applications, including streaming and interpersonal communications. Beyond those use cases, NetInf also supports the integration of the real world and related information and services in the Internet, which we call the *Augmented Internet* scenario. Details about those scenarios can be found in our recent publication by Ahlgren *et al.* [2].

NetInf extends the concept of the identifier/locator split with another level of indirection in order to decouple objects from their storage location(s). As hosts take a secondary role and information ascends into center stage, the information has to become self-certifiable, unlike today’s Internet where information is assumed to be valid because the sender appears legitimate. This information-centric security approach enables a much better level of security and trust than currently present in the Internet.

NetInf can be implemented as an overlay on top of the current Internet infrastructure. However, it is important to point out that NetInf is independent from the underlying transport protocol and is, moreover, in a unique position to benefit from other technologies ranging from new transport technologies to virtualization to in-network management, also developed within 4WARD.

II. NETINF COMPONENTS

This section gives a brief overview of the main components that are required to build an information-centric network architecture and related challenges.

A. Overview

Information plays the key role in an information-centric network. Hence, an information model is required that represents information appropriately and supports efficient information dissemination. To make information access independent from storage locations and to benefit from copies that are available in the network, the Network of Information builds upon an identifier/locator split. Therefore, a namespace is required to name information independent of the storage location. Furthermore, a name resolution mechanism is required that maintains and resolves the binding between locators and identifiers. The *information model*, the *naming* aspects, and the *name resolution* will be discussed in more detail in the next subsections.

In addition to those components, the NetInf concept considers several other aspects that will only be mentioned briefly here. First, we evaluate several different *routing* mechanisms to route requests to the naming system and to route the payload data packets. In conjunction with the routing of the

The research leading to these results has received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 216041.




Fig. 1. NetInf Information Model

payload, the NetInf concept also integrates *storage* into the network and makes local copies accessible in the network. This improves the availability of information and supports efficient information dissemination.

Preceding the overall information retrieval is a *search*. Advanced search mechanisms that complement today's dominating full text search can be enabled by the information-centric network approach based on metadata stored in the information model.

B. Information Model

NetInf elevates information to first-class network citizens, in the form of so-called *Information Objects* (IOs). IOs represent information such as audio and video content, Web pages, and email in our information model. Beyond those rather obvious examples, IOs can also represent streams, real-time services, (video-)telephony data, and physical objects thanks to the flexible and generic nature of the information model.

A special kind of IOs are *Data Objects* (DOs). A DO represents a specific *Bit-Level Object* such as a certain MP3 file with a certain encoding, including copies of this specific file. A DO aggregates (some or) all copies of a certain file by storing locators for those copies.

In many cases, however, a user is not really interested in a specific Data Object but in the information a Data Object represents. For example, a user might be interested in a certain song (e.g., Beethoven's 9th symphony), but might not care about the encoding (MP3 with 128 kbps or WMA with 196 kbps). These higher level semantics are expressed through the actual *Information Objects* (Fig. 1). An IO may, e.g., refer to a certain song without specifying the concrete encoding or the performing orchestra. IOs enable users to find content independently of its specific representation and independently of certain characteristics that *might* not be relevant to the user. IOs can be composed of other IOs or can directly point to one or multiple Data Objects that contain the content itself.

Metadata enables us to further express the semantic meaning of Information Objects, e.g., describing its content or its relation to other objects. Existing research in this field provides an excellent starting point for integrating these features into the network layer, particularly with regard to description

languages such as the Resource Description Framework or the ability to establish relations between IOs.¹

Thanks to the very general Information Object model, NetInf is also able to better integrate information access into the user's physical world by representing real-world objects as Information Objects. Those IOs aggregate Information Objects and Data Objects related to the represented physical entity. For example, an IO could represent the Eiffel tower and could point to related Data Objects like pictures, a wiki page, and a service to buy tickets.

C. Naming and Name Resolution

Name resolution (NR) mechanisms resolve IDs into one or more locations. NR should work on a global scale, guaranteeing correct resolution for any globally-available resource, just as the Internet works today. We call this the *NR Global Resolution Property*. NR should also work in an intermittently connected network if a Data Object is locally available. We call this the *Local Resolution Property*. We implement the Local Resolution Property by supporting multiple coexisting NR systems, some of which have global scope and some have local scope. In other words, NR systems that can resolve any ID worldwide can naturally coexist with NR systems that deal with a local ID space (e.g. company-internal). This important feature eliminates the need for permanent global connectivity and facilitates efficient implementations using anycast and locality-aware content distribution.

Besides the identifier/locator binding (DO-locator), the information model described in Section II-B calls for at least two additional bindings, IO-IO bindings and IO-DO bindings, that need resolution. Those bindings also have to be handled by our Name Resolution service.

The choice of an appropriate NR mechanism will be heavily influenced by the characteristics of the NetInf namespaces. Important attributes of the NetInf namespace are *persistence* of names and *contention freeness* [5]. Those attributes can be met by using flat namespaces. But a flat namespace prevents the use of concepts similar to today's DNS, which is based on a hierarchical architecture and requires, accordingly, a hierarchical namespace. Therefore, a new NR approach is required.

For flat names, Distributed Hash Table (DHT) based systems are a promising approach. DHTs are decentralized, highly scalable, and mostly self-organized, limiting the need for administrative entities. There are several compact routing protocols (e.g., Chord, Pastry, Tapestry, CAN, Kademia [6]), typically used in P2P overlay networks, which can route messages in $O(\log N)$ routing steps, with compact routing tables of $O(\log N)$ states, where N is the number of nodes. $O(\log N)$ resolution steps may, however, result in unacceptably large latencies. Recently proposed, promising approaches can guarantee a constant number of hops per lookup [7], [8]. Although convergence time may still be an issue, these approaches can reduce the number of required hops in exchange

¹<http://www.w3.org/RDF/>, <http://www.w3.org/2001/sw/>

for larger routing tables and increased overhead in case of churn. Therefore, they are well suited for NR systems with local scope, e.g., within an ISP network. Here, a limited number of “carrier-grade” NR nodes is deployed, i.e., hundreds or thousands of nodes, that are expected to be highly stable and reliable, with almost no churn.

On a global scope, however, a DHT-based NR system becomes more problematic. Due to the flat namespace and the intrinsically non-cooperative nature of Autonomous Systems (ASs) and other administrative domains, there is an issue with binding placement and control. Scalability and increased churn also need to be taken into account. As we generally believe in the benefits of a flat namespace over a hierarchical one, a possible solution to the dilemma lies in integrating aspects of promising initiatives from the area of routing on “flat” identifiers into the NR system, such as the Late Locator Construction (LLC). It is one of the key challenges of the NR design to find a satisfactory trade-off between scalability (aggregation based on hierarchical names) and the name persistence offered by the flat namespace.

D. Routing

A first approach to routing requests as well as payload could be to use a traditional topology-based routing scheme, based on shortest path algorithms and hierarchical routing, like the ones used in the current Internet (OSPF, IS-IS, BGP), or a topological-based compact routing scheme. But recent results in routing research [9] are not encouraging, since logarithmic scaling cannot be reached over real networks whose topologies are not static; in fact, network dynamics involve communication costs which cannot grow slower than linearly with the number of nodes and often increase at a very high rate [10].

A second alternative to investigate is to use name-based routing which integrates both the resolution path and the retrieval path; this might result in better performance. We currently evaluate name-based routing schemes and related mechanisms like NodeID [11] and LLC [12] for providing an integrated name resolution and routing architecture.

III. CONCLUSION

The Network of Information can significantly improve the availability and dissemination of information. We are currently at the advanced stage of (re-)defining the architecture and testing the architecture via prototyping.

Building a scalable information-centric architecture involves several challenges. This includes the development of an information model and a naming framework which support efficient information dissemination with improved security properties. It also includes the development of a world-wide scalable name resolution mechanism for the new namespace. This new name resolution service should have the same desirable attributes concerning generality, flexibility, and ease of use that made DNS so successful in today’s Internet.

Although NetInf can be implemented on top of today’s TCP/IP stack to allow for easy migration, it can significantly

benefit from being integrated into a clean slate approach such as developed, e.g., in the 4WARD project.

IV. ACKNOWLEDGMENTS

This work has been carried out in the IST 7th Framework Program Integrated Project 4WARD, partly funded by the Commission of the European Union. I thank my colleagues in Work Package 6 for many fruitful discussions. I especially thank my coauthors of the publications that this paper is based upon, namely, Bengt Ahlgren, Matteo D’Ambrosio, Marco Marchisio, Ian Marsh, Börje Ohlman, Kostas Pentikousis, René Rembarz, Eric Renault, Ove Strandberg, Javier Ubillos, and Vinicio Vercellone.

REFERENCES

- [1] V. Jacobson, M. Mosko, D. Smetters, and J. Garcia-Luna-Aceves, “Content-centric networking,” Palo Alto Research Center, Whitepaper, Jan. 2007.
- [2] C. Dannewitz, K. Pentikousis, R. Rembarz, E. Renault, O. Strandberg, and J. Ubillos, “Scenarios and research issues for a Network of Information,” in *Proc. 4th International Mobile Multimedia Communications Conference*, Oulu, Finland, July 2008.
- [3] C. Dannewitz, “Augmented Internet: An information-centric approach for real-world / Internet integration,” in *Proc. International Workshop on the Network of the Future (Future-Net)*, colocated with IEEE ICC, June 2009, to be published.
- [4] B. Ahlgren, M. D’Ambrosio, C. Dannewitz, M. Marchisio, I. Marsh, B. Ohlman, K. Pentikousis, R. Rembarz, O. Strandberg, and V. Vercellone, “Design considerations for a network of information,” in *Proceedings of the First International Workshop on Re-Architecting the Internet (ReArch2008)*, Madrid, Spain, December 2008.
- [5] M. Walfish, H. Balakrishnan, and S. Shenker, “Untangling the Web from DNS,” in *NSDI’04: Proc. 1st Symp. on Networked Systems Design and Implementation*, San Francisco, CA, USA, 2004.
- [6] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, “A survey and comparison of peer-to-peer overlay network schemes,” *IEEE Communications Survey and Tutorial*, vol. 7, 2005.
- [7] V. Ramasubramanian and E. G. Sirer, “Beehive: O(1) lookup performance for power-law query distributions in peer-to-peer overlays,” in *Proc. Networked System Design and Implementation (NSDI)*, San Francisco, CA, Mar. 2004.
- [8] A. T. Mýzrak, Y. Cheng, V. Kumar, and S. Savage, “Structured superpeers: Leveraging heterogeneity to provide constant-time lookup,” in *Proc. 3rd IEEE Workshop on Internet Applications*, Washington, DC, USA, 2003.
- [9] D. Krioukov, K. Claffy, K. Fall, and A. Brady, “On compact routing for the Internet,” *SIGCOMM Computer Communication Review*, vol. 37, no. 3, 2007.
- [10] Y. Afek, E. Gafni, and M. Ricklin, “Upper and lower bounds for routing schemes in dynamic networks,” in *Proc. 30th Symp. on Foundations of Computer Science*, 1989.
- [11] “Connectivity and dynamic internetworking prototype and evaluation,” Deliverable 23, Ambient Networks Phase 2 project, Dec. 2007, fP6-CALL4-027662-AN P2/ D23-E.2.
- [12] A. Eriksson and B. Ohlman, “Dynamic internetworking based on late locator construction,” in *10th IEEE Global Internet Symposium*, May 2007.