Towards an ICN Cloud Architecture for the IoT

Martine Lenders, adviced by Matthias Wählisch
Freie Universität Berlin
{m.lenders, m.waehlisch}@fu-berlin.de

***Abstract*—Cloud services are an essential part of the current Internet of Things (IoT). They provide end-users with an easy to use abstraction to scale, not only in terms of processing the data but also in terms of distributing the content. However, from a content delivery perspective, clouds are not a natural fit to the current Internet as they badly intermingle application (e.g., DNS) and networking layer to enable efficient access to de-localized data. More recently, Information-Centric Networking (ICN) has been introduced, which provides two key services. First, a name-based interface to content on the network layer, and second, an inherent replicative content distribution scheme. In particular in the IoT, these two features help to improve reliability and to reduce resource requirements. The objective of this thesis is to explore potentials of an ICN cloud architecture for the Internet of Things.**

Introduction

The Internet of Things

The Internet of Things (IoT) is an Internet of *very* heterogeneous devices. Most challenging are constrained devices [1], which exhibit very limited processing, memory, and energy resources. Moreover, these devices are typically connected to the Internet via lossy wireless domains. On the other hand, such constraint devices will make a significant portion of the IoT as they allow for very low-cost deployment [2].

The vast amount of data generated by these devices need to be aggregated and processed. The state-of-the-art is to offload these operations to a cloud service.

To natively connect IoT devices to the IP-based Internet and its services, the IETF provides multiple standards. This includes adaption layers for IPv6 to support the heterogeneous landscape of link-layer technologies in the IoT [3, 4, 5], as well as tailored application layer protocols such as a constrained RESTful alternative to HTTP, CoAP.

Information-Centric Networking (ICN) alternatively offers a new perspective on the IoT with its name-based routing: Data is no longer associated to a specific devices but addressed directly. Likewise, provides the caching mechanism of ICN low traffic-rates by design.

Cloud Computing

"The Cloud" includes a range of different technologies. To avoid confusion, we narrow down the term Cloud Computing to the specific use-case of the IoT: This reduces the set of delivery models to Software-as-a-Service, Storage-as-a-Service, and solutions based around Content-Delivery Networks. These solutions have in common that they aim to hide most of the underlying architecture from the user [6]. By these means the user is able to easily set up a data aggregation and distribution point for the great amount of data the IoT produces in form of a service: All they need to know is the API to the cloud, while all the layers below are handled by the cloud provider.

Information-Centric Networking

The Internet was build upon an end-to-end paradigm. Contrary to that, the modern Internet is primarily used to disseminate unlocated information. Content-Delivery Networks (CDNs) currently provide a way to make this possible. ICN tries to solve this issue on the network layer directly.

Named Data Networking (NDN) is an implementation of ICN [7] we use to exemplify how an ICN-based Internet might work. Any of the following three roles can be assumed by a node – if need be even at once – from the perspective of a specific *named data object* (NDO):

- **Publisher:** The node that originally provided the NDO.
- **Subscriber:** A node that is interested in the NDO.
- **Content Router (CR):** A node involved in forwarding the NDO. A CR holds 3 data structures:
  - The *Forwarding Information Base (FIB)* translates a name to a next hop and is populated by a routing protocol currently out of scope.
  - A *Pending Interest Table (PIT)* that stores the name and the requester of a requested NDO.
  - A *Content Store (CS)* that caches already requested and returned NDOs indexed by their names

In NDN every NDO gets a *hierarchical name* under which the Publisher publishes it. A Subscriber then can send out an *Interest* to solicit a *Data message* with containing the NDO from the network. Based on the name the CRs are able to forward the *Interest* to the Publisher using the *FIB*. On the way to the Publisher the *Interest* is noted down in the *PIT*. *Data messages* are cached in the *CS* while being forwarded back to the Subscriber. For this forwarding step the *PIT* created on Interest traversal entry is used and consumed.

Note, that NDN is only *one* approach to ICN. Other approaches e.g. don't come with hierarchical naming.

Research Challenges

My thesis will primarily focus on pushing data from an ICN Publisher in form of a e.g. sensor to the cloud (which joins the ICN as a Subscriber, see Fig. 1). Ignoring the low-level approach of how to bring ICN-traffic to an IP-based cloud for now, a few key questions are open in this scenario:

- How would a unified API, similar to TCP/IP's sockets, (see Fig. 2) look like that grants access to both the cloud and ICN in a transparent way to the user?
- How to find a good naming scheme that includes both the ICN as well as the cloud?

Especially the second question bares some sub-challenges, which not only include solving Internetworking for ICN, but also need to show if prefix-less naming is a possibility. If this is the case we need to determine if it scales and, considering the developer, if it is still handleable by a human (i.e. intuitively human-readable).

Connecting the cloud and the ICN is a secondary challenge: Most cloud services provide communication interfaces based on TCP/IP-based protocols like MQTT [8], so some kind of gateway solution (at CR1 in Fig. 1) is needed.

Further challenges in this scenario include the problem, that the data needs to be *pushed* to the cloud, for which there is no mechanism in ICN at the moment [9].

Research Agenda

Though secondary, the translation between the communication protocols of the cloud and the ICN needs to be addressed for experimentation. Multiple solutions are thinkable:

- **Tunnel ICN via TCP/IP:** While we might be able to provide a portable and flexible solution with this, multiple questions arise with it:
  - Does the cloud need to know about the ICN?
  - On what TCP/IP layer should the ICN packet be transported? The main candidates are the network layer (making ICN itself a transport layer protocol) and the transport layer (making it an application layer protocol).
  - If transported on the transport layer: which transport layer protocol is best suited?
- **Translate ICN to the given communication protocol:** Here we can't find a generalized solution since we always need to adapt for the communication protocol used and/or the API provided by the cloud provider.

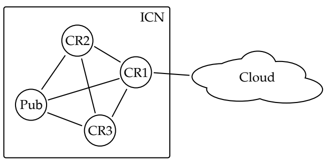

Fig. 1: A cloud as a Subscriber in an ICN ("Pub" denotes a publisher)

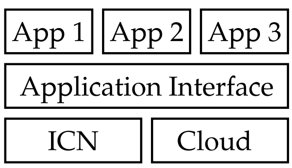

Fig. 2: Sketch of a cloud-to-ICN architecture

Next, a solution for the naming challenge must be found, which will most likely take the most significant portion of my thesis. The name must both allow for scalable routing, be human-readable and ideally must not put any constraints for either cloud-, gateway- and ICN-provider on the choice of their name. Ideally the name may be free to choose with a minimal set of rules to be applied or provided by e.g. a name dissemination mechanism.

Lastly, the transparent API to the ICN and cloud needs to be drafted, implemented and put to test in the real world.

Future Steps

The small fragment of the scenario I want to look in already provided many challenges that need significant research put into them. They might also open doors to new research challenges that could be tackled by my thesis or following ones.

Besides solving the problems of application interfaces, gateway solutions and the naming problem my contribution may also provide first steps into both an ICN-based Internet (as the connection from an ICN to the cloud must be solved), as well as providing a solution for a generic gateway between ICN and other protocol suites beyond the scope of cloud integration.

References

[1] C. Bormann, M. Ersue, and A. Keränen, *Terminology for constrained-node networks*, RFC 7228, https://rfc-editor.org/rfc/rfc7228.txt, May 2014.

[2] Ericsson, "More than 50 billion connected devices," Ericsson, Tech. Rep., 2011.

[3] G. Montenegro, C. Schumacher, and N. Kushalnagar, *Ipv6 over low-power wireless personal area networks (6lowpans): overview, assumptions, problem statement, and goals*, RFC 4919, https://rfc-editor.org/rfc/rfc4919.txt, Aug. 2007.

[4] A. Brandt and J. Buron, *Transmission of ipv6 packets over itu-t g.9959 networks*, RFC 7428, https://rfc-editor.org/rfc/rfc7428.txt, Feb. 2015.

[5] Z. Shelby, J. Nieminen, T. Savolainen, M. Isomaki, B. Patil, and C. Gomez, *Ipv6 over bluetooth(r) low energy*, RFC 7668, https://rfc-editor.org/rfc/rfc7668.txt, Oct. 2015.

[6] T. Erl, R. Puttini, and Z. Mahmood, *Cloud computing: Concepts, technology, & architecture.* Pearson Education, 2013.

[7] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.

[8] Amazon Web Services, Inc., *How the aws iot platform works*, https://aws.amazon.com/iot/how-it-works/, [Accessed: 2016-11-25], 2016.

[9] D. Kutscher, S. Eum, K. Pentikousis, I. Psaras, D. Corujo, D. Saucez, T. C. Schmidt, and M. Wählisch, *Information-centric networking (icn) research challenges*, RFC 7927, https://rfc-editor.org/rfc/rfc7927.txt, Jul. 2016.