

ICNRG
Internet-Draft
Intended status: Informational
Expires: August 22, 2013

D. Corujo
Instituto de Telecomunicacoes
K. Pentikousis
Huawei Technologies
I. Vidal
UC3M
February 18, 2013

ICN Management Considerations
draft-corujo-icn-mgmt-00

Abstract

This document aims to draw the attention of the ICNRG community to network management, an important but hitherto underdeveloped area of research in information-centric networking. We consider that the availability of modern management mechanisms for information-centric networks will foster their deployment in real-world environments. For example, we argue that there is a need for creating basic network management tools early on while ICN is still in the design and experimentation phases that can evolve over time. Perhaps ICN can borrow successful mechanisms from the host-centric paradigm and adapt them to the new network primitives. Alternatively, novel network management schemes can be designed based on ICN primitives. As a discussion starter, this document summarizes recently published approaches for ICN network management. In particular, this first version presents a management framework for named data networking and reviews previous work on NetInf management.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 22, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. NDN Management Considerations	4
2.1. Towards a Management Framework for NDN	5
2.2. NDN Management Operations	7
2.2.1. Discovery Procedure	7
2.2.2. Management Data Exchange	9
2.3. Implementation Experience	10
3. NetInf Management Considerations	11
4. Acknowledgements	11
5. IANA Considerations	12
6. Security Considerations	12
7. Informative References	12
Authors' Addresses	13

1. Introduction

Information-centric networking (ICN) enables new ideas for naming and addressing, privacy, security, and trust, and should also lead us to think new ways for deploying, operating and managing networks in the future. By default, users, programs, information objects and hosts are in general untrustworthy and mobile in an information-centric network. This means that many of the assumptions in traditional network management, including all aspects of FCAPS (Fault, Configuration, Accounting, Performance, and Security) need to be rethought. However, despite the different instantiations of ICN architectures, and the plethora of novel research work built on top of them, little attention has been paid to management aspects so far. This includes both enabling "traditional" network management operations (which work well from small networks to large infrastructure networks), and supporting and optimizing intrinsic procedures of the ICN fabric.

This document aims to draw the attention of ICNMG to the importance of network management for real-world deployments. Today, network management is practically an add-on to host-centric deployments. We can do better as we move forward in ICN research considering the full range of deployments from home-office environments to challenged networks to tier-1 networks. To this end, we draft some first management considerations that, on the one hand, capitalize on ICN concepts for defining management procedures and, on the other, explore the possibilities for defining a common management framework irrespective of the ICN approach taken. We reckon that the later is a much more formidable task and we are looking forward to tackling it together with other members of ICNMG. We start this exercise in this first version based on published literature and in particular with a NDN approach.

We argue that addressing management at an early stage is not only important for real-world adoption and the successful future deployment of ICN, but also to deal with scenarios where management can simplify, enhance or optimize ICN network utilization and performance. The subject becomes particularly challenging, as disparate characteristics from different ICN approaches (e.g., in terms of namespace, granularity, routing, and so on) impact the definition and design of these management mechanisms. Section 2 below provides an initial assessment, proposal and evaluation of management mechanisms leveraging NDN intrinsic capabilities based on [NDN-MGMT], while Section 3 briefly summarizes earlier work on self-management for NetInf.

We plan to incrementally develop the draft and incorporate other ICN approaches (e.g., [PURSUIT] and [NetInf]) as well as address other

pertinent aspects as we receive feedback from the research group members.

2. NDN Management Considerations

The Named Data Networking [NDN] ICN architecture provides a new communication framework built on named data. Like other ICN counterparts, such as [NetInf], [PURSUIT] and [DONA], NDN intrinsically supports security, routing/forwarding, reliability, caching and even mobility, aiming at scalable and more efficient content-distribution than today's IP-based approaches. Fostered by an open-source implementation [CCNx], NDN has been at the heart of an active topic with several research contributions evaluating its deployment feasibility and performance in a number of scenarios [ICN-Scenarios].

NDN relies on a hierarchical, human-readable namespace to address named data objects, where the naming scheme is simultaneously used to both name information and to route it. It relies on content requesters sending an Interest packet with a Content Name, where the prefix can provide information for global and organizational routing, while the suffix indicates versioning and segmentation details. When a node receives an Interest packet asking for content which matches what is already available at the node, it responds with a matching Data packet carrying back the content.

Each NDN node comes with a set of supporting data structures which enable the coordination between the transmission of Interest packets with the reception of the corresponding Data packets. These structures include:

1. Content Store: maintains an indication of locally available content, according to name, and is used for Interest packet matching. If the content is available at the node, the Interest packet is consumed, and a Data packet with the respective content is sent towards the request origin.
2. Pending Interest Table (PIT): keeps track of Interest packets seen previously by the node, on their way to locate matching content. Interest packets in the PIT were not matched to content available in the node. Basically, PIT maintains a degree of state regarding Interest packets, mapping them to a corresponding egress network interface.
3. Forward Information Base (FIB): associates named data to potential holders of the content. A routing protocol can populate the FIB (although this is outside the scope of NDN) or

it can be populated through registration in a local NDN store.

NDN introduces the concept of a Strategy Layer, which can control Interest packet forwarding behavior. It basically determines which is the best interface (or set of interfaces) to send an Interest packet. The "strategy" component establishes a pre-configured algorithm for tackling Interest packet decisions, ranging from sending it sequentially on each interface until a Data packet is received, to evaluating which interfaces provide better performance (i.e., lower average RTT) in retrieving certain content (as discussed in [NDN]).

It is important to keep in mind that NDN replaces the commonly used term "interface" with the term "face", since packets can be forwarded over hardware network interfaces as well as between application interfaces, further acknowledging the information dissemination capabilities of ICN. This aspect is considered in [NDN] and [NDN-R], where programs can be associated to the NDN governing structures (like the FIB), defining configurations such as "sendToAll" and "sendToBest" with respect to managing the content reaching process. Corujo et al. [NDN-MGMT] exploit these concepts enabling management mechanisms to be deployed, and steer network operations and NDN operation, as described in the following section.

2.1. Towards a Management Framework for NDN

An important aspect supporting network management procedures is the interaction of network information residing at the network side with information about the network from the perspective of clients connected to it. The former includes, for instance, information stored in the network operator core about user profiles, associated policies, or data collected by the access network equipment, such as current and past traffic load levels, active flows, and maintenance information. Today, such information can be retrieved for management and operation support through dedicated signaling protocols (e.g., [RFC1157], [RFC6733]), or Operation Support Services (OSS) web services. The client point of view of the network includes information that, for example, a wireless terminal can provide, indicating wireless link quality, average return-trip times (RTT) or perceived Quality of Experience (QoE).

Both types of information can be capitalized upon allowing, for example, the network to coordinate network management procedures, considering as input information obtained from other network elements as well as from user nodes. One way to generate management information in network entities and at client nodes, as well as to consume and act upon it (i.e., using the management information exchange as a control channel) is to couple NDN nodes with Management

Agent (MA) entities.

Fig. 1 (redrawn here from [NDN-MGMT] for convenience) illustrates how a MA can be deployed in both network and client entities, interfacing with different operational aspects and protocol layers of an NDN node. By using NDN content reaching and disseminating mechanisms, management information can be consumed by the MA to steer not only the behavior of application processes and network interfaces, but also to interface with NDN supporting structures (i.e. Content Store, FIB, PIT). Effectively, different kinds of information can be conveyed to a network node responsible for managing the network (under different perspectives and processes), and resubmitted back towards client nodes, affecting the way applications interface with network interfaces and the NDN fabric.

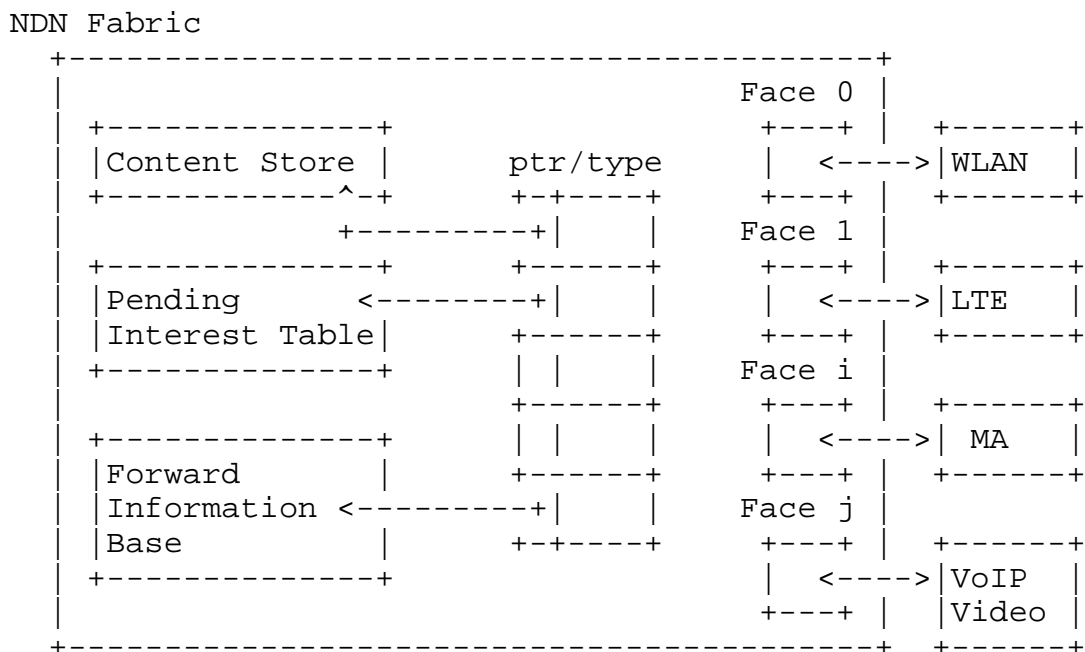


Figure 1. NDN Management Framework

MA can interface with the PIT and FIB structures, acting as a dynamic, application- and/or network-controlled interface to the strategy layer. This could also be used to direct how to forward NDN Interest and Data packets, in a configurable manner. Regarding network interfaces, the MA can interface with them not only to control (i.e., initiate wireless access scanning procedures), but also to collect information (i.e., an informational event regarding detected access points). Finally, the MA can also interface with application processes, drawing out information about the perceived QoS/QoE (e.g., lost packets or delay from a real-time video feed) and also to execute commands, such as selecting a better video codec when

the network commands the video flow to be accessed from a different wireless access interface.

Conversely, MA entities residing in network equipment can provide informational events as well, but related to network-side link layer characteristics (such as number of attached nodes or load), as well as accepting commands from the network (i.e., activate maintenance procedures). Management processes residing in the network core can leverage information collected from applications, client terminals and network equipment, to drive optimization procedures. Such optimization procedures can also tap into other entities, containing complementary information such as policies and subscription information, and use it to produce an overall network decision, which can then be forwarded to multiple client nodes, in a policy enforcing way.

An important consideration from the NDN architecture, is the hierarchical namespace, allowing nodes to request and convey management data, by simply using an appropriate prefix (e.g., `ccn://domain/management/ME`).

By leveraging the NDN information-centric dissemination mechanisms to convey management information and commands as content, these management extensions inherit the intrinsic capabilities of the NDN architecture, including security and reliability, which are fundamental for management procedures.

2.2. NDN Management Operations

In order to implement management operations, besides the interfacing capabilities of the MA entity mentioned in the previous section, a management framework needs other supporting mechanisms in order to provide the envisioned management capabilities, while maintaining the inherent NDN capabilities. Concretely, when nodes connect to the network, the management entities need to become aware of the management capabilities of the newly-connected node. In addition, an asynchronous information exchange capability needs to be provided, allowing not only the request of management information, but also the ability to push information towards a remote node (i.e., sending a command or an informational event).

2.2.1. Discovery Procedure

The discovery procedure is illustrated in Fig. 2 (redrawn from [NDN-MGMT]), and borrows for the procedures described in [NDN-VOIP]. The procedure starts with the newly connected User Equipment (UE) broadcasting an Interest packet (Fig. 2:1) perhaps with a well-known content name (e.g., `ccn://domain/management/mgmt-case/ME`) to its

local network.

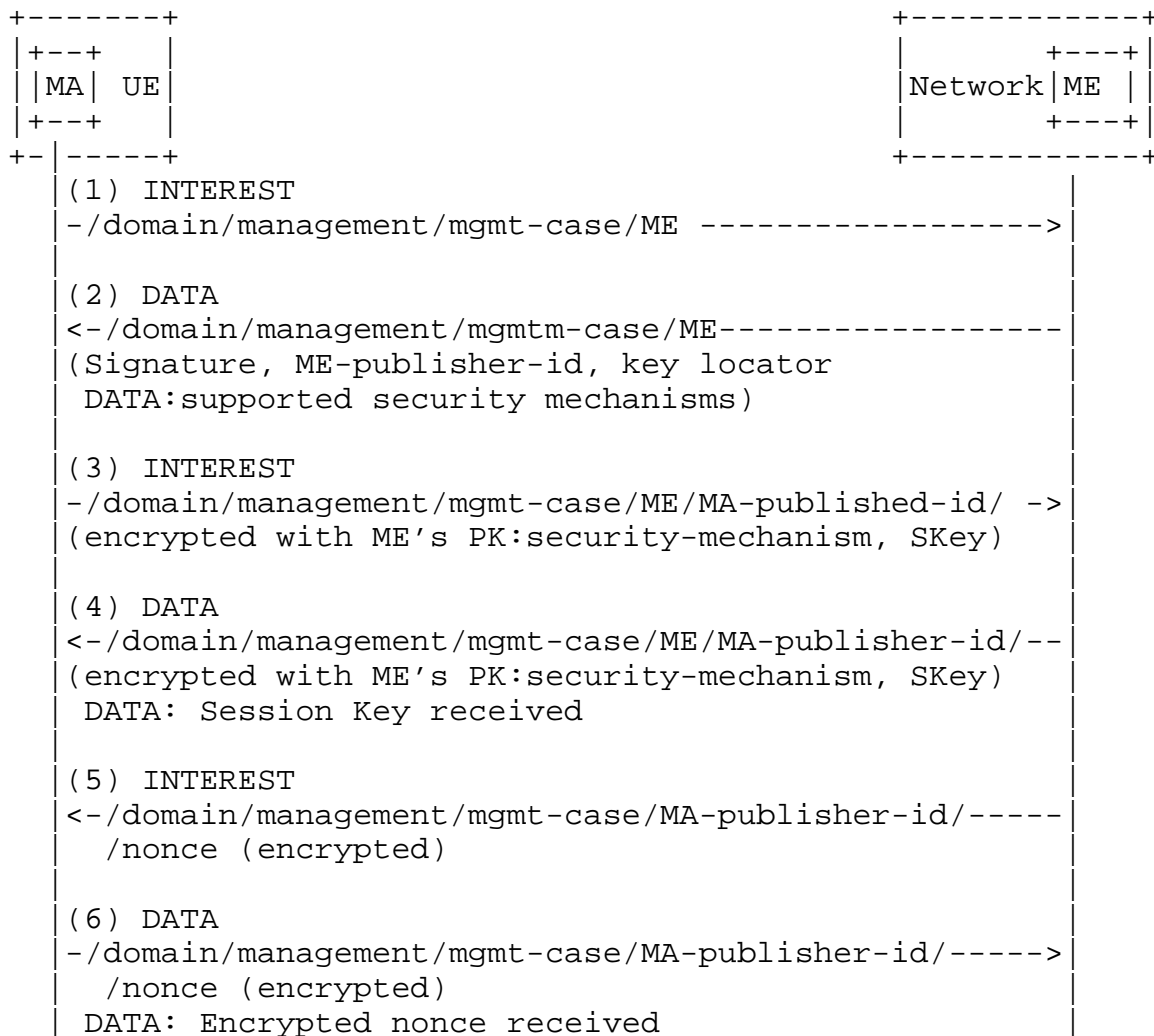


Figure 2. Secure Management Session Establishment

The "mgmt-case" part of the name can be used to select different aspects of management capabilities allowed by a Management Entity (ME) (i.e., a management decision point in the network). The ME then replies to this Interest with a Data packet (Fig. 2:2), providing its shorthand identifier (i.e., ME-publisher-key) and a key locator, indicating how to retrieve its public key (assuming it is authorized by another key trusted by the UE). In this way, the MA at the UE recognizes the ME as a valid signer (and provider) of management content.

A session key, K_s , is generated by the MA, considering an encryption algorithm from the ones indicated by the ME in the Data packet. The

MA then expresses its desire to receive (and reply to) Interests matching a specific NDN name associated with the management service (e.g., `ccn://domain/management/mgmt-case/ME/MA-publisher-id`), where MA-publisher-id uniquely and globally identifies the MA, through a cryptographic digest of its public key. After this, the MA sends an Interest packet (Fig. 2:3) to retrieve management Data from the ME containing the short-hand identifier of the MA (MA-publisher-id), the chosen encryption algorithm and session key (Ks), both encrypted with the public key of the ME. In this way, the confidentiality of the content exchanged between the ME and the MA is guaranteed. The ME responds with a Data packet (Fig. 2:4) signaling the reception of the session Key.

Before the actual exchange of management data begins, the ME generates a challenge (i.e., a nonce) which is sent via an Interest packet (Fig. 2:5) to the MA, indicating through a named data name that it requests the reception of the response to this challenge, sent by the MA using a Data packet (Fig. 2:6). This allows the ME, after verifying the signature of the Data packet, to verify that the encryption algorithm and the session key are valid for the MA, making it ready to exchange information for coordinating management procedures in the network.

2.2.2. Management Data Exchange

After the discovery and security establishment procedures have been finalized, the framework provides the capability for both the MA and the ME to securely obtain management content from one another.

In order to push unsolicited content, a dual Interest/Data procedure can maintain compatibility with the Interest and Data exchange/consumption of the NDN architecture. Fig. 3 (redrawn from Fig.2 of [NDN-MGMT]) illustrates the procedure which is initiated by the MA. In this case, the MA intends to push management information to the ME. It does so via an Interest packet manifesting its interest in receiving management content with a local sequence number. This sequencing allows the recovery of new content over cached content. If the ME is interested in retrieving content from the MA, it answers back with a Data packet, where it indicates that it is willing to receive management content. Then, the ME sends an Interest packet to retrieve the management data with the sequence number provided by the MA, which responds with a Data packet containing the information it wanted to push into the ME.

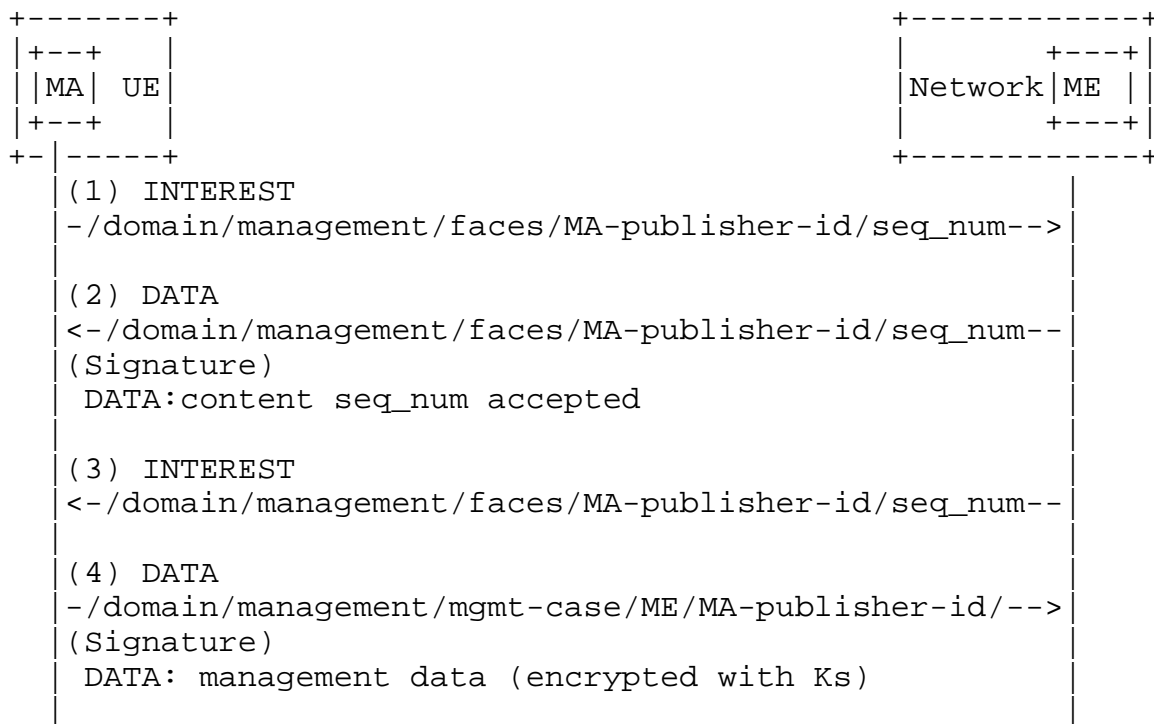


Figure 3. Content Management Push

2.3. Implementation Experience

As a proof-of-concept, a software prototype of the management framework was developed for [NDN-MGMT], using the CCNx Java API [CCNx]. At this early stage, it includes the implementation of an ME and an MA as NDN applications, supporting the NDN management operations outlined in Fig. 3. Thus, the ME and the MA can push unsolicited content to each other, related with management operations.

To validate this basic prototype, [NDN-MGMT] considered a specific use case supported by the framework, i.e., face management. This entails configuring and selecting an appropriate face in a UE to retrieve a given content. Based on the CCNx, an evaluation test-bed was deployed including an NDN UE (featuring an MA and a set of network interfaces), a content server and a network node (featuring an ME). These entities are interconnected by a set of NDN routers. The purpose of the evaluation scenario is to demonstrate feasibility for the protocol exchanges mentioned earlier. Note that the code has been tested in a small-scale environment where the ME is topology-aware and keeps track of conditions of the access networks that are available to the UE. Thus, the ME can provide the MA with management information reporting the appropriate face for content retrieval, or an alternative point of access that could be used to improve the

performance. The MA uses the management information to reconfigure the FIB (and possibly the network interfaces) in the UE, setting the appropriate face to forward subsequent Interests.

For validation purposes, a local application was also implemented at the NDN UE that works similarly to a ping utility, generating periodic Interests that match a given prefix (served by the content server), and computing the Round Trip Time of each Interest/Data exchange. The RTT values obtained by this application in [NDN-MGMT], indicate that the performance of the NDN management framework in the considered evaluation scenario is satisfactory, given the early stage of this work. Further development and testing is ongoing.

3. NetInf Management Considerations

Early-phase work in NetInf management [NetInfSelfX] discussed a two-fold problem. The first question that arises is whether it is possible by adopting a new set of network primitives and in-network storage to usher a new type of network management. In other words, can network management become information-centric while handling often host-centric data? The second question is whether an information-centric network is more suitable for self-management mechanisms than IP-based networks are. In particular with respect to the later, [NetInfSelfX] introduced some design considerations for adding self-management mechanisms in NetInf.

Of interest from this early work are two examples where network management can play a new role. First, network management can get involved in decisions about caching and (re)distribution of content, and not only whether an (inter)face is on or off, or what traffic limits should be enforced. Moreover, network policies can be distributed securely in the same way as other content in the network, removing the need for centralized management, and enabling improved recovery procedures. Second, network management can get involved in more intricate processes such as controlling multiaccess support, intermediating for content adaptation when deemed appropriate, and enabling richer tools for traffic engineering.

4. Acknowledgements

This document has benefited from comments and/or text provided by the following members of ICNRG:

Jaime Garcia-Reinoso (UC3M); Section 2.3

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

TBD

7. Informative References

- [CCNx] PARC, "CCNx Project", 2013, <<http://www.ccnx.org>>.
- [DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", SIGCOMM, ACM , 2007.
- [ICN-Scenarios]
Pentikousis, K., Ohlman, B., Corujo, D., and G. Boggia,
"ICN Baseline Scenarios", draft-pentikousis-icn-scenarios
(work in progress), February 2013.
- [NDN] Jacobson, V., Smetters, D., Thornton, J., Plass, M.,
Briggss, N., and R. Braynard, "Networking Named Content",
CoNEXT 2009, Rome , Dec 2009.
- [NDN-MGMT]
Corujo, D., Vidal, I., Garcia-Reinoso, J., and R. Aguiar,
"A named data networking flexible framework for management
communications", Communications Magazine, IEEE , vol.50,
no.12, pp.36-43 , Dec 2012.
- [NDN-R] Zhang, L. et al., "Named Data Networking (NDN) Project",
NDN Report ndn-0001, Tech Report, PARC , 2010,
<<http://www.named-data.net/techreport/TR001ndn-proj.pdf>>.
- [NDN-VOIP]
Jacobson, V., Smetters, D., Briggss, N., Plass, M.,
Steward, P., and J. Thornton, "VoCCN: Voice Over Content-
Centric Networks", ReARCH 2009, Rome , Dec 2009.
- [NetInf] Ahlgren, B. et al., "Design considerations for a network
of information", CoNEXT, Re-Arch Workshop, ACM , 2008.
- [NetInfSelfX]
Pentikousis, K. et al., "Self-Management for a Network of
Information", IEEE ICC Workshops 2009 , June 2009.

- [PURSUIT] Fotiou, N. et al., "Developing Information Networking Further: From PSIRP to PURSUIT", BROADNETS, ICST , 2010.
- [RFC1157] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol (SNMP)", STD 15, RFC 1157, May 1990.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", RFC 6733, October 2012.

Authors' Addresses

Daniel Corujo
Instituto de Telecomunicacoes
Campus Universitario de Santiago
Aveiro, P-3810-193 Aveiro
Portugal

Phone: +351 234 377 900
Email: dcorujo@av.it.pt

Kostas Pentikousis
Huawei Technologies
Carnotstrasse 4
10587 Berlin
Germany

Email: k.pentikousis@huawei.com

Ivan Vidal
UC3M
Av de la Universidad, 30
28911 Leganes, Madrid
Spain

Email: ivalidal@it.uc3m.es