ICNRG
Internet-Draft
Intended status: Informational
Expires: September 04, 2014

D. Corujo
Instituto de Telecomunicacoes
K. Pentikousis
EICT
I. Vidal
J. Garcia-Reinoso
UC3M
S. Lederer
Alpen-Adria Universitat Klagenfurt
S. Spirou
Intracom Telecom
C. Westphal
Huawei

March 03, 2014

ICN Management Considerations
draft-corujo-icn-mgmt-04

Abstract

Motivated by the need to find and evaluate better ways for reaching on-line content in upcoming Future Internet environments, ICN has been increasingly deployed in an broad range of research and experimental actions. Some deployments even go as far as subjecting ICN to new scenarios beyond content-reaching, exposing the flexibility of ICN core primitives in supporting such mechanisms. In this sense, besides analyzing and discussing the role of network management procedures in ICN environments, this document also analyzes possibilities on how intrinsic core ICN mechanisms can be reutilized for network management. We consider that the availability of management mechanisms for ICN will foster their deployment and, as such, should be tackled still in the design and experimentation phases. Perhaps ICN can adapt successful mechanisms from the host-centric paradigm, or new network management schemes can be designed. Perhaps even both. This document centralizes that discussion, drawing the attention of the ICNRG community to this underdeveloped area of research in ICN.

Table of Contents

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 ICNRG 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 ICNRG. In this document, different ICN research aspects tackled by ICNRG members are analyzed in respect to management possibilities and impact.

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, showcasing considerations on Face Management Section 2.1, Video Adaptation Section 2.2, Content Management Section 2.3, Network Policies Section 2.4 and Cache Management Section 2.5.

We plan to incrementally develop the draft, incorporating considerations on other ICN aspects as well as different approaches (e.g., [PURSUIT] and [NetInf]) as well as address other pertinent aspects as we receive feedback from the research group members.

2. ICN Management Considerations

This section addresses management considerations regarding specific ICN deployments and scenarios, by analyzing the opportunities, requirements and possibilities for management deployment therein. This analysis starts with the proposal of a NDN-based face management framework, followed by considerations from video adaptation, content management and network policies scenarios.

2.1. A Management Framework for NDN - The Face Management Example

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 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.

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 (CS), Forward Information Base (FIB) and Pending Interest Table (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.

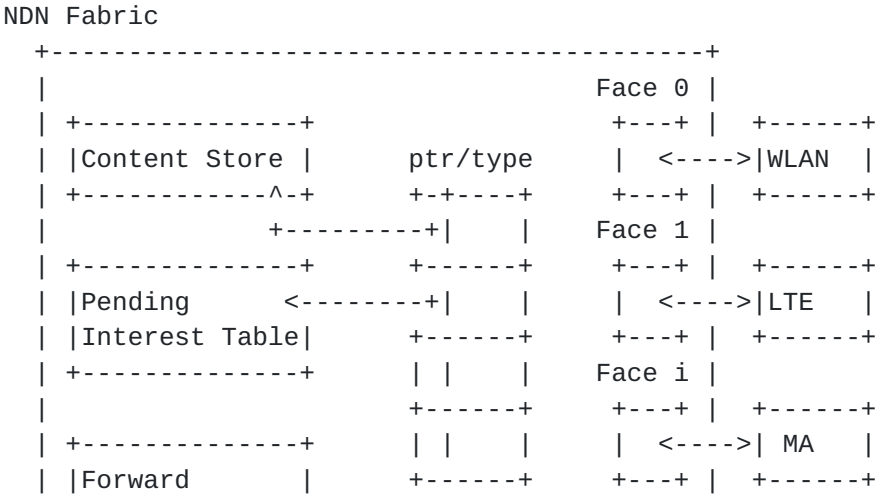

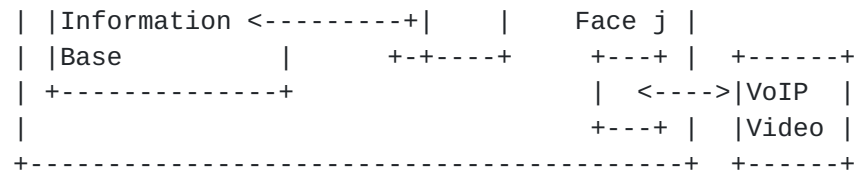

Figure 1. NDN Management Framework

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.

2.1.1. 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).

2.1.1.1. 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.

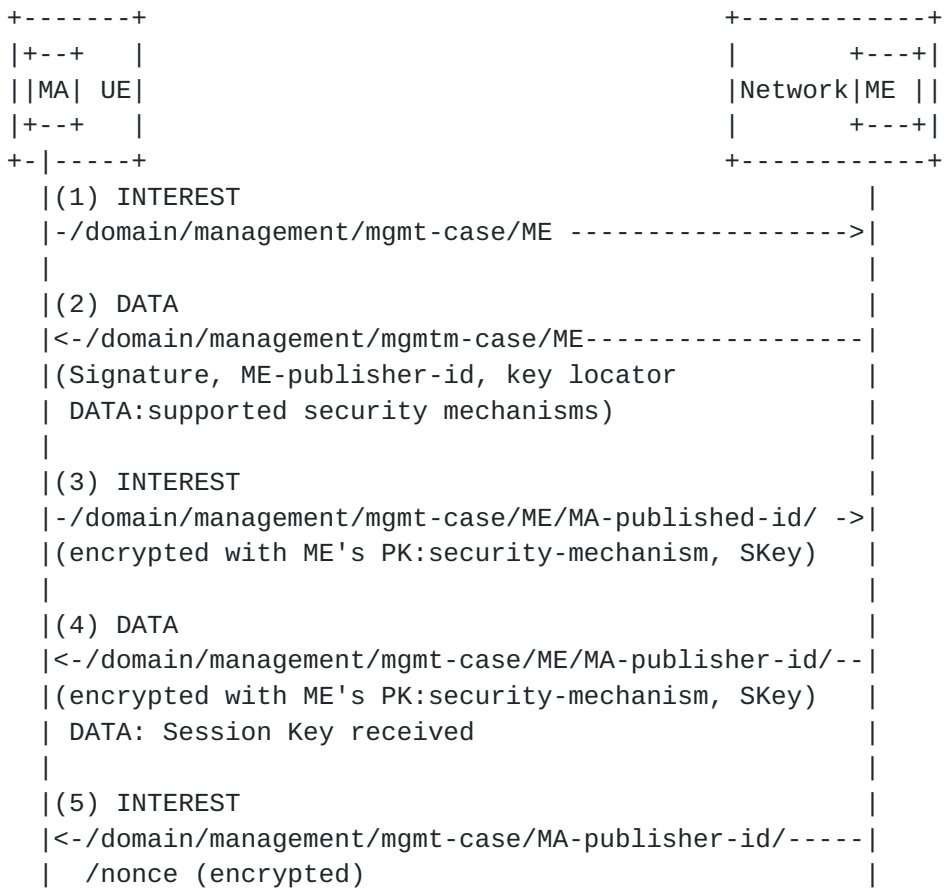

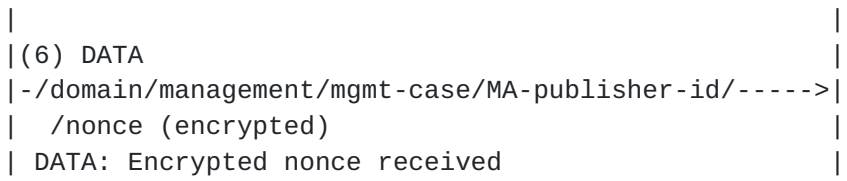

Figure 2. Secure Management Session Establishment

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, Ks, 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.

2.1.1.2. 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.

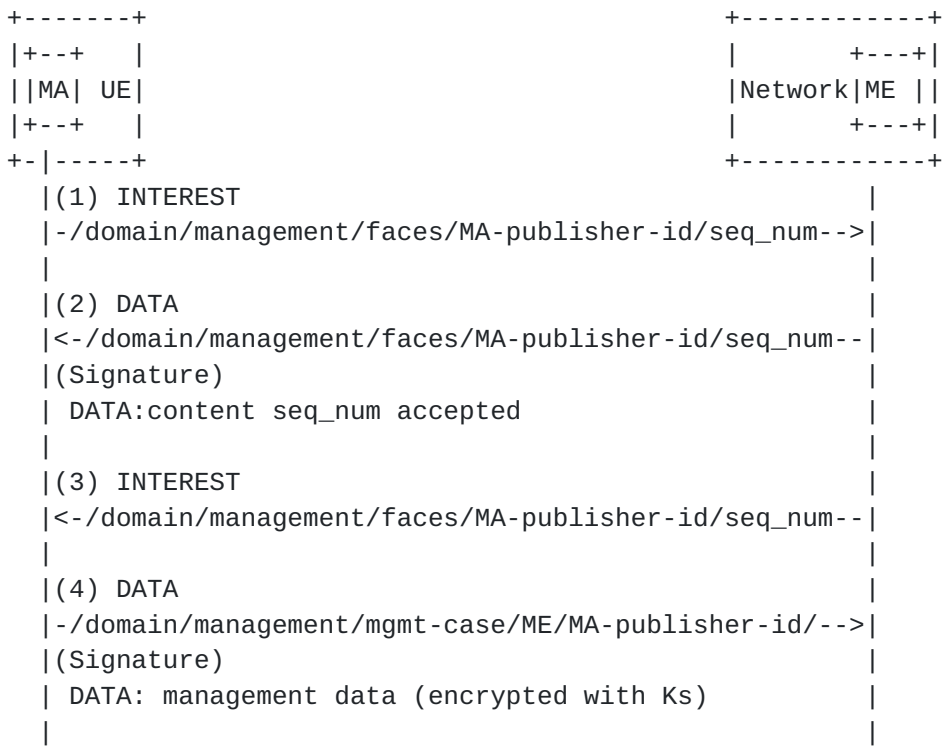

Figure 3. Content Management Push

2.1.2. Implementation Experience

As a proof-of-concept, a software prototype of the management framework, [NDNFlexManager] was developed for [NDN-MGMT], using the

CCNx Java API [CCNx]. At this stage, it includes the support of the discovery procedure described in Section Section 2.1.1.1, as well as the management data exchange operations outlined in Section Section 2.1.1.2. The framework provides an API to easy the development of management applications, i.e., MAs and MEs, which can request and 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 that can be achieved by using the NDN management framework in the considered evaluation scenario is satisfactory, given the early stage of this work. Further development and testing is ongoing.

2.2. Video Adaptation

This section investigates ICN management considerations for the delivery of video data, and especially the adaptive delivery of video. From a content perspective, multimedia is omnipresent in the Internet, e.g., producing 62% of the total Internet traffic in North America's fixed access networks [GIPR2013].

Video, and multimedia content in general, is has specific characteristics, which have to be considered and where network management consideration are necessary. The consumption of multimedia content comes along with timing requirements for the

delivery of the content, for both, live and on-demand consumption. Long startup delays, buffering periods or poor quality, etc. should be avoided to achieve a good Quality of Experience of the consumer of the content. Of course, these requirements are heavily influenced by routing decision and caching, which are central parts of ICN, and which may be leveraged more efficiently by an intelligent network management.

2.2.1. Adaptive Delivery of Multimedia Content in ICN

Today's dominant streaming systems are based on the common approach of leveraging HTTP-based Internet infrastructures, which are consequently based on the Transmission Control Protocol (TCP) and the Internet Protocol (IP). Especially the adaptive multimedia streaming (AMS) via HTTP is gaining more and more momentum and resulted in the standardization of MPEG-DASH [MPEG-DASH], which stands for Dynamic Adaptive Streaming over HTTP. The basic idea of AHS is to split up the media file into segments of equal length, which can be encoded at different resolutions, bitrates, etc. The segments are stored on conventional HTTP Web server and can be accessed through HTTP GET requests from the client. Due to this, the streaming system is pull based and the entire streaming logic is on the client side. This means that the client fully controls the bitrate of the streaming media on a per-segment basis, which has several advantages, e.g., the client knows its bandwidth requirements and capabilities best. As one can see, ICN and adaptive multimedia streaming have several elements in common, such as the client-initiated pull approach, the content being dealt with in pieces as well as the support of efficient replication and distribution of content pieces within the network. As ICN is a promising candidate for the Future Internet (FI) architecture, it is useful to investigate its suitability in combination with AMS systems and standards like MPEG-DASH as shown in [AdaptCCN][InterAdaptCCN], as well as the possibilities and benefits of intelligent network management to improve the performance of AMS in ICN as well as the resulting QoE at the client.

One of the most promising aspects in this context is the possibility of ICN to consume content from different origin nodes as well as over different network links in parallel, which can be seen as an intrinsic error resilience feature w.r.t. the network. This is a useful feature of ICN for adaptive multimedia streaming within mobile environments since most mobile devices are equipped with multiple network links. Here, a focus of ICN management could be in the load balancing of such traffic between the available links. This would increase the effective media throughput of the multimedia content, however, it could potentially lead to high variations of the resulting bandwidth which is available to the client. As DASH is designed for environments with dynamic bandwidth conditions, they can

be compensated in general. However, more conservative adaptation algorithms may prevent too frequent switching between the content's bitrate representations as well as compensate short-term bandwidth drops caused by network link switches more smoothly.

2.3. Content Management

An ICN network aims to facilitate access to, and delivery of, information objects (content and services). Content (in particular, video) access and delivery seems to be the dominant use case in traditional, host-based networks, so ICN networking is forced to adopt content delivery as a minimum requirement. Indeed, virtually all ICN approaches so far target at least content delivery.

From the perspective of a content owner or provider, an ICN network functions essentially as a content delivery network. This creates a set of requirements for ICN. First of all, end-users and content providers alike should be able to Read (consume) a content object available on the ICN network. In addition, content providers need the ability to Create (publish), Update, and Delete content. Finally, Accounting (logging) is necessary to support business models that typically require charging, analytics, and monitoring.

The Read operation has received the lion's share in ICN research. This is expected as content access and delivery is at the heart of ICN. Given a request for a named content object, the ICN network resolves that name to an object replica and proceeds with delivery to the end-user. Of course, different ICN approaches employ different mechanisms to achieve the Read operation. For example, name resolution can be done with a hierarchical system resembling DNS, with DHTs, or with flooding. Similarly, content delivery can be done over normal best-effort paths from the origin server, over dynamically computed provisioned paths, or from caches close to the end-user. Some approaches can even cater to mobile end-users and content hosts. ICN should be able to handle frequent Reads as well as Read spikes (flash crowds). In fact, it seems crucial for ICN's deployment chances to at least match the capabilities of incumbent content delivery systems.

ICN research has not addressed Create as much as Read, but some effort has been expanded on mechanisms for publishing content. Much of this effort has focused on content naming schemes that enable global uniqueness of names and hence allow global addressing of the content objects. It has been difficult to balance human readability of names, efficiency in machine processing, and name aggregation (that can realistically enable request routing by name). Although a fully automated mechanism for (human-readable) name assignment would be desirable, so far it seems that a manual process, similar to that

of domain name registration in DNS, is necessary to allocate at least namespaces. No other restrictions on naming have been seriously considered. The consensus seems to be that with ICN anyone should be able to publish anything. Content semantics are a higher layer issue. This might be a prudent approach when building a transport layer technology, but it could undermine the potential of ICN deployment. A content owner would not want copies of its content published on an ICN network under different names. In any case, once a name has been assigned, the Create operation is mainly about creating an entry in the name resolution system. This is obviously a security risk and furthermore, for highly distributed name resolution systems, it can suffer from considerable lag in availability. Fortunately, Create is a rare operation compared to Read.

Update is an operation that seeks to alter an already created object. A content provider would want to modify the data or the metadata of a published object either to rectify publication errors or to augment the object. It is debatable whether the provider should address the later simply by creating a new object. Another use case for Update comes from the need to rebrand or alias an object when its rights have been sold to another party. Nevertheless, the Update operation has received minimal attention in ICN research. The main problem is one of consistency: once an update has been committed, an ICN network with highly distributed name resolution and content delivery (caching) would host both the old and new versions of the updated content object for some time. Security concerns for the Update and Create operations are similar. Update is normally rarer than Create, but this will not be the case for collaborative media.

Content providers may occasionally need to remove a published object. This is the goal of the Delete operation. An object might be deleted when it was published by mistake, because it's no longer useful or relevant, or because it's illegal. Consistency is a major challenge for the Delete operation as well. The high degree of distribution in ICN can sustain a network state where some data or metadata replicas of an object have been deleted, while others persist. On the other hand, this lag can be beneficial if deletion was initiated erroneously or maliciously. Like with the Update operation, Delete has not been properly investigated in ICN research. Deletes are typically less often than updates.

From the point of view of content providers and end users, an ICN network resembles a content directory and repository, with Create, Read, Update, and Delete as typical operations. As with any database system, the reliability of those operations (or transactions) depends on the properties of atomicity, consistency, isolation, and durability. The challenge for ICN research is to build systems at a massive scale that employ those properties.

2.4. Network Policies

2.4.1. 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.

2.5. Cache Management

Caching is a hot topic research nowadays in ICN. The challenges of caching in ICN are different than those of web caching, mainly because the former has to deal with high line rates and with a huge amount of content. Some ICN works propose to cache content in all ICN routers traversed by the data packet, in an LCE (Leave Copy Everywhere) fashion as in [NDN]. Some studies, like [L4M-ICN], have shown that other cache decision policies, focused to reduce the cache redundancy, may increase the overall caching performance. Some of these decision policies only use the local information available at the ICN routers, but others use the information available at other nodes to cache or not the incoming content. This is known as explicit cache coordination decision, and there are several proposals around this concept [ICN-CACHING]. The idea behind the explicit coordination is to exchange topological information, individual cache's state and content popularity view among a set of ICN routers, in order to coordinate caching decisions.

This way, a given ICN router may forward a request towards another router storing the requested content. In this context, the routing

protocol is affected by the cache's state of surrounding neighbours. For example, in [CATT] the authors propose to distinguish between the source(s) and routers' caches that hold a copy of that content: the former paths are globally advertised, while the latter are only advertised within the router's neighbourhood. In all these cases, the use of a management framework may bring significant advantages, providing standard interfaces that allow the routers to dynamically manage their caches.

2.6. Fine-Grained Management of Resources

While caching has been the focus of much of the attention in ICN, one of the key advantages of the ICN architecture is that it allows a fine grained allocation of content to resources. This has been observed in [CB-TE] and [ICN-TE] for instance. Unlike IP, an ICN packet carries specific, explicit information about the content it carries. Further, this content is uniquely named, and different versions of the content will have different names.

In IP, a flow from a certain source address to a certain destination address can correspond to myriad potential applications: web traffic, video streaming, VoIP call all may use the same port 80 and be hosted by same servers. Therefore, providing appropriate resource to such a flow is a matter of guessing. The simple problem of identifying when a flow terminates is made unnecessarily complex in ICN: a timer is set-up, and when no packets match the flow filter, then the flow is over. Of course, multiple packets from different applications may match the same filter, and flows with different characteristics in terms of inter-arrival times could be broken down into multiple flows with an improper choice of time-out values.

In ICN, there is a unique mapping of the name to the content of the data stream going through the network. If a content object is requested, then it has well defined semantics, and the network management layer can identify exactly when the data stream starts and ends based upon these semantics. Further, a content management layer can also learn the properties of the stream associated with a given identifier. [CB-TE] presented such a mechanism to learn the properties associated with a name, either by counting the bytes on the wire corresponding to this name, or by reading the footprint of the content with this name when stored in a cache. It is therefore possible for the management layer to gather meta-data pertaining to the content that goes through the network, and to use this meta-data to make proper resource allocations.

Of course, the resource manager should only acquire meta-data about content that is likely to be seen again (i.e., popular) or specific in any way (for instance, the name of an elephant flow). This

considerably simplifies the task as the data of interest is concentrated on a few items. One potential usage of this meta-data is to keep track of what content is going through which link. In this scenario, each link keeps an aggregate tally of the amount of data that has been assigned to this link and subtracts the amount that has gone through. The resulting backlog can be then used to allocate new data streams to this or another link.

In [ICN-TE], it was shown that such a policy would significantly reduce the time spent in the network by content streams when considering a WAN topology and its corresponding end-to-end traffic matrix. The network load would stay the same when comparing with a min-MLU policy, but by splitting elephant flows across different paths, the completion time would be reduced. In the simulations of [ICN-TE], min-MLU is roughly 50% slower than a content-based policy.

This is an encouraging result and a step towards a management framework that assigns resource to content in a deterministic and fine-grained manner, unlike the probabilistic allocation of IP. The ICNRG should consider such a management framework, and evaluate the different proposals in light of this opportunity. For instance, an ICN architecture such as [PURSUIT] contains a natural mechanism to perform such allocation of content to paths as it assigns a source route to the content. On the other hand, an ICN architecture such as [NDN] needs to be expanded as the link allocation semantics are, in the current proposal, tied to the content resolution process: the interest homes into the content, and lays the reverse path for the content delivery at the same time. This semantics make the management for multiple link selection more difficult, as multiple interests would have to be sent over multiple links to provide path diversity. However, it could be an area of study for the ICNRG as solving such resource management problem would provide significant benefits to ICN architectures.

3. Acknowledgements

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

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

TBD

6. Informative References

[AdaptCCN] Lederer, S., Mueler, C., Rainer, B., Timmerer, C., and H. Hellwagner, "Adaptive Streaming over Content Centric Networks in Mobile Networks using Multiple Links", Proceedings of the IEEE International Conference on Communication (ICC), Budapest, Hungary , June 2013.

[CATT] Eum, S., Nakauchi, K., Murata, M., Shoji, Y., and N. Nishinaga, "CATT: potential based routing with content caching for ICN", Workshop on Information-centric networking, pp 49-54 , 2012.

[CB-TE] Chanda, A. et al., "Content Based Traffic Engineering in Software Defined Information Centric Networks", IEEE INFOCOM Workshop NOMEN , April 2013.

[CCNx] PARC, "CCNx Project", 2013, <http://www.ccnx.org>.

[DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", SIGCOMM, ACM , 2007.

[GIPR2013] Sandivine, , "Global Internet Phenomena Report 1H 2013", Sandvine Intelligent Broadband Networks , 2013.

[ICN-CACHING] Zhang, G., Li, Y., and T. Lin, "Caching in information centric networking: A survey", Computer Networks, vol. 57, no. 16, pp. 3128-3141, Nov , 2013.

[ICN-Scenarios] Pentikousis, K., Ohlman, B., Corujo, D., and G. Boggia, "ICN Baseline Scenarios", draft-pentikousis-icn-scenarios (work in progress), February 2013.

[ICN-TE] Su, K. et al., "On the Benefit of Information Centric Networks for Traffic Engineering", IEEE ICC , June 2014.

[InterAdaptCCN]

Grandl, R., Su, K., and C. Westphal, "On the Interaction of Adaptive Video Streaming with Content-Centric Networking", Proceedings of the 20th Packet Video Workshop 2013, San Jose, USA , December 2013.

[L4M-ICN] Chai, W. K., He, D., Psaras, I., and G. Pavlou, "Cache "less for more" in information-centric networks", Lecture Notes in Computer Science Vol. 7289, Springer, pp. 27-40 , 2012.

[MPEG-DASH] Sodagar, I., "The MPEG-DASH Standard for Multimedia Streaming Over the Internet", IEEE MultiMedia, IEEE, vol.18, no.4, pp.62-67 , 2011.

[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.K., Briggss, N.H., Plass, M.F., Steward, P., and J.D. Thornton, "VoCCN: Voice Over Content-Centric Networks", ReARCH 2009, Rome , Dec 2009.

[NDNFlexManager] UC3M and ITAV, "Framework for Flexible NDN Management", 2013, <https://github.com/ndnflexmanager/framework>.

[NDN] Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggss, N.H., and R.L. Braynard, "Networking Named Content", CoNEXT 2009, Rome , Dec 2009.

[NetInfSelfX] Pentikousis, K. et al., "Self-Management for a Network of Information", IEEE ICC Workshops 2009 , June 2009.

[NetInf] Ahlgren, B. et al., "Design considerations for a network of information", CoNEXT, Re-Arch Workshop, ACM , 2008.

[PURSUIT] Fotiou, N. et al., "Developing Information Networking Further: From PSIRP to PURSUIT", BROADNETS, ICST , 2010.

[RFC1157] Case, J.D., Fedor, M., Schoffstall, M.L., and J.R. 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
EICT GmbH
Torgauer Strabe 12-15
10829 Berlin
Germany

Email: k.pentikousis@eict.de

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

Email: ividal@it.uc3m.es

Jaime Garcia-Reinoso
UC3M
Av de la Universidad, 30
28911 Leganes, Madrid
Spain

Email: jgr@it.uc3m.es

Stefan Lederer
Alpen-Adria Universitat Klagenfurt
Universitatsstrasse 65-67
Klagenfurt
Austria

Email: stefan.lederer@itec.aau.at

Spiros Spirou
Intracom Telecom
19.7 km Markopoulou Avenue
Peania 19002
Greece

Email: spis@intracom.com

Cedric Westphal
Huawei
2330 Central Expressway
Santa Clara, CA95050
USA

Email: cedric.westphal@huawei.com