Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 5, 2015

D. Kutscher, Ed.
NEC
S. Eum
NICT
K. Pentikousis
EICT
I. Psaras
UCL
D. Corujo
Universidade de Aveiro
D. Saucez
INRIA
T. Schmidt
HAW HAmburg
M. Waehlisch
FU Berlin
February 1, 2015

ICN Research Challenges
draft-irtf-icnrg-challenges-01

Abstract

This memo describes research challenges for Information-Centric Networking. Information-Centric Networking is an approach to evolve the Internet infrastructure to directly support information distribution by introducing uniquely named data as a core Internet principle. Data becomes independent from location, application, storage, and means of transportation, enabling in-network caching and replication. Challenges include naming, security, routing, system scalability, mobility management, wireless networking, transport services, in-network caching, and network 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 5, 2015.

Table of Contents

1. Introduction

Distributing and manipulating named information is a major application in the Internet today. In addition to web-based content distribution, other distribution technologies (such as P2P and CDN) have emerged and are promoting a communication model of accessing data by name, regardless of origin server location.

In order to respond to increasing traffic volume in the current Internet for applications such as mobile video and cloud computing, a set of disparate technologies and distribution services are applied that employ caching, replication and content distribution in different specific ways. These approaches are currently deployed in separate silos -- different CDN providers and P2P applications rely on their own, often proprietary, distribution technologies. It is not possible to uniquely and securely identify named information independently of the distribution channel; and the different distribution approaches are typically implemented as an overlay, potentially leading to unnecessary inefficiency.

For example, creating and sharing multimedia content in a social networking application today, typically requires uploading data objects to centralized service provider platforms, from where it can be accessed individually by other users. Even if content sharing is intended to happen locally, e.g., in a local network or local area, the actual communication will require interactions from any interested user with the service provider. CDNs can alleviate the situation only partly, because, due to organizational and economic reasons, it is not common to deploy CDN gear ubiquitously. Moreover, since CDNs and the respective communication sessions form overlays, the actual communication, i.e., the requests for named content and the actual responses, are largely invisible to the network, i.e., it is not easily possible to optimize efficiency and performance. For example in a wireless access network, it is not possible to leverage the inherent broadcast nature of the medium (and avoid duplicate transmission of the same content) due to limitations from point-to-point and overlay communication.

Information-centric networking (ICN) is an approach to evolve the Internet infrastructure to directly support this use by introducing named data as a core network-layer primitive. Data objects become independent of location, application, storage, and means of transportation, allowing for inexpensive and ubiquitous in-network caching and replication. The expected benefits are improved efficiency, better support for provenance verification and name-content binding validation, better scalability with respect to information/bandwidth demand and better robustness in challenging communication scenarios.

ICN concepts can be deployed by retooling the protocol stack: name-based data access can be implemented on top of the existing IP infrastructure, e.g., by providing resource naming, ubiquitous caching and corresponding transport services, or it can be seen as a packet-level internetworking technology that would cause fundamental changes to Internet routing and forwarding. In summary, ICN is expected to evolve the Internet architecture towards a network model that is more suitable for the current and future usage patterns.

This document presents the ICN research challenges that need to be addressed in order to achieve these goals. These research challenges are seen from a technical perspective, although business relationships between Internet players will also influence developments in this area. We leave business challenges for a separate document, however. The objective of this note is to document the technical challenges and corresponding current approaches and to expose requirements that should be addressed by future research work.

2. Problems with Information Distribution Today

The best current practice to manage the above-mentioned growth in terms of data volume and number of devices is to increase infrastructure investment, employ application-layer overlays such as CDNs, P2P applications, and M2M application platforms that cache content, provide location-independent access to data, and optimize its delivery. In principle, such platforms provide a service model of accessing named data objects (NDOs) (e.g., replicated web resources, M2M data in data centers) instead of a host-to-host packet delivery service model.

However, since this functionality resides in overlays only, the full potential of content distribution and M2M application platforms cannot be leveraged as the network is not aware of data requests and data transmissions. This has the following impact:

- data traffic typically follows sub-optimal paths as it is effectively routed depending on the overlay topology instead of the Internet layer topology;
- network capabilities, such as multicast and broadcast, are largely underutilized or not employed at all. As a result, request and delivery for the same object have to be made multiple times;
- overlays typically require significant infrastructure support, e.g., authentication portals, content storage, and applications servers, making it often impossible to establish local, direct

communication;

- since the network is not aware of the nature of data objects, it is unable to manage access and transmission (without layer violations);
- provenance validation uses host authentication today. As such, even if there are locally cached copies available, it is normally not easily possible to validate their authenticity; and
- many applications follow their own approach to caching, replication, transport, authenticity validation (if at all), although they all share similar models for accessing named data objects in the network.

3. ICN Terminology and Concepts

3.1. Terminology

Information-Centric Networking (ICN): A concept for communicating in a network that provides accessing named data objects as a first order service. See Section 3.2 for details.

Named Data Object (NDO): Addressable data unit in an ICN that can represent a collection of bytes or a piece of information. In ICN, each data object has a name bound to it, and there are typically mechanisms to secure (and validate) this binding. Different ICN approaches have different concepts for how to map NDOs to individual units of transport. Within the context of this document, an NDO is any named object that can be requested from the network.

Requestor: Entitiy in an ICN network that is sending a request for a Named Data Object to the network.

3.2. Concepts

Fundamentally, ICN is providing access to named data as a first-order network service, i.e., the network is able to serve requests to named data natively. That means, network nodes can receive requests for named data and act as necessary, for example, by forwarding the request to a suitable next-hop. Consequently, the network processes requests for named data objects (and corresponding responses) natively. Every network node on a path is enabled to perform forwarding decisions, to cache objects etc. This enables the network to forward such requests on optimal paths, employing the best transmission technologies at every node, e.g., broadcast/multicast

transmission in wireless networks to avoid duplicate transmission of both requests and responses.

In ICN there is a set of common concepts and node requirements beyond this basic service model. Naming data objects is a key concept. In general, ICN names represent neither network nodes nor interfaces -- they represent NDOs independently of their location. Names do play a key role in forwarding decisions and are used for matching requests to responses: In order to provide better support for accessing copies of NDOs regardless of their location, it is important to be able to validate that a response actually delivers the bits that correspond to an original request for named data.

Name-content binding validation is a fundamental security service in ICN, and this is often achieved by establishing a verifiable binding between the object name and the actual object or an identity that has created the object. ICN can support other security services, such as provenance validation, encryption -- depending on the details of naming schemes, object models and assumptions on infrastructure support. Security services such as name-content binding validation are available to any node, i.e., not just the actual receivers. This is an important feature, for enabling ingress gateways to check object authenticity to prevent denial-of-service attacks.

Based on these fundamental properties it is possible to leverage network storage ubiquitously: every node and every device can cache data objects and respond to requests for such objects -- it is not required to validate the authenticity of the node itself since name-content bindings can be validated. Ubiquitous in-network storage can be used for different purposes: it can enable sharing, i.e., the same object copy can be delivered to multiple users/nodes as in today's proxy caches and CDNs. It can also be used to make communication more robust (and perform better) by enabling the network to answer requests from local caches (instead of from origin servers). In case of disruption (message not delivered), a node can re-send the request, and it could be answered by an on-path cache, i.e., on the other side of the disrupted link. The network itself would thus support retransmissions -- enabling shorter round-trip times and offloading origin servers and other parts of the network.

The request/response model and ubiquitous in-network storage also enable new options for implementing transport services, i.e., reliable transmission, flow control, etc. First of all, a request/response model can enable receiver-driven transport regimes, i.e., receivers (the requestors of NDOs) can control message sending rates by regulating the request sending rate (assuming that every response message has to be triggered by a request message). Retransmission would be achieved by re-sending requests, e.g., after a timeout.

Because objects can be replicated, object transmission and transport sessions would not necessarily have end-to-end semantics: requests can be answered by caches, and a node can select one or multiple next-hop destinations for a particular request -- depending on configuration, observed performance or other criteria.

This receiver-driven communication model potentially enables new interconnection and business models: a request for named data can be linked to an interest of a requestor (or requesting network) in data from another peer, which could suggest modeling peering agreements and charging accordingly.

4. ICN Research Challenges

4.1. Naming and data authenticity

Naming data objects is as important for ICN as naming hosts is for today's Internet. Fundamentally, ICN requires unique names for individual NDOs, since names are used for identifying objects independently of their location or container. In addition, since NDOs can be cached anywhere, the origin cannot be trusted anymore hence the importance to establish a verifiable binding between the object and its name (name-data integrity) so that a receiver can be sure that the received bits do correspond to the NDO originally requested (object authenticity). Information about an object's provenance, i.e., who generated or published it, is also useful to associate to the name.

The above functions are fundamentally required for the information-centric network to work reliably, otherwise neither network elements nor receivers can trust object authenticity. Lack of this trust enables several attacks including DoS attacks by injecting spoofed content into the network. There are different ways to use names and cryptography to achieve the desired functions [ICNNAMING] [ICNSURVEY], and there are different ways to manage namespaces correspondingly.

Two types of naming schemes have been proposed in the ICN literature: hierarchical and flat namespaces. For example, a hierarchical scheme may have a structure similar to current URIs, where the hierarchy is rooted in a publisher prefix. Such hierarchy enables aggregation of routing information, improving scalability of the routing system. In some cases, names are human-readable, which makes it possible for users to manually type in names, reuse, and, to some extent, map the name to user intent.

The second general class of naming schemes follows a "self-

certifying" approach, meaning that the object's name-data integrity can be verified without requiring a public key infrastructure (PKI) or other third party to first establish trust in the key. Self-certification is achieved, e.g., by binding the hash of the NDO content to the object's name. For instance, this can be done by directly embedding the hash of the content in the name. Another option is an indirect binding, which embeds the public key of the publisher in the name and signs the hash of the content with the corresponding secret key. The resulting names are typically non-hierarchical, or flat, although the publisher field could be employed to create a structure which could facilitate route aggregation. There are several design trade-offs for ICN naming, which affect routing and security. Self-certifying names are not human readable nor hierarchical. They can however provide some structure for aggregation, for instance, a name part corresponding to a publisher.

Research challenges specific to naming include:

- naming static data objects can be performed by using content hashes as part of object names, so that publishers calculate the hash over existing data objects and receivers (or any ICN node) can validate the name-content binding by re-calculating the hash and comparing it to the name (component). [RFC6920] specifies a concrete naming format for this.

- naming dynamic objects refers to use cases where the name has to be generated before the object is created. For example, this could be the case for live streaming, when a publisher wants to make the stream available by registering stream chunk names in the network. One approach to this can be self-certified names as described above.

- requestor privacy protection can be a challenge in ICN as a direct consequence of the accessing-named-data-objects paradigm: if the network can "see" requests and responses, it can also log request history for network segments or individual users, which can be undesirable, especially since names are typically expected to be long-lived. That is, even if the name itself does not reveal much information, the assumption is that the name can be used to retrieve the corresponding data objects in the future.

- Updating and versioning NDO can be challenging because it can contradict fundamental ICN assumptions: if an NDO can be replicated and stored in in-network storage for later retrieval, names have to be long-lived -- and the name-content binding must not change: updating an object (i.e., changing the content without generating a new name) is not possible. Versioning is one possible solution, but requires a naming scheme that supports it

(and a way for requestors to learn about versions).

- Managing accessibility: whereas in ICN the general assumption is to enable ubiquitous access to NDOs, there can be relevant use cases where access to objects should be restricted, for example to a specific user group. There are different approaches for this, such as object encryption (requiring key distribution and related mechanisms) or the concept of scopes, e.g., based on names that can only be used/resolved under some constraints.

4.2. Security

Security can take many different forms in ICN and instead of discussing specific attacks or technical details, we propose here the most important security challenges that come from the shift to information-centric communications. Some challenges are well-understood, and there are (sometimes multiple different) approaches to address them, whereas other challenges are active research and engineering topics.

4.2.1. Data Object Authentication

As mentioned in section Section 4.1, data object authentication is an important ICN feature, since ICN data objects are retrieved not only from an original copy holder but also from any caching point. Hence, the communication channel endpoints to retrieve NDOs are not trustable anymore and solutions widely used today such as TLS [RFC5246] cannot be used as a general solution. Since data objects can be maliciously modified ICN should provide users with a security mechanism to verify the origin and integrity of the data object, and there are different ways to achieve this.

An efficient approach for static NDOs is providing a name-content-binding by hashing an NDO and using the hash as a part of the object’s name. [RFC6920] provides a mechanism and a format for representing a digest algorithm and the actual digest in a name (amongst other information).

For dynamic objects (where it is desirable to refer to an NDO by name before the object has been created), public-key cryptography is often applied, i.e., every NDO would be authenticated by means of a signature performed by the data object publisher so that any data object consumer can verify the validity of the data object based on the signature. However, in order to verify the signature of an object, the consumer must know the public key of the entity that signed the object.

One research challenge is then to support a mechanism to distribute

the publisher's public keys to the consumers of data objects. There are two main approaches to achieve this; one is based on an external third party authority such as hierarchical Public Key Infrastructure (PKI) [RFC5280] and the other is to adapt a self-certifying scheme. The former, as the name implies, depends on an external third party authority to distribute the public key of the publisher for the consumers. In a self-certifying scheme, the public key (or a hash of it) would be used as part of the name -- which is sufficient to validate the object's authenticity.

In cases where information about the origin of a data object is not available by other means, the object itself would have to incorporate the necessary information to determine the object publisher, for example with a certificate, that can be validated through the PKI. Once the certificate is authenticated, its public key can be used to authenticate the signed data object itself.

4.2.2. Binding NDOs to Real-World Identities

In addition to validating NDO authenticity, it is still important to bind real-world identities, e.g., a publisher identity, to objects, so that a requestor can verify that a received object was actually published by a certain source.

With hash-based and self-certifying names, real-world-identity bindings are not intrinsically established: the name provides the hash of the NDO or of the public key that has been used to sign the NDO. There needs to be another binding to a real-world-identity if that feature is requested.

If the object name directly provides the publisher name and if that name is protected by a certificate that links to PKI-like trust chain, the object name itself can provide an intrinsic binding to a real-world identity.

Binding between NDOs and Real-World Identities is essential but there is no universal way to achieve it as it is all intrinsic to a particular ICN approach.

4.2.3. Traffic aggregation and filtering

One request message to retrieve a data object can actually aggregate requests coming from several consumers. This aggregation of requests reduces the overall traffic but makes per-requestor filtering harder. The challenge in this case is to provide a mechanism that allows request aggregation and per-requestor filtering. A possible solution is to indicate the set of requestors in the aggregated request such that the response can indicate the subset of requestors allowed to

access the data object. However, this solution requires collaboration from other nodes in the network and is not suitable for caching. Another possible solution is to encrypt data objects and ensure that only authorised consumers can decrypt them. This solution does not preclude caching and does not require collaboration from the network. However, it implies a mechanism to generate group keys (e.g., different private keys can be used to decrypt the same encrypted data object) [Chaum].

4.2.4. State overloading

ICN solutions that implement state on intermediate routers for request routing or forwarding (e.g., CCN [CCN]) are subject to denial of service attacks from overloading or superseding the internal state of a router (e.g., 'interest flooding' [BACKSCATTER]). Additionally, stateful forwarding can enable attack vectors such as resource exhaustion or complexity attacks to the routing infrastructure. The challenge is then to provision routers and construct internal state in a way that alleviates sensibility to such attacks. The problem becomes even harder, if the protocol does not provide information about the origin of messages. Without origin, it is a particular challenge to distinguish between regular (intense) use and misuse of the infrastructure.

4.2.5. Delivering data objects from replicas

A common capability of ICN solutions is data replication and in-network storage. Delivering replicated data objects from caches decouples content consumption from data sources, which leads to a loss of control on (1) content access, and (2) content dissemination. In a widely distributed, decentralized environment like the Internet, this raises several challenges.

One group of challenges is related to content management. Without access control, a content provider looses the means to count and survey content consumption, to limit access scopes, to control or know about the number of copies of its data in the network, or to withdraw publication reliably. Any non-cooperative or desynchronized data cache may hinder an effective content management policy.

Another group of challenges arises from potential traffic amplifications in the decoupled environment. ICN solutions that attempt to retrieve content from several replicas in parallel, or decorrelated network routing states, but also distributed attackers may simultaneously initiate the transmission of content from multiple replicas towards the same destination (e.g., 'initiated overloads' or 'blockades' [BACKSCATTER]). Methods for mitigating such threats need rigorous forwarding checks that require alignment with caching

procedures (e.g., on-path or off-path).

4.2.6. Cryptographic robustness

Content producers sign their content to ensure the integrity of data and to allow for data object authentication. This is a fundamental requirement in ICN due to distributed caching. Publishers, who (a) massively sign content, which is (b) long-lived, offer time and data to an attacker for comprising cryptographic credentials. Signing large amount of data eases common attacks that try to breach the key of the publisher. Based on this observation, the following research challenges appear. To which extent does the content publication model conflict with cryptographic limitations? How can we achieve a transparent re-signing without introducing additional cryptographical weaknesses or key management overhead?

4.2.7. Routing and forwarding information bases

In ICN networks, one attack vector is to increase the size of routing and forwarding information bases at ICN nodes, i.e., attacking routing scalability in networks that rely on routing by name. This is an intrinsic ICN security issue: possible mitigation approaches include combining routing information authenticity validation with filtering (e.g., maximum deaggregation level whenever applicable, black lists, etc.).

4.3. Routing and Resolution System Scalability

ICN routing is a process that finds a data object based on its name initially provided by a requestor. ICN routing may comprise three steps: a name resolution step, a discovery step, and a delivery step. The name resolution step translates the name of the requested data object into its locator. The discovery step routes the request to data object based on its name or locator. The last step (delivery) routes the data object back to the requestor. Depending on how these steps are combined, ICN routing schemes can be categorized as: Route-By-Name Routing (RBNR), Lookup-By-Name Routing (LBNR), and Hybrid Routing (HR).

4.3.1. Route-By-Name Routing (RBNR)

RBNR omits the first name resolution step. The name of data object is directly used to route the request to the data object. Therefore, routing information for each data object has to be maintained in the routing table. Since the number of data objects is very large (estimated as 10^11 back in 2007 [DONA] but this may be significantly larger than that, e.g., 10^15 to 10^22), the size of routing tables becomes a concern, as it can be proportional to the number of data

object unless an aggregation mechanism is introduced. On the other hand, RBNR reduces overall latency and simplifies the routing process due to the omission of the resolution process. For the delivery step, RBNR needs another identifier (ID) of either host or location to forward the requested data object back to the requestor. Otherwise, an additional routing mechanism has to be introduced, such as bread-crumbs routing [BREADCRUMBS], in which each request leaves behind a trail of breadcrumbs along its forwarding path, and then the response is forwarded back to the requestor consuming the trail. Specific challenges include:

- How to aggregate the names of data objects to reduce the number of routing entries?

- How does a user learn the name which is designed for aggregation by provider? (For example, although we name our contribution as "ICN research challenge", IRTF (provider) may want to change the name to "/IETF/IRTF/ICN/Research challenge" for aggregation. In this case, how does a user learn the name "/IETF/IRTF/ICN/Research challenge" to retrieve the contribution initially named "ICN research challenge" without any resolution process?)

- Without introducing the name aggregation scheme, can we still achieve scalable routing by taking advantage of topological structure and distributed copies? For example, employing compact routing [COMPACT], random walk [RANDOM] or greedy routing [GREEDY].

- How to incorporate copies of a data object in in-network caches in this routing scheme?

4.3.2. Lookup-By-Name Routing (LBNR)

LBNR uses the first name resolution step to translate the name of requesting data object into its locator. Then, the second discovery step is carried out based on the locator. Since IP addresses could be used as locators, the discovery step can depend on the current IP infrastructure. The delivery step can be implemented similarly to IP routing. The locator of the requestor is included in the request message, and then the requested data object is delivered to the requestor based on the locator. A specific instantiation of such a system is [MDHT]. Specific challenges include:

- How to build a scalable resolution system which provides

  * Fast lookup: mapping the name of data object to its locators (copies as well).

    * Fast update: the location of data object is expected to change frequently. Also, multiple data objects may change their locations at the same time, e.g., data objects in a laptop.

- How to incorporate copies of a data object in in-network caches in this routing scheme?

4.3.3. Hybrid Routing (HR)

HR combines RBNR and LBNR to benefit from their advantages. For instance, within a single administrative domain, e.g., an ISP, where scalability issues can be addressed with network planning, RBNR can be adopted to reduce overall latency by omitting the resolution process. On the other hand, LBNR can be used to route between domains which have their own prefix (locator). A specific challenge here is:

- How to design a scalable mapping system which, given the name of data object, it should return a destination domain locator so that a user request can be encapsulated and forwarded to the domain?

4.4. Mobility Management

Mobility management has been an active field in host-centric communications for more than two decades. In IETF in particular, starting with [RFC2002], a multitude of enhancements to IP have been standardized aiming to "allow transparent routing of IP datagrams to mobile nodes in the Internet" [RFC5944]. In a nutshell, mobility management for IP networks is locator-oriented and relies on the concept of a mobility anchor as a foundation for providing always-on connectivity to mobile nodes. Other standards organizations, such as 3GPP, have followed similar anchor-based approaches. Traffic to and from the mobile node must flow through the mobility anchor, typically using a set of tunnels, enabling the mobile node to remain reachable while changing its point of attachment to the network.

Needless to say, an IP network which supports node mobility is more complex than one that does not, as specialized network entities must be introduced in the network architecture. This is reflected in the control plane as well, which carries mobility-related signaling messages, establishes and tears down tunnels and so on. While mobile connectivity was an afterthought in IP, in ICN this is considered a primary deployment environment. Most, if not all, ICN proposals consider mobility from the very beginning, although at varying levels of architectural and protocol detail. That said, no solution has so far come forward with a definite answer on how to handle mobility in ICN using native primitives. In fact, we observe that mobility appears to be addressed on ICN proposal specific basis. That is,

there is no single paradigm solution, akin to tunneling through a mobility anchor in host-centric networking, that can be applied across different ICN proposals. For instance, although widely-deployed mobile network architectures typically come with their own network entities and associated protocols, they follow the same line of design with respect to managing mobility. This design thinking, which calls for incorporating mobility anchors, permeates in the ICN literature too.

However, employing mobility anchors and tunneling is probably not the best way forward in ICN research for mobile networking. Fundamentally this approach is anything but information-centric and location-indepedent. In addition, as argued in [SEEN], current mobility management schemes anchor information retrieval not only at a specific network gateway (e.g., home agent in Mobile IP) but due to the end-to-end nature of host-centric communication also at a specific correspondent node. However, once a change in the point of attachment occurs, information retrieval from the original "correspondent node" may be no longer optimal. This was shown in [MANI], for example, where a simple mechanism that triggers the discovery of new retrieval providers for the same data object, following a change in the point of attachment, clearly outperforms a tunnel-based approach like Mobile IP in terms of object download times. The challenge here is how to capitalize on location information while facilitating the use of ICN primitives which natively support multicast and anycast.

ICN naming and name resolution, as well as the security features that come along should natively support mobility. For example, CCN [CCN] does not have the restriction of spanning tree routing, so it is able to take advantage of multiple interfaces or adapt to the changes produced by rapid mobility (i.e., there is no need to bind a layer 3 address with a layer 2 address). In fact, client mobility can be simplified by allowing requests for new content to normally flow from different interfaces, or through newly connected points of attachment to the network. However, when the node moving is the (only) content source, it appears that more complex network support might be necessary, including forwarding updates and cache rebuilding. A case in point is a conversation network service, such as a voice or video call between two parties. The requirements in this case are more stringent when support for seamless mobility is required, esp. when compared to content dissemination that is amenable to buffering. Another parameter that needs to be paid attention to is the impact of using different wireless access interfaces based on different technologies, where the performance and link conditions can vary widely depending of numerous factors.

In host-centric networking, mobility management mechanisms ensure

optimal handovers and (ideally) seamless transition from one point of attachment to another. In ICN, however, the traditional meaning of "point of attachment" no longer applies as communication is not restrained by location-based access to data objects. Therefore, a "seamless transition" in ICN ensures that content reception continues without any perceptible change from the point of view of the ICN application receiving that content. Moreover, this transition needs to be executed in parallel with ICN content identification and reaching mechanisms enabling scenarios, such as, preparation of the content reaching process at the target connectivity point, prior to the handover (to reduce link switch disturbances). Finally, these mobility aspects can also be tightly coupled with network management aspects, in respect to policy enforcement, link control and other parameters necessary for establishing the node’s link to the network.

In summary, the following research challenges on ICN mobility management can be derived:

- How can mobility management take full advantage of native ICN primitives?
- How do we avoid the need for mobility anchors in a network that by design supports multicast, anycast and location-indepedent information retrieval?
- How can content retrieval mechanisms interface with specific link operations, such as identifying which links are available for certain content?
- How can mobility be offered as a service, which is only activated when the specific user/content/conditions require it?
- How can mobility management be coordinated between the node and the network for optimization and policing procedures?
- How do we ensure that managing mobility does not introduce scalability issues in ICN?
- How will the name resolution process be affected by rapid topological changes, when the content source itself is mobile?

4.5. Wireless Networking

Today, all layer 2 wireless network radio access technologies (L2) are developed with a clear assumption in mind: the waist of the protocol stack is IP and it will be so for the foreseeable future. By fixing the protocol stack waist, engineers can answer a large set of questions, including how to handle conversational traffic (e.g.,

voice calls) vs. web access to online resources, how to support multicast (the IP flavor), and so on, in a rather straightforward manner. Broadcast, on the other hand, which is inherent in wireless communication is not fully taken advantage of. On the contrary, researchers are often more concerned about introducing mechanisms that ensure that "broadcast storms" do not take down a network. The question of how broadcast can serve ICN needs better has yet to be thoroughly investigated.

Wireless networking is often intertwined with mobility but this is not always the case. In fact, empirical measurements often indicate that many users tend to connect (and remain connected) to a single Wi-Fi access point for considerable amounts of time. A case in point, which is frequently cited in different variations in the ICN literature, is access to a document repository during a meeting. For instance, in a typical IETF working group meeting, a scribe takes notes which are uploaded to a centralized repository (see Figure 1). Subsequently, each meeting participant obtains a copy of the document on their own devices for local use, annotation, and sharing with colleagues that are not present at the meeting. Note that in this example there is no node mobility and that it is not important whether the document with the notes is uploaded in one go at the end of the session or in a streaming-like fashion as is typical today with online (cloud-based) document processing.

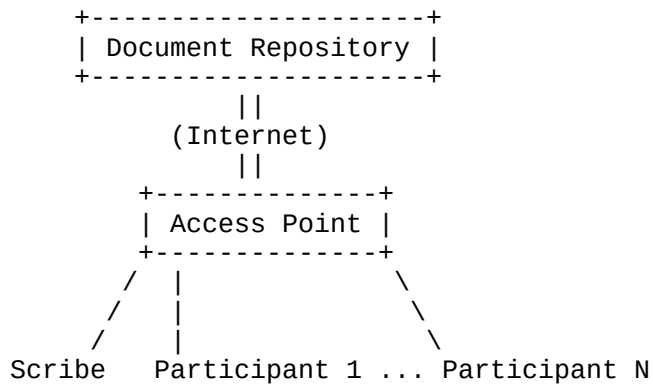

Figure 1: Document sharing during a meeting

In this scenario we observe that the same data object bits (corresponding to the meeting notes) need to traverse the wireless medium at least N+1 times, where N is the number of meeting participants obtaining a copy of the notes. In effect, a broadcast medium is shoehorned into N+1 virtual unicast channels. One could argue that wireless local connectivity is inexpensive, but this is

not the critical factor in this example. The actual information exchange wastes N times the available network capacity, no matter what is the spectral efficiency (or the economics) underlying the wireless technology. This waste is a direct result of extending the remote access paradigm from wired to wireless communication, irrespective of the special characteristics of the latter.

It goes without saying that an ICN approach that does not take into consideration the wireless nature of an interface will waste the same amount of resources as a host-centric paradigm. In-network caching at the wireless access point could reduce the amount of data carried over the backhaul link but, if there is no change in the use of the wireless medium, the NDO will still be carried over the wireless ether N+1 times. Intelligent caching strategies, replica placement cooperation and so on simply cannot alleviate this. On the other hand, promiscuous interface operation and opportunistic caching would maximize wireless network capacity utilization in this example.

Arguably, if one designs a future wireless access technology with an information-centric "layer 3" in mind, many of the design choices that are obvious in an all-IP architecture may no longer be valid. Although this is clearly outside the scope of this document, a few research challenges that the wider community may be interested in include:

- Can we use wireless resources more frugally with the information-centric paradigm than what is possible today in all-IP wireless networks?
- In the context of wireless access, how can we leverage the broadcast nature of the medium in an information-centric network?
- Would a wireless-oriented ICN protocol stack deliver significant performance gains? How different would it be from a wired-oriented ICN protocol stack?
- Is it possible that by changing the network paradigm to ICN we can in practice increase the spectral efficiency (bits/s/Hz) of a wireless network beyond what would be possible with today’s host-centric approaches? What would be the impact of doing so with respect to energy consumption?
- Can wireless interface promiscuous operation coupled with opportunistic caching increase ICN performance, and if so, by how much?
- How can a conversational service be supported at least as efficiently as today’s state-of-the-art wireless networks deliver?

- What are the benefits from combining ICN with network coding in wireless networks?

- How can MIMO and Coordinated Multipoint Transmission (CoMP) be natively combined with ICN primitives in future cellular networks?

4.6. Transport Services

ICN's receiver-driven communication model as described above creates new options for transport protocol design, as it does not rely solely on end-to-end communication from a sender to a receiver. A requested object can be accessible in multiple different network locations. A node can thus decide how to utilize multiple sources, e.g., by sending parallel requests for the same object or by switching sources (or next hops) in a suitable schedule for a series of requests.

In this model, the requestor would control the data rate by regulating its request sending rate and next by performing source/ next-hop selections. Specific challenges depend on the specific ICN approach, but general challenges for receiver-driven transport protocols (or mechanisms, since dedicated protocols might not be required) include flow and congestion control, fairness, network utilization, stability (of data rates under stable conditions) etc. [HRICP] describes a sample request rate control protocol and corresponding design challenges.

As mentioned above, the ICN communication paradigm does not depend strictly on end-to-end flows, as contents might be received from mid-network caches. The traditional concept of a flow is then somewhat cancelled as sub-flows, or flowlets might be formed on the fly, when fractions of an NDO are transmitted from in-network caches. For a transport layer protocol this is challenging, as any measurement related to this flow, as traditionally done by transport protocols such as TCP, will be hugely misleading. For example, false RTT measurements will lead to largely variable average and smooth RTT values, which in turn will trigger false timeout expirations.

Furthermore, out-of-order delivery is expected to be common in a scenario where parts of a content file are retrieved from in-network caches, rather than from the origin server. Several techniques for dealing with out-of-order delivery have been proposed in the past for TCP, some of which could potentially be modified and re-used in the context of ICN. Further research is needed on this direction though to i) choose the right technique and ii) adjust it according to the requirements of the ICN architecture and transport protocol in use.

ICN offers routers the possibility to aggregate requests and can use several paths, meaning that there is no such thing as a (dedicated)

end-to-end communication path, e.g., a router that receives two requests for the same content at the same time only sends one request to its neighbor. The aggregation of requests has a general impact on transport service design.

Achieving fairness for requestors can be one challenge as it is not possible to identify the number of requestors behind one particular request. A second problem related to request aggregation is the management of request retransmissions. Generally, it is assumed that a router will not transmit a request if it transmitted an identical request recently and because there is no information about the requestor, the router cannot distinguish the initial request from a client from a retransmission from the same client. In such a situation, how routers can adapt their timers to use the best of the communication paths. Finally, aggregation of requests has an impact on the server (producer) side. This last has no way to determine the number of clients actually consuming the content it is producing. This shift of model influences the business model of the server, e.g., how to implement pay-per-click.

NDOs can represent content used in different types of applications with different QoS requirements, for example, interactive real-time applications, media streaming, file download. Each of these applications imposes different QoS requirements on different elements in an information-centric network, e.g., regarding cache placement, request-to-cache/source routing. Achieving the necessary quality-of-service levels in a shared network is an active ICN research topic.

4.7. In-Network Caching

Explicitly named data objects allow for caching at virtually any network element, including routers, proxy caches and end-host machines. In-network caching can therefore improve network performance by fetching content from nodes geographically placed closer to the end-user. Several issues that need further investigation have been identified with respect to in-network caching. Here we list some of the most important challenges that relate to the properties of the new ubiquitous caching system.

4.7.1. Cache Placement

The declining cost of fast memory gives the opportunity to deploy caches in network routers and take advantage of explicitly named cached contents. There exist two approaches to in-network caching, namely, on-path and off-path caching. Both approaches have to consider the issue of cache location. Off-path caching is similar to traditional proxy-caching or CDN server placement. Retrieval of contents from off-path caches requires redirection of requests and,

therefore, is closely related to the Request-to-Cache Routing problem discussed later. Off-path caches have to be placed in strategic points within a network in order to reduce the redirection delays and the number of detour hops to retrieve cached contents. Previous research on proxy-caching and CDN deployment is helpful in this case.

On the other hand, on-path caching requires less network intervention and fits more neatly in ICN. However, on-path caching requires line-speed operation, which places more constraints on the design and operation of in-network caching elements. Furthermore, the gain of such a system of on-path in-network caches relies on opportunistic cache hits and has therefore been considered of limited benefit, given the huge amount of contents hosted in the Internet. For this reason, network operators might initially consider only a limited number of network elements to be upgraded to in-network caching elements. The decision on which nodes should be equipped with caches is an open issue and might be based, among others, on topological criteria, or traffic characteristics. These challenges relate to both the Content Placement Problem and the Request-to-Cache Routing Problem discussed below.

In call cases, however, the driver for the implementation, deployment and operation of in-network caches will be its cost. Operating caches at line speed inevitably requires faster memories, which increase the implementation cost. Based on the capital to be invested, ISPs will need to make strategic decisions on the cache placement, which can be driven by several factors, such as: avoid inter-domain/expensive links, centrality of nodes, size of domain and the corresponding spatial locality of users, traffic patterns in a specific part of the network (e.g., university vs. business vs. fashion district of a city).

4.7.2. Content Placement -- Content-to-Cache Distribution

Given a number of (on-path or off-path) in-network caching elements, content-to-cache distribution will affect both the dynamics of the system, in terms of request redirections (mainly in case of off-path caches) and the gain of the system in terms of cache hits. A straightforward approach to content placement is on-path placement of contents as they travel from source to destination. This approach reduces the computation and communication overhead of placing content within the network but, on the other hand, might reduce the chances of hitting cached contents. This relates to the Request-to-Cache Routing problem discussed next.

Furthermore, the number of replicas held in the system brings up resource management issues in terms of cache allocation. For example, continuously replicating data objects in all network

elements results in redundant copies of the same objects. The issue of redundant replication has been investigated in the past for hierarchical web caches. However, in hierarchical web-caching, overlay systems coordination between the data and the control plane can guarantee increased performance in terms of cache hits. Line-speed, on-path in-network caching poses different requirements and therefore, new techniques need to be investigated. In this direction, there already exist some studies that attempt to reduce redundancy of cached copies. However, the issue of coordinated content placement in on-path caches still remains open.

The Content-to-Cache Allocation problem relates also to the characteristics of the content to be cached. Popular content might need to be placed where it is going to be requested next. Furthermore, issues of "expected content popularity" or temporal locality need to be taken into account in designing in-network caching algorithms in order for some contents to be given priority (e.g., popular content vs. one-timers). The criteria as to which contents should be given priority in in-network content caches relate also to the business relationships between content providers and network operators. Business model issues will drive some of these decisions on content-to-cache distribution, but such issues are outside the scope of this note and are not discussed here further.

4.7.3. Request-to-Cache Routing

In order to take advantage of cached contents, requests have to be forwarded to the nodes that temporarily host (cache) the corresponding contents. This challenge relates to name-based routing, discussed before. Requests should ideally follow the path to the cached content. However, instructions as to which content is cached where cannot be broadcast throughout the network. Therefore, the knowledge of a content’s location at the time of the request might either not exist, or it might not be accurate (i.e., contents might have been removed by the time a request is redirected to a specific node).

Coordination between the data and the control planes to update information of cached contents has been considered, but in this case scalability issues arise. We therefore, have two options. We either have to rely on opportunistic caching, where requests are forwarded to a server and in case the content is found on the path, then the content is fetched from this node (instead of the original server); or we employ cache-aware routing techniques. Cache-aware routing can either involve both the control and the data plane, or only one of them. Furthermore, cache-aware routing can be done in a domain-wide scale or can involve more than one individual Autonomous System (AS). In the latter case, business relationships between ASes might need to

be exploited in order to build a scalable model.

4.7.4. Staleness Detection of Cached NDOs

Due to the largely distributed copies of NDOs in in-network caches, ICN should be able to provide a staleness verification algorithm which provides synchronization of NDOs located at their providers and in-network caching points. Two types of approaches can be considered for this problem, namely direct and indirect approaches.

In the direct approach, each cache looks up certain information in the name of NDO, e.g., time stamp which directly indicates its staleness. This approach is well applicable to some NDOs that come from machine-to-machine and Internet of Things scenarios, whose base operation relies on obtaining the latest version of that NDO (i.e., a soil sensor in a farm providing different continuous parameters that are sent to a display or green-house regulation system) [freshness].

In the indirect approach, each cache consults the publisher of the cached NDO about its staleness before serving it. This approach assumes that the name of NDO includes the publisher information which can be used to reach to the publisher. It is suitable for the NDO whose expiring time is difficult to be set in advance, e.g., a webpage which contains main text (that stays the same ever after) and the interactive section such as comments or ads (that is updated irregularly).

It is often argued that ignoring stale NDOs in caches and simply providing new names for updated NDOs might be sufficient rather than using a staleness verification algorithm to manage them. However, notifying the new names of updated NDOs to users is not a trivial task. Unless the update is informed to entire users at the same time, some users would use the old acquainted name by intending to retrieve the updated NDO.

One research challenge is how to design consistency and coherence models for caching NDOs along with their revision handling and updating protocols in a scalable manner.

4.8. Network Management

Managing networks has been a core craft in the IP-based host-centric paradigm ever since the technology was introduced in production networks. However, at the onset of IP, management was considered primarily as an add-on. Essential tools that are used daily by networkers, such as ping and traceroute, did not become widely available until more than a decade or so after IP was first introduced. Management protocols, such as SNMP, also became

available much later than the original introduction of IP and many still consider them insufficient despite the years of experience we have running host-centric networks. Today, when new networks are deployed network management is considered a key aspect for any operator, a major challenge which is directly reflected in higher OPEX if not done well. If we want ICN to be deployed in infrastructure networks, development of management tools and mechanisms must go hand-in-hand with the rest of the architecture design.

Although defining an FCAPS model for ICN is clearly outside the scope of this document, there is a need for creating basic tools early on while ICN is still in the design and experimentation phases that can evolve over time and help network operations centers (NOC) to define policies, validate that they are indeed used in practice, be notified early on about failures, determine and resolve configuration problems. AAA as well as performance management, from a NOC perspective, will also need to be considered. Given the expectations for a large number of nodes and unprecedented traffic volumes, automating tasks, or even better employing self-management mechanisms is preferred. The main challenge here is that all tools we have at our disposal today are node-centric, end-to-end oriented, or assuming a packet-stream communication environment. Rethinking reachability and operational availability, for example, can yield significant insights into how information-centric networks will be managed in the future.

With respect to network management we see three different aspects. First, any operator needs to manage all resources available in the network, which can range from node connectivity to network bandwidth availability to in-network storage to multi-access support. In ICN, users will also bring into the network significant resources in terms of network coverage extension, storage, and processing capabilities. DTN characteristics should also be considered to the degree that this is possible (e.g., content dissemination through data mules). Secondly, given that nodes and links are not at the center of an information-centric network, network management should capitalize on native ICN mechanisms. For example, in-network storage and name resolution can be used for monitoring, while native publish/subscribe functionality can be used for triggering notifications. Finally, management is also at the core of network controlling capabilities by allowing operating actions to be mediated and decided, triggering and activating networking procedures in an optimized way. For example, monitoring aspects can be conjugated with different management actions in a coordinated way, allowing network operations to flow in a concertated way.

However, the considerations on leveraging intrinsic ICN mechanisms

and capabilities to support management operations go beyond a simple mapping exercise. In fact, not only it raises a series of challenges on its own, but also opens up new possibilities for both ICN and "network management" as a concept. For instance, naming mechanisms are central to ICN intrinsic operations, which are used to identify and reach content under different aspects (e.g., hierarchically structured vs. 'flattish' names). In this way, ICN is decoupled from host-centric aspects on which traditional networking management schemes rely upon. As such, questions are raised which can directly be translated into challenges for network management capability, such as, for example how to address a node or a network segment in a ICN naming paradigm, how to identify which node is connected "where", and if there is a host-centric protocol running from which the management process can also leverage upon.

But, on the other hand, these same inherent ICN characteristics also allow us to look into network management through a new perspective. By centering its operations around content, one can conceive new management operations addressing, for example, per-content management or access control, as well as analyzing performance per content name instead of per link or node. Moreover, such considerations can also be used to manage operational aspects of ICN mechanisms themselves. For example, [NDN-MGMT] re-utilizes inherent content-centric capabilities of CCN to manage optimal link connectivity for nodes, in concert with a network optimization process. Conversely, how these content-centric aspects can otherwise influence and impact management in other areas (e.g., security, resilience) is also important, as exemplified by in [ccn-access], where access control mechanisms are integrated into a prototype of the [PURSUIT] architecture.

In this way, a set of core research challenges on ICN management can be derived as:

- Manage and control content reception at the destination
- Coordination of management information exchange and control between ICN nodes and ICN network control points
- Identification of management and controlling actions and items through information naming
- Relationship between NDOs and host entities identification (i.e., how to identify a particular link, interface or flow that need to be managed)

4.9. Application Development

ICN can be applied to different application domains and is expected to provide benefits for application developers by providing a more suitable interface for application developers (in addition to the other ICN benefits described above). [I-D.irtf-icnrg-scenarios] provides an overview of relevant application domains at large. This section discusses opportunities and challenges for selected application types.

4.9.1. Web Applications

Intuitively the ICN request/response communication style seems to be directly mappable to web communication (HTTP). NDO names could be the equivalent of URIs in today's web, proprietary and transparent caching could be obsoleted by ICN in-network caching, and developers could directly use an ICN request/response API to build applications.

Research effort such as [icn2014-web-ndn] have analzed real-world web applications and ways to implement them in ICN. The most significant insight is that, REST-style web communication heavily relies on transmitting user/application context information in HTTP GET requests, which would have to be mapped to corresponding ICN messages. The challenge in ICN would be how to exactly achieve that mapping -- this could be done to some extent by extending name formats or by extending message structure to include cookies and similar context information. The design decisions would need to consider overhead in routers (if larger GET/Interest messages would have to be stored in corresponding tables on routers, for exampe).

Other challenges include the ability to return different results based on client-specific processing in the presence on immutable objects (and name-object bindings) in ICN and the ability for efficient bidirectional communication, which would require some mechanism to name and reach client applications.

4.9.2. Video Streaming and Download

One of ICN's prime application area is video streaming and download where accessing named data, object-level security and in-network storage can fulfill requirements for both video streaming and download. The applicability and benefits of ICN to video has been demonstrated by several prototype developments [icn2014-ahlgren-video-demo].

[I-D.irtf-icnrg-videostreaming] discusses the opportunities and challenges for implementing todays' video services such as DASH-based streaming and download over ICN, considering performance

requirements, relationship to Peer-to-Peer live streaming, IPTV and DRM.

In addition to just porting today's video application from HTTP to ICN there are also promising opportunities to leverage the ICN network services for redesigning and thus significantly enhancing video access and distribution [icnrg-2015-01-westphal]. For example, ICN's store and forward capability could be leveraged for rate adaptation to achieve maximum throughput and optimal QoE in scenarios with varying link properties, if capacity information is fed back to rate selection algorithms at senders. Other optimizations such as more aggressive pre-fetching could be performed in the network by leveraging visibility of chunk NDO names and NDO meta data in the network. Moreover, multi-source rate adaption in combination with network coding could enable better quality of experience, for example in multi-interface/access scenarios where multiple paths from client to upstream caches exist.

4.9.3. Internet of Things

The essence of ICN lies in the name based routing that enables users to retrieve NDOs by the names regardless of their locations. By the definition, ICN is suitable well for IoT applications, where users consume data generated from IoTs without maintaining secure connections to them. The basic put/get style APIs of ICN enable developers to build IoT applications in a simple and fast manner.

On-going efforts such as [I-D.lindgren-icnrg-efficientiot], [I-D.zhang-iot-icn-challenges] have addressed the requirements and challenges of ICN for IoT. For instance, many IoT applications depend on a PUSH model where data transmission is initiated by the publisher, and so they can support various real-time applications: emergency alarm, etc. However, ICN does not support the PUSH model in a native manner due to its inherent receiver-driven data transmission mechanism. The challenge would be how to efficiently support the PUSH model in ICN, and so it provides publish/subscribe style APIs for IoT application developers. This could be done by introducing other types of identifiers such as a device identifier or by extending the current request/response communication style, which may result in heavy overhead in ICN routers.

Moreover, key characteristics of the ICN underlying operation also impact important aspects of IoT, such as the caching in content storage of network forwarding entities (which raise issues, e.g., concerning the freshness of the information received from the cache in contrast to the last value generated by a sensor) as well as pushing content to specific nodes (e.g., for controlling them), which requires individual addressing for identification.

5. Security Considerations

Security related questions related to ICN are discussed in Section 4.2.

6. Informative References

[BACKSCATTER] Waehlisch, M., Schmidt, TC., and M. Vahlenkamp, "Backscatter from the Data Plane - Threats to Stability and Security in Information-Centric Network Infrastructure", Computer Networks Vol 57, No. 16, pp. 3192-3206, November 2013.

[BREADCRUMBS] Rosensweig, E. and J. Kurose, "Breadcrumbs: Efficient, Best-Effort Content Location in Cache Networks", In Proceedings of the IEEE INFOCOM 2009, April 2009.

[CCN] Jacobson, K, D, F, H, and L, "Networking Named Content", CoNEXT 2009 , December 2009.

[COMPACT] Cowen, L., "Compact routing with minimum stretch", In Journal of Algorithms, vol. 38, pp. 170--183, 2001.

[Chaum] Chaum, D. and E. van Heijst, "Group signatures", In Proceedings of EUROCRYPT, 1991.

[DONA] Koponen, T., Ermolinskiy, A., Chawla, M., Kim, K., gon Chun, B., and S. Shenker, "A Data-Oriented (and Beyond) Network Architecture", In Proceedings of SIGCOMM 2007, August 2007.

[GREEDY] Papadopoulos, F., Krioukov, D., Boguna, M., and A. Vahdat, "Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces", In Proceedings of the IEEE INFOCOM, San Diego, USA, 2010.

[HRICP] Carofiglio, G., Gallo, M., and L. Muscariello, "Joint hop-by-hop and receiver-driven interest control protocol for content-centric networks", In Proceedings of ACM SIGCOMM ICN 2012, DOI 10.1145/2342488.2342497, 2012.

[I-D.irtf-icnrg-scenarios] Pentikousis, K., Ohlman, B., Corujo, D., Boggia, G., Tyson, G., Davies, E., Molinaro, A., and S. Eum, "Information-centric Networking: Baseline Scenarios",

draft-irtf-icnrg-scenarios-03 (work in progress), August 2014.

[I-D.irtf-icnrg-videostreaming]
Lederer, S., cedric.westphal@huawei.com, c., Mueller, C., Detti, A., Corujo, D., aytav.azgin, a., Posch, D., and C. Timmerer, "Adaptive Video Streaming over ICN", draft-irtf-icnrg-videostreaming-02 (work in progress), November 2014.

[I-D.lindgren-icnrg-efficientiot]
Lindgren, A., Abdesslem, F., Ahlgren, B., Schelen, O., and A. Malik, "Applicability and Tradeoffs of Information-Centric Networking for Efficient IoT", draft-lindgren-icnrg-efficientiot-02 (work in progress), January 2015.

[I-D.zhang-iot-icn-challenges]
Zhang, Y., Raychadhuri, D., Grieco, L., Baccelli, E., Burke, J., Ravindran, R., and G. Wang, "ICN based Architecture for IoT - Requirements and Challenges", draft-zhang-iot-icn-challenges-01 (work in progress), December 2014.

[ICNNAMING]
Ghodsi, A., Koponen, T., Rajahalme, J., Sarolahti, P., and S. Shenker, "Naming in Content-Oriented Architectures", In Proceedings ACM SIGCOMM Workshop on Information-Centric Networking (ICN), 2011.

[ICNSURVEY]
Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutscher, D., and B. Ohlman, "A Survey of Information-Centric Networking", In Communications Magazine, IEEE , vol.50, no.7, pp.26-36, DOI 10.1109/MCOM.2012.6231276, 2012.

[MANI] Pentikousis, K. and T. Rautio, "A multiaccess Network of Information", WoWMoM 2010, IEEE , June 2010.

[MDHT] D’Ambrosio, M., Dannewitz, C., Karl, H., and V. Vercellone, "MDHT: A hierarchical name resolution service for information-centric networks", ACM SIGCOMM workshop on Information-centric networking Toronto, Canada, 2011, August 2011.

[NDN-MGMT]
Corujo, D., Aguiar, R., Vidal, I., and J. Garcia-Reinoso, "A named data networking flexible framework for management

communications", Communications Magazine, IEEE , vol.50, no.12, pp.36-43 , December 2012.

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

[RANDOM] Gkantsidis, C., Mihail, M., and A. Saberi, "Random walks in peer-to-peer networks: algorithms and evaluation", In Perform. Eval., vol. 63, pp. 241--263, 2006.

[RFC2002] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

[RFC5944] Perkins, C., "IP Mobility Support for IPv4, Revised", RFC 5944, November 2010.

[RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.

[SEEN] Pentikousis, K., "In search of energy-efficient mobile networking", Communications Magazine, IEEE, vol. 48, no. 1, pp.95-103 , January 2010.

[ccn-access] Fotiou, N., Marias, G., and G. Polyzos, "Access control enforcement delegation for information-centric networking architectures", In Proceedings of the second edition of the ICN workshop on Information-centric networking (ICN '12). ACM, New York, NY, USA, 85-90., 2012.

[freshness] Quevedo, J., Corujo, D., and R. Aguiar, "Consumer Driven Information Freshness Approach for Content Centric Networking", IEEE INFOCOM Workshop on Name-Oriented Mobility Toronto, Canada, 2014, May 2014.

[icn2014-ahlgren-video-demo] Ahlgren, B., Jonasson, A., and B. Ohlman, "Demo Overview:

HTTP Live Streaming over NetInf Transport", ACM SIGCOMM Information-Centric Networking Conference Paris, France, 2014, September 2014.

[icn2014-web-ndn]
Moiseenko, I., Stapp, M., and D. Oran, "Communication Patterns for Web Interaction in Named Data Networking", ACM SIGCOMM Information-Centric Networking Conference Paris, France, 2014, September 2014.

[icnrg-2015-01-westphal]
Westphal, C., "Video over ICN", IRTF ICNRG Meeting Cambridge, Massachusetts, USA, 2015, URI http://www.ietf.org/proceedings/interim/2015/01/13/icnrg/slides/slides-interim-2015-icnrg-1-0.pptx, January 2015.

Appendix A. Acknowledgments

The authors would like to thank Georgios Karagiannis for providing suggestions on QoS research challenges and Dimitri Papadimitriou for feedback on the routing section.

Appendix B. Changes

draft-icnrg-challenges-01

* changes to the routing section based on comments by Dimitri Papadimitriou
* Added new section on Application Development (4.9)

draft-icnrg-challenges-00

* added paragraph on QoS to transport services section
* rewrote text 4.7.4 (Staleness Detection of Cached NDOs)

Authors' Addresses

Dirk Kutscher (editor)
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: kutscher@neclab.eu

Suyong Eum
National Institute of Information and Communications Technology
4-2-1, Nukui Kitamachi, Koganei
Tokyo 184-8795
Japan

Phone: +81-42-327-6582
Email: suyong@nict.go.jp

Kostas Pentikousis
EICT GmbH
Torgauer Strasse 12-15
Berlin 10829
Germany

Email: k.pentikousis@eict.de

Ioannis Psaras
University College London, Dept. of E.E. Eng.
Torrington Place
London WC1E 7JE
United Kingdom

Email: i.psaras@ucl.ac.uk

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

Email: dcorujo@av.it.pt

Damien Saucez
INRIA
2004 route des Lucioles - BP 93
Sophia Antipolis 06902 Cedex
France

Email: damien.saucez@inria.fr

Thomas C. Schmidt
HAW HAmburg
Berliner Tor 7
Hamburg 20099
Germany

Email: t.schmidt@ieee.org

Matthias Waehlisch
FU Berlin
Takustr. 9
Berlin 14195
Germany

Email: waehlisch@ieee.org

ICNRG
Internet Draft
Intended status: Informational
Expires: August 31, 2015

S. Lederer
D. Posch
C. Timmerer
Alpen-Adria University Klagenfurt
C. Westphal, Ed.
A. Azgin
S. Liu
Huawei
C. Mueller
Bitmovin
A.Detti
University of Rome Tor Vergata
D. Corujo
University of Aveiro

February 23, 2015

Adaptive Video Streaming over ICN
draft-irtf-icnrg-videostreaming-03.txt

Status of this Memo

Copyright Notice

Abstract

This document considers the consequences of moving the underlying network architecture to an Information-Centric Network (ICN) architecture on video distribution. As most of the traffic in future networks is expected to be video, we consider how to modify the existing video streaming mechanisms. Several important topics related to video distribution over ICN are presented, covering a wide range of scenarios: we look at how to evolve DASH to work over ICN, and leverage the recent ISO/IEC MPEG Dynamic Adaptive Streaming over HTTP (DASH) standard; we consider layered encoding over ICN; P2P mechanisms introduce distinct requirements for video and we look at how to adapt PPSP for ICN; IPTV adds delay constraints, and this will create more stringent requirements over ICN as well. As part of the discussion on video, we discuss DRMs in ICN. Finally, in addition to consider how existing mechanisms would be impacted by

ICN, this document lists some research issues to design ICN specific video streaming mechanisms.

Table of Contents

1. Introduction

The unprecedented growth of video traffic has triggered a rethinking of how content is distributed, both in terms of the underlying Internet architecture and in terms of the streaming mechanisms to deliver video objects.

In particular, the IRTF ICN working group has been chartered to study new architectures centered upon information; the main contributor to Internet traffic (and information dissemination) is video, and this is expected to stay the same in the short- to mid-term future. If ICN is expected to become prominent, it will have to support video streaming efficiently.

As such, it is necessary to discuss along two directions:

- Can the current video streaming mechanisms be leveraged and adapted to an ICN architecture?

- Can (and should) new, ICN-specific video streaming mechanisms be designed to fully take advantage of the new abstractions exposed by the ICN architecture?

This document intends to focus on the first question, in an attempt to define the use cases for video streaming and some requirements.

This document focuses on a few scenarios, namely Netflix-like video streaming, peer-to-peer video sharing and IPTV, and identifies how the existing protocols can be adapted to an ICN architecture. In doing so, it also identifies the main issues with these protocols in this ICN context.

Some documents have started to consider the ICN-specific requirements of dynamic adaptive streaming [2][3][4][6].

In this document, we give a brief overview of the existing solutions for the selected scenarios. We then consider the interactions of such existing mechanisms with the ICN architecture and list some of the interactions any video streaming mechanism will have to consider. We then identify some areas for future research.

2. Conventions used in this document

In examples, "C:" and "S:" indicate lines sent by the client and server respectively.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

In this document, the characters ">>" preceding an indented line(s) indicates a compliance requirement statement using the key words listed above. This convention aids reviewers in quickly identifying or finding the explicit compliance requirements of this RFC.

3. Use case scenarios for ICN and Video Streaming

For ICN specific descriptions, we refer to the other working group documents. For our purpose, we assume here that ICN means an architecture where content is retrieved by name and with no binding of content to a specific network location.

The consumption of multimedia content comes along with timing requirements for the delivery of the content, for both, live and on-demand consumption. Additionally, real-time use cases such as audio-/video conferencing [7], game streaming, etc., come along with more strict timing requirements. Long startup delays, buffering periods or poor quality, etc., should be avoided to achieve a good Quality of Experience (QoE) to the consumer of the content. Of course, these requirements are heavily influenced by routing decisions and caching, which are central parts of ICN and which have to be considered when streaming video in such infrastructures.

Due to this range of requirements, we find it useful to narrow the focus on four scenarios (more can be included later):

- a video delivery architecture similar to that if itune, where the whole file is being downloaded to the client and can be replayed there multiple times;
- a video streaming architecture for playing back movies; this is relevant for the naming and caching aspects of ICN, as well as the interaction with the rate adaptation mechanism necessary to deliver the best QoE to the end-user;
- a peer-to-peer architecture for sharing videos; this introduces more stringent routing requirements in terms of locating copies of the content, as the location of the peers evolves and peers join and leave the swarm they use to exchange video chunks;
- IPTV; this introduces requirements for multicasting and adds stronger delay constraints.

Other scenarios, such as video-conferencing and real-time video communications are not explicitly discussed in this document, while they are in scope. Also, events of mass-media distribution, such as a large crowd in a live event, are also adding new requirements to be included in later version.

We discuss how the current state-of-the-art protocols in an IP context can be modified for the ICN architecture. The remainder of this document is organized as follows. In the next section, we consider video download. Then in Section 5, we briefly describe DASH [1], and Layered Encoding (MDC, SVC). P2P is the focus of Section 6, where we describe PPSP. Section 7 highlights the requirements of IPTV, while Section 8 describes the issues of DRM. Section 9 lists some research issues to be solved for ICN-specific video delivery mechanisms.

This research items include videoconferencing and real-time video communications, which will be detailed more in future versions of this document; as well as the mass distribution of content at live large-scale events (stadium, concert hall, etc) for which there is no clearly adopted existing protocol.

4. Video download

Video download, namely the fetching of a video file from a server or a cache down to the user’s local storage, is a natural application of ICN. It should be supported natively without requiring any specific considerations.

This is supported now by a host of protocols (say, scp, ftp, or over http), which would need to be replaced by the protocols to retrieve content in ICNs.

However, current mechanisms are built atop existing transport protocol. Some ICN proposals (say, CCN or NDN for instance) attempt to leverage the work down upon these transport protocol and it has been proposed to use mechanisms such as the TCP congestion window (and the associated Adaptive Increase, Multiplicative Decrease - AIMD) to decide how many object requests ("interests" in CCN/NDN terminology) should be in flight at any point in time.

It should be noted that ICN intrinsically supports different transport mechanisms, which could achieve better performance than TCP, as they subsume TCP into a special case. For instance, one could imagine a link-by-link transport coupled with caching. This is enabled by the ICN architecture, and would facilitate the point-to-point download of video files.

5. Video streaming and ICN

5.1. Introduction to client-driven streaming and DASH

Media streaming over the hypertext transfer protocol (HTTP) and in a further consequence streaming over the transmission control protocol (TCP) has become omnipresent in today’s Internet. Content providers such as Netflix, Hulu, and Vudu do not deploy their own streaming equipment but use the existing Internet infrastructure as it is and they simply deploy their own services over the top (OTT). This streaming approach works surprisingly well without any particular support from the underlying network due to the use of efficient video compression, content delivery networks (CDNs), and adaptive video players. Earlier video streaming research mostly recommended to use the user datagram protocol (UDP) combined with the real time transport protocol (RTP). It assumed it would not be possible to transfer multimedia data smoothly with TCP, because of its throughput variations and large retransmission delays. This point of view has significantly evolved today. HTTP streaming, and especially its most simple form known as progressive download, has become very popular over the past few years because it has some major benefits compared to RTP streaming. As a consequence of the consistent use of HTTP for this streaming method, the existing Internet infrastructure, consisting of proxies, caches and CDNs, could be used. Originally, this architecture was designed to support best effort delivery of files and not real time transport of multimedia data. Nevertheless, real time streaming based on HTTP could also take advantage of this architecture, in comparison to RTP, which could not leverage any of the aforementioned components. Another benefit that results from the use of HTTP is that the media stream could easily pass firewalls or network address translation (NAT) gateways, which was definitely a key for the success of HTTP

streaming. However, HTTP streaming is not the holy grail of streaming as it also introduces some drawbacks compared to RTP. Nevertheless, in an ICN-based video streaming architecture these aspects also have to be considered.

The basic concept of DASH [1] is to use segments of media content, which can be encoded at different resolutions, bitrates, etc., as so-called representations. These segments are served by conventional HTTP Web servers and can be addressed via HTTP GET requests from the client. As a consequence, the streaming system is pull-based and the entire streaming logic is located on the client, which makes it scalable, and allows to adapt the media stream to the client's capabilities.

In addition to this, the content can be distributed using conventional CDNs and their HTTP infrastructure, which also scales very well. In order to specify the relationship between the contents' media segments and the associated bitrate, resolution, and timeline, the Media Presentation Description (MPD) is used, which is a XML document. The MPD refers to the available media segments using HTTP URLs, which can be used by the client for retrieving them.

5.2. Layered Encoding

Another approach for video streaming consist in using layered encoding. Namely, scalable video coding formats the video stream into different layers: a base layer which can be decoded to provide the lowest bit rate for the specific stream, and enhancement layers which can be transmitted separately if network conditions allow. The higher layers offer higher resolutions and enhancement of the video quality, while the layered approach allows to adapt to the network conditions. This is used in MPEG-4 scalable profile or H.263+. H264SVC is available, but not much deployed. JPEG2000 has a wavelet transform approach for layered encoding, but has not been deployed much either.

It is not clear if the layered approach is fine-grained enough for rate control.

5.3. Interactions of Video Streaming with ICN

5.3.1. Interaction of DASH and ICN

Video streaming, and DASH in particular, have been designed with goals that are aligned with that of most ICN proposals. Namely, it is a client-based mechanism, which requests items (in this case, chunks of a video stream) by name.

ICN and MPEG-DASH [1] have several elements in common:

- the client-initiated pull approach;
- the content being dealt with in pieces (or chunks);
- the support of efficient replication and distribution of content pieces within the network;
- the scalable, session-free nature of the exchange between the client and the server at the streaming layer: the client is free to request any chunk from any location;
- the support for potentially multiple sources.

As ICN is a promising candidate for the Future Internet (FI) architecture, it is useful to investigate its suitability in combination with multimedia streaming standards like MPEG-DASH. In this context, the purpose of this section is to present the usage of ICN instead of HTTP in MPEG-DASH

However, there are some issues that arise from using a dynamic rate adaptation mechanism in an ICN architecture:

o Naming of the data in DASH does not necessarily follow the ICN convention of any of the ICN proposals. Several chunks of the same video stream might currently go by different names that for instance do not share a common prefix. There is a need to harmonize the naming of the chunks in DASH with the naming conventions of the ICN. The naming convention of using a filename/time/encoding format could for instance be made compatible with the convention of CCN.

o While chunks can be retrieved from any server, the rate adaptation mechanism attempts to estimate the available network bandwidth so as to select the proper playback rate and keep its playback buffer at the proper level. Therefore, there is a need to either include some location semantics in the data chunks so as to properly assess the throughput to a specific location; or to design a different mechanism to evaluate the available network bandwidth.

o The typical issue of access control and accounting happens in this context, where chunks can be cached in the network outside of the administrative control of the content publisher. It might be a requirement from the owner of the video stream that access to these data chunks needs to be accounted/billed/monitored.

- Dynamic streaming multiplies the representations of a given video stream, therefore diminishing the effectiveness of caching: namely, to get a hit for a chunk in the cache, it has to be for the same format and encoding values. Alternatively, to get the same hit rate as for a stream using a single encoding, the cache size must be scaled up to include all the possible representations.

- Caching introduces oscillatory dynamics as it may modify the estimation of the available bandwidth between the end user and the repository where it is getting the chunks from. For instance, if an edge cache holds a low resolution representation near the user, the user getting this low resolution chunks will observe a good performance, and will then request higher resolution chunks. If those are hosted on a server with poor performance, then the client would have to switch back to the low representation. This oscillation may be detrimental to the perceived QoE of the user.

- The ICN transport mechanism needs to be compatible to some extent with DASH. To take a CCN example, the rate at which interests are issued should be such that the chunks received in return arrive fast enough and with the proper encoding to keep the playback buffer above some threshold.

- The usage of multiple network interfaces is possible in ICN, enabling a seamless handover between them. For the combination with DASH, an intelligent strategy which should focus on traffic load balancing between the available links may be necessary. This would increase the effective media throughput of DASH by leveraging the combined available bandwidth of all links, however, it could potentially lead to high variations of the media throughput.

- DASH does not define how the MPD is retrieved; hence, this is compatible with CCN. However, the current profiles defined within MPEG-DASH require the MPD to contain HTTP-URLs (incl. http and https URI schemes) to identify segments. To enable a more integrated approach as described in this document, an additional profile for DASH over CCN has to be defined, enabling ICN/CCN-based URIs to identify and request the media segments.

We describe in Section 5.4 a potential implementation of a dynamic adaptive video stream over ICN, based upon DASH and CCN [5].

5.3.2. Interaction of ICN with Layered Encoding

Issues of interest to an Information-Centric network architecture in the context of layered video streaming include:

- Caching of the multiple layers. The caching priority should go to the base layer, and defining caching policy to decide when to cache enhancement layers;
- Synchronization of multiple content streams, as the multiple layers may come from different sources in the network (for instance, the base layer might be cached locally while the enhancement layers may be stored in the origin server);
- Naming of the different layers: when the client requests an object, the request can be satisfied with the base layer alone, aggregated with enhancement layers. Should one request be sufficient to provide different streams? In a CCN architecture for instance, this would violate a one interest-one data packet principle and the client would need to specify each layer it would like to receive. In a Pub/Sub architecture, the rendezvous point would have to make a decision as to which layers (or which pointer to which layer's location) to return.

5.4. Possible Integration of Video streaming and ICN architecture

5.4.1. DASH over CCN

DASH is intended to enable adaptive streaming, i.e., each content piece can be provided in different qualities, formats, languages, etc., to cope with the diversity of todays' networks and devices. As this is an important requirement for Future Internet proposals like CCN, the combination of those two technologies seems to be obvious. Since those two proposals are located at different protocol layers - DASH at the application and CCN at the network layer - they can be combined very efficiently to leverage the advantages of both and potentially eliminate existing disadvantages. As CCN is not based on classical host-to-host connections, it is possible 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 CCN for adaptive multimedia streaming within mobile environments since most mobile devices are equipped with multiple network links like 3G and WiFi. CCN offers this functionality out of the box which is beneficial when used for DASH-based services. In particular, it is possible to enable adaptive video streaming handling both bandwidth and network link changes. That is, CCN handles the network link decision and DASH is implemented on top of CCN to adapt the video stream to the available bandwidth.

In principle, there are two options to integrate DASH and CCN: a proxy service acting as a broker between HTTP and CCN as proposed in [6], and the DASH client implementing a native CCN interface. The former transforms an HTTP request to a corresponding interest packet as well as a data packet back to an HTTP response, including reliable transport as offered by TCP. This may be a good compromise to implement CCN in a managed network and to support legacy devices. As such a proxy is already described in [6] this draft focuses on a more integrated approach, aiming at fully exploiting the potential of a CCN DASH Client. That is, we describe a native CCN interface within the DASH client, which adopts a CCN naming scheme (CCN URIs) to denote segments in the Media Presentation Description (MPD). In this architecture, only the network access component on the client has to be modified and the segment URIs within MPD have to be updated according to the CCN naming scheme.

Initially, the DASH client retrieves the MPD containing the CCN URIs of the content representations including the media segments. The naming scheme of the segments may reflect intrinsic features of CCN like versioning and segmentation support. Such segmentation support is already compulsory for multimedia streaming in CCN and, thus, can also be leveraged for DASH-based streaming over CCN. The CCN versioning can be adopted in a further step to signal different representations of the DASH-based content, which enables an implicit adaptation of the requested content to the clients' bandwidth conditions. That is, the interest packet already provides the desired characteristics of a segment (such as bit rate, resolution, etc.) within the content name (or potentially within parameters defined as extra types in the packet formats). Additionally, if bandwidth conditions of the corresponding interfaces or routing paths allow so, DASH media segments could be aggregated automatically by the CCN nodes, which reduces the amount of interest packets needed to request the content. However, such approaches need further research, specifically in terms of additional intelligence and processing power needed at the CCN nodes.

After requesting the MPD, the DASH client will start to request particular segments. Therefore, CCN interest packets are generated by the CCN access component and forwarded to the available interfaces. Within the CCN, these interest packets leverage the efficient interest aggregation for, e.g., popular content, as well as the implicit multicast support. Finally, the interest packets are satisfied by the corresponding data packets containing the video segment data, which are stored on the origin server or any CCN node, respectively. With an increasing popularity of the content, it will be distributed across the network resulting in lower transmission

delays and reduced bandwidth requirements for origin servers and content providers respectively.

With the extensive usage of in-network caching, new drawbacks are introduced since the streaming logic is located at the client, i.e., clients are not aware of each other and the network infrastructure and cache states. Furthermore, negative effects are introduced when multiple clients are competing for a bottleneck and when caching is influencing this bandwidth competition. As mentioned above, the clients request individual portions of the content based on available bandwidth which is calculated using throughput estimations. This uncontrolled distribution of the content influences the adaptation process of adaptive streaming clients. The impact of this falsified throughput estimation could be tremendous and leads to a wrong adaptation decision which may impact the Quality of Experience (QoE) at the client, as shown in [8]. In ICN, the client does not have the knowledge from which source the requested content is actually served or how many origin servers of the content are available, as this is transparent and depends on the name-based routing. This introduces the challenge that the adaptation logic of the adaptive streaming client is not aware of the event when the ICN routing decides to switch to a different origin server or content is coming through a different link/interface. As most algorithms implementing the adaption logic are using bandwidth measurements and related heuristics, the adaptation decisions are no longer valid when changing origin servers (or links) and potentially cause playback interruptions and, consequently, stalling. Additionally, ICN supports the usage of multiple interfaces and a seamless handover between them, which again comes together with bandwidth changes, e.g., switching between fixed and wireless, 3G/4G and WiFi networks, etc. Considering these characteristics of ICN, adaptation algorithms merely based on bandwidth measurements are not appropriate anymore, as potentially each segment can be transferred from another ICN node or interface, all with different bandwidth condition. Thus, adaptation algorithms taking into account these intrinsic characteristics of ICN are preferred over algorithms based on mere bandwidth measurements.

5.4.2. Testbed, Open Source Tools, and Dataset

For the evaluations of DASH over CCN, a testbed with open source tools and datasets is provided in [9]. In particular, it provides two client player implementations, (i) a libdash extension for DASH over CCN and (ii) a VLC plugin implementing DASH over CCN. For both implementations the CCNx implementation has been used as a basis.

The general architecture of libdash is organized in modules, so that the library implements a MPD parser and an extensible connection manager. The library provides object-oriented interfaces for these modules to access the MPD and the downloadable segments. These components are extended to support DASH over CCN and available in a separate development branch of the github project available at http://www.github.com/bitmovin/libdash. libdash comes together with a fully featured DASH player with a QT-based frontend, demonstrating the usage of libdash and providing a scientific evaluation platform. As an alternative, patches for the DASH plugin of the VLC player are provided. These patches can be applied to the latest source code checkout of VLC resulting in a DASH over CCN-enabled VLC player.

Finally, a DASH over CCN dataset is provided in form of a CCNx repository. It includes 15 different quality representation of the well-known Big Buck Bunny Movie, ranging from 100 kbps up to 4500 kbps. The content is split into segments of two seconds, and described by an associated MPD using the presented naming scheme in Section 4.1. This repository can be downloaded from [9], and is also provided by a public accessible CCNx node. Associated routing commands for the CCNx namespaces of the content are provided via scripts coming together with the dataset and can be used as a public testbed.

6. P2P video distribution and ICN

Another form of distributing content - and video in particular- which ICNs need to support is Peer-to-Peer distribution (P2P). We see now how an existing protocol such as PPSP can be modified to work in an ICN environment.

6.1. Introduction to PPSP

P2P video Streaming (PPS) is a popular approach to redistribute live media over Internet. The proposed P2PVS solutions can be roughly classified in two classes:

- Push/Tree based

- Pull/Mesh based

The Push/Tree based solution creates an overlay network among peers that has a tree shape. Using a progressive encoding (e.g. Multiple Description Coding or H.264 Scalable Video Coding), multiple trees could be set up to support video rate adaptation. On each tree an enhancement stream is sent. The more the number of stream received,

the higher the video quality. A peer control video rate by fetching or not the streams delivered on the distribution trees.

The Pull/Mesh based solution is inspired by the BitTorrent file sharing mechanism. A Tracker collects information about the state of the swarm (i.e. set of participating peers). A peer forms a mesh overlay network with a subset of peers, and exchange data with them. A peer announces what data items it disposes and requests missing data items that are announced by connected peers. In case of live streaming, the involved data set includes only a recent window of data items published by the source. Also in this case, the use of a progressive encoding can be exploited for video rate adaptation.

Pull/Mesh based P2PVS solutions are the more promising candidate for the ICN deployment, since most of ICN approach provides a pull-based API [5][10][11][12]. In addition, Pull/Mesh based P2PVS are more robust than Push/Tree based one [13] and the Peer to Peer Streaming Protocol (PPSP) working group [14] is also proposing a Pull/Mesh based solution.

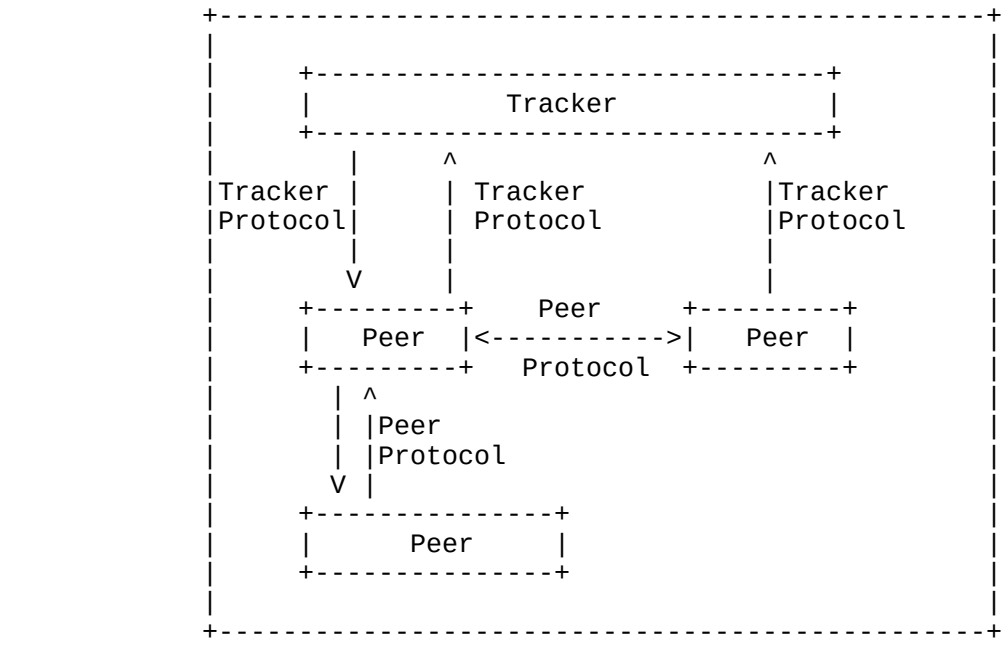

Figure 1: PPSP System Architecture (source [RFC6972])

Figure 1 reports the PPSP architecture presented in [RFC6972]. PEERs announce and share video chunks and a TRACKER maintains a list of PEERs participating in a specific audio/video channel or in the distribution of a streaming file. The tracker functionality may be centralized in a server or distributed over the PEERs. PPSP standardize the Peer and Tracker Protocols, which can run directly over UDP or TCP.

This document discusses some preliminary concepts about the deployment of PPSP on top of an ICN that exposes a pull-based API, meanwhile considering the impact of MPEG DASH streaming format.

6.2. PPSP over ICN: deployment concepts

6.2.1. PPSP short background

PPSP specifies peer protocol (PPSPP) [15] and tracker protocol (PPSP-TP)[16].

Some of the operations carried out by the tracker protocol are the followings. When a peer wishes to join the streaming session it contacts the Tracker (CONNECT message), obtains a PEER_ID and a list of PEER_IDs (and IP addresses) of other peers that are participating to the SWARM and that the tracker has singled out for the requesting peer (this may be a subset of the all peers of the SWARM). In addition to this join operation, a peer may contact the tracker to request to renew the list of participating peers (FIND message), to periodically update its status to the tracker (STAT_REPORT message), etc.

Some of the operations carried out by the peer protocol are the following. Using the list of peers delivered by the tracker, a peer establishes a session with them (HANDSHAKE message). A peer periodically announces to neighboring peers which chunks it has available for download (HAVE message). Using these announcements, a peer requests missing chunks from neighboring peers (REQUEST messages), which will send back them (DATA message).

6.2.2. From PPSP messages to ICN named-data

An ICN provides users with data items exposed by names. The bundle name and data item is usually referred as named-data, named-content, etc. To transfer PPSP messages though an ICN the messages should be be wrapped as named-data items, and receivers should request them by name.

A PPSP entity receives messages from peers and/or tracker. Some operations require gathering the messages generated by another specific host (peer or tracker). For instance, if a peer A wishes to gain information about video chunks available from peer B, the former shall fetch the PPSP HAVE messages specifically generated by the later. We refer to these kinds of named-data as "located-named-data", since they should be gathered from a specific location (e.g. peer B).

For other PPSP operations, like to fetch a DATA message (i.e. a video chunk), what it is relevant for a peer is just to receive the requested content, independently from who is the endpoint that generate the data. We refer this information with the generic term "named-data".

The naming scheme differentiates named-data and located-named-data items. In case of named-data, the naming scheme only includes a content identifier (e.g. the name of the video chunk), without any prefix identifying who provides the content. For instance, a DATA message containing the video chunk n. 1 may be named as "ccnx:/swarmID/chunk/chunkID", where swarmID is a unique identifier of the streaming session, "chunk" is a keyword and chunkID is the chunk identifier (e.g. a integer number).

In case of located-named-data, the naming scheme includes a location-prefix, which uniquely identifies the host generating the data item. This prefix may be the PEER_ID in case the host was a peer or a tracker identifier in case the host was the tracker. For instance, a HAVE message generated by a peer B may be named as "ccnx:/swarmID/peer/PEER_ID/HAVE", where "peer" is a keyword, PEER_ID_B is the identifier of peer B and HAVE is a keyword.

6.2.3. Support of PPSP interaction through a pull-based ICN API

The PPSP procedures are based both on pull and push interactions. For instance, the distribution of chunks availability can be classified as a push-based operation, since a peer sends an "unsolicited" information (HAVE message) to neighboring peers. Conversely the procedure used to receive video chunks can be classified as pull-based, since it is supported by a request/response interaction (i.e. REQUEST, DATA messages).

As we said, we refer to an ICN architecture which provides a pull-based API. Accordingly, the mapping of PPSP pull-based procedure is quite simple. For instance, using the CCN architecture [5] a PPSP

DATA message may be carried by a CCN Data message and a REQUEST message can transferred by a CCN Interest.

Conversely, the support of push-based PPSP operations may be more difficult. We need of an adaptation functionality that carries out a push-based operation using the underlying pull-based service primitives. For instance, a possible approach is to use the request/response (i.e. Interest/Data) four ways handshakes proposed in [7]. Another possibility is that receivers periodically send out request messages of the named-data that neighbors will push and, when available, sender inserts the pushed data within a response message.

6.2.4. Abstract layering for PPSP over ICN

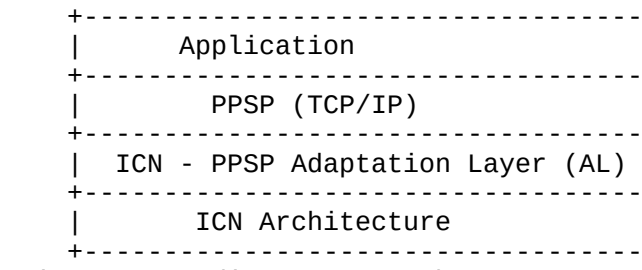

Figure 2: Mediator approach

Figure 2 provides a possible abstract layering for PPSP over ICN. The Adaptation Layer acts as a mediator (proxy) between legacy PPSP entities based on TCP/IP and the ICN architecture. In facts, the role the mediator is to use ICN to transfer PPSP legacy messages.

This approach makes possible to merely reuse TCP/IP P2P applications whose software includes also PPSP functionality. This "all-in-one" development approach may be rather common since the PPSP-Application interface is not going to be specified. Moreover, if the Operating System will provide libraries that expose a PPSP API, these will be initially based on a underlying TCP/IP API. Also in this case, the mediator approach would make possible to easily reuse both the PPSP libraries and the Application on top of an ICN.

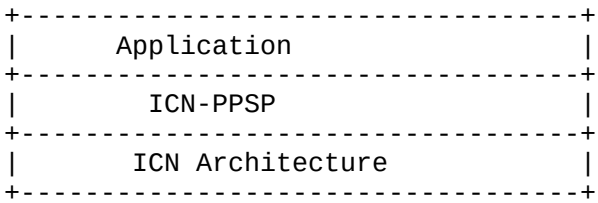

Figure 3: Clean-slate approach

Figure 3 sketches a clean-slate layering approach in which the application directly includes or interacts with a PPSP version based on ICN. Likely such a PPSP_ICN integration could yield a simplier development, also because it does not require implementing a TCP/IP to ICN translation as in the Mediator approach. However, the clean-slate approach requires developing the application (in case of embedded PPSP functionality) or the PPSP library from scratch, without exploiting what might already exist for TCP/IP.

Overall, the Mediator approach may be considered as the first step of a migration path towards ICN native PPSP applications.

6.2.5. PPSP interaction with the ICN routing plane

Upon the ICN API a user (peer) requests a content and the ICN sends it back. The content is gathered by the ICN from any source, which could be the closest peer that disposes of the named-data item, an in-network cache, etc. Actually, "where" to gather the content is controlled by an underlying ICN routing plane, which sets up the ICN forwarding tables (e.g. CCN FIB [5]).

A cross-layer interaction between the ICN routing plane and the PPSP may be required to support a PPSP session. Indeed, ICN shall forward request messages (e.g. CCN Interest) towards the proper peer that can handle them. Depending on the layering approach, this cross-layer interaction is controlled either by the Adaptation Layer or by the ICN-PPSP. For example, if a peer A receives a HAVE message indicating that peer B disposes of the video chunk named "ccnx:/swarmID/chunk/chunkID", then former should insert in its ICN forwarding table an entry for the prefix "ccnx:/swarmID/chunk/chunkID" whose next hop locator (e.g. IP address) is the network address of peer B [17].

6.2.6. ICN deployment for PPSP

The ICN functionality that supports a PPSP session may be "isolated" or "integrated" with the one of a public ICN.

In the isolated case, a PPSP session is supported by an instance of an ICN (e.g. deployed on top of IP), whose functionalities operate only on the limited set of nodes participating to the swarm, i.e. peers and the tracker. This approach resembles the one followed by current P2P application, which usually form an overlay network among peers of a P2P application. And intermediate public IP routers do not carry out P2P functionalities.

In the integrated case, the nodes of a public ICN may be involved in the forwarding and in-network caching procedures. In doing so, the swarm may benefit from the presence of in-network caches so limiting uplink traffic on peers and inter-domain traffic too. These are distinctive advantages of using PPSP over a public ICN, rather than over TCP/IP. In addition, such advantages aren't likely manifested in the case of isolated deployment.

However, the possible interaction between the PPSP and the routing layer of a public ICN may be dramatic, both in terms of explosion of the forwarding tables and in terms of security. These issues specifically take place for those ICN architectures for which the name resolution (i.e. name to next-hop) occurs en-route, like the CCN architecture.

For instance, using the CCN architecture, to fetch a named-data item offered by a peer A the on-path public ICN entities have to route the request messages towards the peer A. This implies that the ICN forwarding tables of public ICN nodes may contain many entries, e.g. one entry per video chunk, and these entries are difficult to be aggregated since peers avail sparse parts of a big content, whose names have a same prefix (e.g. "ccnx:/swarmID"). Another possibility is to wrap all PPSP messages into a located-named-data. In this case the forwarding tables should contain "only" the PEER_ID prefixes (e.g. "ccnx:/swarmID/peer/PEER_ID"), so scaling down the number of entries from number of chunks to number of peers. However, in this case the ICN mechanisms recognize a same video chunk offered by different peers as different contents, so vanishing caching and multicasting ICN benefits. Moreover, in any case routing entries should be updated either the base of the availability of named-data items on peers or on the presence of peers, and these events in a P2P session is rapidly changing so possibly hampering the convergence of the routing plane. Finally, since peers have an impact on the ICN forwarding table of public nodes, this may open obvious security issues.

6.3. Impact of MPEG DASH coding schemes

The introduction of video rate adaptation may valuably decrease the effectiveness of P2P cooperation and of in-network caching, depending of the kind of the video coding used by the MPEG DASH stream.

In case of a MPEG DASH streaming with MPEG AVC encoding, a same video chunk is independently encoded at different rates and the encoding output is a different file for each rate. For instance, in case of a video encoded at three different rates R1,R2,R3, for each

segment S we have three distinct files: S.R1, S.R2, S.R3. These files are independent of each other. To fetch a segment coded at R2 kbps, a peer shall request the specific file S.R2. The estimation of the best coding rate is usually handled by receiver-driven algorithms, implemented by the video client.

The independence among files associated to different encoding rates and the heterogeneity of peer bandwidths, may dramatically reduce the interaction among peers, the effectiveness of in-network caching (in case of integrated deployment), and consequently the ability of PPSP to offload the video server (i.e. a seeder peer). Indeed, a peer A may select a coding rate (e.g. R1) different from the one selected by a peer B (e.g. R2) and this prevents the former to fetch video chunks from the later, since peer B avails of chunks coded at a rate different from the ones needed by A. To overcome this issue, a common distributed rate selection algorithm could force peers to select the same coding rate [17]; nevertheless this approach may be not feasible in the in case of many peers.

The use of SVC encoding (Annex G extension of the H.264/MPEG-4 AVC video compression standard) should make rate adaptation possible, meanwhile neither reducing peer collaborations nor the in-network caching effectiveness. For a single video chunk, a SVC encoder produces different files for the different rates (roughly "layers"), and these files are progressively related each other. Starting from a base-layer which provides the minimum rate encoding, the next rates are encoded as an "enhancement layer" of the previous one. For instance, in case the video is coded with three rates R1 (base-layer), R2 (enhancement-layer n.1), R3 (enhancement-layer n.2), then for each DASH segment we have three files S.R1, S.R2 and S.R3. The file S.R1 is the segment coded at the minimum rate (base-layer). The file S.R2 enhances S.R1, so as S.R1 and S.R2 can be combined to obtain a segment coded at rate R2. To get a segment coded at rate R2, a peer shall fetch both S.R1 and S.R2. This progressive dependence among files that encode a same segment at different rates makes peer cooperation possible, also in case peers player have autonomously selected different coding rates. For instance, if peer A has selected the rate R1, the downloaded files S.R1 are useful also for a peer B that has selected the rate R2, and vice versa.

7. IPTV and ICN

7.1. IPTV challenges

IPTV refers to the delivery of quality content broadcast over the Internet, and is typically associated with strict quality

requirements, i.e., with a perceived latency of less than 500 ms and a packet loss rate that is multiple orders lower than the current loss rates experienced in the most commonly used access networks. We can summarize the major challenges for the delivery of IPTV service as follows.

Channel change latency represents a major concern for the IPTV service. Perceived latency during channel change should be less than 500ms. To achieve this objective over the IP infrastructure, we have multiple choices:

(i) receiving fast unicast streams from a dedicated server (most effective but not resource efficient);
(ii) connecting to other peers in the network (efficiency depends on peer support, effective and resource efficient, if also supported with a dedicated server);
(iii) connecting to multiple multicast sessions at once (effective but not resource efficient, and depends on the accuracy of the prediction model used to track user activity).

The second major challenge is the error recovery. Typical IPTV service requirements dictate the mean time between artifacts to be approximately 2 hours. This suggests the perceived loss rate to be around or less than 10^-7. Current IP-based solutions rely on the following proactive and reactive recovery techniques: (i) joining the FEC multicast stream corresponding to the perceived packet loss rate (not efficient as the recovery strength is chosen based on worst-case loss scenarios), (ii) making unicast recovery requests to dedicated servers (requires active support from the service provider), (iii) probing peers to acquire repair packets (finding matching peers and enabling their cooperation is another challenge).

7.2. ICN benefits for IPTV delivery

ICN presents significant advantages for the delivery of IPTV traffic. For instance, ICN inherently supports multicast and allows for quick recovery from packet losses (with the help of in-network caching). Similarly, peer support is also provided in the shape of in-network caches that typically act as the middleman between two peers, enabling therefore earlier access to IPTV content.

However, despite these advantages, delivery of IPTV service over Information Centric Networks brings forth new challenges. We can list some of these challenges as follows:

. Messaging overhead: ICN is a pull-based architecture and relies on a unique balance between requests and responses. A user needs to make a request for each data packet. In the case of IPTV, with rates up to, and likely to be, above 15Mbps, we observe significant traffic upstream to bring those streams. As the number of streams increase (including the same session at different quality levels), so as the burden on the routers. Even if the majority of requests are aggregated at the core, routers close to the edge (where we observe the biggest divergence in user requests) will experience a significant increase in overhead to process these requests. The same is true at the user side, as the uplink usage multiplies in the number of sessions a user requests (for instance, to minimize the impact of bandwidth fluctuations).
. Cache control: As the IPTV content expires at a rapid rate (with a likely expiry threshold of 1s), we need solutions to effectively flush out such content to also prevent degradatory impact on other cached content, with the help of intelligently chosen naming conventions. However, to allow for fast recovery and optimize access time to sessions (from current or new users), the timing of such expirations needs to be adaptive to network load and user demand. However, we also need to support quick access to earlier content, whenever needed, for instance, when the user accesses the rewind feature (note that in-network caches will not be of significant help in such scenarios due to overhead required to maintain such content).
. Access accuracy: To receive the up-to-date session data, users need to be aware of such information at the time of their request. Unlike IP multicast, since the users join a session indirectly, session information is critical to minimize buffering delays and reduce the startup latency. Without such information, and without any active cooperation from the intermediate routers, stale data can seriously undermine the efficiency of content delivery. Furthermore, finding a cache does not necessarily equate to joining a session, as the look-ahead latency for the initial content access point may have a shorter lifetime than originally intended. For instance, if the

user that has initiated the indirect multicast leaves the session early, the requests from the remaining users need to experience an additional latency of one RTT as they travel towards the content source. If the startup latency is chosen depending on the closeness to the intermediate router, going to the content source in-session can lead to undesired pauses.

8. Digital Rights Managements in ICN

This section discusses the need for Digital Rights Management (DRM) functionalities for multimedia streaming over ICN. It focuses on two possible approaches: modifying AAA to support DRM in ICN, and using Broadcast Encryption.

It is assumed that ICN will be used heavily for digital content dissemination. It is vital to consider DRM for digital content distribution. In today’s Internet there are two predominant classes of business models for on-demand video streaming. The first model is based on advertising revenues. Non-copyright protected (usually user-generated content, UGC) is offered by large infrastructure providers like Google (YouTube) at no charge. The infrastructure is financed by spliced advertisements into the content. In this context DRM considerations may not be required, since producers of UGC may only strive for the maximum possible dissemination. Some producers of UGC are mainly interested to share content with their families, friends, colleges or others and have no intention to make profit. However, the second class of business models requires DRM, because they are primarily profit oriented. For example, large on-demand streaming platforms like Netflix establish business models based on subscriptions. Consumers may have to pay a monthly fee in order to get access to copyright protected content like TV series, movies or music. This model may be ad-supported and free to the content consumer, like YouTube Channels or Spotify. But the creater of the content expects some remuneration for his work. From the perspective of the service providers and the copyright owners, only clients that pay the fee (explicitly or implicitly through ad placement) should be able to access and consume the content. Anyway, the challenge is to find an efficient and scalable way of access control to digital content, which is distributed in information-centric networks.

8.1. Broadcast Encryption for DRM in ICN

The section discusses Broadcast Encryption (BE) as a suitable basis for DRM functionalities in conformance to the ICN communication paradigm. Especially when network inherent caching is considered the advantage of BE will be highlighted.

In ICN, data packets can be cached inherently in the network and any network participant can request a copy of these packets. This makes it very difficult to implement an access control for content that is distributed via ICN. A naive approach is to encrypt the transmitted data for each consumer with a distinct key. This prohibits everyone other than the intended consumers to decrypt and consume the data. However, this approach is not suitable for ICN's communication paradigm since it would reduce the benefits gained from the inherent network caching. Even if multiple consumers request the same content the requested data for each consumer would differ using this approach. A better but still insufficient idea is to use a single key for all consumers. This does not destruct the benefits of ICN's caching ability. The drawback is that if one of the consumers illegally distributes the key, the system is broken and any entity in the network can access the data. Changing the key after such an event is useless since the provider has no possibility to identify the illegal distributer. Therefore this person cannot be stopped from distributing the new key again. In addition to this issue other challenges have to be considered. Subscriptions expire after a certain time and then it has to be ensured that these consumers cannot access the content anymore. For a provider that serves millions of daily consumers (e.g. Netflix) there could be a significant number of expiring subscriptions per day. Publishing a new key every time a subscription expires would require an unsuitable amount of computational power just to re-encrypt the collection of audio-visual content.

A possible approach to solve these challenges is Broadcast Encryption (BE) [22] as proposed in [23]. From this point on, this section will focus only on BE as an enabler for DRM functionality in the use case of ICN video streaming. This subsection continues with the explanation of how BE works and shows how BE can be used to implement an access control scheme in the context of content distribution in ICN.

BE actually carries a misleading name. One might expect a concrete encryption scheme. However, it belongs to the family of key-management schemes (KMS). KMS are responsible for the generation, exchange, storage and replacement of cryptographic keys. The most

interesting characteristics of Broadcast Encryption Schemes (BES) are:

- A BES typically uses a global trusted entity called the licensing agent (LA), which is responsible for spreading a set of pre-generated secrets among all participants. Each participant gets a distinct subset of secrets assigned from the LA.
- The participants can agree on a common session key, which is chosen by the LA. The LA broadcasts an encrypted message that includes the key. Participants with a valid set of secrets can derive the session-key from this message.
- The number of participants in the system can change dynamically. Entities may join or leave the communication group at any time. If a new entity joins the LA passes on a valid set of secrets to that entity. If an entity leaves (or is forced to leave) the LA revokes the entity's subset of keys, which means that it cannot derive the correct session key anymore when the LA distributes a new key.
- -Traitors (entities that reveal their secrets) can be traced and excluded from ongoing communication. The algorithms and preconditions to identify a traitor vary between concrete BES.

This listing already illustrates why BE is suitable to control the access to data that is distributed via an information-centric network. BE enables the usage of a single session key for confidential data transmission between a dynamically changing subset or network participants. ICN caches can be utilized since the data is encrypted only with a single key known by all legitimate clients. Furthermore, traitors can be identified and removed from the system. The issue of re-encryption still exists, because the LA will eventually update the session key when a participant should be excluded. However, this disadvantage can be relaxed in some way if the following points are considered:

- The updates of the session key can be delayed until a set of compromised secretes has been gathered. Note that secrets may become compromised because of two reasons. First, if the secret has been illegally revealed by a traitor. Second, if the subscription of an entity expires. Delayed revocation temporarily enables some non-legitimate entities to consume content. However, this should not be a severe problem in home entertainment scenarios. Updating the session key in regular (not too short) intervals is a good tradeoff. The longer the interval last the less computational resources are required for content re-encryption and the better the cache utilization in the ICN will be. To evict old data from ICN caches that has

been encrypted with the prior session key the publisher could indicate a lifetime for transmitted packets.

. Content should be re-encrypted dynamically at request time. This has the benefit that untapped content is not re-encrypted if the content is not requested during two session key updates and therefore no resources are wasted. Furthermore, if the updates are triggered in non-peak times the maximum amount of resource needed at one point in time can be lowered effectively, since in peak times generally more diverse content is requested.

. Since the amount of required computational resources may vary strongly from time to time it would be beneficial for any streaming provider to use cloud-based services to be able to dynamically adapt the required resources to the current needs. Regarding to a lack of computation time or bandwidth the cloud service could be used to scale up to overcome shortages.

Figure 4 show the potential usage of BE in a multimedia delivery frameworks that builds upon ICN infrastructure and uses the concept of dynamic adaptive streaming, e.g., DASH. BE would be implemented on the top to have an efficient and scalable way of access control to the multimedia content.

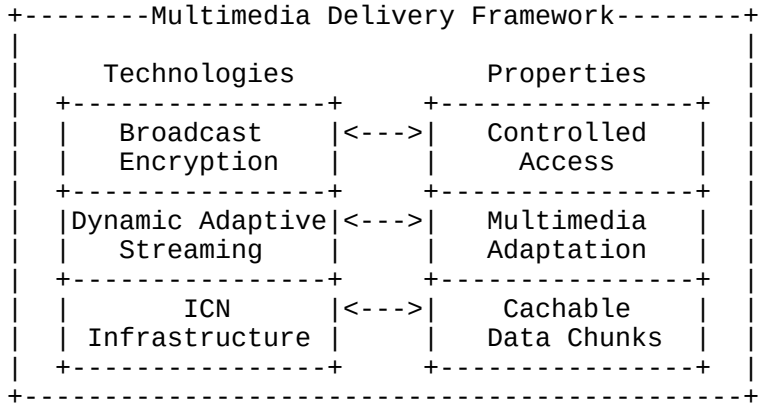

Figure 4: A potential multimedia framework using BE.

8.2 . AAA Based DRM for ICN Networks

8.2.1. Overview

Recently, a novel approach to Digital Rights Management (DRM) has emerged to link DRM to usual network management operations, hence linking DRM to authentication, authorization, and accounting (AAA) services. ICN provides the abstraction of an architecture where content is requested by name and could be served from anywhere. In DRM, the content provider (the origin of the content) allows the destination (the end user account) to use the content. The content provider and content storage/cache are at two different entities in ICC and for traditional DRM only source and destination count and not the intermediate storage. The proposed solution allows the provider of the caching to be involved in the DRM policies using well known AAA mechanisms. It is important to note that this solution is compatible with the proposes the Broadcast Encryption (BE) proposed earlier in this draft. The BE proposes a technology as this solution is more operational.

8.2.2. Implementation

With the proposed AAA-based DRM, when a content is requested by name from a specific destination, the request could link back to both the content provider and the cahing provider via traditional AAA mechanisms, and trigger the appropriate DRM policy independently from where the content is stored. In this approach the caching, DRM and AAA remain independent entities but can work together through ICN mechanisms. The proposed solution enables extending the traditional DRM done by the content provider to jointly being done by content provider and network/caching provider.

The solution is based on the concept of a "token". The content provider authenticates the end user and issues an encrypted token to authenticate the a named content ID or IDs that the user can access. The token will be shared with the network provider and used as the interface to the AAA protocols. At this point all content access is under the control of the network provider and the ICN. The controllers and switches can manage the content requests and handle mobility. The content can be accessed from anywhere as long as the token remains valid or the content is available in the network. In such a scheme the content provider does not need to be contacted every time a named content is requested. This reduces the load of the content provider network and creates a DRM mechanism that is much more appropriate for the distributed caching and peer-to-peer

storage characteristic of ICN networks. In particular, the content requested by name can be served from anywhere under the only condition that the storage/cache can verify that the token is valid for content access.

The solution is also fully customizable to both content and network provider's needs as the tokens can be issued based on user accounts, location and hardware (MAC address for example) linking it naturally to legacy authentication mechanisms. In addition, since both content and network providers are involved in DRM policies pollution attacks and other illegal requests for the content can be more easily detected. The proposed AAA-based DRM is currently under full development.

9. Future Steps for Video in ICN

The explosion of online video services, along with their increased consumption by mobile wireless terminals, further exacerbates the challenges of Video Adaptation leveraging ICN mechanisms. The following sections present a series of research items derived from these challenges, further introducing next steps for the subject.

9.1. Large Scale Live Events

An active area of investigation and a potential use case where ICN would provide significant benefits, is that of distributing content, and video in particular, using local communications in large scale events such as sports event in a stadium, a concert or a large demonstration.

Such use-case involves locating content that is generated on the fly and requires discovery mechanisms in addition to sharing mechanisms. The scalability of the distribution becomes important as well.

9.2. Video Conferencing and Real-Time Communications

Current protocols for video-conferencing have been designed, and this document needs to take input from them to identify the key research issues. Real-time communication add timing constraints (both in terms of delay and in terms of synchronization) to the scenario discussed above.

9.3. Store-and-Forward Optimized Rate Adaptation

One of the benefits of ICN is to allow the network to insert caching in the middle of the data transfer. This can be used to reduce the overall bandwidth demands over the network by caching content for

future re-use. But it provides more opportunities for optimizing video streams.

Consider for instance the following scenario: a client is connected via an ICN network to a server. Let’s say the client is connected wirelessly to a node that has a caching capability, which is connected through a WAN to the server. Assume further that the capacity of each of the links (both the wireless and the WAN logical links) vary with time.

If the rate adaptation is provided in an end-to-end manner, as in current mechanisms like DASH, then the maximal rate that can be supported at the client is that of the minimal bandwidth on each link.

For instance, if during time period 1, the wireless capacity is 1 and the wired capacity is 2, and during time period 2, the wireless is 2 due to some hotspot, and the wired is 1 due to some congestion in the network, then the best end-to-end rate that can be achieved is 1 during each period.

However, if the cache is used during time period 1 to pre-fetch 2 units of data, then during period 2, there is 1 unit of data at the cache, and another unit of data, which can be streamed from the server, and the rate that can be achieved is therefore 2 units of data. In this case, the average bandwidth rises from 1 to 1.5 over the 2 periods.

This straw man example illustrate a) the benefit of ICN for increasing the throughput of the network, and b) the need for the special rate adaptation mechanisms to be designed so as to take advantage of this gain. End-to-end rate adaptation can not take advantage of the cache availability.

9.4. Heterogeneous Wireless Environment Dynamics

With the ever-growing increase in online services being accessed by mobile devices, operators have been deploying different overlapping wireless access networking technologies. In this way, in the same area, user terminals are within range of different cellular, Wi-Fi or even WiMAX networks. Moreover, with the advent of the Internet of Things (e.g., surveillance cameras feeding video footage), this list can be further complemented with more specific short-range technologies, such as Bluetooth or ZigBee.

In order to leverage from this plethora of connectivity opportunities, user terminals are coming equipped with different

wireless access interfaces, providing them with extended connectivity opportunities. In this way, such devices become able to select the type of access which best suits them according to different criteria, such as available bandwidth, battery consumption, access do different link conditions according to the user profile or even access to different content. Ultimately, these aspects contribute to the Quality of Experience perceived by the end-user, which is of utmost importance when it comes to video content.

However, the fact that these users are mobile and using wireless technologies, also provides a very dynamic setting, where the current optimal link conditions at a specific moment might not last or be maintained while the user moves. These aspects have been amply analyzed in recently finished projects such as FP7 MEDIEVAL [18], where link events reporting on wireless conditions and available alternative connection points were combined with vide requirements and traffic optimization mechanisms, towards the production of a joint network and mobile terminal mobility management decision. Concretely, in [19] link information about the deterioration of the wireless signal was sent towards a mobility management controller in the network. This input was combined with information about the user profile, as well as of the current video service requirements, and used to trigger the decrease or increase of scalable video layers, adjusting the video to the ongoing link conditions. Incrementally, the video could also be adjusted when a new better connectivity opportunity presents itself.

In this way, regarding Video Adaptation, ICN mechanisms can leverage from their intrinsic multiple source support capability and go beyond the monitoring of the status of the current link, thus exploiting the availability of different connectivity possibilities (e.g., different "interfaces"). Moreover, information obtained from the mobile terminal’s point of view of its network link, as well as information from the network itself (i.e., load, policies, and others), can generate scenarios where such information is combined in a joint optimization procedure allowing the content to be forward to users using the best available connectivity option (e.g., exploiting management capabilities supported by ICN intrinsic mechanisms as in [20]).

In fact, ICN base mechanisms can further be exploited in enabling new deployment scenarios such as preparing the network for mass requests from users attending a large multimedia event (i.e., concert, sports), allowing video to be adapted according to content, user and network requirements and operation capabilities in a dynamic way.

The enablement of such scenarios require further research, with the main points highlighted as follows:

. Development of a generic video services (and obviously content) interface allowing the definition and mapping of their requirements (and characteristics) into the current capabilities of the network;

. How to define a scalable mechanism allowing either the video application at the terminal, or some kind of network management entity, to adapt the video content in a dynamic way;

. How to develop the previous research items using intrinsic ICN mechanisms (i.e., naming and strategy layers);

. Leverage intelligent pre-caching of content to prevent stalls and poor quality phases, which lead to bad Quality of Experience of the user. This includes in particular the usage in mobile environments, which are characterized by severe bandwidth changes as well as connection outages, as shown in [21].

9.5. Network Coding for Video Distribution in ICN

An interesting research area for combining heterogeneous sources is to use network coding [24]. Network coding allows to asynchronously combine multiple sources by having each of them send information that is not duplicated by the other but can be combined to retrieve the video stream.

However, this creates issues in ICN in terms of defining the proper rate adaptation for the video stream; securing the encoded data; caching the encoded data; timeliness of the encoded data; overhead of the network coding operations both in network resources and in added buffering delay, etc.

10. Security Considerations

This is informational. Security considerations are TBD.

11. IANA Considerations

This is informational. IANA considerations are TBD.

12. Conclusions

This draft proposed adaptive video streaming for ICN, identified potential problems and presented the combination of CCN with DASH as

a solution. As both concepts, DASH and CCN, maintain several elements in common, like, e.g., the content in different versions being dealt with in segments, combination of both technologies seems useful. Thus, adaptive streaming over CCN can leverage advantages such as, e.g., efficient caching and intrinsic multicast support of CCN, routing based on named data URIs, intrinsic multi-link and multi-source support, etc.

In this context, the usage of CCN with DASH in mobile environments comes together with advantages compared to today's solutions, especially for devices equipped with multiple network interfaces. The retrieval of data over multiple links in parallel is a useful feature, specifically for adaptive multimedia streaming, since it offers the possibility to dynamically switch between the available links depending on their bandwidth capabilities, transparent to the actual DASH client.

13. References

13.1. Normative References

[RFC6972] Y. Zhang, N. Zong, "Problem Statement and Requirements of the Peer-to-Peer Streaming Protocol (PPSP)", RFC6972, July 2013

13.2. Informative References

[1] ISO/IEC DIS 23009-1.2, Information technology - Dynamic adaptive streaming over HTTP (DASH) - Part 1: Media presentation description and segment formats

[2] Lederer, S., Mueller, C., Rainer, B., Timmerer, C., Hellwagner, H., "An Experimental Analysis of Dynamic Adaptive Streaming over HTTP in Content Centric Networks", in Proceedings of the IEEE International Conference on Multimedia and Expo 2013, San Jose, USA, July, 2013

[3] Liu, Y., Geurts, J., Point, J., Lederer, S., Rainer, B., Mueller, C., Timmerer, C., Hellwagner, H., "Dynamic Adaptive Streaming over CCN: A Caching and Overhead Analysis", in Proceedings of the IEEE international Conference on Communication (ICC) 2013 - Next-Generation Networking Symposium, Budapest, Hungary, June, 2013

[4] Grandl, R., Su, K., Westphal, C., "On the Interaction of Adaptive Video Streaming with Content-Centric Networks", eprint arXiv:1307.0794, July 2013.

[5] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs and R. Braynard, "Networking named content", in Proc. of the 5th int. Conf. on Emerging Networking Experiments and Technologies (CoNEXT '09). ACM, New York, NY, USA, 2009, pp. 1-12.

[6] A. Detti, M. Pomposini, N. Blefari-Melazzi, S. Salsano and A. Bragagnini, "Offloading cellular networks with Information-Centric Networking: The case of video streaming", In Proc. of the Int. Symp. on a World of Wireless, Mobile and Multimedia Networks (WoWMoM '12), IEEE, San Francisco, CA, USA, 1-3, 2012.

[7] V. Jacobson, D. K. Smetters, N. H. Briggs, M. F. Plass, P. Stewart, J. D. Thornton, and R. L. Braynard, "VoCCN: Voice over content-centric networks," in ACM ReArch Workshop, 2009

[8] Christopher Mueller, Stefan Lederer and Christian Timmerer, A proxy effect analysis and fair adaptation algorithm for multiple competing dynamic adaptive streaming over HTTP clients, In Proceedings of the Conference on Visual Communications and Image Processing (VCIP) 2012, San Diego, USA, November 27-30, 2012.

[9] DASH Research at the Institute of Information Technology, Multimedia Communication Group, Alpen-Adria Universitaet Klagenfurt, URL: http://dash.itec.aau.at

[10] A. Detti, N. Blefari-Melazzi, S. Salsano, and M. Pomposini, "CONET: A content centric inter-networking architecture," in ACM Workshop on Information-Centric Networking (ICN), 2011.

[11] W. K. Chai, N. Wang, I. Psaras, G. Pavlou, C. Wang, G. C. de Blas, F. Ramon-Salguero, L. Liang, S. Spirou, A. Beben, and E. Hadjioannou, "CURLING: Content-ubiquitous resolution and delivery infrastructure for next-generation services," IEEE Communications Magazine, vol. 49, no. 3, pp. 112-120, March 2011

[12] NetInf project Website http://www.netinf.org

[13] N. Magharei, R. Rejaie, Yang Guo, "Mesh or Multiple-Tree: A Comparative Study of Live P2P Streaming Approaches," INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE , vol., no., pp.1424,1432, 6-12 May 2007

[14] PPSP WG Website https://datatracker.ietf.org/wg/ppsp/

[15] A. Bakker, R. Petrocco, V. Grishchenko, "Peer-to-Peer Streaming Peer Protocol (PPSPP)", draft-ietf-ppsp-peer-protocol-08

[16] Rui S. Cruz, Mario S. Nunes, Yingjie Gu, Jinwei Xia, Joao P. Taveira, Deng Lingli, "PPSP Tracker Protocol-Base Protocol (PPSP-TP/1.0)", draft-ietf-ppsp-base-tracker-protocol-02

[17] A.Detti, B. Ricci, N. Blefari-Melazzi,"Peer-To-Peer Live Adaptive Video Streaming for Information Centric Cellular Networks", IEEE PIMRC 2013,London, UK, 8-11 September 2013

[18] http://www.ict-medieval.eu

[19] B. Fu, G. Kunzmann, M. Wetterwald, D. Corujo, R. Costa, "QoE-aware Traffic Management for Mobile Video Delivery", Proc. 2013 IEEE ICC, Workshop on Immersive & Interactive Multimedia Communications over the Future Internet (IIMC), Budapest, Hungary, Jun 2013.

[20] Corujo D., Vidal I., Garcia-Reinoso J., Aguiar R., "A Named Data Networking Flexible Framework for Management Communications", IEEE Communications Magazine, Vol. 50, no. 12, pp. 36-43, Dec 2012

[21] Crabtree B., Stevens T., Allan B., Lederer S., Posch D., Mueller C., Timmerer C., Video Adaptation in Limited or Zero Network Coverage, CCNxConn 2013,PARC, Palo Alto, pp. 1-2, 2013

[22] Fiat A., Naor M., "Broadcast Encryption", in Advances in Cryptology (Crypto'93), volume 773 of Lecture Notes in Computer Science, pages 480-491. Springer Berlin / Heidelberg, 1994.

[23] Posch D., Hellwagner H., Schartner P., "On-Demand Video Streaming based on Dynamic Adaptive Encrypted Content Chunks", in Proceedings of the 8th International Workshop on Secure Network Protocols (NPSec' 13), Los Alamitos, IEEE Computer Society Press, October, 2013.

[24] Montpetit M.J., Westphal C., Trossen D., "Network Coding Meets Information Centric Networks," in Proceedings of the workshop on Name-Oriented Mobility (NOM), jointly with ACM MobiHoc 2013, Hilton Head, SC, June 2013.

14. Authors' Addresses

Stefan Lederer, Christian Timmerer, Daniel Posch
Alpen-Adria University Klagenfurt
Universitaetsstrasse 65-67, 9020 Klagenfurt, Austria

Email: {firstname.lastname}@itec.aau.at

Cedric Westphal, Aytac Azgin, Sucheng (Will) Liu
Huawei
2330 Central Expressway, Santa Clara, CA95050, USA

Email: {cedric.westphal, aytac.azgin, liushucheng}@huawei.com

Christopher Mueller
bitmovin GmbH
Lakeside B01, 9020 Klagenfurt, Austria

Email: christopher.mueller@bitmovin.net

Andrea Detti
Electronic Engineering Dept.
University of Rome Tor Vergata
Via del Politecnico 1, Rome, Italy

Email: andrea.detti@uniroma2.it

Daniel Corujo,
Advanced Telecommunications and Networks Group
Instituto de Telecomunicacoes
Campus Universitario de Santiago
P-3810-193 Aveiro, Portugal

Email: dcorujo@av.it.pt

15. Acknowledgements

This work was supported in part by the EC in the context of the SocialSensor (FP7-ICT-287975) project and partly performed in the Lakeside Labs research cluster at AAU. SocialSensor receives research funding from the European Community's Seventh Framework Programme. The work for this document was also partially performed in the context of the FP7/NICT EU-JAPAN GreenICN project, http://www.greenicn.org. Apart from this, the European Commission has no responsibility for the content of this draft. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability. The authors would like to Dr. Jianping Wang (City University Hong Kong) and Marie-Jose Montpetit of MIT for their help in writing the AAA for DRM section.

ICNRG
Internet-Draft
Intended status: Informational
Expires: September 6, 2015

M. Arumaithurai
J. Chen
X. Fu
University of Goettingen
K. Ramakrishnan
University of California, Riverside
J. Seedorf
NEC

March 5, 2015

Enabling Publish/Subscribe in ICN
draft-jiachen-icn-pubsub-01

Abstract

Information-Centric Networks (ICN) provide substantial flexibility for users to obtain information without regard to the source of the information or its current location. Publish/subscribe (pub/sub) systems have gained popularity in society to provide the convenience of removing the temporal dependency of the user having to indicate an interest each time he or she wants to receive a particular piece of related information. Such an "information-centric" communication model should be supported in the new ICN network paradigm. This document outlines some research directions for ICN with respect to enhancing the inherently pull-based ICN approaches for achieving efficient pub/sub capability.

Table of Contents

1. Introduction

This document points out the need to support publish/subscribe (pub/sub) capabilities in ICN and the problems with the existing solutions. Further, the document discusses potential directions for enhancing Information Centric Networking (ICN) to achieve efficient pub/sub.

Section 2 describes the pub/sub systems and the challenges of such systems to the current Internet. Section 3 demonstrates the use of pub/sub systems in different scenarios. Section 4 outlines the requirements of an efficient pub/sub architecture and Section 5 discusses the related works and some possible shortcomings. In Section 6 we brief our standardisation considerations.

2. Pub/Sub Communication

Users increasingly desire access to information, ranging from news, financial markets, healthcare, to disaster relief and beyond, independent of who published it, where it is located, and often, when it was published. Typical representation of these usages are microblogs, RSS feed, social network, search engines, etc. A consumer may not wish (or it may even be infeasible) to receive all of the "channels" belonging to a myriad of information providers that disseminate items of interest, either on demand (such as web, twitter, blogs and social networks), or tune to a broadcast channel (e.g., television, radio, newspaper). In these cases, the consumer would rather prefer obtaining the data based on Content Descriptors (CD) such as a keyword, a tag, or a property of the content (publisher identity, published date etc.).

Publish/subscribe (pub/sub) systems are particularly suited for such kind of large scale content-oriented information dissemination, and provide the exibility for users to subscribe to information of interest, without being intimately tied to when that information is made available by publishers. With the use of an appropriate interface, users can select and filter the information desired so that they receive only what they are interested in, often irrespective of the publisher.

Intelligent end-systems and information aggregators (e.g., Google News and Yahoo! News, cable and satellite providers) have increasingly adapted their interfaces to provide a content-oriented pub/sub-based delivery method. However, these mechanisms are built on top of a centralized server based framework and can also result in a waste of network resources as shown in [Ramasubramanian2006][Katsaros2011], since the Internet protocol suite is focused on end-to-end delivery of data. Furthermore, issues of "coverage" and "timeliness" still exist in such forms of dissemination, where the aggregator may be selective in what information is made available.

Information-Centric Networks (ICN) is a new network paradigm that intendeds to achieve large scale data delivery with greater ease for users, greater scalability in terms of the amount of information disseminated as well as number of producers and consumers of

information, and greater efficiency in terms of network and server resource utilization.

It is also desirable for such a network to assist the pub/sub communication model that delivers the information from any of the producers to all subscribers. Moreover, it is desirable for the network to assist in delivering fine-grained information to the subscriber.

Recently, works such as [Schmidt2012],[Carzaniga2011],[Chen2011],[Chen2012] have also highlighted the need for ICN to support a pub/sub like communication model.

3. Scenarios of Pub/Sub Architecture

In this section, we list several use cases of pub/sub architectures in ICN. They help us to understand the requirements of an efficient pub/sub architecture and why the existing solutions fall short.

3.1. Online Social Networks and RSS Feeds

Online social networks (e.g., Twitter, Facebook, etc.) and Rich Site Summary (RSS) feeds are typical use cases for a content-centric pub/sub system. In such systems, the receivers receive messages either from friends, followees, or from some information aggregators. They do not care which exact machine is sending the message (content-centric), nor do they know when and what is the name of the next message they are going to receive (temporal separation).

To prevent the receivers from polling all the possible providers, existing systems use web servers as rendezvous points: the publishers send new messages to the servers and the receivers/subscribers poll the server periodically. This still causes great wastage for the (HTTP) servers answering "304 - Not Modified" repeatedly since the message update frequency is usually lower than the polling frequency.

3.2. Online Gaming and Audio/Video Conferencing

Massively multiplayer online role-playing games (MMORPGs, e.g., Counter-Strike, Quake, World of Warcraft, etc.) and audio/video conferencing (e.g., Skype meeting, Web Whiteboard, Etherpad, etc.) is another kind of content-centric pub/sub systems. Similar to the social network scenario, users in such systems only care about the content, either the area of interest (AoI) or the conference partners, and they do not know when and from where the next message will come. But different from the previous scenario, such systems

require real-time update (message) delivery and these messages are usually smaller in size compared to the online social networks.

Many of these systems choose to use HTTPS or direct TCP connection between the server and the users to enable the capability of server "pushing" the updates to the user. But maintaining such links are costly. MMORPGs usually limit the number of players in a same game which greatly reduces the interesting of these games.

3.3. Notification Systems in Disaster

Disasters have often disrupted communications because of damages to critical infrastructure. For instance in the aftermath of the Japanese Earthquake in 2011, approximately 1,200,000 fixed telephone lines and 15,000 base-stations were not functioning. On average, 22% (with peaks up to 65% in some areas) of the base-stations had to shut down due to the lack of power or damages to the infrastructure.

Contradictory to the loss of available hardware capacity, during and in the aftermath of a disaster, there is a substantial increase in the amount of traffic generated because of the natural anxiety and panic among people and the need to organize rescue and emergency services. Many of these traffic are in the form of a pub/sub communication model, e.g., the government needs to publish some notifications (recovery status, new shelter locations, etc.), the refugees need to notify their friends about their safety, or people needs to ask for help from ambulances or fire brigade. In the Japanese case, the congestion caused by such traffic resulted in restrictions in voice traffic up to 95%, including emergency priority calls.

4. Requirements of an Efficient Pub/Sub Architecture

Given a pub/sub communication model as described in Section 2, on a high-level one can derive the following (incomplete) list of basic requirements:

- Decouple publishers and subscribers: In an ideal pub/sub environment, publishers only focus on their core task of publishing while not having to maintain membership status, and subscribers receive content from a multitude of sources without having to worry about maintaining a list of publishers and frequently polling them for the availability of fresh data. Moreover, a consumer may not wish (and it may even be infeasible) to subscribe to all of the channels belonging to a myriad of information providers that disseminate items of interest, either on demand (such as web, twitter, blogs and social networks), or tune to a broadcast channel (e.g., television, radio, newspaper).

In these cases, support should be provided to the consumer who would prefer obtaining the data based on descriptors such as keywords, tags, or other properties of the published data.

- Push enabled dissemination: The ability to exploit push-based delivery is a key to achieving timeliness and to avoid wasting server and network resources because of redundant polls. Therefore, an efficient pub/sub architecture must provide the capability for publishers to push information to online subscribers interested in it. Such timely dissemination is necessary in many scenarios such as disaster (e.g., Tsunami) warnings, stock market information, news and gaming.

- Scalability: The target architecture should be able to accommodate a large number of subscribers as well as publishers (often subscribers are also publishers as user-generated content becomes common). Therefore, it should minimize the amount of states maintained in the network, ensure the load on the publisher grows slowly (sublinearly) with the number of subscribers. The load on the subscribers should also grow slowly with the number of publishers (e.g., dealing with the burden of duplicate elimination). Importantly, the load on the network should not grow significantly with the growth in the number of publishers and subscribers. There is also a need to accommodate a very large range in the amount of information that may be disseminated, and the need for all elements of the pub/sub framework in a content-centric environment to scale in a manageable way.

- Efficiency: The architecture should enable a nearly unlimited amount of information being generated by publishers, allow for delivery of information related to subscriptions independent of the frequency at which that information is generated by publishers. The architecture must utilize network and server resources efficiently. It is desirable that content is not transmitted multiple times by a server or on a link. Furthermore, the overhead on publisher and subscriber end-points to query unnecessarily for information must be minimized.

- Dynamicity: The architecture should be able to deal with the substantial churn in subscription state, allowing a large number of users to join, leave and frequently change their subscriptions. The topics of interest may change frequently as well (e.g., in a Twitter-like publishing environment, where the popular topics change frequently).

Additionally, to support a full-fledge pub-sub environment, it is desirable that the target system support the following additional features:

- Support hierarchies and context in naming content: It is desirable to be able to exploit both context and hierarchies in identifying content. Hierarchical naming has been recognized by NDN as well. Exploiting context enables a richer identification of content (in both subscriptions and published information), as noted in the database community.

- Support two-step dissemination for policy control and user interest: There is a need for pub/sub environments to support a two-step dissemination process both for reasons of policy and access control at the publisher as well as managing delivery of large volume content. In such a scenario, the pub/sub framework would be designed to publish only a snippet of the data (containing a description of the content and the method how to obtain it) to subscribers. The subscribers then request for the content based on their interest and allowance.

- Subscriber offline support: Another typical characteristic of pub-sub environments is that subscribers could be offline at the time the data is published. There is clearly a need for asynchronous delivery of information in a pub/sub environment in an efficient, seamless and scalable manner. The system needs to allow users who were online to retrieve the data that they have missed. It should also allow new subscribers to retrieve previously published content that they are interested in. We envisage a server that stores all the content published.

- Prevent Spam/DoS: Spam and DoS attacks are security issues that concern push based pub/sub mechanisms. Efforts to mitigate this at the network layer as well as at the application layer should be considered.

Additionally, it will be desirable to have the following features to support a (limited) pub-sub environment in a disaster affected scneario:

- TBD

- TBD

5. Related Work

5.1. IP/Overlay Multicast

IP multicast [RFC1112] is a candidate solution for efficiently delivering content to multiple receivers. A sender sends data to a multicast group address that subscribers could join. Multicast routing protocols such as PIM-SM [RFC4601] construct and maintain a

tree from each sender to all receivers of a multicast group. However, IP multicast isn't an efficient pub/sub delivery mechanism for several reasons: 1) IP multicast is designed for delivery of packets to connected end-points. Dealing with disconnected operation (when subscribers are online) would have to be an application layer issue. Overlay multicast solutions such as [Jannotti2000][Chu2002][Banerjee2002] are agnostic of the underlying network topology, usually relying on multiple unicasts in the underlay path and are therefore also inefficient as a pub/sub delivery mechanism. 2) The somewhat limited multicast group address space makes it difficult to support a direct mapping of CDs to IP multicast addresses. 3) Current IP multicast is not able to exploit relationships between information elements, such as CDs. CDs may be hierarchical or may have a contextual relationship, which enables multiple CDs to be mapped to a group. For example, consider a publisher that sends a message to all the subscribers interested in football, and subscribers who are interested in receiving messages about all sports. The message from the publisher will have to be sent to two distinct IP multicast groups. If there happens to be a subscriber of messages on sports and football, (s)he will receive the same message twice and will have to perform redundancy elimination in the application layer. The result is a waste in network traffic and processing at both ends.

5.2. Named-Data Networking (NDN)

NDN has limited intrinsic support for pub/sub systems, a critical need in a content centric environment. The aggregation of pending Interests at routers achieves efficient dissemination of information from NDN nodes. But this aggregation is similar to a cache hit in a content distribution network (CDN) cache, which occurs only if subscribers send their Interests with some temporal locality. Thus it avoids multiple Interest queries having to be processed directly by the content provider. Note however that this is still a pull-based information delivery method and depends both on temporal locality of interests and a large enough cache to achieve effective caching in the (content centric) network. On the other hand, native multicast support allows for a much more scalable push-based pub/sub environment, since it is not sensitive to issues such as the cycling of the cache when a large amount of information is disseminated.

TBD: Update it based on recent modifications by the NDN team

5.3. CCN

TBD: Update it based on recent modifications by the CCN team

5.4. Content-Oriented Publish/Subscribe(COPSS)

COPSS enhances CCN/NDN with a push-based delivery mechanism using multicast in a content-centric framework. It is designed to satisfy the requirements mentioned above, especially to provide temporal separation between subscription (or expression of Interest) and publication. At the content-centric network layer, COPSS uses a multiple-sender, multiple-receiver multicast capability, in much the same manner as PIM-SM.

5.5. PSIRP Project

TBD

5.6. NetInf Project (http://www.netinf.org)

TBD

5.7. Pursuit Project (http://www.fp7-pursuit.eu/)

TBD

5.8. Other Related Works

Here we list the other related works we are considering. The list might not be complete and we intend to add to it based on feedback received in further revisions.

- A. Carzaniga, M. Rutherford, A. Wolf, A routing scheme for content-based networking, in: INFOCOM, 2004.
- B. Segall, D. Arnold, J. Boot, M. Henderson, T. Phelps, Content Based Routing with Elvin, in: AUUG2K, 2000.
- C. Esteve, F. Verdi, M. Magalhaes, Towards a new generation of information-oriented Internetworking architectures, in: ReArch, 2008.
- G. Chockler, R. Melamed, Y. Tock, R. Vitenberg, SpiderCast: a scalable interest-aware overlay for topic-based pub/sub communication, in: DEBS, 2007.
- H. Eriksson, Mbone: the multicast backbone, Commun. ACM 37 (8) (1994) 54-60.
- M. Ott, L. French, R. Mago, D. Makwana, Xml-based semantic multicast routing: an overlay network architecture for future information services, in: GLOBECOM, 2004.

- P. T. Eugster, P. A. Felber, R. Guerraoui, A.-M. Kermarrec, The many faces of publish/subscribe, ACM Comput. Surv. 35 (2) (2003) 114-131.

- R. Baldoni, R. Beraldi, V. Quema, L. Querzoni, S. Tucci-Piergiovanni, TERA: topic-based event routing for peer-to-peer architectures, in: DEBS, 2007.

- R. V. Renesse, K. P. Birman, W. Vogels, Astrolabe: A Robust and Scalable Technology for Distributed System Monitoring, Management, and Data Mining, ACM TOCS 21 (2001) 66-85.

- S. Voulgaris, E. Riviere, A.-M. Kermarrec, M. Van Steen, Sub-2-Sub: Self-Organizing Content-Based Publish and Subscribe for Dynamic and Large Scale Collaborative Networks, Research report, INRIA (December 2005).

- T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, I. Stoica, A data-oriented (and beyond) network architecture, in: SIGCOMM, 2007.

- V. Ramasubramanian, R. Peterson, E. G. Sirer, Corona: a high performance publish-subscribe system for the world wide web, in: NSDI, 2006.

- V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard, Networking Named Content, in: CoNEXT, 2009.

- Y. Cui, B. Li, K. Nahrstedt, ostream: asynchronous streaming multicast in application-layer overlay networks, JSAC 22 (1) (2004) 91-106.

- Y. Diao, S. Rizvi, M. J. Franklin, Towards an internet-scale XML dissemination service, in: VLDB, 2004.

6. Standardisation Considerations

Future versions of this document will outline a concrete protocol specification for pub/sub support for ICN. Below some initial standardisation considerations are outlined.

An initial list of details that need to be specified is the following:

- Pub/Sub related interfaces/APIs

- Pub/Sub related data structure modification to existing ICN proposals

We are also considering to write a survey paper that accumulates all the Pub/sub related work.

7. References

7.1. Normative References

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

7.2. Informative References

[Banerjee2002] Banerjee, S., Bhattacharjee, B., and C. Kommareddy, "Scalable application layer multicast", SIGCOMM, 2002, .

[Carzaniga2011] Carzaniga, A., Papalini, M., and A. Wolf, "Content-based Publish/Subscribe Networking and Information-centric Networking", Proceedings of the ACM SIGCOMM workshop on Information-centric networking, ACM, 2011, .

[Chen2011] Chen, J., Arumaithurai, M., Fu, X., and K. Ramakrishnan, "COPSS: An Efficient Content Oriented Publish/Subscribe System", ACM/IEEE 7th Symposium on Architectures for Networking and Communications Systems (ANCS), 2011, .

[Chen2012] Chen, J., Arumaithurai, M., Fu, X., and K. Ramakrishnan, "G-COPSS: A Content Centric Communication Infrastructure for Gaming Applications", IEEE 32nd International Conference on Distributed Computing Systems (ICDCS), 2012, .

[Chu2002] Chu, Y., Rao, S., Seshan, S., and H. Zhang, "A case for end system multicast", IEEE Journal on Selected Areas in Communications 20, no. 8 (2002): 1456-1471, .

[Fenner2005]
Fenner, W., Rabinovich, M., Ramakrishnan, K., Srivastava, D., and Y. Zhang, "XTreeNet: Scalable overlay networks for XML content dissemination and querying (synopsis)", 10th International Workshop on Web Content Caching and Distribution (WCW), 2005, .

[Jannotti2000]
Jannotti, J., Gifford, D., Johnson, K., and M. Kaashoek, "Overcast: reliable multicasting with on overlay network", Proceedings of the 4th conference on Symposium on Operating System Design & Implementation-Volume 4, pp. 14-14. USENIX Association, 2000, .

[Katsaros2011]
Katsaros, K., Xylomenos, G., and G. Polyzos, "MultiCache: An overlay architecture for information-centric networking", Computer Networks 55.4 (2011): 936-947, .

[Ramasubramanian2006]
Ramasubramanian, V., Peterson, R., and E. Sirer, "Corona: A High Performance Publish-Subscribe System for the World Wide Web", NSDI. Vol. 6. 2006, .

[Schmidt2012]
Schmidt, T. and M. Waehlisch, "Why We Shouldn t Forget Multicast in Name-oriented Publish/Subscribe", arXiv preprint arXiv:1201.0349 (2012), .

Appendix A. Acknowledgment

This document has been supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Authors' Addresses

Mayutan Arumaithurai
University of Goettingen
Goldschmidt Str. 7
Goettingen 37077
Germany

Phone: +49 551 39 172046
Fax: +49 551 39 14416
Email: arumaithurai@informatik.uni-goettingen.de

Jiachen Chen
University of Goettingen
Goldschmidt Str. 7
Goettingen 37077
Germany

Phone: +49 551 39 172051
Fax: +49 551 39 14416
Email: jiachen@informatik.uni-goettingen.de

Xiaoming Fu
University of Goettingen
Goldschmidt Str. 7
Goettingen 37077
Germany

Phone: +49 551 39 172023
Fax: +49 551 39 14416
Email: fu@informatik.uni-goettingen.de

K. K. Ramakrishnan
University of California, Riverside
900 University Ave
Riverside CA 92521
USA

Email: kkramakrishnan@yahoo.com

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

ICNRG
Internet-Draft
Intended status: Informational
Expires: September 4, 2015

J. Seedorf
NEC
M. Arumaithurai
University of Goettingen
A. Tagami
KDDI R&D Labs
K. Ramakrishnan
University of California
N. Blefari Melazzi
University Tor Vergata

March 3, 2015

Using ICN in disaster scenarios
draft-seedorf-icn-disaster-03

Abstract

Information Centric Networking is a new paradigm where the network provides users with named content, instead of communication channels between hosts. This document outlines some research directions for Information Centric Networking (ICN) with respect to applying ICN approaches for coping with natural or human-generated, large-scale disasters.

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 September 4, 2015.

Copyright Notice

Table of Contents

1. Introduction

This document summarizes some research challenges for coping with natural or human-generated, large-scale disasters. Further, the document discusses potential directions for applying Information Centric Networking (ICN) to address these challenges.

Section 2 gives some examples of what can be considered a large-scale disaster and what the effects of such disasters on communication networks are. Section 3 outlines why ICN can be beneficial in such scenarios and provides a high-level overview on corresponding research challenges. Section 4 describes some concrete use cases and requirements for disaster scenarios. In Section 5, some concrete ICN-based solutions approaches are outlined. Related research activities are ongoing in the GreenICN research project; Section 6 provides an overview of this project.

2. Disaster Scenarios

An enormous earthquake hit Northeastern Japan (Tohoku areas) on March 11, 2011, and caused extensive damages including blackouts, fires, tsunamis and a nuclear crisis. The lack of information and means of communication caused the isolation of several Japanese cities. This impacted the safety and well-being of residents, and affected rescue work, evacuation activities, and the supply chain for food and other essential items. Even in the Tokyo area that is 300km away from the Tohoku area, more than 100,000 people became 'returner' refugees, who could not reach their homes because they had no means of public transportation (the Japanese government has estimated that more than 6.5 million people would become returner refugees if such a catastrophic disaster were to hit the Tokyo area).

That earthquake in Japan also showed that the current network is vulnerable against disasters and that mobile phones have become the lifelines for communication including safety confirmation. The aftermath of a disaster puts a high strain on available resources due to the need for communication by everyone. Authorities such as the President/Prime-Minister, local authorities, Police, fire brigades, and rescue and medical personnel would like to inform the citizens of possible shelters, food, or even of impending danger. Relatives would like to communicate with each other and be informed about their wellbeing. Affected citizens would like to make enquiries of food distribution centres, shelters or report trapped, missing people to the authorities. Moreover, damage to communication equipment, in addition to the already existing heavy demand for communication highlights the issue of fault-tolerance and energy efficiency.

Additionally, disasters caused by humans such as a terrorist attack may need to be considered, i.e. disasters that are caused deliberately and willfully and have the element of human intent. In such cases, the perpetrators could be actively harming the network by launching a Denial-of-Service attack or by monitoring the network passively to obtain information exchanged, even after the main disaster itself has taken place. Unlike some natural disasters that are predictable using weather forecasting technologies and have a slower onset and occur in known geographical regions and seasons, terrorist attacks may occur suddenly without any advance warning. Nevertheless, there exist many commonalities between natural and human-induced disasters, particularly relating to response and recovery, communication, search and rescue, and coordination of volunteers.

The timely dissemination of information generated and requested by all the affected parties during and the immediate aftermath of a disaster is difficult to provide within the current context of global

information aggregators (such as Google, Yahoo, Bing etc.) that need to index the vast amounts of specialized information related to the disaster. Specialized coverage of the situation and timely dissemination are key to successfully managing disaster situations. We believe that network infrastructure capability provided by Information Centric Networks can be suitable, in conjunction with application and middleware assistance.

3. Research Challenges and Benefits of ICN

3.1. High-Level Research Challenges

Given a disaster scenario as described in Section 2, on a high-level one can derive the following (incomplete) list of corresponding technical challenges:

- Enabling usage of functional parts of the infrastructure, even when these are disconnected from the rest of the network: Assuming that parts of the network infrastructure (i.e. cables/links, routers, mobile bases stations, ...) are functional after a disaster has taken place, it is desirable to be able to continue using such components for communication as much as possible. This is challenging when these components are disconnected from the backhaul, thus forming fragmented networks. This is especially true for today's mobile networks which are comprised of a centralised architecture, mandating connectivity to central entities (which are located in the core of the mobile network) for communication. But also in fixed networks, access to a name resolution service is often necessary to access some given content.

- Decentralised authentication: In mobile networks, users are authenticated via central entities. In order to communicate in fragmented or disconnected parts of a mobile network, the challenge of decentralising such user authentication arises. Independently of the network being fixed or mobile, data origin authentication of content retrieved from the network is challenging when being 'offline' (e.g. disconnected from servers of a security infrastructure such as a PKI).

- Delivering/obtaining information in congested networks: Due to broken cables, failed routers, etc., it is likely that in a disaster scenario the communication network has much less overall capacity for handling traffic. Thus, significant congestion can be expected in parts of the infrastructure. It is therefore a challenge to guarantee message delivery in such a scenario. This is even more important as in the case of a disaster aftermath, it

may be crucial to deliver certain information to recipients (e.g. warnings to citizens).

- Delay/Disruption Tolerant Approach: Fragmented networks makes it difficult to support end-to-end communication. However, communication in general and especially during disaster can tolerate some form of delay. E.g. in order to know if his/her relatives are safe or a 'SOS' call need not be supported in an end-to-end manner. It is sufficient to improve communication resilience in order to deliver such important messages.

- Energy Efficiency: Long-lasting power outages may lead to batteries of communication devices running out, so designing energy-efficient solutions is very important in order to maintain a usable communication infrastructure.

The list above is most likely incomplete; future revisions of this document intend to add additional challenges to the list.

3.2. How ICN can be Beneficial

Several aspects of ICN make related approaches attractive candidates for addressing the challenges described in Section 3.1. Below is an (incomplete) list of considerations why ICN approaches can be beneficial to address these challenges:

- Routing-by-name: ICN protocols natively route by named data objects and can identify objects by names, effectively moving the process of name resolution from the application layer to the network layer. This functionality is very handy in a fragmented network where reference to location-based, fixed addresses may not work as a consequence of disruptions. For instance, name resolution with ICN does not necessarily rely on the reachability of application-layer servers (e.g. DNS resolvers). In highly decentralised scenarios (e.g. in infrastructureless, opportunistic environments) the ICN routing-by-name paradigm effectively may lead to a 'replication-by-name' approach, where content is replicated depending on its name.

- Authentication of named data objects: ICN is built around the concept of named data objects. Several proposals exist for integrating the concept of 'self-certifying data' into a naming scheme (see e.g. [RFC6920]). With such approaches, the origin of data retrieved from the network can be authenticated without relying on a trusted third party or PKI.

- Content-based access control: ICN can regulate access to data objects (e.g. only to a specific user or class of users) by means

of content-based security; this functionality could facilitate trusted communications among peer users in isolated areas of the network.

- Caching: Caching content along a delivery path is an inherent concept in ICN. Caching helps in handling huge amounts of traffic, and can help to avoid congestion in the network (e.g. congestion in backhaul links can be avoided by delivering content from caches at access nodes).

- Sessionless: ICN does not require full end-to-end connectivity. This feature facilitates a seemless aggregation between a normal network and a fragmented network, which needs DTN-like message forwarding.

The list above is most likely incomplete; future revisions of this document intend to add more considerations to the list and to argue in more detail why ICN is suitable for addressing the aforementioned research challenges.

4. Use Cases and Requirements

This Section describes some use cases for the aforementioned disaster scenario (as outlined in Section 2) and discusses the corresponding technical requirements for enabling these use cases.

- Delivering Messages to Relatives/Friends: After a disaster strikes, citizens want to confirm to each other that they are safe. For instance, shortly after a large disaster (e.g., Earthquake, Tornado), people have moved to different refugee shelters. The mobile network is not fully recovered and is fragmented, but some base stations are functional. This use case imposes the following high-level requirements: a) People must be able to communicate with others in the same network fragment, b) people must be able to communicate with others that are located in different fragmented parts of the overall network. More concretely, the following requirements are needed to enable the use case: a) a mechanism for scalable message forwarding scheme that dynamically adapts to changing conditions in disconnected networks, b) DTN-like mechanisms for getting information from disconnected island to another disconnected island, and c) data origin authentication so that users can confirm that the messages they receive are indeed from their relatives or friends.

- Spreading Crucial Information to Citizens: State authorities want to be able to convey important information (e.g. warnings, or information on where to go or how to behave) to citizens. These kinds of information shall reach as many citizens as possible.

i.e. Crucial content from legal authorities shall potentially reach all users in time. The technical requirements that can be derived from this use case are: a) Data origin authentication, such that citizens can confrim the authenticity of messages sent by authorities, b) mechanisms that guarantee the timeliness and loss-free delivery of such information, which may include techniques for prioritizing certain messages in the network depending on who sent them, and c) DTN-like mechanisms for getting information from disconnected island to another disconnected island.

It can be observed that different key use cases for disaster scenarios imply overlapping and similar technical requirements for fulfilling them. As discussed in Section 3.2, ICN approaches are envisioned to be very suitable for addressing these requirements with actual technical solutions.

5. Solution Design

This Section outlines some ICN-based approaches that aim at fulfilling the previously mentioned use cases and requirements.

- ICN ’data mules’: To facilitate the exchange of messages between different network fragments, mobile entitites can act as ICN ’data mules’ which are equipped with storage space and move around the disaster-stricken area gathering information to be disseminated. As the mules move around, they deliver messages to other individuals or points of attachment to different fragments of the network. These ’data mules’ could have a pre-determined path (an ambulance going to and fro from a hospital), a fixed path (drone/ robot assigned specifically to do so) or a completely random path (doctors moving from one camp to another).

- Priority dependent Name-based replication: By allowing spatial and temporal scoping of named messages, priority based replication depending on the scope of a given message is possible. Clearly, spreading information in disaster cases involves space and time factors that have to be taken into account as messages spread. A concrete approach for such scope-based prioritisation of ICN messages in disasters, called ’NREP’, has been proposed [Psaras2014], where ICN messages have attributes such as user-defined priority, space, and temporal-validity. These attributes are then taken into account when prioritizing messages. In [Psaras2014], evaluations show how this approach can be applied to the use case ’Delivering Messages to Relatives/Friends’ decribed in Section 4

- Energy Efficiency: A large-scale disaster causes a large-scale blackout and thus a number of base stations (BSs) will be operated by their batteries. Capacities of such batteries are not large enough to provide cellular communication for several days after the disaster. In order to prolong the batteries' life from one day to several days, different techniques need to be explored: Priority control, cell-zooming, and collaborative upload. Cell zooming switches-off some of the BSs because switching-off is the only way to reduce power consumed at the idle time. In cell zooming, areas covered by such inactive BSs are covered by the active BSs. Collaborative communication is complementary to cell zooming and reduces power proportional to a load of a BS. The load represents cellular frequency resources. In collaborative communication, end-devices delegate sending and receiving messages to and from a base station to a representative end-device of which radio propagation quality is better. The design of an ICN-based publish/subscribe protocol that incorporates collaborative upload is ongoing work. In particular, the integration of collaborative upload techniques into the COPSS (Content Oriented Publish/ Subscribe System)} framework is envisioned [COPSS2011].

- Data-centric confidentiality and access control: In ICN, the requested content is not anymore associated to a trusted server or an endpoint location, but it can be retrieved from any network cache or a replica server. This call for 'data-centric' security, where security relies on information exclusively contained in the message itself, or, if extra information provided by trusted entities is needed, this should be gathered through offline, asynchronous, and non interactive communication, rather than from an explicit online interactive handshake with trusted servers. The ability to guarantee security without any online entities is particularly important in disaster scenarios with fragmented networks. One concrete cryptographic technique is 'Ciphertext-Policy Attribute Based Encryption' (CP-ABE), allowing a party to encrypt a content specifying a policy, which consists in a Boolean expression over attributes, that must be satisfied by those who want to decrypt such content. Such encryption schemes tie confidentiality and access-control to the transferred data, which can be transmitted also in an unsecured channel, enabling the source to specify the set of nodes allowed to decrypt.

- Decentralised authentication of messages: Self-certifying names provide the property that any entity in a distributed system can verify the binding between a corresponding public key and the self-certifying name without relying on a trusted third party. Self-certifying names thus provide a decentralized form of data origin authentication. However, self-certifying names lack a binding with a corresponding real-world identity. Given the

decentralised nature of a disaster scenario, a PKI-based approach for binding self-certifying names with real-world identities is not feasible. Instead, a Web-of-Trust can be used to provide this binding. Not only are the cryptograohic signatures used within a Web-of-Trust independent of any central authority; there are also technical means for making the inherent trust relationships of a Web-of-Trust available to network entities in a decentralised, 'offline' fashion, such that information received can be assessed based on these trust relationships. A concrete scheme for such an approach has been published in [Seedorf2014], where also concrete examples for fulfilling the use case 'Delivering Messages to Relatives/Friends' with this approach are given.

6. The GreenICN Project

This section provides a brief overview of the GreenICN project. You can find more information at the project web site http://www.greenicn.org/

The recently formed GreenICN project, funded by the EU and Japan, aims to accelerate the practical deployment of ICN, addressing how ICN networks and devices can operate in a highly scalable and energy-efficient way. The project will exploit the designed infrastructure to support multiple applications including the following two broad exemplary scenarios: 1) The aftermath of a disaster, e.g. hurricane, earthquake, tsunami, or a human-generated network breakdown when energy and communication resources are at a premium and it is critical to efficiently distribute disaster notification and critical rescue information. Key to this is the ability to exploit fragmented networks with only intermittent connectivity, the potential exploitation of multiple modalities of communication and use of query/response and pub/sub approaches; 2) Scalable, efficient pub/sub video delivery, a key requirement in both normal and disaster situations.

GreenICN will expose a functionality-rich API to spur the creation of new applications and services expected to drive industry and consumers, with special focus on the EU and Japanese environments, into ICN adoption. Our team, comprising researchers with diverse expertise, system and network equipment manufacturers, device vendors, a startup, and mobile telecommunications operators, is very well positioned to design, prototype and deploy GreenICN technology, and validate usability and performance of real-world GreenICN applications, contributing to create a new, low-energy, Information-Centric global communications infrastructure. We also plan to make contributions to standards bodies to further the adoption of ICN technologies.

7. Conclusion

This document outlines some research directions for Information Centric Networking (ICN) with respect to applying ICN approaches for coping with natural or human-generated, large-scale disasters. The document describes high-level research challenges as well as a general rationale why ICN approaches could be beneficial to address these challenges. One main objective of this document is to gather feedback from the ICN community within the IETF and IRTF regarding how ICN approaches can be suitable to solve the presented research challenges. Future revisions of this draft intend to include additional research challenges and to discuss what implications this research area has regarding related, future IETF standardisation.

8. References

8.1. Normative References

[RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.

8.2. Informative References

[COPSS2011]
Chen, J., Arumaithurai, M., Jiao, L., Fu, X., and K. Ramakrishnan, "COPSS: An Efficient Content Oriented Publish/Subscribe System", Seventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS), 2011, .

[Psaras2014]
Psaras, I., Saino, L., Arumaithurai, M., Ramakrishnan, K., and G. Pavlou, "Name-Based Replication Priorities in Disaster Cases", 2nd Workshop on Name Oriented Mobility (NOM), 2014, .

[Seedorf2014]
Seedorf, J., Kutscher, D., and F. Schneider, "Decentralised Binding of Self-Certifying Names to Real-World Identities for Assessment of Third-Party Messages in Fragmented Mobile Networks", 2nd Workshop on Name Oriented Mobility (NOM), 2014, .

Appendix A. Acknowledgment

The authors would like to thank Ioannis Psaras for useful comments.

This document has been supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Authors' Addresses

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

Mayutan Arumaithurai
University of Goettingen
Goldschmidt Str. 7
Goettingen 37077
Germany

Phone: +49 551 39 172046
Fax: +49 551 39 14416
Email: arumaithurai@informatik.uni-goettingen.de

Atsushi Tagami
KDDI R&D Labs
2-1-15 Ohara
Fujimino, Saitama 356-85025
Japan

Phone: +81 49 278 73651
Fax: +81 49 278 7510
Email: tagami@kddilabs.jp

K. K. Ramakrishnan
University of California
Riverside CA
USA

Email: kkramakrishnan@yahoo.com

Nicola Blefari Melazzi
University Tor Vergata
Via del Politecnico, 1
Roma 00133
Italy

Phone: +39 06 7259 7501
Fax: +39 06 7259 7435
Email: blefari@uniroma2.it

ICNRG
Internet-Draft
Intended status: Informational
Expires: September 6, 2015

J. Seedorf
NEC

March 5, 2015

Binding Self-certifying Names to Real-World Identities with a Web-of-Trust
draft-seedorf-icn-wot-selfcertifying-01

Abstract

Self-certifying names are one way of binding a given public key to a certain name in Information Centric Networking. However, an additional binding of a self-certifying name to a Real-World identity is needed in most cases, so that a recipient of some information cannot only verify that the publisher was in possession of the correct corressponding private key for the requested name, but that in addition the name itself is the intended one. This draft specifies how such a binding of Real-World identities with self-certifying ICN names can be done, taking existing IETF specifications into account.

Table of Contents

1. Introduction

Self-certifying names provide the useful property that any entity in a distributed system can verify the binding between a corresponding public key and the self-certifying name without relying on a trusted third party [Aura2003]. Self-certifying names thus provide a decentralized form of data origin authentication. This feature makes self-certifying names a prime candidate for addressing the security requirements in Information Centric Networking (ICN) (which are inherently different from IP networks): a source can digitally sign data associated with a self-certifying name, and any intermediate entity (e.g. ICN-router/Cache) or receiving entity (i.e. issuer of a request for the name) can verify the signature, without the need to verify the identity of the host that caches the object, nor relying on a trusted third party, or a Public Key Infrastructure (PKI). However, as noted in [Ghodsi2011] and elsewhere, self-certifying names lack a binding with a corresponding real-world identity (RWI): the concept enables to verify that whoever signed some data was in possession of the private key associated with the self-certifying name, but it does not provide any means to verify what real-world identity corresponds to the public key, i.e. who actually signed the data [Ghodsi2011] [Nom2014].

In principle, this binding between a public key and an RWI could be provided by a PKI, or alternatively by a Web-of-Trust (WoT)

[Ghodsi2011]. Several ICN approaches use a PKI [Survey] . However, until recently, there have not been concrete proposals for a WoT-based approach for binding a public key (or a self-certifying name) with an RWI in content-oriented architectures. A concrete approach on how this can be done has been proposed in [Nom2014]. This document has the objective of providing the corresponding necessary standards specification to enable this approach (or similar ones) in principle in an interoperable way.

2. High-Level Design

On a high level, binding of self-certifying names and a Web-of-Trust can be achieved in the following way (see [Nom2014] for a detailed example of such an approach): The WoT key-ID is equivalent to the self-certifying name part used in the naming scheme. This ties the self-certifying name with the ID of the corresponding public key in the WoT.

For instance, in the existing PGP Web-of-Trust, the V4 key ID is the lower 64 bits of the fingerprint of the public key, where the fingerprint is essentially the 160-bit SHA-1 hash of the public key [RFC2440]. So if a self-certifying name would be based on the same lower 64-bits of the fingerprint of a given public key, this public key would be tied to the self-certifying name and at the same time be tied to the real-world identity used in the WoT, e.g. an email-address or the real (i.e. non-self-certifying) name of a given ICN publisher.

Thus, if a user requests the content for a self-certifying name in a given ICN architecture, he/she would retrieve the content which contains a digital signature and the corresponding public key for the self-certifying name. The user can then verify that the content retrieved indeed belongs to the name by first hashing the public key and confirm that the hash (or part of it) matches the requested name, and second using the public key to verify the signature over the content. This is in principle the general way of using self-certifying names for data origin authentication in distributed systems. If, in addition, (part of) the self-certifying name is equivalent to a WoT key-ID, the user can use any WoT infrastructure (e.g. PGP keyservers) to retrieve certificates for the key ID that contain/confirm the binding between the corresponding (to the WoT key ID) public key with a real-world identity, such as an email address. This binding provides the requesting user with assurance that the self-certifying name indeed is owned by the intended publisher, i.e. is the correct, intended name from the requestor’s perpective.

The current PGP specification [RFC2440] considers only a bitlength of 64-bit for forming the key-ID, which is not very collision-resistant

(collision-resistance among different key-IDs was not a design goal for PGP [RFC2440]). For securely binding a self-certifying name to a WoT key-ID, collision-resistance is a design goal, because otherwise attckaers could potentially forge a binding of their public key with a given self-certifying name. Thus, either a longer bitlength of the hash of the public key (or its fingerprint) must be used, or hash extension techniques [Aura] must be used, which effectively make collision attacks harder for constant bitlengths at the price of the time needed to create a public/private key pair. Future versions of this document will take these design considerations into account.

3. Standardisation Considerations

Future versions of this document will outline a concrete protocol specification for binding self-certifying names to a Web-of-Trust as outlined on a high level in the previous Section. Below some initial standardisation considerations are highlighted, as well as an assessment of existing IETF standards that could be used as building blocks. Also, future versions of this document will look in more detail into existing IETF specifications, e.g. regarding ICN naming ([RFC6920]) and Web-of-Trust ([RFC2440]), and inspect to what extend such existing specifications can be used directly or in a modified form.

3.1. High-Level Considerations

An initial list of details that need to be specified is the following:

- (List of) Asymmetric cryptography algorithm(s) and corresponding bit-length(s)
- (List of) Hash algorithm(s) and corresponding bit-length(s)
- Rules that define what part of the hash is used for forming the self-certifying part of the name, i.e. the Web-of-Trust Key-ID
- Rules for forming a self-certifying name based on a public key
- Semantics of a signature in the Web-of-Trust
- Defintion of how many bits are used in case of hash extension techniques [Aura][RFC3972]

3.2. Existing Information-Centric Naming Schemes in the IETF

RFC 6920 ’Naming Things with Hashes’ defines a standard for correctly identifying data ’using the output from a hash function’ [RFC6920]. In particular, it specifies a ’(ni) URI Format’ (see [RFC6920], Section 3) and a ’Named Information Hash Algorithm Registry’ (see [RFC6920], Section 9.4). These building blocks allow to specify a format for self-certifying names as hashes of WoT public keys, as outlined above,. In particular, truncated hash formats are clearly defined which can be used to form a self-certifying name from a Web-of-Trust public key by defining what part of the hash is used for forming the WoT key-ID self-certifying part of the name (e.g. ’sha-256-64’ for a truncated SHA-256 hash to 64 bits).

3.3. Existing Web-of-Trust Standards in the IETF

RFC 2440 asymmetric cryptography algorithms and corresponding bit-length for usage in a Web-of-Trust [RFC2440]. Thus, there is an existing IETF specification that provides this building block needed for binding Self-certifying Names to Real-World Identities with a Web-of-Trust.

3.4. Hash Extension Techniques

RFC 3972 discusses hash extension techniques, i.e. approaches that ’increase the cost of both address generation and brute-force attacks by the same parameterized factor while keeping the cost of address use and verification constant’ [RFC3972]. This can be a building block for using hash extension techniques for binding Self-certifying Names to Real-World Identities with a Web-of-Trust.

4. Conclusion

One option for binding self-certifying names to real-world identities is using a Web-of-Trust. This document aims at a concrete specification for providing such a binding, taking existing IETF specification into account. An inspection of existing Web-of-Trust and Naming Scheme standards in the IETF reveal that the basic building blocks for the intended specification for binding Self-certifying Names to Real-World Identities with a Web-of-Trust are already available as IETF standards. Future versions of this document will provide a more detailed specification.

5. References

5.1. Normative References

[RFC2440] Callas, J., Donnerhacke, L., Finney, H., and R. Thayer, "OpenPGP Message Format", RFC 2440, November 1998.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.

5.2. Informative References

[Aura] Aura, T. and M. Roe, "Strengthening Short Hash Values", http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.145.7681, .

[Aura2003] Aura, T., "Cryptographically Generated Addresses (CGA)", 6th International Conference on Information Security (ISC), 2003, .

[Ghodsi2011] Ghodsi, A., Koponen, T., Rajahalme, J., Sarolahti, P., and S. Shenker, "Naming in Content-oriented Architectures", ACM SIGCOMM Workshop on Information-centric Networking, 2011, .

[I-D.seedorf-icn-disaster] Seedorf, J., Arumaithurai, M., Tagami, A., Ramakrishnan, K., and N. Blefari-Melazzi, "Using ICN in disaster scenarios", draft-seedorf-icn-disaster-03 (work in progress), March 2015.

[Nom2014] Seedorf, J., Kutscher, D., and F. Schneider, "Decentralised Binding of Self-Certifying Names to Real-World Identities for Assessment of Third-Party Messages in Fragmented Mobile Networks", 2nd Workshop on Name Oriented Mobility (NOM), 2014, .

[Survey] Xylomenos, G., Ververidis, C., Siris, V., Fotiou, N., Tsilopoulos, C., Vasilakos, X., Katsaros, K., and G. Polyzos, "A Survey of Information-Centric Networking Research", IEEE Communications Surveys and Tutorials, Vol. 16, No. 2, pp 1024-1049, 2014, .

Appendix A. Acknowledgment

This document has been supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Author’s Address

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

ICN Research Group
Internet Draft
Intended status: Informational
Expires: September 4, 2015

P. Truong
Orange
K. Satzke
Alcatel-Lucent
B. Mathieu
Orange
E. Stephan
Orange

March 4, 2015

Named data networking for social network content delivery

draft-truong-icnrg-ndn-osn-00.txt

Abstract

Online Social Networking (OSN) applications have attracted millions of people over the last few years. Their traffic represents a large part of the traffic of the Internet. For instance, Facebook represents near 25 percent of the Internet traffic [14][15], and a part of this traffic is exchanged amongst groups of end-users which are located in the same geographic area. In this document, we introduce a Named Data Networking (NDN) architecture to improve the delivery of OSNs contents requested by end-users in the neighbourhood of the publishers: Having the knowledge of the social network graph and the end-users network location, a SDN-based NDN controller dynamically configures the NDN routers to route the interest requests directly between the end-users.

Table of Contents

1. Introduction

Internet usage has rapidly evolved over the last decade, moving beyond simple one-to-one connections toward more complex interconnections between hosts. Online Social Networking (OSN) services are distributed platforms through which people sharing the same interests, activities, background and so on can interact. The most prominent OSNs are Facebook, Twitter and LinkedIn [1].

Analysing the OSN end-users behaviour is done by several research teams, wishing to better understand the social relationships between users, their habits, their way of using and consuming OSNs, etc. From those papers, such as [2][3][4], we can say that locality plays an important role in OSN applications. People are very frequently connected to other people that are in the same town, same region, in short in a close vicinity (e.g. tweets are distributed locally to local followers, users often send their tweets from the same location, etc.), except for very popular accounts (e.g. a Twitter account having millions of followers). However, the current networking behaviour of the OSN applications does not take this into account. Messages are always transferred toward remote centralised servers, while the real destinations of the messages can be very close[3].

We then defined a framework of NDN for OSN Delivery (NOD), which is an name-based forwarding scheme designed to optimize the delivery of OSN data, based on the social relationships. Our NOD framework focuses on four key features:

- The separation of the NDN control plane from forwarding plane

- A centralized NDN controller and view of the network topology

- Interface between the NDN controller and the OSN server to share social network graph information

- Interface between the NDN controller and the NDN routers to dynamically configure the forwarding plane.

This migration of control, formerly tightly bound in individual network devices, enables the underlying infrastructure to be abstracted for applications and network services, which can treat the network as a logical or virtual entity.

The centralized view and the separation of the control plane and the data plane mean that the controller can create and maintain a topology of how the forwarding nodes under his control are connected and, based on some combination algorithms, can create paths through the network. That allows the controller to better manage traffic flows across the entire network and to react to changes quicker and more intelligently.

The examples given in the draft configure the forwarding plane using SDN interfaces to highlight how Information-Centric Networking (ICN) concepts can help to deliver OSN contents more efficiently. These interfaces and underlying networking data transport could be done using other content naming schemes or routing protocols.

2. Analysis of OSN Networking Behaviour

OSN applications work on graph structures that define the social relationships between end-users. The study [4] of the graph of the user communities of Twitter and Facebook shows that they are very similar:

. A majority of users have friends in the immediate vicinity (city, district or country).
. Only very popular users have "friends" (or followers) distributed almost everywhere around the world.

The delivery of contents amongst OSN users can then be categorized into two groups:

a) OSN users having a mainly local (e.g., national) follower group, their content is consumed locally

b) OSN users having non-local (i.e., international) follower group, their content is consumed worldwide.

The study of OSN network traffic shows that the flows do not reflect the end-users graphs. OSN applications currently forward end-users' requests toward either the OSN remote servers or some CDN caches irrespectively of the end-users' relationships. For example, the exchange of OSN content between two users, who are friends on the OSN and located in the same access ISP network, is directed outside the ISP network toward the OSN servers.

The delivery of OSN content locally represents an important field of resources saving (transit cost, quality of experience, ...).

Our study proposes a new architectural solution for optimizing content delivery in OSNs. It is based on a routing scheme that takes into account end-users' co-locality.

This draft proposes a network model based on Named Data Networking (NDN) and the knowledge about the co-locality of end-users: The NDN controller configures NDN routing states in the network based on its knowledge about the locations of NDN forwarding nodes and OSN content, as well as the social graph from the OSN.

3. NDN-based Naming Scheme

While our approach leverages name-based routing principles, it does not depend on a particular naming scheme. We propose to adopt a hierarchical naming scheme allowing addressing OSN applications, end-users of a specific OSN and the contents produced by end-users. As an example, for a particular OSN application the following hierarchical naming scheme is proposed:

/OSNapp/UserID/contentID

The name prefix /OSNapp denotes a specific OSN application. Then, the prefix /OSNapp/UserID identifies an end-user in this OSN "OSNapp". Finally, /OSNapp/UserID/contentID refers to a content produced by the end-user "UserID" of the OSN "OSNapp".

4. Locality-aware Name-based Routing

Based on the analysis of the OSN networking behaviour and the end-users behavior as described in Section 2, we propose a centralised controller-based NDN forwarding scheme that differentiates end-users based on their OSN graph, in particular on the properties of the list of followers of an OSN user.

Popular end-users, whose content is consumed worldwide, should have a different way of working than non-popular local end-users, whose content will be locally consumed. For popular end-users the current networking behaviour of OSNs (e.g.; requests toward the centralized server and data given back by the server) can be retained. It enables OSN providers to keep knowledge of their clients and possibly adapt some processing for very popular end-users. Furthermore, it can be an added value for them (and something they can monetize) for popularity measurements, targeted advertisement (e.g. announcement of a live of a popular singer that is followed by millions of people, etc.).

For local end-users, i.e. for users that have mainly friends in the immediate vicinity, we adopt a modified architecture to alleviate the load on the OSN server and to reduce the network load between the users' access networks and the OSN server, as according to our research the majority of the traffic stays locally. As such, we suggest to route interest requests for local content toward the local end-users themselves who will provide the content.

To make local routing possible, we first define the notion of locality between OSN end-users by using the network routing hop. We say that two users are local if there are separated by two routing hops (or any other value depending on the design configuration).

We assume now that all end-users (local end-users but also non-local ones and popular ones) can announce their prefix name in the network.

Following the hierarchical naming approach described in the previous section, locality-aware name-based routing is possible since NDN allows for the routing of interests based on the longest-prefix matching on the name. Indeed, let us suppose that locality is defined by 2 routing hops, and that the user Bob requests an Interest for Alice's content who name is "/OSNapp/Alice/content245". We then have the two following possible situations:

- If Bob is located more than 2 routing hops far from Alice, his Interest will be routed toward the OSN server based on the OSN application name prefix "/OSNapp". (Note that this forwarding process toward the OSN server is also very likely to be applied to popular end-users who have friends/followers located in numerous places around the world.)
- On the opposite side, if Bob is located within the network region delimited by 2 routing hops, his Interest will be directly routed toward Alice using the name prefix "/OSNapp/Alice". Alice is then required to serve Bob's Interest by send the requested content in a Data message to Bob

As we can see, in addition to have the route for the prefix "/OSNapp", we should then have different entries "/OSNapp/UserID" (for locally serving content requests) in the Forwarding Information Base (FIB) of the NDN routers. So in order to avoid a scalability issue when storing those routes for local users in the FIBs, we advocate the use of a logically-centralized NDN controller to dynamically configure the routes toward end-users taking into account information from their OSN application social graph required, as described in the following section.

5. SDN-based routing configuration employing OSN information

We advocate for a dynamic and temporary configuration of the FIB tables with the help of the NDN controller because is not realistic to think that every forwarding node will keep knowledge of all end-users in their FIB table for scalability reasons.

For ensuring better scalability in the control plane with user's prefix announcements, we adopt a dynamical approach, based a Software-Defined Networking (SDN)-like architecture between a centralized controller and the NDN routers, which allows to dynamically configure the NDN forwarding tables for applying our local routing scheme, based on the social interactions in the OSN (e.g. add route in the local NDN routers according to the network operator's requirements and policies for a certain network region and/or a particular group of contents/users). In a general context of the workflow, when the controller receives an update message from the OSN server containing information about OSN contents or users' locations, it can decide to re-calculate routes using topological and OSN social graph information. If a newly calculated route requires modifications to one or more NDN forwarding elements, the

controller communicates the changes to the registered NDN forwarding nodes, respectively.

A threshold value can be defined for each local user to estimate his popularity which takes into account the number of followers in an OSN application.

If the user's number of followers exceeds the predefined threshold value, the OSN application can instruct the NDN controller to modify the FIBs of forwarding nodes located in close vicinity to the user location as discussed above by adding new prefixes and targets.

On the other hand, if the number of followers of the local user is below a predefined threshold, the OSN application can instruct the NDN controller to modify the FIBs of the previously selected forwarding nodes to remove the additional FIB entries for the particular local user.

This helps scaling the name-based routing approach by avoiding the accumulation of unnecessary FIB entries in the forwarding nodes.

We describe in the following the main types of notification which can be sent to our NDN controller.

5.1. Notification to the controller for setting routes to the OSN server

At the startup, the OSN provider is responsible for announcing its name prefix '/OSNapp' in the network: this allows to set up routes in the network toward the OSN server. The announcement can be registered by a service running on the NDN controller to calculate routes for interest requests in the network to the OSN application server.

Subsequently the controller populates the Forwarding Information base (FIB) of the registered NDN forwarding nodes. The FIB of each forwarding node will then contain at least one entry matching '/OSNapp' to forward all requests for OSNapp content to the respective application server if no other FIB entry has a longer prefix match. It is presumed that the OSN application servers are stable and always on so that there will no requirements for frequent updates of this particular FIB entry.

5.2. Notification to the controller for OSN users' location

When an user, let us say Joe, is getting online, the OSN server notifies the controller that Joe is now available. In the notification message, the OSN server provides the controller with information about Joe, such as:

- the location of the NDN router to which Joe is connected
- the list of Joe's friends/followers who are currently online, along with their location (i.e. the location of the NDN routers to which Joe's online friends are connected)
- the list of users Joe follows (friendship in OSNs is not necessarily two-way, reciprocal.)
- additionally, some other optional meta-information, which can be especially useful when Joe defines different access rights to his contents for his friends. For example, Joe can restrict the access to his personal photos to only his family members.

Based on those elements from the OSN server and on its knowledge of the network topology and policies, the controller can configure the network with new routes toward Joe, by adding the entry "/OSNapp/Joe" into the FIB of the NDN nodes that are located within the local network region around Joe (i.e. located at most 2 routing hops far from the Joe, supposing that we have set the locality value to 2 - the value depends on network topology and configuration). Note that this route is added to a local NDN router only if it is the access router of Joe or there are Joe's online friends attached to the router.

After the controller has configured the local NDN forwarding nodes for Joe's reachability, local friends can retrieve Joe's contents from himself: all related content requests will thus forwarding to Joe who is responsible for serving the demand.

Now, when Joe is disconnecting, the OSN server notifies the controller again, which in turn can check the list of Joe's followers to evaluate with the help of a pre-defined threshold value if the number of followers of Joe still exceeds the threshold or if the FIB of a NDN forwarding node can be cleaned from the entries corresponding to "/OSNapp/Joe" name prefix. And also in the reverse operation, the controller can check the list of users that Joe himself follows in order to check if it can remove other FIB entries to those users to improve scalability of the approach.

6. Application Call Flows

This section describes in details the call flows for our local-aware NDN-based architecture with dynamic routing configuration.

Let us consider the network configuration as illustrated by Figure 1. We represent the different NDN routers in the Figure by their FIB (F1, F2, etc). All NDN routers have been registered with the controller. Therefore, as described in Section 5.1, all the FIBs, F1, F2,...,F6 in the Figure, contain, as a minimum, the entry "/OSNapp" for the route to the OSN server.

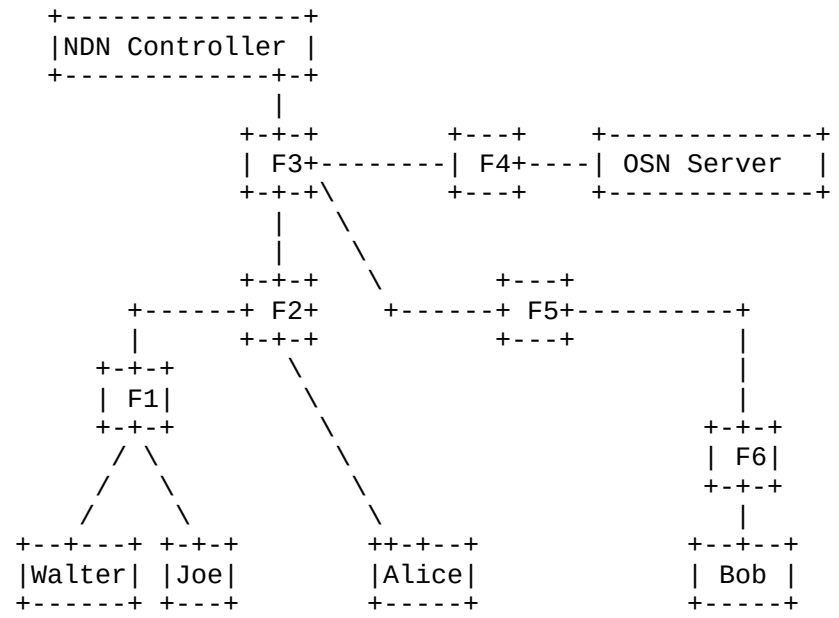

Figure 1: Social- and Local-aware network routing configuration

6.1. Publication of Tweets

For local-aware routing, we set the locality as network regions delimited by two routing hops. As a consequence, Alice, Joe, Walter are local users, i.e. located in the same local area.

1. Joe is getting online: he connects to the OSN server using an application client on his smartphone or computer.
2. The OSN server registers Joe’s client and evaluates the number of Joe’s followers
3. In case the number of followers exceeds a predefined threshold value, the OSN server can decide to inform the NDN controller with a notification message containing meta information related to Joe
4. The NDN controller configures the local NDN routers (located at most two routing hops from Joe) by adding the route "/OSNapp/Joe" for Joe’s reachability in the related FIBs (F1 and F2 in Figure 1). (Note that in Figure 1, we suppose Alice and Walter follow Joe. However, if Alice does not follow Joe, and there are no other users following Joe under the same router, the controller does not add the route "/OSNapp/Joe" to the FIB.)
5. The NDN router FIBs F1 and F2 contain now an entry "/OSNapp/Joe" for the route to Joe’s device.
6. Joe can now publish his new content: the content object is stored in the OSN server.

6.2. Retrieval of Tweets for Local Users

We keep the same network configuration as in 6.1 (Figure 1). As we set locality to two routing hops, Alice and Walter are local users compared to Joe’s location. They are located in the same local network region. Alice and Walter can then benefits from local-aware routing.

1. Alice wants to retrieve Joe’s content. She then expresses an Interest for "/OSNapp/Joe/Video10".
2. Thanks to the previous routing configuration by the controller, (the router FIBs F1 and F2 in Figure 1 contain a route for "/OSNapp/Joe"), the network knows how to forward the Interest("/OSNapp/Joe/Video10"), which will finally directed to Joe.
3. Joe, receiving the Interest, returns a Data message for serving the requested content /OSNapp/Joe/Video10.
4. While forwarding the Data message on the reverse path (taken by the Interest), Joe’s content is also cached in the Content Store of the traversed routers.
5. Alice can then enjoy Joe’s content.

6. Now, when Walter requests Joe’s content /OSNapp/Joe/Video10, he will get it directly from the cache of his NDN access router.

6.3. Retrieval of Tweets for Non-Local Users

Now, we suppose that Bob wants to retrieve Joe’s content Video10 (Figure 1). In this case, as Bob is a non-local user, i.e. located too far from Joe, his Interest for "/OSNapp/Joe/Video10" will be forwarded using the prefix "/OSNapp". The OSN server is then responsible for serving the request, and the sent content will be cached in the content store of the different traversed NDN routers on the reverse path.

7. Security Considerations

This document does not impact the security of the Internet.

8. IANA Considerations

This document presents no IANA considerations.

9. Informative References

[1] http://www.alexa.com/topsites/category/Computers/Internet/On_the_Web/Online_Communities/Social_Networking

[2] M. P. Wittie, V. Pejovic, L. Deek, K. C. Almeroth, and B. Y. Zhao, "Exploiting locality of interest in online social networks," in ACM CoNEXT ’10, New York, USA, 2010, pp. 25:1-25:12.

[3] R. Cuevas, R. Gonzalez, A. Cuevas, and C. Guerrero, "Understanding the locality effect in twitter: measurement and analysis," Personal and Ubiquitous Computing, pp. 1-15, 2013.

[4] eCousin Deliverable D3.1, "Measurement, Modelling, and Prediction of Social-Content Interdependencies": www.ict-ecousin.eu/public-deliverables-dissemination/public-deliverables/ecousin-deliverable-d3.1-v1.2-final.pdf/at_download/file

[5] W. Wongyai and L. Charoenwatana, "Examining the network traffic of facebook homepage retrieval: An end user perspective," in Computer Science and Software Engineering (JCSSE), 2012 International Joint Conference on, 2012, pp. 77-81.

[6] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," Communications Magazine, IEEE, vol. vol.50, July 2012.

[7] L. Wang, M. A. Hoque, C. Yi, A. Alyyan, and B. Zhang, "OSPFN: An OSPF Based Routing Protocol for Named Data Networking," Name Data Networking, Tech. Rep. NDN-003, July 2012.

[8] V. Jacobson, D. K. Smetters, N. H. Briggs, M. F. Plass, P. Stewart, J. D. Thornton, and R. L. Braynard, "VoNDN: voice-over content-centric networks," in Proceedings of the 2009 workshop on Re-architecting the internet, ser. ReArch ’09. New York, NY, USA: ACM, 2009, pp. 1-6.

[9] A. K. M. M. Hoque, S. O. Amin, A. Alyyan, B. Zhang, L. Wang, and L. Zhang, "NLSR: Named-data link state routing protocol", In SIGCOMM 2013 ICN Workshop, 2013

[10] M. Almishari, P. Gasti, N. Nathan, and G. Tsudik, "Optimizing bi-directional low-latency communication in named data networking", SIGCOMM Comput. Commun. Rev. 44, 1 (December 2013), 13-19.

[14] http://www.adweek.com/socialtimes/in-q2-facebook-drove-nearly-a-quarter-of-web-traffic/300175

[15] http://www.digitaltrends.com/mobile/facebook-25-pct-of-u-s-traffic-and-100-million-app-downloads/

10. Acknowledgments

The research leading to these results has received funding from the European Union’s Seventh Framework Programme ([FP7/2007-2013]) under grant agreement n 18398, project eCOUSIN.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the eCOUSIN project or the European Commission.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Patrick Truong

Orange Labs

2 Av. Pierre Marzin

22300 Lannion

France

Email: patrick.truong@orange.com

Klaus Satzke

Alcatel-Lucent Bell Labs

Lorenzstrasse 10

70435 Stuttgart

Email: Klaus.Satzke@alcatel-lucent.com

Bertrand Mathieu

Orange Labs

2 Av. Pierre Marzin

22300 Lannion

France

Email: bertrand2.mathieu@orange.com

Emile Stephan

Orange Labs

2 Av. Pierre Marzin

22300 Lannion

France

Email: emile.stephan@orange.com