

Internet Research Task Force (IRTF)
Request for Comments: 7927
Category: Informational
ISSN: 2070-1721

D. Kutscher, Ed.
NEC
S. Eum
Osaka University
K. Pentikousis
Traveling
I. Psaras
UCL
D. Corujo
Universidade de Aveiro
D. Saucez
INRIA
T. Schmidt
HAW Hamburg
M. Waehlich
FU Berlin
July 2016

Information-Centric Networking (ICN) Research Challenges

Abstract

This memo describes research challenges for Information-Centric Networking (ICN), 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 or enhancing a number of desirable features, such as security, user mobility, multicast, and in-network caching. Mechanisms for realizing these benefits is the subject of ongoing research in the IRTF and elsewhere. This document describes current research challenges in ICN, including naming, security, routing, system scalability, mobility management, wireless networking, transport services, in-network caching, and network management.

This document is a product of the IRTF Information-Centric Networking Research Group (ICNRG).

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Information-Centric Networking Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7927>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	4
2. Problems with Host-Centric Communications	4
3. ICN Terminology and Concepts	6
3.1. Terminology	6
3.2. Concepts	6
4. ICN Research Challenges	8
4.1. Naming, Data Integrity, and Data Origin Authentication	8
4.2. Security	10
4.2.1. Data Integrity and Origin Authentication	10
4.2.2. Binding NDOs to Real-World Identities	11
4.2.3. Access Control and Authorization	12
4.2.4. Encryption	13
4.2.5. Traffic Aggregation and Filtering	13
4.2.6. State Overloading	13
4.2.7. Delivering Data Objects from Replicas	14
4.2.8. Cryptographic Robustness	14
4.2.9. Routing and Forwarding Information Bases	15
4.3. Routing and Resolution System Scalability	15
4.3.1. Route-By-Name Routing	15
4.3.2. Lookup-By-Name Routing	16
4.3.3. Hybrid Routing	17
4.4. Mobility Management	18
4.5. Wireless Networking	20
4.6. Rate and Congestion Control	22
4.7. In-Network Caching	24
4.7.1. Cache Placement	24
4.7.2. Content Placement: Content-to-Cache Distribution	25
4.7.3. Request-to-Cache Routing	26
4.7.4. Staleness Detection of Cached NDOs	26
4.7.5. Cache Sharing by Multiple Applications	27
4.8. Network Management	27
4.9. ICN Applications	29
4.9.1. Web Applications	30
4.9.2. Video Streaming and Download	30
4.9.3. Internet of Things	31
5. Security Considerations	32
6. Informative References	32
Acknowledgments	37
Authors' Addresses	37

1. Introduction

Information-Centric Networking (ICN) is an approach to evolve the Internet infrastructure to directly support accessing Named Data Objects (NDOs) as a first-order network service. 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 and security, 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 allowing for named data structures, 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 can evolve the Internet architecture towards a network model based on named data with different properties and different services.

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 memo is to document the technical challenges and corresponding current approaches and to expose requirements that should be addressed by future research work.

This document has been reviewed, commented on, and discussed extensively for nearly two years by the vast majority of ICNRG members, which certainly exceeds 100 individuals. It is the consensus of ICNRG that the research challenges described in this document should be published in the IRTF stream of the RFC series. This document does not constitute a standard.

2. Problems with Host-Centric Communications

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 that cache content such as Content Distribution Networks (CDNs) and Peer-to-Peer (P2P) applications, 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 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 platforms cannot be leveraged as the network is not aware of data requests and data transmissions. This has the following impact:

- o Data traffic typically follows sub-optimal paths as it is effectively routed, depending on the overlay topology instead of the Internet-layer topology.
- o 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.
- o Overlays typically require significant infrastructure support, e.g., authentication portals, content storage, and applications servers, making it often impossible to establish local, direct communication.
- o The forwarding layer cannot cooperate with transport-layer functions, so sometimes useful functionality such as local retransmission and local rate control have to be implemented with TCP proxies or other intermediaries.
- o 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.
- o Many applications follow their own approach to caching, replication, transport, and authenticity validation (if at all), although they all share similar models for accessing named data objects in the network.

Host-centric communication systems restrict applications to data transfer between end-hosts only. Naming data directly provides a powerful "hook" for applications to exploit and natively support multi-party communication, e.g., multi-source/multi-destination communication and a ubiquitous information ecosystem that is not restricted to end-host addresses.

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 information-centric network 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, e.g., chunks and segments. Sometimes smaller units may be represented by NDOs themselves. Within the context of this document, an NDO is any named data object that can be requested from the network, and we do not consider sub-units below the NDO level. In this document, we often use the terms "NDO" and "data object" interchangeably.

Requestor: Entity in an ICN network that is sending a request for a named data object to the network.

Publisher: Entity in an ICN network that publishes an NDO to the network, so that corresponding requests can reach the publisher. The publisher does not need to be identical to the actual creator, for example, a publisher could provide the service of hosting NDOs on behalf of the actual creators/owners.

3.2. Concepts

Fundamentally, ICN provides 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, 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 and 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 requestors. 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 ICN node 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 resend 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 be able to send local retransmissions, which enables shorter round-trip times and the offloading of origin servers and other parts of the network.

ICN potentially retrieves segments of NDOs from multiple data sources, so only a requestor can determine the completion of a retrieval process, i.e., the retrieval of NDOs or individual segments is typically controlled by a requestor. For this reason, ICN transport protocols are typically based on a receiver-driven mechanism: requestors 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 resending 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, Data Integrity, and Data Origin Authentication

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 of establishing a verifiable binding between the object and its name (name-data binding validation) so that a requestor can be sure that the received bits do correspond to the NDO originally requested (data integrity). Data origin authentication is a different security service that can be related to naming, i.e., verifying that an NDO has indeed been published by a publisher (that could be identified by a name prefix).

The above functions are fundamentally required for the information-centric network to work reliably; otherwise, neither network elements nor requestors 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 the user intent.

The second general class of naming schemes enables verifying the object's name-data integrity without requiring a Public Key Infrastructure (PKI) or other third party to first establish trust in the key. This 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 private key. The resulting names are typically non-hierarchical, or flat, although the publisher field could be employed to create a structure that could facilitate route aggregation.

There are several design trade-offs for ICN naming that affect routing and security. Hash-based names are not human readable nor hierarchical. They can, however, provide some structure for aggregation, for instance, a name part corresponding to a publisher. In hash-based names with indirect binding, the key of the publisher is bound to the name of NDO, so when a user receives, e.g., the triplet, namely (data, key, signature), the receiving entity can verify that the NDO has been generated by the possessor of the private/public key pair and that the NDO has not been changed in transit (data integrity). This can be done by cryptographically hashing the received key and the name of the NDO, and comparing it with the received hashed key. Then, the key can be used to verify the signature.

Data origin authentication can be achieved by validating signatures based on public key cryptography about an NDO's name and content. In order to ascertain data integrity and origin authenticity with such an approach, a PKI-like system is required that would allow linking the corresponding public key to a trust chain.

Research challenges specific to naming include:

- o Naming static data objects can be performed by using content hashes as part of object names, so that publishers can calculate the hash over existing data objects and requestors, and any ICN node can validate the name-content binding by recalculating the hash and comparing it to the name (component). [RFC6920] specifies a concrete naming format for this.
- o 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 hash-based names with indirect binding as described above.
- o 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.

- o Updating and versioning NDOs 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 newer and older versions).
- o Managing accessibility can also be a challenge. In ICN, the general assumption is to enable ubiquitous access to NDOs, but 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 is an active research field in ICN. This section provides an overview of important security features and corresponding challenges that are related to shifting 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 Integrity and Origin Authentication

As mentioned in Section 4.1, data integrity verification is an important ICN feature, since NDOs 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 Transport Layer Security (TLS) [RFC5246] cannot be used as a general solution. Since data objects can be maliciously modified, ICN should provide receivers with a security mechanism to verify the 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.

Data origin authentication, i.e., verifying that an NDO has indeed been published by a publisher, requires a secure binding of an NDO name to a publisher identity -- this is also typically implemented using public key cryptography, i.e., by requiring a receiver to verify digital signatures that are part of received messages.

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) (see [RFC5280] for a description of hierarchical PKI), and the other is to adapt a hash-based scheme with indirect binding. 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 hash-based scheme with indirect binding, the public key (or a hash of it) would be used as part of the name -- which is sufficient to validate the data integrity.

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 names, real-world identity bindings are not intrinsically established: the name provides the hash of the NDO or of the public key that was 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 a 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. Access Control and Authorization

Access control and authorization is a challenge in ICN, because of the lack of user-to-server authentication in the fundamental communication model based on named data.

All ICN entities are capable of delivering NDOs on demand due to their in-network caching function. In such an environment, traditional access control schemes based on Access Control List (ACL) are ill-suited since widely distributed ICN entities have to maintain an identical control policy over NDOs for each consumer, which is prohibited due to computational overhead and privacy issues. There are two complementary approaches to address the issues:

1. Separated approach: access control service from a third party that is independent from ICN entities. Due to the clear separation, ICN entities are free from computational overhead to determine the accessibility of NDOs by consumers; also, consumers can secure their privacy through the independent authorization entity [ACCESS-CTL-DEL]. Relevant challenges to this approach include reducing the authorization delay (when communicating to the access control provider) and currency and consistency of access control information (when access control lists are distributed).
2. Integrated approach: access control service from ICN entities. This mechanism is often based on content encryption and key distribution [ENCRYPTION-AC]. As mentioned previously, this approach suffers from prohibitive overhead for ICN entities due to the process of key verification. While key distribution is challenging per se, this approach is beneficial in a way that NDOs can be retrieved without the help of an external access control provider. Challenges to this approach include:

1. applying an access control mechanism for dynamic NDOs in in-network caches in a timely manner;
2. providing consumers with the different levels of accessibility to individual NDOs in a scalable manner; and
3. managing key revocation and similar PKI management functions.

4.2.4. Encryption

In ICN, NDOs can be encrypted to implement access control (only consumers in possession of corresponding decryption keys can access the content) or privacy (same approach). Distributing and managing the corresponding keys as well as providing usable interfaces to applications and human users are challenges and the subject of ongoing work.

In principle, the challenges are similar to those of broadcast/media distribution, and similar approaches (combing symmetric with public key cryptography) are being investigated [NDN-CTL-SHARING].

4.2.5. 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 authorized 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].

4.2.6. State Overloading

ICN solutions that implement state on intermediate routers for request routing or forwarding (e.g., Content-Centric Networking (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.

4.2.7. 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 loses 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 earlier publications 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).

4.2.8. 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 massively sign content, which is long-lived, offer time and data to an attacker for comprising cryptographic credentials. Signing a large amount of data eases common attacks that try to breach the key of the publisher. Based on this observation, the following research challenges emerge:

- o To which extent does the content publication model conflict with the cryptographic limitations?
- o How can we achieve transparent re-signing without introducing additional cryptographic weaknesses or key management overhead?

In general, ICN implementations should be designed considering the guidelines provided by [RFC7696], especially regarding cryptographic algorithm agility, for example, [RFC6920] specifies a naming scheme for hash-based names that was designed to support algorithm agility.

4.2.9. Routing and Forwarding Information Bases

In information-centric 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 de-aggregation level whenever applicable, blacklists, etc.,).

4.3. Routing and Resolution System Scalability

ICN routing is a process that finds an NDO based on its name initially provided by a requestor. ICN routing may comprise three steps: (1) name resolution, (2) discovery, and (3) delivery. The name resolution step translates the name of the requested NDO into its locator. The discovery step routes the request to the 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) as discussed in the following subsections.

4.3.1. Route-By-Name Routing

RBNR omits the first name resolution step as the name of the NDO 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 objects 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 breadcrumbs 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.

Challenges specific to RBNR include:

- o How can we aggregate the names of data objects to reduce the number of routing entries?
- o How does a user learn the name that is designed for aggregation by the provider? For example, although we name our contribution as "ICN research challenges", the IRTF (provider) may want to change the name to "/IETF/IRTF/ICN/Research challenges" for aggregation. In this case, how does a user learn the name "/IETF/IRTF/ICN/Research challenges" to retrieve the contribution initially named "ICN research challenges" without any resolution process?
- o Without introducing the name aggregation scheme, can we still achieve scalable routing by taking advantage of topological structure and distributed copies? For example, would employing compact routing [COMPACT], random walk [RANDOM], or greedy routing [GREEDY] work at the Internet scale?
- o How can we incorporate copies of a data object in in-network caches in this routing scheme?
- o Breadcrumbs routing implies a symmetric path for ICN request and response delivery. Some network configurations and link types prohibit symmetric path forwarding, so it would be challenging to interconnect such networks to an infrastructure based on breadcrumbs routing. For example, certain forwarding strategies in Delay-Tolerant Networking (DTN) [RFC4838] are employing opportunistic forwarding where responses cannot be assumed to travel the same path as requests.

4.3.2. Lookup-By-Name Routing

LBNR uses the first name resolution step to translate the name of the 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. An instantiation of LBNR is [MDHT].

Challenges specific to LBNR include:

- o How can we build a scalable resolution system that provides:
 - * Fast lookup: Mapping the name of a data object to its locators (copies as well).
 - * Fast update: The location of a 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.
- o How can we incorporate copies of a data object in in-network caches in this routing scheme?

4.3.3. Hybrid Routing

HR combines RBNR and LBNR to benefit from their advantages. 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 that have their own prefix (locator).

For instance, a request message initially includes the name of the NDO for the operation of RBNR and is forwarded to a cached copy of the NDO or the original server. When the request message fails to find a routing entry in the router, a name resolution step kicks in to translate the name into its locator before forwarding the request message based on the retrieved locator.

Challenges specific to HR are:

- o How can we design a scalable mapping system that, given the name of the NDO, should return a destination domain locator so that a user request can be encapsulated and forwarded to the domain?
- o How can the mapping information be secured to prevent a malicious router from hijacking the request message by chaining its locator?
- o How can the bind between the name and the content of the NDO be maintained for the verification of its origin and integrity when the name changes due to the retrieved locator?

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 (see [MMIN]). 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 that 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 an 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-independent. 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 also at a specific correspondent node due to the end-to-end nature of host-centric communication. However, once a change in the point of attachment occurs, information retrieval from the original "correspondent node" may no longer be 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, especially 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 delivery mechanisms, enabling scenarios such as preparation of the content delivery 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 for ICN mobility management can be derived:

- o How can mobility management take full advantage of native ICN primitives?

- o How do we avoid the need for mobility anchors in a network that by design supports multicast, anycast, and location-independent information retrieval?
- o How can content retrieval mechanisms interface with specific link operations, such as identifying which links are available for certain content?
- o How can mobility be offered as a service that is only activated when the specific user/content/conditions require it?
- o How can mobility management be coordinated between the node and the network for optimization and policing procedures?
- o How do we ensure that managing mobility does not introduce scalability issues in ICN?
- o 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 (L2) wireless network radio access technologies 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 traffic, how to support multicast, 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 can broadcast better serve ICN needs 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.

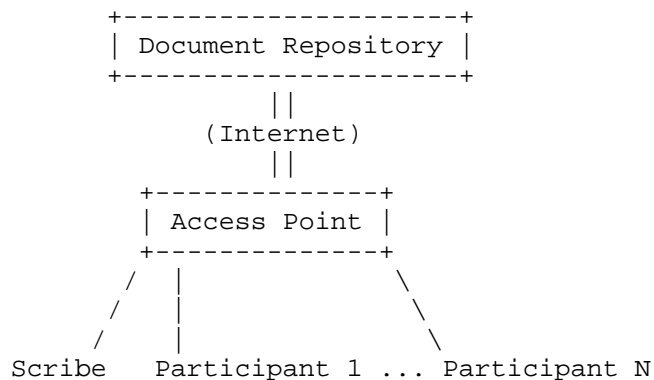


Figure 1: Document Sharing During a Meeting

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 the spectral efficiency (or the economics) underlying the wireless technology is. 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:

- o Can we use wireless resources more frugally with the information-centric paradigm than what is possible today in all-IP wireless networks?
- o In the context of wireless access, how can we leverage the broadcast nature of the medium in an information-centric network?
- o Would a wireless-oriented ICN protocol stack deliver significant performance gains? How different would it be from a wired-oriented ICN protocol stack?
- o 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?
- o Can promiscuous wireless interface operation coupled with opportunistic caching increase ICN performance, and if so, by how much?
- o How can a conversational service be supported at least as efficiently as today's state-of-the-art wireless networks deliver?
- o What are the benefits of combining ICN with network coding in wireless networks?
- o How can Multiple-Input Multiple-Output (MIMO) and Coordinated Multipoint Transmission (CoMP) be natively combined with ICN primitives in future cellular networks?

4.6. Rate and Congestion Control

ICN's receiver-driven communication model as described above creates new opportunities for transport protocol design, as it does not rely solely on end-to-end communication from a sender to a requestor. A requested data 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 NDO 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] and [CONTUG] describe request rate control protocols 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 in-network caches. The traditional concept of a flow is then somewhat not valid 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, can often prove misleading. For example, false Round-Trip Time (RTT) measurements will lead to largely variable average and smoothed 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 data object 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 reused in the context of ICN. Further research is needed in this direction though to choose the right technique and 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 protocol design and offers new options for employing per-node forwarding strategies and for rethinking in-network resource sharing [RESOURCE-POOL].

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, routers can adapt their timers to use the best of the communication paths.

4.7. In-Network Caching

Explicitly named data objects allow for caching at virtually any network element, including routers, proxy caches, and end-user devices. Therefore, in-network caching can improve network performance by fetching content from nodes that are geographically placed closer to the end user. Several issues that need further investigation have been identified with respect to in-network caching. In this section, we list 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 to take advantage of cached NDOs. We identify 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 below. 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 most 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 memory,

which increases 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 avoidance of inter-domain/expensive links, centrality of nodes, size of domain and the corresponding spatial locality of users, and 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; therefore, new techniques need to be investigated. In this direction, reducing the redundancy of cached copies is a study item. However, the issue of coordinated content placement in on-path caches 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 relates 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 cache the corresponding contents. This challenge relates to name-based routing, discussed earlier. Requests should ideally follow the path to the cached NDO. However, instructions as to which content is cached where cannot be broadcast throughout the network. Therefore, the knowledge of an NDO location at the time of the request either might not exist or 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 NDO is found on the path, then the content is fetched from this node instead of the origin server, or we employ cache-aware routing techniques. Cache-aware routing can involve either 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 that 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 NDO's name, e.g., the timestamp, that directly indicates its staleness. This approach is 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 greenhouse 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 NDO includes the publisher information, which can be used to reach the publisher. It is suitable for the NDO whose expiration time is difficult to be set in advance, e.g., a web page

that contains the main text (which stays the same ever after) and the interactive sections such as comments or ads (which are 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 all users at the same time, some users would use the old name although they intended 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.

4.7.5. Cache Sharing by Multiple Applications

When ICN is deployed as a general, application-independent network and cache infrastructure, multiple consumers and producers (representing different applications) would communicate over the same infrastructure. With universal naming schemes or sufficiently unique hash-based identifiers, different application could also share identical NDOs in a transparent way.

Depending on the naming, data integrity, and data origin authentication approaches, there may be technical and business challenges to share caches across different applications, for example, content protection, avoiding cache poisoning, ensuring performance isolation, etc. As ICN research matures, these challenges should be addressed more specifically in a dedicated document.

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 that is directly reflected in higher operational cost 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 (Fault, Configuration, Accounting, Performance, and Security) [ISO/IEC-7498-4] management 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 (NOCs) to define policies, validate that they are indeed used in practice, be notified early on about failures, and determine and resolve configuration problems. Authentication, Authorization, and Accounting (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 are preferred. The main challenge here is that all tools we have at our disposal today are node-centric, are end-to-end oriented, or assume 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. Delay Tolerant Networking (DTN) characteristics should also be considered to the degree that this is possible (e.g., content dissemination through data mules). Second, 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 concerted manner.

However, the considerations on leveraging intrinsic ICN mechanisms and capabilities to support management operations go beyond a simple mapping exercise. In fact, it not only 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 network management schemes rely. As such, questions are raised that can directly be translated into challenges for network management capability, such as, for example, how to address a node or a network segment in an ICN naming paradigm, how to identify which node is connected "where", how to be aware of the node capabilities (i.e., high or low-powered machine-to-machine (M2M) node), and if there is a host-centric protocol running where the management process can also leverage.

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 NDOs, one can conceive new management operations addressing, for example, per-content management or access control, as well as analyzing performance per NDO 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] reutilizes 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 and resilience) is also important, as exemplified in [CCN-ACCESS], where access control mechanisms are integrated into a prototype of the [PURSUIT] architecture.

The set of core research challenges for ICN management includes:

- o Management and control of NDO reception at the requestor
- o Coordination of management information exchange and control between ICN nodes and ICN network control points
- o Identification of management and controlling actions and items through information naming
- o Relationship between NDOs and host entities identification, i.e., how to identify a particular link, interface, or flow that needs to be managed

4.9. ICN Applications

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). [RFC7476] 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 over 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 efforts such as [ICN2014-WEB-NDN] have analyzed real-world web applications and ways to implement them in ICN. The most significant insight is that REST-style (Representational State Transfer) web communication relies heavily 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 degree 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 (for example, if larger GET/Interest messages would have to be stored in corresponding tables on routers).

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

4.9.2. Video Streaming and Download

One of ICN's prime application areas 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].

[VIDEO-STREAMING] discusses the opportunities and challenges of implementing today's video services such as DASH-based (Dynamic Adaptive Streaming over HTTP) streaming and download over ICN, considering performance requirements, relationship to peer-to-peer live streaming, IPTV, and Digital Rights Management (DRM).

In addition to just porting today's video application from a host-centric paradigm 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 store and forward could be leveraged for rate adaptation to achieve maximum throughput and optimal Quality of Experience (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 prefetching could be performed in the network by leveraging visibility of chunk NDO names and NDO metadata in the network. Moreover, multi-source rate adaptation in combination with network coding could enable better QoE, for example, in multi-interface/access scenarios where multiple paths from client to upstream caches exist [RFC7476].

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 definition, ICN is well suited for IoT applications, where users consume data generated from IoTs without maintaining secure connections to them. The basic request/response style APIs of ICN enable developers to build IoT applications in a simple and fast manner.

Ongoing efforts such as [ICN-FOR-IOT], [ICN-ARCH], and [ICN2014-NDNWILD] 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, 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, 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. This allows the simplification of ICN-based IoT application development. Since the network is able to act on named content, generic names provide a way to address content independently of the underlying device (and access) technology, and bandwidth consumption is optimized due to the availability of cached content. However, these aspects raise

challenges themselves 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. In addition, due to the heterogeneous nature of IoT nodes, their processing capabilities might not be able to handle the necessary content signing verification procedures.

5. Security Considerations

This document does not impact the security of the Internet. Security questions related to ICN are discussed in Section 4.2.

6. Informative References

[ACCESS-CTL-DEL]

Fotiou, N., Marias, G., and G. Polyzos, "Access control enforcement delegation for information-centric networking architectures", Proceedings of the second edition of the ICN workshop on Information-centric networking (ICN '12) Helsinki, Finland, DOI 10.1145/2342488.2342507, 2012.

[BACKSCATTER]

Waehlich, M., Schmidt, T., 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, DOI 10.1016/j.comnet.2013.07.009, November 2013.

[BREADCRUMBS]

Rosensweig, E. and J. Kurose, "Breadcrumbs: Efficient, Best-Effort Content Location in Cache Networks", In Proceedings of the IEEE INFOCOM 2009, DOI 10.1109/INFOCOM.2009.5062201, April 2009.

[CCN]

Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggs, N., and R. Braynard, "Networking Named Content", CoNEXT 2009, DOI 10.1145/1658939.1658941, December 2009.

[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, DOI 10.1145/2342488.2342507, 2012.

- [CHAUM] Chaum, D. and E. van Heijst, "Group signatures", In Proceedings of EUROCRYPT, DOI 10.1007/3-540-46416-6_22, 1991.
- [COMPACT] Cowen, L., "Compact routing with minimum stretch", In Journal of Algorithms, vol. 38, pp. 170-183, DOI 10.1006/jagm.2000.1134, 2001.
- [CONTUG] Arianfar, S., Nikander, P., Eggert, L., Ott, J., and W. Wong, "ConTug: A Receiver-Driven Transport Protocol for Content-Centric Networks", Technical Report Aalto University Comnet, 2011.
- [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, DOI 10.1145/1282427.1282402, August 2007.
- [ENCRYPTION-AC] Kurihara, J., Uzun, E., and C. Wood, "An Encryption-Based Access Control Framework for Content-Centric Networking", IFIP Networking 2015, Toulouse, France, DOI 10.1109/IFIPNetworking.2015.7145300, September 2015.
- [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, DOI 10.1109/INFCOMW.2014.6849279, May 2014.
- [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, DOI 10.1109/INFCOM.2010.5462131, 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.
- [ICN-ARCH] Zhang, Y., Raychadhuri, D., Grieco, L., Baccelli, E., Burke, J., Ravindran, R., Ed., and G. Wang, "ICN based Architecture for IoT - Requirements and Challenges", Work in Progress, draft-zhang-iot-icn-challenges-02, August 2015.

[ICN-FOR-IOT]

Lindgren, A., Ben Abdesslem, F., Ahlgren, B., Schelen, O., and A. Malik, "Applicability and Tradeoffs of Information-Centric Networking for Efficient IoT", Work in Progress, draft-lindgren-icnrg-efficientiot-03, July 2015.

[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, DOI 10.1145/2660129.2660136, September 2014.

[ICN2014-NDNWILD]

Baccelli, E., Mehlis, C., Hahm, O., Schmidt, T., and M. Waehlich, "Information Centric Networking in the IoT: Experiments with NDN in the Wild", ACM SIGCOMM Information-Centric Networking Conference Paris, France, DOI 10.1145/2660129.2660144, 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, DOI 10.1145/2660129.2660152, September 2014.

[ICNNAMING]

Ghods, 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), DOI 10.1145/2018584.2018586, 2011.

[ICN2015-01-WESTPHAL]

Westphal, C., "Video over ICN", IRTF ICN2015 Meeting Cambridge, Massachusetts, USA, January 2015, <<http://www.ietf.org/proceedings/interim/2015/01/13/icnrg/slides/slides-interim-2015-icnrg-1-0.pptx>>.

[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.

- [ISOIEC-7498-4]
ISO, "Information Processing Systems -- Open Systems Interconnection -- Basic Reference Model -- Part 4: Management Framework", November 1989,
<[http://standards.iso.org/ittf/PubliclyAvailableStandards/s014258_ISO_IEC_7498-4_1989\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s014258_ISO_IEC_7498-4_1989(E).zip)>.
- [MANI] Pentikousis, K. and T. Rautio, "A multiaccess Network of Information", WoWMoM 2010 IEEE,
DOI 10.1109/WOWMOM.2010.5534922, 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,
DOI 10.1145/2018584.2018587, August 2011.
- [MMIN] Pentikousis, K. and P. Bertin, "Mobility management in infrastructure networks", Internet Computing, IEEE vol. 17, no. 5, pp. 74-79, DOI 10.1109/MIC.2013.98, October 2013.
- [NDN-CTL-SHARING]
Yu, Y., "Controlled Sharing of Sensitive Content", IRTF ICNRG Meeting San Francisco, USA, October 2015,
<<https://www.ietf.org/proceedings/interim/2015/10/03/icnrg/slides/slides-interim-2015-icnrg-4-8.pdf>>.
- [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, DOI 10.1109/MCOM.2012.6384449, December 2012.
- [PURSUIT] Fotiou et al., N., "Developing Information Networking Further: From PSIRP to PURSUIT", In Proceedings of Proc. BROADNETS. ICST, DOI 10.1007/978-3-642-30376-0_1, 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,
DOI 10.1016/j.peva.2005.01.002, 2006.
- [RESOURCE-POOL]
Psaras, I., Saino, L., and G. Pavlou, "Revisiting Resource Pooling: The case of In-Network Resource Sharing", ACM HotNets Los Angeles, USA, DOI 10.1145/2670518.2673875, October 2014.

- [RFC2002] Perkins, C., Ed., "IP Mobility Support", RFC 2002, DOI 10.17487/RFC2002, October 1996, <<http://www.rfc-editor.org/info/rfc2002>>.
- [RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", RFC 4838, DOI 10.17487/RFC4838, April 2007, <<http://www.rfc-editor.org/info/rfc4838>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [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, DOI 10.17487/RFC5280, May 2008, <<http://www.rfc-editor.org/info/rfc5280>>.
- [RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, DOI 10.17487/RFC5944, November 2010, <<http://www.rfc-editor.org/info/rfc5944>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, DOI 10.17487/RFC6920, April 2013, <<http://www.rfc-editor.org/info/rfc6920>>.
- [RFC7476] Pentikousis, K., Ed., Ohlman, B., Corujo, D., Boggia, G., Tyson, G., Davies, E., Molinaro, A., and S. Eum, "Information-Centric Networking: Baseline Scenarios", RFC 7476, DOI 10.17487/RFC7476, March 2015, <<http://www.rfc-editor.org/info/rfc7476>>.
- [RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<http://www.rfc-editor.org/info/rfc7696>>.
- [SEEN] Pentikousis, K., "In search of energy-efficient mobile networking", Communications Magazine, IEEE vol. 48 no. 1, pp. 95-103, DOI 10.1109/MCOM.2010.5394036, January 2010.

[VIDEO-STREAMING]

Westphal, C., Ed., Lederer, S., Posch, D., Timmerer, C.,
Azgin, A., Liu, S., Mueller, C., Detti, A., Corujo, D.,
Wang, J., Montpetit, M., Murray, N., Azgin, A., and S.
Liu, "Adaptive Video Streaming over ICN", Work in
Progress, draft-irtf-icnrg-videostreaming-08, April 2016.

Acknowledgments

The authors would like to thank Georgios Karagiannis for providing suggestions on QoS research challenges, Dimitri Papadimitriou for feedback on the routing section, and Joerg Ott and Stephen Farrell for reviewing the whole document.

Authors' Addresses

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

Email: kutscher@neclab.eu

Suyong Eum
Osaka University, School of Information Science and Technology
1-5 Yamadaoka, Suita
Osaka 565-0871
Japan

Phone: +81-6-6879-4571
Email: suyong@ist.osaka-u.ac.jp

Kostas Pentikousis
Travelping
Koernerstr. 7-10
Berlin 10785
Germany

Email: k.pentikousis@travelping.com

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@haw-hamburg.de

Matthias Waehlisch
FU Berlin
Takustr. 9
Berlin 14195
Germany

Email: waehlisch@ieee.org

