

ICN Research Group  
Internet-Draft  
Intended status: Experimental  
Expires: July 12, 2015

A. Lindgren  
F. Ben Abdesslem  
B. Ahlgren  
SICS  
O. Schelen  
Lulea University of Technology  
A. Malik  
Ericsson  
January 8, 2015

**Applicability and Tradeoffs of Information-Centric Networking for  
Efficient IoT  
draft-lindgren-icnrg-efficientiot-02**

**Abstract**

This document outlines the tradeoffs involved in utilizing Information Centric Networking (ICN) for the Internet of Things (IoT) scenarios. It describes the contexts and applications where the IoT would benefit from ICN, and where a host-centric approach would be better. The requirements imposed by the heterogeneous nature of IoT networks are discussed (e.g., in terms of connectivity, power availability, computational and storage capacity). Design choices are then proposed for an IoT architecture to handle these requirements, while providing efficiency and scalability. An objective is to not require any IoT specific changes of the ICN architecture per se, but we do indicate some potential modifications of ICN that would improve efficiency and scalability for IoT and other applications.

This document mainly serves as a problem statement and will not present a conclusive architecture design. It can be used as a basis for further discussion and to design architectures for the IoT.

**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 July 12, 2015.

#### Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

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



## Table of Contents

<a href="#">1.</a>	<a href="#">Motivation . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Advantages of using ICN for IoT . . . . .</a>	<a href="#">5</a>
<a href="#">2.1.</a>	<a href="#">Naming of Devices, Data and Services . . . . .</a>	<a href="#">5</a>
<a href="#">2.2.</a>	<a href="#">Distributed Caching . . . . .</a>	<a href="#">5</a>
<a href="#">2.3.</a>	<a href="#">Decoupling between Sender and Receiver . . . . .</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">Design Challenges of IoT over ICN . . . . .</a>	<a href="#">6</a>
<a href="#">3.1.</a>	<a href="#">Naming of Devices, Data and Services . . . . .</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Efficiency of Distributed Caching . . . . .</a>	<a href="#">7</a>
<a href="#">3.3.</a>	<a href="#">Decoupling between Sender and Receiver . . . . .</a>	<a href="#">8</a>
<a href="#">4.</a>	<a href="#">Proposed Design Choices for IoT over ICN . . . . .</a>	<a href="#">9</a>
<a href="#">4.1.</a>	<a href="#">Relationship to existing Internet protocols . . . . .</a>	<a href="#">9</a>
<a href="#">4.2.</a>	<a href="#">Data naming, format and composition . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.</a>	<a href="#">Immutable atomic data units . . . . .</a>	<a href="#">10</a>
<a href="#">4.4.</a>	<a href="#">Data naming in streams of immutable data units . . . . .</a>	<a href="#">10</a>
<a href="#">4.5.</a>	<a href="#">The importance of time . . . . .</a>	<a href="#">11</a>
<a href="#">4.6.</a>	<a href="#">Decoupling and roles of senders and receivers . . . . .</a>	<a href="#">12</a>
<a href="#">4.7.</a>	<a href="#">Combination of PULL/PUSH model . . . . .</a>	<a href="#">12</a>
<a href="#">4.8.</a>	<a href="#">Capability advertisements . . . . .</a>	<a href="#">13</a>
4.9.	<a href="#">Name-based routing vs name resolution + 1-step vs 2-step . . . . .</a>	<a href="#">14</a>
<a href="#">4.10.</a>	<a href="#">What's naming and what's searching . . . . .</a>	<a href="#">14</a>
<a href="#">4.11.</a>	<a href="#">Tagging/tracing of data, and partial data . . . . .</a>	<a href="#">14</a>
<a href="#">4.12.</a>	<a href="#">Handling actuators in the ICN model . . . . .</a>	<a href="#">15</a>
<a href="#">4.13.</a>	<a href="#">Role of constrained IoT devices as ICN nodes . . . . .</a>	<a href="#">15</a>
<a href="#">5.</a>	<a href="#">Other Issues . . . . .</a>	<a href="#">17</a>
<a href="#">5.1.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">17</a>
<a href="#">5.1.1.</a>	<a href="#">Retrieving trusted content from untrusted caches . . . . .</a>	<a href="#">17</a>
5.1.2.	<a href="#">Enabling application-layer processing in untrusted intermediaries . . . . .</a>	<a href="#">18</a>
<a href="#">5.1.3.</a>	<a href="#">Energy efficiency of cryptographic mechanisms . . . . .</a>	<a href="#">18</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">19</a>



## **1. Motivation**

Information Centric Networking (ICN) has been shown to efficiently meet current usage demands of computer networks, where users consume content from the network instead of communicating with specific hosts. The applications and usage of the Internet of Things (IoT) often imply information centric usage patterns, where users or devices consume IoT generated content from the network instead of communicating with specific hosts or devices.

However, while the IoT shares many characteristics with typical information centric applications, it differs because of the high heterogeneity of connected devices (including sensors and actuators), the very high rate of new information being generated, and the heterogeneity in requirements from applications regarding information retrieval and dynamic actuation. Because of these differences, using an Information Centric Network to design an architecture of the IoT is often, but not always, beneficial. Depending on the context, the IoT architecture may benefit from using an ICN or a host-centric network (HCN). In practice, the right approach is a complex tradeoff that depends on the applications and usage of the IoT network.

This document describes some advantages and inconveniences of using an ICN architecture for the IoT, and helps finding the right tradeoff between using an ICN or an HCN, depending on the context. In this, we explore how to represent and model IoT on top of existing ICN solutions, without requiring IoT specific functionality in the ICN. We discuss this in terms of effectiveness, efficiency and scalability. However, in some cases we do invite for discussion on tentative additions of functionality to ICN in order to make the overall IoT solution more efficient and scalable.



## **2. Advantages of using ICN for IoT**

A key concept of ICN is the ability to name data independently from the current location at which it is stored, which simplifies caching and enables decoupling of sender and receiver. Using ICN to design an architecture for IoT networks potentially provides such advantages compared to using traditional host-centric networks. This section highlights general benefits that ICN could provide to IoT networks.

### **2.1. Naming of Devices, Data and Services**

The heterogeneity of both network equipment deployed and services offered by IoT networks leads to a large variety of data, services and devices. While using a traditional host-centric architecture, only devices or their network interfaces are named at the network level, leaving to the application layer the task to name data and services. In many common applications of IoT networks, data and services are the main goal, and specific communication between two devices is secondary. The network distributes content and provides a service, instead of establishing a communication link between two devices. In this context, data content and services can be provided by several devices, or group of devices, hence naming data and services is often more important than naming the devices.

### **2.2. Distributed Caching**

While caching mechanisms are already used by other types of overlay networks, IoT networks can potentially benefit even more from caching systems, because of their resource constraints. Wireless bandwidth and power supply can be limited for multiple devices sharing a communication channel, and for small mobile devices powered by batteries. In this case, avoiding unnecessary transmissions with IoT devices to retrieve and distribute IoT data to multiple places is important, and storing such content in the network can save wireless bandwidth and battery power. Moreover, as for other types of networks, applications for IoT networks requiring shorter delays can benefit from local caches to reduce delays between content request and delivery.

### **2.3. Decoupling between Sender and Receiver**

IoT devices may be mobile and face intermittent network connectivity. When specific data is requested, such data can often be delivered by ICN without any consistent direct connectivity between devices. Apart from using structured caching systems as described previously, information can also be spread by forwarding data opportunistically.





### **3. Design Challenges of IoT over ICN**

As outlined in [Section 2](#), there are potential benefits from using ICN to implement IoT communication architectures. However, in order to obtain a scalable and efficient architecture there are some aspects of ICN that must be specifically considered in making the right design choices for IoT. In fact, using an ICN may not be beneficial in all desired sub-functions and scenarios. This section outlines some of the ICN specific challenges that must be considered and describes some of the trade offs that will be involved. We will address these challenges in our proposed design choices later in [Section 4](#).

#### **3.1. Naming of Devices, Data and Services**

Naming devices is a common element of both ICN and host-centric approaches. However, naming devices in the IoT raises different challenges that have to be addressed if an ICN approach is adopted. The ICN approach also provides a means to name data and services independently of the location they are hosted at. Whether they are named in the network layer or the layers above it depends on the specific ICN protocol in question. In any case naming devices, data and services should be done carefully depending on the context.

- o Naming of devices: Naming devices is often important in an IoT network. The presence of actuators requires clients to act specifically on a device, e.g. to switch it off. Also, managing and monitoring the devices for administration purposes requires devices to have a specific name allowing to identify them uniquely. There are multiple ways to achieve device naming, even in systems that are data centric by nature. For example, in systems that are addressable or searchable based on metadata or sensor content, the device identifier can be included as a special kind of metadata or sensor reading.
- o Size of data/service name: In information centric applications, the size of the data is often larger than its name. For the IoT, sensors and actuators are very common, and they can generate or use data as small as a short integer containing a temperature value, or a one-byte instruction to switch off an actuator. The name of the content for each of these pieces of data has to uniquely identify the content. For this reason, many existing naming schemes have long names that are likely to be longer than the actual data content for many types of IoT applications. Furthermore, naming schemes that have self certifying properties (e.g., by creating the name based on a hash of the content), suffer from the problem that the object can only be requested when the object has been created and the content is already known, thus



requiring some form of indexing service. While this is an acceptable overhead for larger data objects, it is infeasible for use when the object size is on the order of a few bytes.

- o Hash-based content name: Hash algorithms are commonly used to name content in order to verify that the content is the one requested. This is only possible in contexts where the requested object is already existing, and where there is a directory service to look up names. This approach is suitable for systems with large data objects where it is important to verify the content.
- o Metadata-based content name: Relying on metadata allows to generate a name for an object before it is created. However this mechanism requires metadata matching semantics.
- o Naming of services: Similarly to naming of devices or data, services can be referred to with a unique identifier, provided by a specific device or by someone assigned by a central authority as the service provider. It can however also be a service provided by anyone meeting some certain metadata conditions. Example of services include content retrieval, that takes a content name/description as input and returns the value of that content, and actuation, that takes an actuation command as input and possibly returns a status code afterwards.

### **3.2. Efficiency of Distributed Caching**

Distributed caching is a key opportunity with ICN. However, an IoT framework must be carefully designed to reap the maximum benefits of ICN caching. When content popularity is heterogeneous, some content is often requested repeatedly. In that case, the network can benefit from caching. Another case where caching would be beneficial is when devices with low duty cycle are present in the network and when access to the cloud infrastructure is limited.

However, using distributed caching mechanisms in the network is not useful when each object is only requested at most once, as a cache hit can only occur for the second request and later. It may also be less useful to have caches distributed throughout ICN nodes in cases when there are overlays of distributed repositories, e.g., a cloud or a Content Distribution Network (CDN), from which all clients can retrieve the data. Using ICN to retrieve data from such services is beneficial, but in case of dense occurrence of overlay CDN servers the additional benefit of caching in ICN nodes would be lower. Another example is when the name of the data has a different meaning depending on the context, or if the name refers to an object with variable content/state. For example, when the last value for a sensor reading is requested, the returned data should change every



time the sensor reading is updated. In that case, ICN caching may increase the risk that cache inconsistencies result in old data being returned.

### **3.3. Decoupling between Sender and Receiver**

Decoupling the sender and receiver is useful mechanism offered by the ICN approach, especially for content retrieval with duty cycling devices or devices with intermittent connectivity. However, in order to efficiently retrieve data it must be possible for requestors (receivers) to easily deduce the name of the data to request, without any direct contact with the responder (sender).

Nevertheless, de-coupling is a challenge when authentication is needed for management and actuation, or when real-time interaction between devices is necessary. Solutions for object security supporting decoupled authentication (e.g., similar to signing by proxy), and solutions for pushing data to decoupled entities must be explored.



#### **4. Proposed Design Choices for IoT over ICN**

This section describes some fundamental design choices and trade-offs to allow for effective, efficient and scalable handling of IoT data in an ICN network. An objective with these choices is to facilitate that an ICN network can be used without requiring additions of IoT application specific functionality in the ICN network. However, in some cases we do invite for discussion on tentative additions of functionality to ICN in order to make the overall IoT solution more efficient and scalable.

##### **4.1. Relationship to existing Internet protocols**

IoT devices can have a role as content generators (e.g., sensors) in where an ICN paradigm should be effective for data retrieval and dissemination. However, IoT devices may also have roles as actuators in which such devices shall be accessed for control purposes. The use of an ICN network may be less natural when actuation and control of specific devices is the key objective. As ICN networks are likely to coexist with existing Internet protocols in most situation, often being deployed as overlay networks, we will consider that there may be situations where a host centric addressing is more suitable for IoT. Thus, to facilitate support of IoT for both data generation and control/actuation, we assume that ICN routing should therefore work in concert with existing Internet protocols. However, we will also investigate the possibility of utilizing ICN network primitives for actuation as well to see what the tradeoffs are, as can be seen in [Section 4.12](#).

##### **4.2. Data naming, format and composition**

The data served by ICN may be aggregated from smaller components. Although IoT data components in many cases are small and simple, a general challenge in defining ICN applications is to decide how to compose (i.e. group) the data so that it can be effectively named and requested. Requesting partial data inside a composition may become a challenge. Indeed, if data is composed and sub components are requested, which are not directly namable by the requestor, finding such a subset will resemble a database query which may require processing to resolve. The ICN network should not have to support such complexity.

A design choice regarding IoT data is therefore to keep the ICN network free from supporting any advanced queries and instead only support directly addressable (i.e., named) data units. Any advanced composition (hierarchical, graph-based, hyperlink, etc.) of IoT data, and related searching for sub-components, would be handled in servers/endpoints instead of inside the ICN network. The issue of





structure and searching is for further study. For effective ICN interoperability, only the structure of the atomic addressable data units must be agreed. There are several advantages of this design choice. First, the size of the directly addressable units can be kept fairly small to avoid that unwanted bulk data is pulled over resource constrained networks or spread over various caches in the ICN network. This results in better resource utilization, better localization of desired data, and ultimately better scalability. There is however one tradeoff in that smaller data units results in a larger overhead. Second, the computational requirement is kept low in the ICN network, essentially limiting it to deciding whether there is a cache hit or not. Third, few new requirements are put on ICN data dissemination. Existing methods will be sufficient. Fourth, this simplification means that a flat address space would be sufficient, but for practical reasons a hierarchical address space may be preferred. There is flexibility in the choice of exact addressing scheme and it may depend on which existing ICN framework that is used for IoT data.

#### **4.3. Immutable atomic data units**

The number of IoT devices as well as the amount of data produced by these devices may potentially be very large, and the data may be spread over very large ICN networks. The potential problem of cache inconsistencies in an ICN network may therefore be large if we allow for data to be mutable objects. To support scalability and horizontal distribution it is essential to define data properties that facilitate independency and consistency, while minimizing the need for dynamic global synchronization.

A key design choice is therefore to mandate that IoT only uses immutable atomic data units. This supports large scale distribution by ensuring that there is no stale data in the ICN domain. A hit is always a clean hit. A trade-off from this is that dynamic data must be modeled as a stream of immutable data units, potentially consuming more resources. However, this challenge can be resolved by smart caching strategies where old data is dropped.

#### **4.4. Data naming in streams of immutable data units**

To support immutable streamed data efficiently, we recommend that names of data can include a sequence number. When data can be named with sequence number, any request may or may not include such a sequence number. If no number is included in the request, the nearest cache hit will result in a response. If a sequence number is included in the request, only an exact cache match will result in a response. A client that wants the "latest" reading can according to our previously mentioned design choice, in [Section 4.2](#), not ask the



ICN network such a high level query, instead it must ask for the specific (version of) information. To avoid complicated searching in the ICN nodes, there is intentionally no way to ask the network for the "latest" reading, or any other "range" of sequence numbers.

Should a client want the latest reading from a sensor, a method for this is to make a subscription for the pushed stream of data, as described in [Section 4.7](#). The confirmation of that subscription will contain the latest reading, and then obviously the normal stream will be received. The reason for including the latest reading in the response is to immediately provide the "state" of sensors that generate new data infrequently.

It is for further study whether any extensions are needed to the ICN paradigm to support sequence numbers as part of naming, or if it with some tradeoffs can be supported with, e.g. clever use of metadata, namespace, and search functionality. It may also depend on the particular flavor of ICN. The naming scheme of CCN/NDN may here provide an advantage.

#### **[4.5](#). The importance of time**

We want to emphasize that time almost always is a very important property of IoT data, and especially so for data that change over time. When modeling dynamic IoT data with a stream of immutable data, it is often the case that a certain IoT data object is a sensor reading at a particular point in time, and the next object in the stream is the next reading. Thus, dynamic data is in this case dynamic over time, with well defined (immutable) values for particular points in time. We therefore argue that it is important to find a way to represent these time-related streams of immutable data. It should be possible to request data from a certain time, and to infer/find the name of the latest, most current, data.

There are several methods for finding readings from a certain time, or the latest reading, for example through a high level request from a server/endpoint, or by using a naming scheme where the name can be directly inferred, e.g., if an IoT device has advertised under which conditions it produces data and how it is named.

To represent absolute time so that it can be directly inferred, one method is that the producer of data in its capability advertisements ([Section 4.8](#)) provide a mapping function between sequence number and time. Thereby also readings on the time axis are immutable while it is still possible to efficiently find the latest reading, as described in [Section 4.4](#). It should be noted that sequence numbers then may have gaps in order to cater for triggered non periodic data, etc. In addition, meta data may include information on absolute



time. Using this mapping scheme data from the current second can be efficiently requested (provided that clock synchronisation is accurate enough, which is out of scope of this document). These methods are based on the request/pull method. For continuous real-time update (most accurate info), there is an option to use dissemination based on the push model as described later in [Section 4.7](#).

We also note that time is also important for other applications, in particular for live streaming video. Live video also produces a time-related stream of immutable objects, and would in the same way benefit from such support in the ICN service.

#### **[4.6. Decoupling and roles of senders and receivers](#)**

Since ICN networks essentially support a request/response model of interaction, we denote the receivers of information as requestors, and the senders of information as responders. The ICN network in itself provides decoupling of requestors and responders, but it does not (and should not) provide any transformation or aggregation of data. The IoT dissemination architecture should therefore allow for any number of intermediate processing nodes. An intermediate node will be an endpoint in the ICN network that can act as both requestor and responder. Such a node may perform aggregation, filtering, selection, etc. The instantiation of such nodes may for example form a directed (acyclic) graph between ultimate responders (IoT devices) and ultimate requestors (the final applications). It is for further study how to define such an architecture.

It is a design choice to keep the IoT dissemination and aggregation functionality outside of the ICN domain. That architecture would be an overlay that may have intricate structure, and put the ICN usage in a new context, where content from ultimate requestors to ultimate responders may go through many IoT processing nodes that collect, process and re-publish data through an ICN for various purposes.

#### **[4.7. Combination of PULL/PUSH model](#)**

A critical decision regarding IoT data is whether to use a PULL model, a PUSH model, or both. In this document, we define a PULL model as a system where data is only sent when explicitly requested, while a PUSH model indicate that data transmission is initiated by the source based on some trigger (either periodic, for each new object, or based on some condition on the generated data). There are some intrinsic trade offs between these models. The PULL model is for example resource efficient when there is an abundant amount of IoT information, potentially redundant from many devices, and the clients only occasionally or partially are interested in the



information. The PUSH model is for example efficient when there is real-time information and the clients are interested in all information from specific devices all the time.

A design decision in the IoT domain is to support both PULL and PUSH. The base model should be PULL, meaning that requestors must always start by sending a request. If the request is for some specific data, it can be resolved by returning the data (if it exists). The pull model can be supported efficiently and scalably by an ICN network. A request can however also include triggers, which means that data will be returned (pushed) when triggers are fulfilled, which may be immediately, or in the future at one or several occasions. This can be used to select alarm conditions, to request continuous or periodic push, etc. The trigger conditions can be set by the requestor, or be pre-defined by the responder. The former is more flexible but may have performance/scalability issues. The latter is more scalable since there will be a predefined and finite number of trigger conditions. Our recommended choice, at least for the initial phase, is to go for a simple and scalable solution and therefore adopt the model where available trigger conditions are defined and advertised by the responder. The ICN would be apt for supporting such capability advertisements, given that they are fairly static.

We recommend to have a discussion on whether an ICN network can or should provide an option to effectively support a push model of data. Such support can make real-time IoT data dissemination more efficient and scalable as previously mentioned in [Section 4.3](#). However, since we assume that the ICN works with existing IP protocols, such functionality can be provided without ICN, by using traditional unicast or multicast communication. We finally note that an ICN supported push service model would make the ICN network more like a publish/subscribe system.

#### **[4.8](#). Capability advertisements**

Capability advertisements and discovery can be used by requestors to discover which responders to connect to. In a deployment with large numbers of responders, the functionality of automatic advertisement and discovery becomes a critical factor to support scaling. Responders should advertise their methods (inputs, outputs, parameters, triggers, etc) and provide relevant metadata. Such capability advertisements should be conservative with resources, which suggests that new advertisements should be posted with reasonably low frequency. This implies that an ICN network can be used for providing capability advertisements. The advertisements should be provided as a stream of immutable objects, or alternatively the IoT system should be tolerant to stale caches. Should there be a





need real-time awareness of dynamic changes, a subscription/push model of capability advertisements could be used as earlier described in [Section 4.7](#).

#### **[4.9](#). Name-based routing vs name resolution + 1-step vs 2-step**

As described in [Section 4.2](#), the IoT framework should be defined so that new functionality in the ICN is not needed. For data that is frequently generated and regenerated, it makes sense to keep simple structures and provide directly inferable naming/addressing of data objects, so that requestors can directly address the data. For more complex data, such as pre-processed, aggregated and structured data a two-step resolution model is recommended. The IoT devices can provide a higher level resolution based on for example queries and searching, resulting in a number of concrete directly addressable ICN objects. This is similar to what web servers do when they return URLs that requestors can use, but in this case it is named content that is returned.

Consequently, the IoT framework should have no requirement that the ICN network itself should support 2-step addressing (although such 2-step methods may exist in some ICNs)

#### **[4.10](#). What's naming and what's searching**

As described in [Section 4.2](#), the IoT framework should be defined so that no new functionality is required in the ICN for searching data or subcomponents of data. The ICN network supports just naming of atomic data objects, while any searching is provided by the IoT framework, which in itself may be constituted by a highly distributed set of nodes that provide processing, analysis and aggregation of IoT data.

#### **[4.11](#). Tagging/tracing of data, and partial data**

IoT data may be tagged with metadata to tell where it originates from. Tagging is made at the level above the ICN network and may for example be a list of strings. It can be added/changed by the originating node (or a node that assigns the originating ID), and added/changed/deleted by any node that processes the data. The tag can in some cases be used to trace data back to origins. For the ICN network, the metadata units are just black-box data that is to be conveyed, and therefore are not to modify the tags. However, in some cases it makes no sense to transmit any metadata. For efficiency reasons the ICN network should have support for optional delivery of metadata. This is to be conservative with scarce resources, for example when a wireless node requests data which is cached in the ICN network, it would be beneficial if the requestor could tell that it



is desirable to not receive any metadata. There should be a discussion whether there should be just one, or more than one, piece of optional information in ICN content to be future proof.

#### **4.12. Handling actuators in the ICN model**

If actuators should be controlled using the ICN communication model, we need to map the functionality of the actuator to named data and/or the requesting of named data. We see two main models with some variants as described in the following paragraphs.

In the first model, the state of the actuator is represented by a named data object. The actuator periodically requests its state using the name of its designated state object. There then has to be a producer of that state data that responds with the current state. When the actuator receives the response, it sets that new state, invoking its actuation function.

A variant of this first model is that a requester first requests the state of the actuator. The requester supplies additional information with the request including the name of the new state data it will produce. The actuator responds with its state, and then requests its new state using the name that was supplied with the additional information in the first request.

In the second model, the actuator invokes its actuation function as a side-effect of receiving a particular request. There are several plausible models. The new state could be encoded in the name of the requested data in the request, or could be supplied as additional information with the request. Regardless, the actuator acts on the new state information as a side effect, and responds with data, possibly its state, to the requester.

To reap the advantages of caching, it should be possible to cache the state of the actuator in both the aforementioned models. However, we think that caching is not as relevant for actuation as it is for other IoT use cases.

#### **4.13. Role of constrained IoT devices as ICN nodes**

Typical ICN nodes such as routers and gateways are deemed to be rich in resources like energy, processing, bandwidth and storage. IoT devices, on the other hand, are quite constrained in such resources. It is also worth noticing that some resources are more crucial than others. In most cases energy, processing and bandwidth are quite expensive for constrained IoT devices. In contrast, storage has shown a considerably rapid decreasing trend in prices over the past few years. A similar trend is expected in the future with increasing



ubiquity of cloud networks and information-centric networks.

A design decision in this regard is that we logically separate IoT functionality (such as sensing and transmitting IoT data) and ICN functionality (such as routing and caching data generated by other devices). A resource constrained device may choose to only implement IoT functionality and not act as intermediate ICN nodes. However, since storage is not as expensive as other resources, IoT devices should be able to cache their own content and, in essence, act as sources to ICN.

## **5. Other Issues**

### **5.1. Security Considerations**

The ICN paradigm is content-centric as opposed to state-of-the-art host-centric internet. Besides aspects like naming, content retrieval and caching this also has security implications. ICN advocates the model of trust in content rather than trust in network hosts. This brings in the concept of Object Security which is contrary to session-based security mechanisms such as TLS/DTLS prevalent in the current host-centric internet.

Object Security is based on the idea of securing information objects unlike session-based security mechanisms which secure the communication channel between a pair of nodes. This reinforces an inherent characteristic of ICN networks i.e. to decouple senders and receivers. In the context of IoT, the Object Security model has several concrete advantages. As discussed earlier in [Section 2.1](#), in many IoT applications data and services are the main goal and specific communication between two devices is secondary. Therefore it makes more sense to secure IoT objects instead of securing the session between communicating endpoints.

It is important that while security mechanisms complement the ICN architecture in a coherent fashion, they do so without laying down any strict requirements or constraints. Therefore, the decision of what security mechanisms are employed should be handled at a layer above ICN, in this case within the IoT framework. This facilitates flexibility and allows IoT applications more freedom to decide what security mechanism suits them best (session-based security, object security or a hybrid). Though the idea of Object Security is very much inline with the ICN concept, there can still be some use cases where Object Security does not add much e.g. a Pub/Sub interaction where a client is expected to interact more or less with the same server node (a session-based security protocol should suffice here) or use cases where application layer headers should also be secured (which can be achieved by TLS/DTLS). We, therefore, effectively imply that there is no need to modify typical ICN standards to accommodate Object Security.

The following sub-sections discuss some advantages of using Object Security in IoT applications.

#### **5.1.1. Retrieving trusted content from untrusted caches**

When functioning in an ICN network, an IoT client is expected to rely on the network to deliver the requested content in an optimal fashion without concerning itself with where the content actually lies. This





could potentially mean that each individual object within a stream of immutable objects is retrieved from a different source. Having a trust relationship with each of these different sources is not realistic. This gives rise to the need of retrieving trusted content from untrusted nodes/caches in an ICN network. Object security is ideal in such use cases because it relieves an IoT client application from the hassle of having to establish trust with each node that can potentially cache an IoT object. This also means that a requesting client can make use of more caches in the network, hence resulting in better throughput and latency.

#### **5.1.2. Enabling application-layer processing in untrusted intermediaries**

Object Security ensures that objects in application-layer payload are secure e.g. XML, JSON objects. However, the application-layer header is unencrypted and available for processing. Securing content at the object level means greater granularity. This facilitates application-layer processing in untrusted intermediary nodes (e.g. proxies and caches) without compromising security. An example use case is untrusted caching nodes that should have the ability to cache individual encrypted objects without being able to see what is there in those objects. In this case there is a need for the caching nodes to identify the object URI which can be done by looking into the application-layer header. But the object is still encrypted and unknown to the caching nodes.

#### **5.1.3. Energy efficiency of cryptographic mechanisms**

Session-based security protocols rely on the exchange of several messages before a secure session is established between a pair of nodes. Use of such protocols in constrained IoT devices can have serious consequences in terms of power efficiency because in most cases transmission and reception of messages is more costly than the cryptographic operations. This is especially true for wireless devices.

The problem is amplified proportionally with the number of nodes the constrained device has to interact with because a secure session would have to be established with every node. If the constrained device is acting as a consumer of data this would mean setting up secure sessions with every caching node that the device retrieves data from. When acting as a producer of data the constrained device would have to setup secure sessions with all the consumers. The Object Security model eliminates this problem because the content is readily available in a secure state in the network. IoT devices producing data can secure it w.r.t. all the intended consumers and start transmitting it right away.



## Authors' Addresses

Anders F. Lindgren  
SICS Swedish ICT  
Box 1263  
Kista SE-164 29  
SE

Phone: +46707177269  
Email: andersl@sics.se  
URI: <http://www.sics.se/~andersl>

Fehmi Ben Abdesslem  
SICS Swedish ICT  
Box 1263  
Kista SE-164 29  
SE

Phone: +46705470642  
Email: fehmi@sics.se  
URI: <http://www.sics.se/~fehmi>

Bengt Ahlgren  
SICS Swedish ICT  
Box 1263  
Kista SE-164 29  
SE

Phone: +46703141562  
Email: bengta@sics.se  
URI: <http://www.sics.se/people/bengt-ahlgren>

Olov Schelen  
Lulea University of Technology  
Lulea SE-971 87  
SE

Phone:  
Email: olov.schelen@ltu.se  
URI:



Adeel Mohammad Malik  
Ericsson  
Kista SE-164 80  
SE

Phone: +46725074492  
Email: adeel.mohammad.malik@ericsson.com  
URI: