

Opportunities and Challenges for Named Data Networking to Increase the Agility of Military Coalitions

Christopher Gibson,* Pablo Bermell-Garcia,† Kevin Chan,‡ Bongjun Ko,§ Alex Afanasyev,¶ and Lixia Zhang¶

*IBM UK

†Airbus Group UK

‡U.S. Army Research Laboratory(kevin.s.chan.civ@mail.mil)

§IBM T. J. Watson Research Center

¶University of California, Los Angeles (aa,lixia@cs.ucla.edu)

Abstract—The fundamental aim of this paper is to position the opportunities and challenges for adopting Named Data Networking (NDN) in the specific context of military coalition operations and tactical networks. The characteristic properties of tactical networks include high dynamics in multiple dimensions: bandwidth, network congestion, frequent topological changes, geographical mobility of assets, as well as dynamic changes in information access policies. Furthermore, coalition networks must provide secure and efficient communication across coalition boundaries and mitigate the impact of adversarial entities attempting to obstruct the mission. In this paper, we elaborate on the basic NDN architecture characteristics, including robust data discovery and retrieval over ad hoc and intermittent connectivity, inherent security, efficient content distribution, and automatic in-network caching; we also articulate how the above properties can all be utilized to enable resilient and secure data collection, improve the analytics capacity of the network, and to speed up and improve the quality of distributed decision making in challenging coalition environments.

I. INTRODUCTION

The success of military coalitions depends, to a large extent, on the ability to exchange relevant information between their assets in a secure and timely manner. While digitalisation in recent years has undoubtedly increased our ability to capture battlefield information in multimedia forms with fine granularity, it has also raised the requirements for collecting, processing, and analyzing data to drive the most informed decision making available and to share information with coalition partners under strict policy control. Given that military forces operate in the very challenging network environments of the battlefield, the federated nature of coalitions brings the additional challenges of controlling information flow across various boundaries, including problems of information discovery, sharing, and managing access privileges to individual data pieces.

Although today’s operational TCP/IP networking architecture, developed 40 years ago, still works effectively in meeting society’s needs for large-scale data dissemination over well engineered network infrastructures, it has shown its age in meeting the specific requirements of highly dynamic tactical networks, particularly security. What is needed are new

technologies and architectures that can exploit the increased capacities in data capture and processing to increase the agility of distributed analytics in coalition operations.

We posit that the Named Data Networking (NDN) design [1] can help achieve this goal effectively. Initially funded under the U.S. National Science Foundation’s Future Internet Architecture Program, the NDN architecture changes the level of abstraction in networking. In NDN, the focus of networking is shifted from forwarding packets between adjacent nodes identified by IP addresses to directly accessing secured data by name. This effectively moves the web semantics—fetching data by name (URLs)—from the application layer to the network layer. This fundamental change in communication abstraction decouples information from its containers and, more importantly, removes the information access dependency on the lower layer naming (e.g., IP addresses) and the information security dependency on the lower layer transport (e.g., IPsec and TLS channels). Named and secured data become the focus of all network operations, allowing the network to utilize all available resources to satisfy applications’ needs.

More specifically, because all communicated data in NDN are identified by hierarchically structured and semantically meaningful names and are secured at the network layer, communication can happen over any available media (potentially in parallel). An application can directly request data by name, without requiring any translation from the name to end point addresses. Regardless how or from where the information is retrieved, it can always be validated. Moreover, all information is secured (signed and, if needed, encrypted) not only during transport across the network, but also when it is stored, e.g., at in-network caches.

All of the above properties are critically important in challenged network environments and to the mission of military coalition operations. To make our case, below we first present a motivating scenario to articulate the required system properties and functions to be supported for battlefield information analytics. We then highlight the technical underpinning that makes NDN the most suitable network architecture candidate to meet the challenges in coalition missions, as well as highlight




Fig. 1. An example scenario of battlefield information capture and collection from a combination of statically placed and mobile multimedia sensors.

research directions to bring the NDN vision into reality.

II. MOTIVATING SCENARIO

To illustrate the communication requirements and to identify the challenges in providing secure, resilient data collection and analytics in military coalition environments, we use a field trial scenario (Figure 1), sponsored by the Coalition Warfare Program (CWP) of the Office of the Secretary of Defense (OSD) [2]. The goal of this project was to demonstrate the seamless integration of a disparate set of coalition assets into a single data-to-decision (D2D) solution. The trial aimed to decouple the assets from their original bespoke systems, and use a microservice architecture to rapidly assembly and deploy a flexible integrated solution with highly granular, policy-controlled data sharing and dissemination.

The scenario describes a wide-area surveillance system that delivers imagery of sites of interest around a fixed location, specifically the sources of acoustic events around a military base. The coalition assets consist of a set of acoustic unattended ground sensors (UGS), a fixed camera with pan-tilt-zoom functionality, a mobile camera mounted on an autonomous unmanned aerial vehicle (UAV), a second set of low-power acoustic UGS, and several software services providing analytics functionality. Each of the sensor assets, with the exception of the low-power UGS, connects wirelessly to the base network. The data from the low-power UGS are relayed via a second UAV tasked to fly between the sensors harvesting their telemetry.

The assets are collectively provided by three coalition partners, each with varying sharing policies. In operation the acoustic sensors publish lines of bearing (LOB) to the sources of acoustic events around the base. The LOBs are discovered and consumed by an analytics service that localizes and publishes the location of each event. Additional services consume this location information to cue the fixed camera and task the UAV to deliver imagery of the source.

This scenario raises the following requirements to the supporting communication systems:

- **Information Discovery:** The system must support automatic information discovery of disparate information sources deployed across networks of different communication technologies. This is challenging because rapid deployments and mobile ad hoc connectivity (e.g., the UAV flying between sensors to harvest their telemetry) are in sharp contrast with the established operational practice of IP which relies on pre-planning, a stable environment, and configuration.
- **Composition of Services:** The system must provide support for automatic composition of ad hoc distributed analytics services, which requires the discovery and best placement of processing functions to be done in a secure and decentralized fashion, to fully and efficiently utilize all available resources. It is difficult to achieve this goal in an IP network because neither the functions nor the data to be processed are visible at the network layer

which possesses the detailed knowledge about network resources.

- **Cross-Boundary Data Sharing:** To support distributed data sharing among coalition partners according to defined policies and coalition command and control (including cross-platform tipping and cueing of the sensing devices), the system must ensure reliable and decentralized authentication of partner devices and an effective means to deliver data in the absence of established secure channels (e.g., data may have to be carried by data mules). Such a requirement is again in sharp contrast with today’s practice of centralized (cloud-based) authentication and authorization solutions.
- **Policy Maintenance:** Data sharing policies may change over time, therefore the system must provide mechanisms to discover and disseminate the up-to-date policies. Conventional means of policy updates yet again rely on consulting centralized authority servers, connectivity to which may not be available at all times.

III. NDN AS THE ENABLER

Since a network’s basic task is to deliver packets, a fundamental question is how each packet is labeled. An IP network sees the world being made of channels connecting up individual nodes, and labels packets with IP addresses that name the locations/nodes where the packets should be delivered. An NDN network views any and all physical connectivities, as well as storage and processing capabilities, as a means to supply requested data, so it labels a packet with the data’s name, and secures the authenticity and integrity of the content of the packet by a cryptographic signature. This new network abstraction opens profound potential to realize resilient and secure data sharing and analytics in challenging coalition environments as we elaborate below.

An analytics application’s need to acquire field data (e.g., video frames at a specific angle, the identification of objects in an image, or acoustic signals at a specific position, etc.) can be expressed, at the network layer, as a sequence of NDN “Interest packets”, each specifying the name of the desired “Data packet” that may either already exist (e.g., chunks of a captured video frame), or can be produced in response to an Interest (e.g., video frame analysis to identify objects).

Forwarding functions are performed on all NDN-enabled nodes to route Interest packets towards the source of the data. Based on the data name carried in an Interest packet, a forwarder will send it toward the closest location where the Data may reside or can be produced; it then records the incoming and outgoing interfaces of this Interest packet in its *Pending Interest Table (PIT)* (Figure 2). When multiple requests for the same data are received, the forwarder forwards the first one and simply records the incoming interfaces for the rest. This state in the PIT is used to deliver the returned Data packet, allowing the Data packet to traverse the Interest path in reverse to get back to the original consumer(s). The Data packet can also be cached in the forwarder’s *Content Store (CS)* to satisfy future requests for the same data. With this

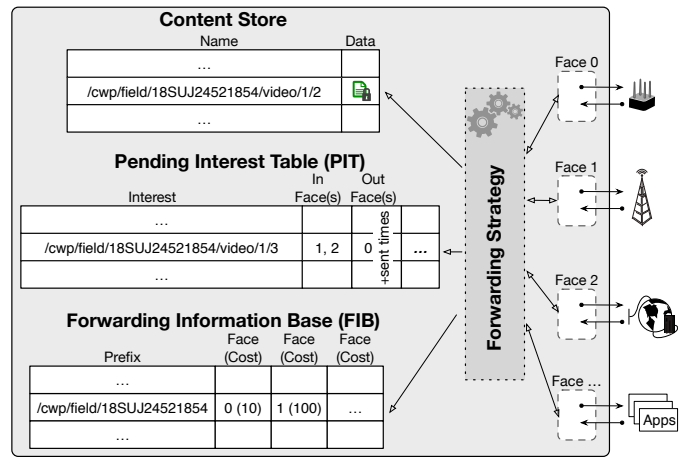


Fig. 2. NDN Forwarder Model

design, NDN natively supports multicast delivery and scalable content distribution.

Maintaining the PIT makes NDN’s data plane stateful, which is explored by the *forwarding strategy* module residing at each NDN node. The recorded state ensures that Interest packets cannot loop in the network and, therefore, the forwarding strategy on each NDN node is free to choose any available interface to forward the Interest. The closed loop between the expressed Interest and the corresponding Data packet (or lack of it) gives a per-packet assessment of the selected forwarding decision. This assessment, coupled with the information from the lower layers (e.g., signal quality/congestion level of a particular interface) and higher-level policies (e.g., desire to prioritize retrieval of data with certain given prefixes), can be used to adjust the decision of where the subsequent Interests will be forwarded.

This intelligence of forwarding strategy enables *information discovery* within a local environment: a wireless channel is broadcast by nature, all nodes within one hop can hear an expressed Interest packet; a node can respond if it has the requested data, or help further forward the request (with a specified hop limit). If the requested data is located in the vicinity, the returned Data packet informs the forwarders of the successful path, so that these forwarders need only broadcast the first Interest; all the future Interests under the same name prefix will be guided by the retrieval path of the first Data packet.

To assist efficient Interest forwarding at large scale, data producers can announce their name prefixes via a routing protocol, which constructs the Forwarding Information Base (FIB) with a set of optimal (shortest or fastest) paths to retrieve Data for each announced prefix. Note that routing announcements are themselves named, secured data that can be verified by the receivers, therefore an NDN network has the routing security automatically built in.

To ensure the authenticity and integrity of the Data regardless of how it may be delivered to consumers, producers cryptographically sign each Data packet, binding its name to

the content. The name of the signing key is included in the Data packet in the dedicated “KeyLocator” field, and the key itself is simply another Data packet. Therefore, each Data packet is associated with a chain of certificates. Whenever needed, Data packets (as well as lower parts of the names) can be encrypted to provide further secrecy properties. The link layer over which NDN packets are communicated can also be secured, providing an additional level of confidentiality and communication privacy.

A. Information Discovery via Named Data

Information naming in NDN is defined by applications, independently from network connectivity (or absence of it) and the underlying communication technologies. Therefore all nodes and applications can communicate via any channel as soon as a channel comes into existence. This makes information discovery in NDN a straightforward operation: the requesting application expresses an Interest packet with the name of required piece of information, the network takes responsibility for fetching the requested information.

However, before an analytic application can request video and acoustic sensor information, for example, it must be able to determine how the information is named. The application can obtain this knowledge by leveraging several mechanisms.

1) *Naming Conventions*: Constructing a well-defined set of naming conventions can provide consumers in a given context the ability to construct names of required data. For example, the “./video/1/1” convention can define the format for raw data frames at specific observation angles and “./SLAM/(f=./video/1)” can define the processing results of feature extraction and annotation. In other words, one can leverage NDN’s hierarchical naming to organize data in a structured way, enabling analytics applications to effectively discover, retrieve, and share sensed, captured, or processed data, as well as to request the production of new data. The dimensions of the namespace design in the coalition environment can reflect the nature of the captured information (location, position, angle) as well as additional meta-parameters of the capture itself (e.g., velocity of the camera during the capture, wind speed, temperature, humidity of the environment, etc.). Given a namespace, the data can be retrieved from the original capture device directly, from nearby devices, or from dedicated storage servers deployed by the mission (centralized or ad hoc).

2) *In-Network Name Discovery*: NDN supports information retrieval using partial names (name prefixes). Analytics applications are assumed to know the prefixes of the data context they require, such as location context, via a well-known name “/cwp/field/18SUJ24521854”, but may not know the name of the latest instance of the data (reflected in the version number or timestamp in the name). When an application uses a given prefix to retrieve a piece of data which carries its full name, the application can then follow established naming conventions to construct the full name of desired Data packets, or further refine its request.

3) *Metadata*: It is possible to devise a metadata-based mechanism to aid name discovery, remotely resembling the

distributed filesystem structure: each level of the namespace is associated with metadata that contains information regarding the names/content below it. In this way, given a namespace prefix (“/cwp/field”), the available data names under this prefix (e.g., “18SUJ24521854”, “11SLT66787062”, etc.) can be learned by requesting the metadata using an established naming convention for metadata.

B. Seamless Integration with Processing via Named Data

Data names can identify not only the existing data, but also data to be produced, effectively integrating network connectivity, storage, and processing engines under the same architectural umbrella. Using names as the enabling technique, NDN offers a simple yet elegant paradigm of *named functions* to solve the problem of rendezvous between data and analytics engines [3].

Distributed analytics in coalition networks can be viewed as a transient step in data fetching: computing engines (“processing producers”) can act as consumers for the captured input data. Through the naming rules computing engines can automatically derive the name for the final data, whilst the computing engines themselves can announce their processing capabilities. For example, UAV and HMMWV nodes may advertise that they have simultaneous localization and mapping (SLAM) processing capabilities by announcing (e.g., “/FUN:/SLAM”) prefixes. An analytics application can request localization and mapping over the captured data by expressing Interest that includes a function’s prefix and arguments of data names to be processed. The forwarding strategy module of NDN will ensure that these Interests are forwarded to the closest or the most available/highest fidelity engine, or the engine that is closest to the processed data. For example, the engine can evaluate the received Interests, retrieve the required raw data packets (e.g., “/cwp/field/18SUJ24521854/video/1/1”, ..., “/cwp/field/18SUJ24521854/video/1/5”) or request the further downstream processing, perform processing, and package the results in one or multiple named and signed data packets to satisfy all original and subsequent Interests for the same “named function processed” data.

The function itself can also be considered as (executable) data and can be packaged in the form of NDN data packets. This way, functions can be easily moved to run at whatever locations that can minimize the overhead of moving data and maximize the processing performance.

C. Security via Named Data

1) *Name-Based Trust*: An architectural difference between NDN and the TCP/IP-based Internet is that in NDN every named piece of content (data packet) must be signed. This ensures that the data can be authenticated regardless of how/where it is retrieved. Besides the signature, each data packet also carries additional metadata including the signing key name. To authenticate a data packet, one needs a trust model that defines which keys are authorized to sign which data (trust rules) and one or more trusted keys to bootstrap the trust (trust anchors). Any entity—applications, dedicated

network storage elements, and even network routers—that learns the trust model for a given piece of content can verify its authenticity and perform necessary actions when the authentication fails (e.g., discard the packet, or try an alternative path or channels to retrieve). Keys in NDN are just another type of data, thus they also have unique names and can be authenticated in the same way as other data packets.

The power of names allows NDN to define general trust model rules as a set of relationships between names of data packets and names of the keys authorized to sign those data packets, e.g., both must share the same prefix, share the same suffix, and/or have specific name components at certain positions of the names. The resulting trust schema [4] is effectively a tool to facilitate application and enforcement of the trust models in a completely distributed fashion. Specifically, a set of strict trust rules captured in the trust schema defines what is (are) legitimate key(s) for each data packet that the application produces or consumes. Given a trust schema that correctly reflects the trust model of the application, data producers can select (and if necessary generate) the right keys to sign the produced data automatically, and consumers can properly authenticate each retrieved data packet.

2) *Name-Based Access Control*: NDN supports content confidentiality by encrypting the content at the time of production. The encrypted data packets can be stored anywhere in the network and delivered to requesters through any path without revealing the payload to any intermediary. Only the authorized parties are given the correct decryption keys to access the original content. Securely distributing coalition data to authenticated parties using fine-grained policies overcomes the current limitation of requiring coalition partners and data to rendezvous at secured data centers via end-to-end secure channels. NDN can also support requester authentication by requiring Interest packets be signed [5].

To support multi-point secret communication, which is most likely required by analytics applications in coalition settings, NDN can apply conventional and novel encryption methods, including shared secret keys, broadcast encryption, public-key based group encryption, and attribute-based encryption. In particular, the ongoing development includes the name-based access control (NAC) mechanism [6], allowing management of the fine-grained access control leveraging hierarchical namespace structure. NAC leverages principles of attribute-based encryption (ABE) [7] to decouple the access control from data production by introducing an intermediate entities, *namespace managers*. A namespace manager performs two important functions: (1) defining the namespace granularity for data encryption and supplying producers with the encryption keys (i.e., which key to use and how long it can be used); and (2) controlling read access to the encrypted data by securely distributing the decryption keys to authorized individual users or groups of users.

D. Policy Management via Named Data

Actors with different authorities can define security rules that control the dissemination and sharing of data between

coalition partners. These security rules can then be packaged as named, secured bags of bits and safely disseminated throughout the network to authorized sensors and analytics components, regardless the channels they come from.

Through naming conventions and pre-defined protocol procedures highlighted in Section III-A, policies can be easily discovered by all interested parties. Given all changeable NDN data requires an explicit name versioning [8], all policy changes are automatically reflected in the name, allowing discovery and retrieval of the latest version of the policy. For example, U.S. policies on information sharing with coalition partners could take a naming convention like “/US:<country-i>/<info-j>/_v=<N>”, where “<country-i>” is a specific country’s name with which the U.S. is sharing information, “<info-j>” represents the specific class of the shared information, and “<N>” is the automatically discovered version number of this policy. This pre-defined naming convention can be shared among all coalition parties to ease policy dissemination.

Since coalition policies may change over time, it is important that these changes can reach the parties who need to execute the policies. In an NDN-based system, consider a U.K. UAV enters the field as depicted in Figure 1: it could fetch the policy “/US:UK/location-1/acoustic” and learn whether it has access to the acoustic data collected by the U.S. at location-1, and if so, how to obtain the access key. Performing such tasks with today’s TCP/IP would require establishing secure channels back to a centralized policy server. In NDN, policies are simply specific types of data and can be fetched as any other data in a secure and resilient way.

The above example also touches upon an important question of where to store the policy data. As of today, they are generated and stored in centralized secure servers. In NDN, one can do the same, but with the generated policies stored as named and secured NDN data packets with versioning. One can also distribute these named and secured policy data into distributed policy repositories (by simply instructing the repos to fetch the latest policies from the central policy generator).

Security, flexibility, and availability are critical factors in successful coalition policy management. We believe that the above examples illustrate the potential help that NDN could offer to policy management.

IV. RESEARCH QUESTIONS

So far we have illustrated various potential advantages from applying NDN to address the challenges that arise from tactical networks and coalition operations. Below we also identify several research topics to be investigated in order to turn those identified advantages to an operational reality.

We note that the following list only identifies some immediate questions to address; this list is by no means exhaustive.

A. Naming Conventions

NDN relies on well defined naming conventions to enable each data consumer to automatically construct the proper name to fetch the data it requires. As we highlighted in

Section III-B, naming conventions and standards are critical for seamless integration of networking, storage, and processing in an efficient and decentralized way. In addition to being used at the network layer for packet forwarding, data names are also used for security purposes (Section III-C). Thus, how best to design the namespace standards remains an open question. We are still in the early stage of developing general guidelines and hope to gain more insights by experimentation with a larger number of diverse application scenarios.

B. Information Discovery

Because tactical coalition environments are highly heterogeneous and highly dynamic, establishing naming conventions may not eliminate cases where consumers express interests using names that do not match to the names of the content being produced. For example, this can happen when the content is directly relevant but not an exact match to the query. It is important to investigate and validate novel methods and tools to enable NDN Interest forwarding to accommodate various approximate name matching schemes, possibly through iterative learning from consumers' reception or rejection of the retrieved data. Such adaptive learning can be useful in handling the uncertainty of namespaces across coalition networks.

C. Name Confidentiality

As we have shown in Section III-C2, NDN enables powerful mechanisms of fine-grained data-centric confidentiality, i.e., the content stays confidential and can only be accessed by the parties who are authorized. However, battlefield communication may also require name confidentiality: a mechanism to prevent unauthorized observers and adversaries from understanding the nature of the requests and retrieved data. One preliminary idea to address this problem is to encrypt all communications among coalition entities to obscure all packet exchanges. Given all the entities in an NDN network must possess valid certificates, we can explore designs that utilize naming conventions to facilitate/automate encryption key management. In other words, a plain-text communication with "signed interests" [5] can be used to discover the encryption key, which is then used to completely conceal (encrypt) all future Interest/Data exchanges, potentially multi-point in nature.

D. Policy Management

Section III-D illustrates how NDN can facilitate policy management. At the same time, trust and confidentiality policies and policy management are not yet an active research area in NDN, but are critical in realizing resilient and secure collection and processing of data in coalition environments. Therefore, it is vital to investigate novel methods to (a) use policy to secure and manage access to named data, (b) to use NDN to distribute policy updates across the network, and (c) to apply the ideas to specific deployment scenarios to verify the relevance and feasibility of the identified ideas.

V. CONCLUSION

Although Named Data Networking (NDN) as a new Internet architecture is still under active research, multiple research initiatives—including the DARPA SHARE Program (Sharing Battlefield Information at Multiple Classification Levels) [9], and joint NSF/Intel ICN-WEN Program (ICN at Wireless Edge Networks) [10]—have been launched recently to explore its applicability and advantages in highly dynamic, highly heterogeneous environments. NDN is expected to excel in such environments because of its resilient delivery of named, secured data, independently from individual nodes, channels, or locations.

This paper aims to motivate the research to exploit the use of NDN to improve the capability of distributed coalition systems. Since this is a new field of study, as a next step we must choose appropriate methodologies to validate and evaluate the perceived advantages of NDN and thoroughly examine any remaining issues.

ACKNOWLEDGMENTS

This work is partially supported by the U.S. Army Research Laboratory and the U.K. Ministry of Defense under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defense or the U.K. Government.

REFERENCES

- [1] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, K. Claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named Data Networking," *ACM Computer Communication Reviews*, June 2014.
- [2] F. Bergamaschi and D. Conway-Jones, "ITA/CWP and ICB technology demonstrator: a practical integration of disparate ISR/ISTAR assets and technologies," in *Proceedings of SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2012.
- [3] C. Tschudin and M. Sifalakis, "Named functions and cached computations," in *Proceedings of 11th Consumer Communications and Networking Conference (CCNC)*, Jan 2014.
- [4] Y. Yu, A. Afanasyev, D. Clark, K. Claffy, V. Jacobson, and L. Zhang, "Schematizing trust in Named Data Networking," in *Proceedings of 2nd ACM Conference on Information-Centric Networking*, September 2015.
- [5] NDN Team, "Signed Interest," Online: <https://named-data.net/doc/ndn-cxx/current/specs/signed-interest.html>, 2016.
- [6] Y. Yu, A. Afanasyev, and L. Zhang, "Name-based access control," NDN, Tech. Report NDN-0034, January 2016.
- [7] V. Goyal, O. Pandey, A. Sahai, and B. Waters, "Attribute-based encryption for fine-grained access control of encrypted data," in *Proceedings of the 13th ACM conference on Computer and communications security*, 2006, pp. 89–98.
- [8] NDN Project Team, "Ndn technical memo: Naming conventions," NDN, Technical Report NDN-0022, 2014.
- [9] "Sharing Battlefield Information at Multiple Classification Levels via Mobile Handheld Devices," Online: <http://www.darpa.mil/news-events/2017-01-10>, January 2017.
- [10] "NSF/Intel Partnership on Information-Centric Networking in Wireless Edge Networks (ICN-WEN)," Online: <https://www.nsf.gov/pubs/2016/nsf16586/nsf16586.htm>, September 2016.