ICN Research Group
Internet-Draft
Intended status: Informational
Expires: January 18, 2018

R. Ravindran
A. Chakraborti
A. Azgin
Huawei Technologies

July 17, 2017

Forwarding-Label support in CCN Protocol
draft-ravi-icnrg-ccn-forwarding-label-01

Abstract

The objective of this proposal is to enable ID and Locator namespace split in the CCN protocol that has several applications such as towards Interest routing optimization, seamless mobility and providing mobility as a service, conversational session support, handling indirections in manifests, and routing scalability. We enable this through the notion of a forwarding-label object (FLO), which is an optional hop-by-hop payload in the Interest message with a topological name, which identifies a network domain, a router or a host. FLO can be inserted within the Interest message by the applications or by the network. FLO can be interpreted by the network in multiple ways, for instance, to terminate the message or to swap it with a new FLO based on the service context. Furthermore, depending on the application and trust context, FLO can be subjected to policy based actions by the forwarders, such as invoking security verification or enabling other service-centric FLO management actions.

Table of Contents

1. ID-Locator Namespace Split in CCN

In the context of ICN/CCN, we define Identifier and Locator as follows:

- o Identifier (ID) is a persistent secure or non-secure flat-ID or a hierarchical name assigned to a content, device or service. If the ID is secure, then there is a direct relationship between the ID and the key of the principal. Otherwise, a binding is provided by a third party using the certificate mechanism or the web-of-

trust model. Generally the identifier space is managed by services and applications.

- o Locator ID (LID), on the other hand, is a topological name assigned to a network entity, which can be a network/domain, a router, a host or an interface. Generally, locator space is assigned and managed by the network administrators. Within the context of this document, Locator and Locator ID carry the same meaning and refer to a topological identifier.

We discuss here the motivations behind the need for separating the persistent name (ID) and the locator ID (LID) in the Interest message within the context of CCN, and present a proposal to achieve this.

The advantages of ID/Locator split have been extensively studied in the literature and it has been part of many host-centric protocols such as HIP [6], ILNP [7], and LISP [8]. However, in this document, we refer to these terms within the context of Information-centric networks, where the network layer uses name(s) to resolve a requested content (or service or host) to a given location, while providing mechanisms to cache or apply computation on the named (data) objects depending on the context (such as the requirements indicated within the Interest message).

Within this ID/Locator split context, ICN architectures such as MobilityFirst [23], NetInf [12] assume an explicit representation for Identifiers and Locators within their architectures, considering their use of non-aggregate-able flat IDs; while the CCN/NDN architecture assumes the aggregate-ability of names within its architecture, thereby applying its use on routing and forwarding. We have argued in [1], the problems associated with the use of name prefixes for routing, which include the challenges related to scalability, loss of name aggregate-ability when data and services are replicated, handling of mobility, and situations where conversational sessions are required for service level authentication and authorization [2]. These issues have also been argued quantitatively in [14], including the scenario, where there is an explosion in the namespace when there are many different ways of naming an entity because of the rich context associated with it. Therefore, providing the ID-LID separation distinctly in the protocol offers the following advantages, which are also discussed in detail in [1]:

- o CCN applications request persistent IDs for contents, services or hosts; while their resolution is handled through per-hop name-based routing by a CCN forwarder using unicast, anycast, or broadcast mechanisms, routing scalability is handled using name prefixes. This model can introduce problems when the named entity

is mobile, migrated, or replicated; as the names have to be announced in the routing control plane, which can in turn introduce routing convergence issues and scalability challenges. Introducing an ID/Locator namespace split within the architecture, which uses a name resolution service, shall address the routing challenges due to dynamic entities, while also improving the routing scalability (due to limiting the state in the core Internet routers to the set of topological names).

- o ID and Locator namespaces are managed by different entities. IDs are managed by applications and services, hence are relevant in the service layer; while Locators are assigned to the networked entity, hence are managed by network administrators. Locators map to network domains or specific network elements, through which the named entity is locally/globally reachable. The relationship between IDs and Locators is established during namespace registration (or namespace publishing) phase, and managed by a separate name resolution service. The distinction between ID and Locator in CCN allows an application to manage its own namespace and to not be restricted by the naming rules imposed by the network.

- o Allowing the representation of IDs and Locators within an Interest message offers many advantages in CCN, especially when a centralized control is applied, such as using a service orchestration framework [13]. This enables efficiency and flexibility through service-centric name resolution, routing, and mobility, service chaining [4] and routing scalability.

Considering the above requirements, we propose a Forwarding-label object (FLO), which includes a locator along with the encapsulated security bindings; and provides the flexibility to forward Interests on a name other than the one provided within the original Interest message, with the ability to terminate or swap it in the network during forwarding. Note that, to effectively handle the ID/Locator mapping, we require a control plane infrastructure and appropriate network layer security functions to prevent malicious usage. Specific control plane or security mechanisms to support secure ID/Locator mapping is out of the scope of this document, as many techniques can be used towards achieving this objective. This draft specifically presents various considerations towards the management of FLOs, such as: (i) FLO insertion/modification/deletion, (ii) processing of FLO by a CCN forwarder, (iii) PIT/CS implications for Interests carrying FLOs, (iv) packet format for the FLO Interest, and (v) security/trust considerations. We also provide a discussion on the various application scenarios for the use of FLOs.

2. Forwarding-Label Object Proposal

In the following, we discuss various aspects to FLO in regards to its semantics and management.

2.1. FLO Naming

We assume the FLO to consist of three components: LID, service specific metadata, and security attributes for authentication. LIDs are hierarchically structured topological names, where the names follow the format defined in [3]. Security attributes are optional, and may include validation payload and algorithm as discussed in [3].

2.2. FLO Insertion

Insertion of FLOs within an Interest message can be done either by the application that requests the named entity or by the network routers that forward the Interest messages.

In some specific scenarios, application logic may use --within the Interest message-- both an FLO and an ID. Such action may also be triggered by a feedback from the network, for instance, due to a routing failure of an Interest using solely IDs, as explained in [15]. In such situations, forwarders, which process the traffic generated by applications outside their trust domain, would require a way to validate the FLOs. One possible approach to ensure trust in such situations is discussed in [15], where a trust binding is provided between an ID and a LID as a link object, which can be validated by the forwarder. To avoid the possibility of an FLO misuse, a default policy for the network may be to ignore the inserted FLOs from untrusted applications and to choose to route only by the content IDs. Another possible policy to implement in this case is for the network to insert FLOs, which is explained next.

Another possibility would be to insert the FLOs by the network, for which the FLO insertion would be triggered at the ingress (or service) routers of the network domain. Network-insertion of an FLO to an incoming Interest message may be triggered based on several considerations, including: (i) the interface, over which the Interest message is received; (ii) whether the Interest message satisfies the flow service profiles that are imposed by the network administrator at the ingress routers; (iii) the default behavior by the network, if it chooses to route only on LIDs. Accordingly, service profile matching actions may include matching an Interest name to a set of service prefixes that are triggered by certain markings or metadata carried within the Interest message, such as the context-IDs. Here, context IDs may refer to service, network, trust, or location related

metrics or information. Note that, a FLO that is inserted within the trust domain may not require security validation.

2.3. FLO Swapping

An FLO can be swapped by another FLO within the network, in the context of a given service, at the designated network points, which typically include the service edge routers of the network.

Future considerations also include the case, where FLOs are potentially stacked based on the semantics of the current FLO.

2.4. FLO Termination

FLOs are terminated by a forwarder, when the LID carried within matches the forwarder's own LID. Here we assume a forwarder to possibly have multiple LIDs that correspond to domain-IDs, router-IDs or Interface-IDs. For example, a forwarder (in a domain) identified as /company/geoloc can process an FLO, for which the LID is set to this router's domain name or to a forwarder ID such as /company/geoloc/PoP*. Whenever an FLO is terminated by the forwarder, it can attach a new FLO depending on the service context, or conduct additional processing (such as re-resolution of the ID to a new FLO) based on the Interest parameters (and again depending on the service context). An FLO can also include optional policy metadata, based on which FLOs can be swapped in the network.

3. FLO Message Format

As the FLOs are swappable in the network, FLO is proposed as an hop-by-hop field within the optional body of the fixed header as shown in Figure 1. The optional FLO includes an attribute of type FL-Object, which contains a name TLV identifying the LID (Figure 2). LID (or FLO-LID), here, is a hierarchically structured variable length name as defined in [5]. In addition to the LID, the optional FLO metadata includes contextual information for the application or the service to aid the network in invoking an appropriate FLO processing, such as trust validation for the FLO. Optional security attributes, such as the authentication information, can be included depending on the specific use case scenarios, which may include secure name delegation information as discussed in [15], or the signature of the consumer.

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    CCN Fixed Header                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              <Optional Hop-by-Hop TLVs>                       |
/                                                               /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Type = FL-Object    |      Length                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        LID TLV                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
/                      Optional FLO Metadata                    /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
/               Optional FLO Security Attributes                /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
/                       Interest Message Body                   /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Interest message with hop-by-hop Optional Forwarding-Label TLV

Forwarding-Label	Meaning	Value
LID TLV	Identifies an AS-ID/Gateway-ID/Host-ID/Interface-ID	Name TLV

Figure 2: Locator-ID (LID) Definition

4. FLO Processing Rules

The following discussion is based on the assumption that all forwarders must process the optional header fields. Within the context of CCN packet processing, FLO is only relevant when the decision to forward the Interest message is to be made. In this draft, the default policy to be applied by all the CCN routers is to route on FLOs, if an FLO exists in the Interest message. Based on this policy, the following considerations may apply:

- o When an Interest message with an FLO arrives at a CCN router, if the FLO is trusted and the lookup based on LID succeeds, the router first checks if the LID matches any of the network names associated with the receiving router. If there is match, received Interest is treated as one that terminates flow, in which case, a

name based lookup is conducted. The lookup at the router may return another FLO, or may result in a forwarding face for the Interest message towards the next hop. If there is no match, which suggests a non-terminating flow case, the FIB lookup on LID would result in a next hop towards which the Interest needs to be forwarded.

- o If the lookup based on LID fails, then the router can try to forward the Interest based on the Interest ID. In the case, where routing based on the Interest ID fails, then the router can raise an error condition and feedback the error message to the previous hop with the appropriate error code.

- o FLO-validation depends on the trust context, which is indicated by the information inserted by the ingress domain router within the FLO. In trusted forwarding scenarios, where the applications and the network are managed by the same authority, the ingress and the core routers can bypass the FLO validation. In untrusted forwarding scenarios, the edge router may only validate the FLO that is inserted by the sender and avoid re-validations by the successive forwarders.

5. Implications on PIT Processing

As FLO is considered as a routing directive, its presence shouldn't affect the functionality of the PIT and its processing under normal situations. However, including FLO within the Interest message could give rise to new questions, which need to be addressed, based on application or network requirements. One such scenario is when there is an implicit binding between the ID and the LID (i.e., multiple Interests arrive at a router carrying the same ID but with different LIDs). One possible approach to handle such case is to treat each such combination of (ID,LID) differently, thereby saving them both in the PIT as separate entries. However, this also requires the content object to piggyback the LID to ensure proper matching with the stored PIT entry. In the case, when the FLO is swapped by the intermediate routers, the PIT should save both the incoming and the outgoing FLOs, and also the content object should be updated with the appropriate FLO to ensure matching of the stored PIT entries along the previous hops. These considerations are similar to those elaborated in [21].

6. Implications on Caching

Caching function shouldn't be affected by this draft proposal. Even if the FLOs are included in the content object as discussed earlier, routers are expected to remove it before caching the content.

7. Multiple Domain Scenario

In wide area networking scenarios, Interest message can cross multiple domains. If an FLO is only trusted within the domain boundaries, then the FLO is removed before the Interest message is forwarded to the next domain, which then, upon entry, inserts a new FLO with the associated security attributes at the ingress of the next domain. However, in the case when there exists trust between the domains, such as one through a trusted third party (validated based on the FLO security bindings), to use the FLO inserted by the previous domain, the intermediate domains can avoid further FLO processing and use the FLO passed on by the previous domains.

8. FLO Security

FLO security is related to the purpose it is used for and the control plane mechanism used to manage it. Depending on the use case scenario for the FLO, appropriate security mechanisms should be applied to secure the control and data plane functionalities and operations to avoid exploitation of this feature.

Generally, the major threat against the FLO approach is to manipulate the relationship between the name and the FLO. Such manipulations can happen in various scenarios, some of which are listed as follows: (i) a malicious interceptor (acting as a publisher) may intentionally inject an incorrect mapping into the name resolution system; (ii) a malicious interceptor (between the edge router and the resolution server) may manipulate the mapping sent back from the name resolution system when the edge router queries the mapping system; (iii) a compromised intermediate router may maliciously change the FLO, for instance, with a wrong FLO or an out-dated FLO; and (iv) an untrusted application may inject an invalid FLO into the Interest message.

To achieve network level FLO security, appropriate mechanisms should be applied to provide mapping provenance and mapping integrity and to prevent replay attack to address these issues. The security mechanisms applicable to the above discussed scenarios (i) and (ii) are similar to ones applied to secure other mapping systems such as LISP [9] and DNS [11]. Scenario (iii) requires new security mechanisms, one such way is to enable a domain level trust infrastructure so that the mapping between the name and the FLO can be authenticated by the successive routers.

In untrusted environments, when an FLO is inserted within the Interest message by the end hosts, appropriate authentication information should also be included in the FLO to allow ingress routers to optionally validate the delegation of the Interest ID to

LID [15]. Furthermore, additional security policies can be enabled by the network to handle FLOs outside its trust domain.

9. Use Case Scenarios

Here we provide the discussions related to using FLOs in different scenarios.

9.1. Handling Producer Mobility

In the literature, we can classify the techniques to handle producer mobility into two main categories as application-based approaches and network-based approaches.

- In application-based approach, the application takes the responsibility for announcing its reachability to the network and triggering a network state change to enable Interest message routing towards the mobile producer. Most of the current proposals fall under this category, and these include the following two approaches: (i) Kite proposal [20], which implements an anchor based approach where consumers and producers agree on an anchor point based on external mechanisms and uses application initiated (traced and tracing) Interests to handle the producer mobility; (ii) Anchor-less proposal [22], which is another application-based approach wherein enhanced name-based routing and forwarding is used to track the mobile producer. While these approaches allow consumers and producers to work on a single name space, their use may raise scalability concerns with increasing number of mobile nodes and number of applications signaling into the network. As these approaches introduce additional signaling in the network, the operational efficiency of packet forwarding is negatively affected due to state changes that have to be applied to ensure optimized Interest routing to the mobile producer. Another potential issue with these approaches is security, as they can be prone to flooding attacks by malicious applications targeting specific application names and impeding their normal operation.

- In network-based approach, late-binding technique [18][16] is used, wherein the reachability of the mobile node is handled by the network. In this case, applications can explicitly request mobility for a given name space [16], with the network handling the mobility by tracking its latest location in the network through a name resolution system. At the same time, through coordination of the old and new point-of-attachments (PoA) and in-band signaling, one can achieve minimal loss for a given Interest flow. The late-binding technique uses ID/Locator split that is only applied at the PoAs, thereby avoiding challenges related to

routing convergence in the network due to producer mobility, while offering better scalability performance as the number of mobile producers increases. In this approach, a user entity (UE) registers a name prefix that requires mobility support with its current point-of-attachment (PoA). The PoA then registers the mapping between the name prefix and the PoA's locator in its local name resolution system. The domain then updates the UE's home domain name resolution system with its current domain LID. When a correspondent nodes expresses Interest for the name, it is first resolved to the current UE domain by the home domain. When the Interest enters the domain offering mobility service, it is resolved again to the UE's current location. Furthermore PoA-to-PoA signaling can be enabled to offer seamless forwarding of Interests whenever an UE changes its PoA. In addition to correcting the path stretch, the Interests re-routed from the old PoA can be marked and re-routed to the new PoA with the new FL. On the return path, the content objects are also marked, and this in-band marking used by the ingress PoA at the consumer's end results in re-resolving the mobile prefix to a new forwarding locater that would correct the path stretch.

9.2. Manifests

The FLOs can also be used to support the retrieval of nameless objects [17]. Using the current manifest proposal [10] a consumer receives a manifest with the ContentObjectHashIDs and their respective locator information. A consuming application uses the locator as a routable content name, while the ContentObjectHashID is used as a HashID restriction parameter. Multiplexing the Interest name field as an ID and also as a LID has the following consequences: (i) a forwarder cannot distinguish between Interest messages containing ID or LID in the name field, as the protocol does not differentiate between these two names; (ii) it complicates Interest message processing, where two different Interest processing logics need to be applied on Interests (with or without the hash-id). In this situation, routers should first check for the presence of ContentObjectHashID and uses it to index the Interest based on it, rather than using the locator name; (iii) more complications may arise, if an Interest packet arrives with two IDs, for example, a ContentObjectHashID as the hash restriction and the ID as the content name, in which case, one of them should seek precedence over the other.

The above issues can be avoided through the use of ContentObjectHashID as the content name and the locator as LID with the FLO. In this case, a forwarder will always index the pending Interest table or the cache as expected on the content name. The routing decision would then be based on the FLO, depending on the

routing policy implemented by the forwarder. This also avoids the situation of dealing with two IDs in the Interest message, where the application has to choose either the ID or the ContentObjectHashID as the content ID.

A possible high level forwarding logic for the edge/core router to support nameless objects based on the above discussion is presented in Figure 3. Here, edge router can also be a gateway node.

```
          Begin

          if Edge_Router
                 If Interest arrives on a face with a flat-ID
                        Then check for the presence of FLO
                          If FLO is present, use the LID in the FLO for
Interest forwarding

                        If there no FLO
                          If policy allows, resolve the flat-ID with a NRS to
obtain an FLO
                                Use the FLO to route the Interest

                 End

                If the Interest arrives with a routeable ID
                  If there is FLO
                        Then use the ID for forwarding and Remove the FLO

                  If there is no FLO
                        Match Interest ID with name policy for e.g. mobility or
interest routing optimization
                          If a name policy for resolution exists
                                Then network resolution service is invoked on
the ID  which returns an FLO
                                 Use the FLO to direct the Interest to the
appropriate next hop
                End
          End

          if Core_Router
                  if Interest arrives with an FLO
                        Use the LID for forwarding
                  Else if Interest is with a Routeable ID
                        Use the name for forwarding
          End
```

Figure 3: Forwarding logic to support flat-ID and

routeable ID at the edge router

We next discuss the security implications of using IDs and FLOs within the Interest message, depending on the ID type, as follows:

- Case 1 - ContentObjectHashId with FLO: This use case is a straightforward simplification of what is being proposed in [17]. Here, the locator is included within the FLO and the ContentObjectHashId is used as the name, so this shouldn't introduce any new security concerns. This holds true for both application and network based FLO insertions.

- Case 2 - Routable ID with FLO: For UE based FLO insertion, this scenario can cause cache poisoning in the absence of signature check enforcement at the forwarders. For the case, when FLO is managed within a trusted domain, the security implications are discussed in Section 8.

9.3. Support for Conversational Sessions

FLO can be used to bind a service ID to a given location in the network, so that the consumer's session is correctly directed to the service instance keeping state of the conversation. An example for this is discussed in [2]. Using an ID-based anycast routing cannot guarantee this as the name prefix state used for forwarding would treat all possible instances equally. One way to mitigate this, which may compromise efficiency, would be to introduce a load balancer, through which all such Interest flows are routed.

9.4. Interest Routing Optimization

A network domain, which hosts its own or third party contents/ services, can benefit from the ability to handle Interest routing logic within its domain opportunistically. When an Interest message seeking a specific content or service enters a network domain, the ingress router can redirect the Interest to the closest cache point or service location.

9.5. Routing Scalability

As discussed in [15], locator-based routing can address the routing scalability, as the number of ASs are many orders less than the number of information objects. Doing so would reduce the forwarding table in the DFZ to the order of the number of ASs in the Internet. In addition, unlike [15], this proposal offers the features of swapping FLOs and late-binding, which may lead to more flexibility and efficiency towards scalability and routing optimization.

10. Informative References

[1] Azgin, A. and R. Ravindran, "Enabling Network Identifier (NI) in Information Centric Networks to Support Optimized Forwarding", draft-azgin-icnrg-ni-01 (work in progress), May 2017.

[2] Mosko, M., Uzun, E., and C. Wood, "CCNx Key Exchange Protocol Version 1.0", draft-wood-icnrg-ccnxkeyexchange-02 (work in progress), July 2017.

[3] Mosko, M., Solis, I., and C. Wood, "CCNx Messages in TLV Format", draft-irtf-icnrg-ccnxmessages-04 (work in progress), March 2017.

[4] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <http://www.rfc-editor.org/info/rfc7665>.

[5] CCN Wire format, CCNX1., "http://www.ietf.org/id/draft-mosko-icnrg-ccnxmessages-00.txt.", 2013.

[6] Nikander, P., Gurtov, A., and T. Henderson, "Host identity protocol (HIP): Connectivity, mobility, multi-homing, security, and privacy over IPv4 and IPv6 networks", IEEE Communications Surveys and Tutorials, pp: 186-204, 2010.

[7] Atkinson, R., "An Overview of the Identifier-Locator Network Protocol (ILNP)", Technical Report, University College London, 2005.

[8] LISP, RFC6380., "https://tools.ietf.org/html/draft-ietf-lisp-sec-07.", 2014.

[9] LISP-Security, LISP-SEC., "https://tools.ietf.org/html/draft-ietf-lisp-sec-07.", 2014.

[10] CCNx, Manifest., "http://www.ccnx.org/pubs/draft-wood-icnrg-ccnxmanifests-00.html.", 2015.

[11] DNS-SEC, RFC4033., "DNS Security Introduction and Requirements.", 2005.

[12] FP7-ICT-2009-5-257448/D.B.3, "SAIL: Scalable and Adaptable Internet Solutions", 2013, <http://www.sail-project.eu/wp-content/uploads/2013/01/SAIL-DB3-v1.1-final-public.pdf>.

[13] Ravindran, R., Chakraborti, A., Amin, S., Azgin, A., and GQ. Wang, "5G-ICN : Delivering ICN Services over 5G using Network Slicing.", IEEE Communication Magazine, May, 2017.

[14] Adhatarao, S., Chen, J., Arumaithurai, M., Fu, X., and K. Ramakrishnan, "Comparison of Naming Schema in ICN.", IEEE LANMAN, 2016.

[15] Afanasyev, A., "Map-and-Encap for Scaling NDN Routing.", NDN Technical Report ndn-004-02, 2015.

[16] Azgin, A., Ravidran, R., Chakraborti, A., and G. Wang, "Seamless Producer Mobility as a Service in Information Centric Networks.", 5G/ICN Workshop, ACM ICN Sigcomm 2016, 2016.

[17] Mosko, M., "Nameless Objects.", IETF/ICNRG, Paris Interim 2016, 2016.

[18] Azgin, A., Ravindran, R., and G. Wang, "A Scalable Mobility-Centric Architecture for Named Data Networking.", ICCCN (Scene Workshop) , 2014.

[19] Cisco System Inc., CISCO., "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update", 2016-2021.

[20] Zhang, Y., Zhang, H., and L. Zhang, "Kite: A Mobility Support Scheme for NDN.", NDN, Technical Report NDN-0020 , 2014.

[21] CCNx Label Forwarding, CCNLF., "http://www.ccnx.org/pubs/ccnx-mosko-labelforwarding-01.txt.", 2013.

[22] Auge, J., Carofiglio, G., Grassi, G., Muscariello, L., Pau, G., and X. Zeng, "Anchor-less Producer Mobility in ICN.", ICN, Sigcomm, 2015 , 2015.

[23] NSF FIA project, MobilityFirst., "http://www.nets-fia.net/", 2010.

Authors' Addresses

Ravishankar Ravindran
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: ravi.ravindran@huawei.com

Asit Chakraborti
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: asit.chakraborti@huawei.com

Aytac Azgin
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: aytac.azgin@huawei.com