

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

A. Azgin
R. Ravindran
Huawei Technologies
July 17, 2017

Enabling Network Identifier (NI) in Information Centric Networks to
Support Optimized Forwarding
draft-azgin-icnrg-ni-02

Abstract

The objective of this proposal is to introduce the notion of network identifier (NI) in the ICN architecture. This is in addition to the existing names (i.e., content identifiers, CIs, or application identifiers, AIs, in general) that are currently used for both naming and routing/forwarding purposes. Network identifiers are needed considering the requirements on future networking architectures such as: (i) to support persistent names (or persistently named objects) and large-scale and high-speed mobility of any network entity (i.e., devices, services, and content), (ii) to accommodate different types of Internet of Things (IoT) services, many of which require low-latency performance, and enabling edge computing to support service virtualization, which will require support for large scale migration and replication of named resources, and (iii) to scale the ICN architecture to future Internet scale considering the exponentially increasing named entities. If information on AI-to-NI mappings are not directly accessible to the consumers, for instance, using specific datasets like manifests, these considerations may require enabling a name resolution service, which can be network based or application driven, to support efficient and scalable routing. Current document do not impose any restrictions on the name resolution architecture, regarding its scope.

In the current draft, we begin by highlighting the issues associated with ICN networking when utilizing only the AIs, which include persistently named content, services, and devices. Next we discuss the function NI serves, and provide a discussion on the two current NI-based proposals, along with their scope and functionalities. This is with the objective of having a single NI construct for ICN that is flexible enough to adapt to different networking contexts.

Requirements Language

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

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 January 18, 2018.

Copyright Notice

Copyright (c) 2017 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

1. Introduction	3
2. Application Identifier (AI) vs. Network Identifier (NI) in ICN	3
3. NI based ICN Forwarding	7
3.1. Label based ICN forwarding	8
3.2. Link-object based ICN forwarding	9
3.3. Link Object vs. Forwarding Label	10
4. Name Resolution System Considerations	12
5. Differences with respect to Existing IP-based Proposals . . .	13
6. References	14
6.1. Normative References	14
6.2. Informative References	14
Appendix A. Additional Stuff	16
Authors' Addresses	17

1. Introduction

Information centric networking (ICN) is proposed as a future Internet architecture to evolve the current host-centric design of Internet towards a content-oriented one, where the named object becomes the principle entity in networking. In doing so, contents, services, and devices become disentangled from any particular host (or hosts) allowing for efficient use of the distributed in-network caches and compute resources with more flexible and dynamic packet forwarding techniques. ICN is expected to offer a scalable and secure networking solution to address many challenges of the current IP architecture. Towards this, we propose to formalize the notion of network identifier (NI) in ICN protocol, that is separate from content name or application identifier (CI/AI, or simply AI) used to both name resources and route user requests.

2. Application Identifier (AI) vs. Network Identifier (NI) in ICN

AI represents the names of services, contents, or devices assigned by the application providers or device manufacturers, and which can be validated through appropriate security mechanisms. AIs are specifically used for content request and distribution. ICN should provide flexibility in accommodating a broad set of identifiers, within which the two well-known classes include hierarchical and flat identifiers. While a hierarchical identifier provides contextual richness for the names, a self-certified flat identifier offers a fixed predictable overhead and context-based security features (as they can be hash of the content object or hash of the public key). Today, this identifier set is already in the order of billions (with hundreds of millions domain name registrations across all top-level domain names [VRSN]). As tens of billions of devices are expected to join the network, this identifier set will be further augmented with the corresponding data objects significantly expanding its size. To decouple applications from the underlying network dynamics, identifiers are expected to be persistent within the scope of the application and its deployment.

NI provides a binding for the AI to the network, at a location and in a topology relevant manner. In more specific cases, such as with anycast use of NIs or for multi-source-homing scenarios, binding can target a set of locations rather than a single location. NI is managed by the network provider to name the routers, point of attachments, servers and end devices. In addition to ICN names, in an overlay deployment, NI could assume names associated with the underlay network as well, such as IP or Ethernet addresses, in which case the NIs would be carried within the underlying protocol headers, potentially with further address translations occurring at the ICN-enabled routers, hosts, or devices. The growth in the NI space is

proportional to the rate of growth in domain topology, the total number of AS(s), and the end points (if they are managed by the network), whereas the growth of the AI space is proportional to the rate of growth of the named resources within it. Considering the potential use cases for ICN, the growth in AI space can be expected to be much faster than that of NI space. Furthermore, as NI is a routing construct, which can be modified en-route using per-hop name resolution at the domain boundaries, the forwarding table sizes at the core routers can be limited to the number of AS(s) instead of the size for the set of end points. Hence if the objective is to limit the size of the forwarding table and scale control plane, it is desirable to route requests on NIs, with the mapping between AI and NI is achieved in a scalable manner using one of many ways including but not limited to using a network based name resolution system or using a manifest-driven information database system to provide such mappings.

Content-centric design used by ICN allows end hosts to make requests using any type of name supported by the applications, including hierarchical (human-readable or hash-based) identifiers (as considered by CCN, NDN[CCN] for both the client application use and the network use-for routing-), or fixed flat identifiers (as considered by MobilityFirst[MFRST] in the network for routing). We refer to an ICN architecture that supports any application naming format (i.e., human-readable or flat) within the network for routing as a non-restricted ICN architecture (as in CCN/NDN), whereas an ICN architecture with a fixed naming format for routing within the network as a restricted ICN architecture (as in MobilityFirst).

As packet forwarding in ICN utilizes names or identifiers (associated with contents, hosts, or services) which are typically managed by applications, thereby of 'mostly' persistent nature, using such names in packet forwarding introduces certain challenges in regards to routing scalability and forwarding efficiency [NAMES]. Depending on the application context, it is possible for an application to support the use of non-persistent names, for instance, in the case of real-time multimedia services, with further challenges towards achieving scalable routing and efficient forwarding. We list the most critical challenges with respect to use of names in routing as follows.

- o Using AI for Routing/Forwarding: Assigning dual functions to an application identifier to include using it as a locator may, in certain scenarios and depending on the ICN implementation, lead to unstable 'routing control and forwarding plane' operations, particularly when replication and mobility of content or end points are taken into consideration. Specifically, if application identifiers are used in routing, we can express the update overhead to be proportional to the multiplication of update-reach

(i.e., the level of reachability of the update in terms of the number of routers that need to update their routing/forwarding tables), update-frequency, updated-object-count, which can easily reach unmanageable levels, with shift towards more mobile communications or higher level of content replication.

- o Applications typically construct names and replicate contents or services to optimize their delivery without any consideration towards network scalability or efficiency. Accordingly, name aggregation may not help with scaling the routing and forwarding as typically considered [RFC7927], and the cost of this would be quite significant in real world scenarios, as discussed in more detail in [NCMP]. Furthermore, it is also observed in [QCMP] that, in certain scenarios (such as content mobility), name-based forwarding approaches can operate more efficiently, if used in conjunction with address-assisted schemes such as DNS or anchor point based approaches like Mobile IP [RFC3220]. Additionally, when names are used for network reachability, more practical problems such as name-suffix hole may arise, as the content requests are forwarded towards non-existent caches [MDHT].
- o Routing/Forwarding Scalability: Routing scalability is typically achieved by designing routing state with aggregate-able property, which is the case for the current IP architecture. However, having such feature in a non-restricted ICN architecture would lead to relinquishing the persistency of the names, along with its security binding such as trust, as the names would involve a topological component for scalability, which can also suggest resources to be renamed depending on, for instance, network or business specs or characteristics. Note that, routing on names with aggregate-able property would mean that names need to change as the location for name changes, for instance, with publisher/producer mobility. As we assume that trust relationships are established based on names, changing names would mean updating security bindings accordingly, dynamically as requests are pushed towards the content source.
- o When content names or application identifiers use a hierarchical identifier format, we observe scalability problems in control and data plane operations [SFWD]. Such problems are caused by various factors. For instance, the explosive growth observed in namespaces can lead to a similar growth in routing/forwarding information base or table sizes [AFWD][SPIT][WPIT], even when namespace aggregation is enabled, to significantly limit the forwarding efficiency and forwarding capacity. If ICN routing with hierarchical naming is the accepted form of naming, name-aggregation is highly unlikely to achieve any practical scalability. This is because, naming ontology and assignment

typically consider application objectives of contextualizing names, service and content placement and replication to better suit the consumers' needs without considering any network objectives on control and data plane efficiency and scalability.

- o Handling Mobility, Migration, and Replication: The impact of namespace expansion on routing/forwarding performance is typically exacerbated with content mobility, or the use of multi-homing and resource replication due to diminished aggregate-ability [NCMP]. The authors in [QCMP] concludes that, as more than 20% of end hosts make more than 10 network address transitions every day, thereby suggesting that mobility should be considered as the norm rather than the exception. Furthermore, to achieve location independent routing based on AIs, each mobility event associated with a device or a popular content may trigger updates on up to 14% of Internet routers.

For the above reasons, restructuring the identifier to directly or indirectly contain a globally routable component becomes an important requirement, especially, to handle mobility at the network layer for architectures that do not restrict names or identifiers to any specific format. We can refer to such operation as the Application and Network identifier split (where the NI represents the globally routable component, and the AI represents the persistent name/identifier) which enables splitting of the namespace to support routable, persistent, and human-friendly names or identifiers. In such a framework, names would be divided accordingly, i.e., based on application binding (offering persistent names) vs. advertised network entities (in routing plane) to provide a more scalable routing architecture. For instance, a persistent name or identifier /Provider/Type/Name, which would be used to create secure content objects, can be published by multiple content distributors, where it would be mapped to different NIs, such as /Distributor/Region/Zone/Storage, to resolve content names or identifiers to specific infrastructure entities. The fundamental requirements with this form of splitting is no different than that of MobilityFirst [MFRST] or LISP [RFC6830], which is the requirement of a network based name resolution system to map the two namespaces.

So far, various approaches have been proposed to support the use of NI in ICN-based networking architectures, depending on how this information is structured and where it is placed within the Interest (which may also determine the structuring of Data packets). Next, we discuss these solutions by specifically focusing on label-based ICN forwarding [FWLDR][FWLRP][MAAS] and ICN-based Map-and-Encap [MPNCP][SNAMP] to provide a general guidance on the use of NI in information centric networks.

3. NI based ICN Forwarding

AI based routing is a feasible solution within certain contexts such as: (i) when resources are static and routing is limited to local area networks or local domains, such as access networks within the scalability considerations of the control and forwarding plane; (ii) in ad hoc situations where AI can be combined with suitable suffix filters to seek content of interest for the applications; (iii) in infrastructure-less scenarios with limited scalability requirements.

On the other hand, the use of NI becomes important in the following situations: (i) when the Interest packet goes outside the local domain, where routing on AI is optionally supported (i.e., as routing scalability and efficiency seeks precedence, forwarders can choose to exclude certain AIs from FIB processing, which may limit the forwarding of requests carrying such AIs); (ii) when the Interest enters a local domain, and the domain has specific knowledge of an NI associated with the resource inside its domain (as the use of NIs would address routing efficiency through exact matching on the NI rather than performing longest prefix matching on AIs). The first situation can limit routability of requests if information on how to reach an AI is not carried in all domains, whereas the second situation, for various reasons, can help with efficient forwarding of requests by routing on NIs rather than on AIs even though it is supported (e.g., NI lookup may be performed on a more performance efficient state table using exact match rather than longest prefix matching).

With the above considerations, with respect to end-to-end networking, using NI for routing is not a mandatory feature, but an optional one. However, as significant amount of user traffic fetches resources outside the requesting host's local domain, it becomes crucial to provide architectural support for routing on NIs in an ICN protocol. Here, we consider NI as an implemented feature for communication among static network components (e.g., as router identifiers) or cross domains (e.g., as domain identifiers or global identifiers), and can be designed using locally or globally defined policies, which for the latter case may require globally agreed semantics for trust management to validate bindings between NIs and AIs.

So far, two solutions for NI in ICN, overall with the same objectives but serving different purposes, have been proposed. These include the forwarding-label proposal [FWLDR] and the Link Object described in [SNAMP]. We next summarize these proposals and discuss their differences.

3.1. Label based ICN forwarding

Label-based ICN forwarding provides NI capability by encoding a network address along with (optional) security binding attributes within an Interest packet to guide it towards a content source (which can be the Producer, a content repository or a cache). We refer to this label that provide the NI capability as the forwarding label [FWLDR], which can be offered as part of an ICN network service (such as a name resolution service with ICN APIs to register and resolve names). Security binding attributes are considered optional for forwarding labels as their scope can be limited to use within a domain, and within the boundaries of a domain, with an established trust among forwarding hosts (i.e., network routers) such bindings may not be needed. On the other hand, to ensure cross-domain validity of forwarding labels, in the absence of prior established trust relationships, security binding attributes are considered as mandatory, and enforced at domain boundaries to ensure that end-to-end NI-based packet forwarding is supported. There may be exceptions to the above scenarios depending on how the NIs are utilized and updated.

For the forwarding label, we have the following important considerations: (i) forwarding label, if present in the Interest packet, takes precedence (over AI) for routing to ensure consistency in packet forwarding whenever its used is triggered (for instance, to avoid the emergence of loops which can occur as a consequence of alternating between routing on AIs and routing on forwarding labels), (ii) forwarding label is mutable in the sense that it can be swapped or removed by intermediate network elements in the network based on routing considerations within its domain. Note that, to ensure compatibility with future potential use cases, label based ICN forwarding can also utilize dynamic precedence, for instance to prevent routers from unnecessarily dropping requests. Here, forwarding labels are not limited to only the ICN names, but, in an overlay mode, they can also represent names from other transport layers as well, for instance, an IP address or a MAC address.

Forwarding label consists of multiple components, one of which is the NI that represents the locator information. If the AI namespace supports the use of an NI to reach a specific destination, forwarding label is embedded within the Interest message at the edge router or the end point within certain trust considerations. For security reasons, edge routers can validate the label, which is inserted by the end hosts, based on the trust context. For instance, if the inserted label cannot be validated, edge router can ignore the label inserted by the end-host and swap it with a new one depending on the feedback from the name resolution system; or if forwarding on labels is not supported, edge router can ignore any such label inserted by

an ICN forwarder at the end hosts, by simply removing the inserted label. Such an approach requires no trust relationship among different domains, as each domain is capable of resolving content namespace to a target domain, and swapping the received label with one to which it resolves.

Forwarding label support for a namespace can be offered at a global scale (i.e., supported by all the domains) or a local scale (supported by a subset of the existing domains). For instance, some autonomous systems can prioritize forwarding solely based on the content names (or offer limited support for label-based forwarding on specific namespaces). In such case, forwarding labels can include additional service tag (or information on the associated service, for which the use of forwarding label might be supported in certain domains, such as towards mobility service) for routing packets on the supported domains. In doing so, we can strategically forward requests over domains that support such service to provide more deterministic service guarantees.

If forwarding label use is supported (or permitted) within a domain, by default, forwarding label is given preference over content identifiers for packet forwarding. In such case, to maximize the forwarding efficiency, additional mapping tables can be implemented at the edge or border ICN routers for quick longest-prefix matching (LPM) lookup on content names to determine a (or the) matching forwarding label(s), which can then be used by the router to perform LPM lookup on the FIB. As forwarding label typically represents a target domain or router, a single LPM lookup on the FIB may suffice to find the outgoing interface for the received Interest. This state can also be software-defined based on application requirements using an SDN based control plane.

3.2. Link-object based ICN forwarding

ICN-based Map-and-Encap [SNAMP] utilizes link objects, which include information on how to retrieve content objects. For instance, link objects can represent domains that host the content object, or direction towards which the requests need to be forwarded to find a matching content object. Link objects consist of two optional headers: (i) a link header, which includes the potential directives that can be used for forwarding and is signed by the Producer to validate its authenticity during forwarding, and (ii) a delegation header, which is used to represent the link choice utilized by the previous forwarder. Since delegations may change at consecutive hops depending on the view of forwarders' network state and forwarding strategy, delegation header represents a variable component that can be altered during packet forwarding.

The role of link objects is mainly for guidance, to provide global routing support on locally defined or routable content identifiers. Hence, if link objects are implemented, they are consulted by the ICN enabled routers only when forwarding lookup on content identifiers returns no match on the forwarding information base.

3.3. Link Object vs. Forwarding Label

Next we list the major differences between a link object and a forwarding label.

- o Link objects are set by the end host's forwarding daemon with certain level of trust associated to it, restricting the link component to be immutable during forwarding. Forwarding labels are initially set by the ICN edge routers or the end-host applications and allow mutable operation through late-binding during Interest forwarding. In doing so, forwarding label offers the ability of network based management during Interest forwarding, which allows each domain to perform packet forwarding according to its administrative and service policies. Note that, it is possible for link objects to trigger network based management operations, however their impact would be limited, in the worst case triggering NACKs that may prevent the use of link objects to help with forwarding.
- o Link objects constrain a network operator from overriding a consumer's intent, which, in some cases, may potentially lead to better performance compared to forwarding over native network service provider paths. Forwarding labels require additional mechanisms to support such feature, for instance, to enable the desired path for the consumer, which may necessitate the use of additional forwarding labels within requests.
- o For the link objects, security binding is mandatory and trust relationship is established by default, by putting all the trust assessment at the end hosts. On the other hand, security binding is optional for the forwarding label, which allows the use of trust association to bind AI to the NI depending on the context associated with its use. However, without appropriate support in trust management, forwarding label use may introduce problems such as route hijacking, hence contextual management should be capable of addressing such challenges using, for instance, approaches identified in [FWLDR].
- o Another difference is related to the processing of forwarding labels and link objects at the ICN routers. A link object is processed only if the router cannot find a matching FIB entry for the content identifier limiting routing flexibility. Furthermore,

if no matching FIB entry is found, link objects would trigger additional lookups on the FIB, leading to efficiency issues with frequent occurrences. On the other hand, forwarding label is processed before a content identifier, if its use is enabled, hence presents a more flexible and efficient operation in routing.

- o Link object can be considered as a hint towards where to find the content, and since it is processed after FIB lookup on the content identifier fails, it typically leads to lower computational and bandwidth efficiency. Forwarding label, on the other hand, can be considered as an enabler for faster packet processing at the ICN routers as it allows bypassing content/application identifier based processing at the supported routers, while at the same offering optimized routing towards a content source.
- o Link object is considered as an application driven component and network service agnostic, thereby allowing the network to decide on its use. Forwarding label, on the other hand, can be enabled as part of a service, which limits the use of forwarding label to the supported namespaces, while at the same time requiring its use whenever/wherever supported. For instance, within the context of virtualizing the ICN network among multiple services, where compute, storage, and bandwidth resources are shared among these services, ICN service edge routers can apply rules on namespaces to decide on how to dedicate network resources. One example of such service is the mobility service, which can utilize forwarding labels, to provide stricter service quality guarantees for the end hosts. In such case, if the namespace requires mobility support, forwarding label is used in effect to achieve more efficient forwarding. Note that, service use can be triggered with the use of service tags integrated within forwarding labels, once validated to be used with the corresponding namespace(s).
- o As a link object can encode multiple routing hints, it can direct a request towards multiple identifier locations, giving an ICN router the option to choose any one of them based on the router's forwarding strategy. Even though this selection is shared between consecutive routers, it is not enforced, thereby potentially leading to non-optimal forwarding paths. Forwarding label, on the other hand, is enforced consistently at consecutive hops within a domain whenever/wherever its use is supported and/or enabled. Hence, forwarding label presents the network with the ability to consistently forward packets over optimal paths towards a content source (with respect to routers forwarding the requests towards the same direction, rather than choosing alternating destinations).

4. Name Resolution System Considerations

To manage the AI to NI mapping, we need a name resolution system (NRS). In addition to exposing APIs to application to register its name to the NRS, it should also scale and work efficiently considering the scale of named resources that need to be published, resolved, removed, and updated at high frequency, for instance, corresponding to high-speed mobility scenarios.

The following are the design choices for the NRS:

- o Hierarchical System: Here, AI to NI mapping is managed by the application providers, but similar to DNS, the service has to sync its name reachability information with high level name resolvers. NDN-DNS (NDNS) is an example of such a system [NDNS], which utilizes a zone-based hierarchy (i.e., root level, top-level domain, etc.) and which is queried iteratively at every component level of the content/application identifier (e.g., /tld/sld would trigger iterative requests of /dns/tld/NS and /dns/sld/NS, at the root and tld levels). NDNS also supports recursive queries to scope the route requests for a content/application identifier. Even though the design of NDNS supports forwarding of Interests on content/application identifiers not present in the FIB, its design is typically suitable in cases when resources are static, rather than for highly dynamic systems such as ICN, where replication and mobility will be the norm. Supporting mobility in NDNS may require frequent updates and requests to setup and identify routes towards mobile entities, which may lead to performance-related problems. Also, such system has to scale to resolve not just the end hosts, which represents the current use, but also the information objects.
- o Network-Integrated Flat System: Here, resolution service is integrated within the ICN infrastructure, where the router contributes a part of its compute and storage resources to enable this service. This integration allows multiple ways of designing a generic name resolution service, similar to the overlaid or in-network designs considered for Global Name Resolution Service (GNRS) in MobilityFirst [GNS] [ASPC] [GNRS] allowing for good scalability performance with proven handling of dynamic updates, aided by a separation of entity identifiers from network identifiers. In GNRS, flat names are queried to obtain the corresponding self-certifying identifiers, such as the network address, before forwarding an Interest for the flat globally-unique identifier.
- o Distributed System: Compared to a flat resolution system, this type of architecture preserves the contextual nature of DNS, by

using the context in the content identifier (such as the network or host identifier portion of the name) to identify a resolution server corresponding to the context information, such as home controller, which stores the mappings associated with registered names carrying the controller's context information and where the respective AI to NI mapping can be resolved. Such a system removes the need for the home controllers to sync up with high level resolvers, as for successful resolution it is sufficient for each controller to manage the names registered to or under it. For instance, /company/content-id would be mapped with a local resolver, identified with /company/resolver-id, that manages any namespace registered under its domain identifier (/company). In doing so, content mobility can effectively be handled through localized updates (for intra-domain mobility) or remote updates at the home controller (for inter-domain mobility) with minimal signaling overhead, while maintaining global scalability.

5. Differences with respect to Existing IP-based Proposals

To address persistent identity, routing scalability, multihoming, and mobility limitations of the current IP, various incremental solutions have been proposed, among which identifier/locator split emerged as the key solution to address these challenges [RFC4984]. Here, we specifically focus on three of these solutions: (i) Host Identity Protocol (HIP) [HIP], (ii) Identifier-Locator Network Protocol (ILNP) [ILNP], and Locator/Identifier Separation Protocol (LISP) [RFC6830]. HIP and ILNP achieve ID/locator separation and binding at the host level whereas LISP achieves that at the network level (i.e., at the network edge using service routers).

In HIP, public cryptographic keys are used as host identifiers, which provide the binding to higher layer protocols instead of IP addresses [RFC7401]. ILNP divides IP namespace into two distinct namespaces of identifiers and locators, each of which carrying distinct semantics with identifier representing the non-topological name for the host and locator representing the topologically bound name for the network [RFC6740]. LISP is a map-and-encap type protocol, which achieves id/locator separation by defining (i) endpoint identifiers, which are used for routing at the access network and which represent the IP address for the host, and (ii) routing locators, which are used for routing at the core and which represent the IP address for the egress routers.

These protocols fundamentally differ from ICN's objective to define a new network layer, where name based routing, location independent caching, mobility, multihoming, and multi-path routing are the integral features. More specifically, this draft proposes to enable

AI/NI binding as a network service to allow efficient routing of user requests depending on the application context.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

6.2. Informative References

- [AFWD] Yi, C., Afanasyev, A., Wang, L., Zhang, B., and L. Zhang, "Adaptive Forwarding in Named Data Networking", ACM CCR, Jul 2012.
- [ASPC] Sharma, A., Tie, X., Uppal, H., Venkataramani, A., Westbrook, D., and A. Yadav, "A Global Name Service for a Highly Mobile Internetwork", ACM SIGGCOM, 2014.
- [CCN] Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggs, N., and R. Braynard, "Networking Named Content", ACM CoNEXT, 2009.
- [FWLDR] Ravindran, R., Chakraborti, A., and A. Azgin, "Forwarding Label Support in CCN Protocol", draft-ravi-icnrg-ccn-forwarding-label-01, July, 2017.
- [FWLRP] Azgin, A., Ravindran, R., and G. Wang, "A Scalable Mobility-Centric Architecture for Named Data Networking", IEEE ICCCN Scene Workshop, 2014.
- [GNRS] Hu, Y., Yates, R., and D. Raychaudhuri, "A Hierarchically Aggregated In-Network Global Name Resolution Service for the Mobile Internet".
- [GNS] Venkataramani, A., Sharma, A., Tie, X., Uppal, H., Westbrook, D., Kurose, J., and D. Raychaudhuri, "Design Requirements for a Global Name Service for a Mobility-Centric, Trustworthy Internetwork", IEEE COMSNETS, 2013.
- [HIP] 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.

- [ILNP] Atkinson, R., "An Overview of the Identifier-Locator Network Protocol (ILNP)", Technical Report, University College London, 2005.
- [MAAS] Azgin, A., Ravindran, R., Chakraborti, A., and G. Wang, "Seamless Producer Mobility as a Service in Information Centric Networks", ACM ICN IC5G Workshop, 2016.
- [MDHT] Liu, H., Foy, X., and D. Zhang, "A Multi-level DHT routing Framework with Aggregation", ACM SIGCOMM ICN Workshop, 2012.
- [MFRST] Venkataramani, A., Kurose, J., Raychaudhuri, D., Nagaraja, K., Mao, M., and S. Banerjee, "MobilityFirst: A Mobility-centric and Trustworthy Internet Architecture", ACM SIGCOMM CCR, 2014.
- [MPNCP] Afanasyev, A., Yi, C., Wang, L., Zhang, B., and L. Zhang, "Map-and-Encap for Scaling NDN Routing", NDN Technical Report, ndn-004-02, 2015.
- [NAMES] Baid, A., Vu, T., and D. Raychaudhuri, "Comparing Alternative Approaches for Networking of Named Objects in the Future Internet", IEEE INFOCOM NOMEN Workshop, 2012.
- [NCMP] Adhatarao, S., Chen, J., Arumaithurai, M., Fu, X., and K. Ramakrishnan, "Comparison of Naming Schema in ICN", IEEE LANMAN, 2016.
- [NDNS] Afanasyev, A., "Addressing Operational Challenges in Named Data Networking Through NDNS Distributed Database", 2013.
- [QCMP] Gao, Z., Venkataramani, A., Kurose, J., and S. Heimlicher, "Towards a Quantitative Comparison of Location-Independent Network Architectures", ACM SIGCOMM, 2014.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <<http://www.rfc-editor.org/info/rfc2629>>.
- [RFC3220] Perkins, C., "IP Mobility Support for IPv4", RFC 3220, 2002.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<http://www.rfc-editor.org/info/rfc3552>>.

- [RFC4984] Meyer, D., Zhang, L., and K. Fall, "Report from the IAB Workshop on Routing and Addressing", RFC 4984, 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6740] Atkinson, R. and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, 2012.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, 2013.
- [RFC7401] Moskowitz, R., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, 2015.
- [RFC7927] Kutscher, D., Eum, S., Pentikousis, K., Psaras, I., Corujo, D., Saucez, D., Schmidt, T., and M. Waehlisc, "Information-Centric Networking (ICN) Research Challenges", RFC 7927, 2016.
- [SFWD] Yuan, H., Song, T., and P. Crowley, "Scalable NDN Forwarding: Concepts, Issues and Principles", IEEE ICCCN, 2012.
- [SNAMP] Afanasyev, A., Yi, C., Wang, L., Zhang, B., and L. Zhang, "SNAMP: Secure Namespace Mapping to Scale NDN Forwarding", IEEE Global Internet Symposium, 2015.
- [SPIT] Yuan, H. and P. Crowley, "Scalable Pending Interest Table Design: From Principles to Practice", IEEE INFOCOM, 2014.
- [VRSGN] "Verisign Domain Name Industry Brief", July 2016.
- [WPIT] Varvello, M., Perino, D., and L. Linguaglossa, "On the Design and Implementation of a Wire-speed Pending Interest Table", IEEE INFOCOM NOMEN Workshop, 2013.

Appendix A. Additional Stuff

This becomes an Appendix.

Authors' Addresses

Aytac Azgin
Huawei Technologies
Santa Clara, CA 95050
USA

Email: aytac.azgin@huawei.com

Ravishankar Ravindran
Huawei Technologies
Santa Clara, CA 95050
USA

Email: ravi.ravindran@huawei.com