

ICN Research Group
Internet-Draft
Intended status: Informational
Expires: September 30, 2015

A. Afanasyev
UCLA
R. Ravindran
G. Wang
Huawei Technologies
L. Wang
University of Memphis
B. Zhang
University of Arizona
L. Zhang
UCLA
March 29, 2015

ICN Packet Format Design Requirements
draft-icn-packet-format-requirements-01

Abstract

It is a great challenge to design the right packet format for the Information-Centric Networking (ICN), because this new architecture is still in its research stage. We do not have good prediction regarding either the technology advances or the application needs may involve in next 10 years and beyond. To minimize the danger of premature constraints in the packet format design, we believe it is important to first clearly identify the design requirements, so that we can use the requirements to guide the design effort. In this document we describe our understanding of the design requirements and their tradeoffs.

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 September 30, 2015.

Copyright Notice

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

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

Table of Contents

1.	Introduction	2
2.	Requirements for ICN Packet Format Design	3
2.1.	Universality (Elasticity)	3
2.2.	Flexibility and Extensibility	4
2.3.	Processing Efficiency	5
2.3.1.	Decoding and Encoding Efficiency	5
2.3.2.	Continuity of Security Envelope	6
2.3.3.	Preservation of Network-Level Processing	6
2.4.	Auditability / Robust Design	6
3.	Separating Information-Centric and Multi-Hop Forwarding Element in ICN Packet Format Design	7
4.	ICN Packet Format Success Factors (A Few Lessons from Today's Internet)	8
5.	Security Considerations	9
6.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

The proposed Information-Centric Networking (ICN) architectures (TRIAD [\[4\]](#), DONA [\[5\]](#), NDN [\[6\]](#) [\[7\]](#) [\[8\]](#) [\[3\]](#), CCNx [\[9\]](#), and others) put forward an objective to become a new universal narrow waist for local and global communication. However, all these proposals are still in research stage with many details yet to be specified, many questions yet to be answered, and many more issues yet to be identified and resolved through experimentation.

It is a great challenge to design the right packet format for ICN that can satisfy all desired research goals including both those that have been identified now and those that are yet to be recognized further down the road. This document is our first attempt to

formalize the general requirements for the ICN packet format design by leveraging past experience and following the framework defined in [1]. Clearly identifying the design requirements can help guiding the decision process in evaluating design tradeoffs and avoiding premature constraints. In this draft we do not discuss any specific packet format proposals, instead our focus is on identifying fundamental requirements on the packet format design, and on understanding different design tradeoffs.

Some of the requirements listed in [Section 2](#) may conflict with each other; how to prioritize such conflicting requirements determines the outcome of the packet format design. In this draft we list our identified requirements and put them in a priority order as the best match to the current and future needs of ICN protocols as we can see at this time. We hope that this draft contributes to the ongoing discussion about ICN packet format design.

2. Requirements for ICN Packet Format Design

This section presents an ordered list of identified requirements for the ICN packet format. In each sub-section we briefly describe potential design space that satisfies the requirement and discuss the tradeoffs between different design choices.

2.1. Universality (Elasticity)

As the new narrow waist of the Internet, the packet format must be able to support a diversity of usage scenarios and underlying network technologies: from highly constrained IoT environments [2] to ultra high-speed network channels [13].

A single IP packet format has served us well in the past. However, continued expansion of Internet usage frontier in recent years has shown the constraints of the existing IP packet formats with fixed header format. As one example: due to its fixed length and limited size, the shortage of IPv4 address space called for the design and deployment of IPv6. Although the design of IPv6 has long been done, its deployment is slow in coming. As we discuss later in [Section 4](#), some earlier sketch of IP protocol design proposed variable address fields, but was changed to fixed length to ease implementation; compared to the variable length address design, this ease of implementation has costed both unnecessarily long IP header when the Internet was small in its early days, and the global effort of moving from IPv4 to IPv6 in the last 10 years. As another example, the design of IPv6 format aimed to provide sufficiently large address space for foreseeable future, but it leads to the concerns of excessive overhead when used in emerging low-power constrained networks (IoT), which were not foreseen in early 1990s when the

initial IPv6 design was sketched out. In turn, IETF had to launch 6LoWPAN effort to address this issue together with a number of others caused by the constrained environment.

These observations suggest that the ICN format design should avoid similar issues. We would like to see that the ICN packet format allows applications to choose their packet size in whichever way as they see fit, from a few tens of bytes in constrained environments to jumbograms to take advantage of future high capacity transport and increase in network MTU capability. One way to achieve this goal is to ensure that packet length (as well as the lengths of other packet format fields) uses a variable-length encoding scheme to allow different packet sizes.

2.2. Flexibility and Extensibility

Given the experimental nature of the ongoing ICN architectural research, the packet format should stay flexible in order to allow new elements to be added to the format as they are identified, as well as to remove any existing elements that are no longer deemed necessary. The required fields in the packet format should also be kept as minimal as possible, as our understanding of the requirements are likely to change over time.

BGP [10] is one of the successful protocols that has gone through a series of evolutionary changes. ([TODO: cite other examples]) After going through a few rapid version changes in its early days, BGP4 adopted the flexible protocol encoding format of type-length-value (TLV), which allows BGP to encode a new function by simply introducing a new TLV block; old TLV blocks can be deprecated over time without a flag day for base protocol format alterations.

Thus, the use of TLV encoding can offer the potential to provide protocol flexibility and extensibility (and the existing ICN protocol designs do already).

Observations over past protocol design experience also show that to ease the introduction of new features, one can encode the information about desired actions to be taken by any entity which does not recognize the newly defined feature. For example, optional TLVs in the extension header of IPv6 [11] reserve first two bits of the type to specify the action to be taken by routers when an unrecognized TLV is encountered: skip over and continue processing or discard the packet, and/or send an error report to the source of the packet. Similarly, SCTP design [12] also reserves two bits to indicate the desired processing for packets with unrecognized chunks. ICN packet format may consider to adopt a similar mechanism by reserving some bits from the type field; an alternative scheme, which does not

require dedicated bits for the purpose, is to split the type code space into different bins, and to assign each type code into the bin based on the desired actions to be taken by entities which do not recognize the new TLV type.

2.3. Processing Efficiency

It is desirable to design the ICN packet format in a way that allows efficient packet processing to the extent possible. However making the format for processing efficiency often runs into conflict with elasticity and flexibility requirements. For example, the latter leads the format to contain variable parts, which in turn lead to higher processing cost as compared to formats with fixed length encoding. Similarly, if the packet format contains a fixed part, that part becomes unchangeable, without a global flag day for making the changes.

Furthermore, when we consider processing efficiency, our thinking should not be constrained by either the currently available technology or the currently understood implementation approach. History showed us that technologies can advance quickly, especially when there is a high demand from the market. Similarly, new implementation approaches are also discovered in response to the need, providing efficient solutions to problems that once were considered infeasible. For example, the switch to classless inter-domain routing (CIDR) during 1990's created challenges for line speed router implementations [TODO: cite], and that challenge was quickly resolved.

Therefore, although the processing efficiency requirement on the format design is an important one, we believe it should be put below the elasticity and flexibility considerations, i.e., the design should not sacrifice the long-term protocol flexibility for short-term efficiency gains.

The rest of this section discusses a few specific approaches to improve processing efficiency.

2.3.1. Decoding and Encoding Efficiency

Decoding and encoding of ICN packet format should be as efficient as possible. Packet format should allow partial decoding of the packet, efficient access to individual fields, and efficient skipping over unprocessed (e.g., application-defined or unrecognized) fields. It is also desirable for the key fields needed to make packet forwarding decision (e.g., name) to appear in the beginning of the packet.

2.3.2. Continuity of Security Envelope

Cryptographic operations should be defined over a continuous packet block, containing one or multiple complete TLV elements. In other words, the security envelope must be defined over a continuous sequence of whole TLV elements.

2.3.3. Preservation of Network-Level Processing

Application-defined elements added to the packet format should not increase routers' processing cost. For example, when an application adds custom elements into the packet format, these elements must go to an application-specific group that is always skipped by routers, or should be clearly separated from router-processed elements and can be easily skipped without processing.

2.4. Auditability / Robust Design

Generally speaking, TLV-encoded packet format includes multiple nested TLV blocks. To provide ability to audit packet encoding without requiring full understanding of semantics of each nested TLV level, it is highly desirable to assign unique type codes to each element in the packet format that is processed by routers. In other words, we would like to require unique type code assignment to all network level TLV elements. This requirement may not apply to vendor-specific or application-specific type codes, whether they may, or may not reuse type code space.

One of the advantages of using unique type code assignment is reduction of potential implementation errors. Unique types across multiple TLV levels is also important for network traffic debugging tools similar to tcpdump and wireshark: with unique type codes for all network-level elements, implementation of those tools would be simple with potential of decoding known TLV types at any level of nesting. When the same type codes are reuse (at different TLV nesting levels), such tools also can still be implemented, but potentially with increased complexity, as it would be required to define semantical contexts for each nested TLV block.

The tradeoff of the unique type code assignment is the coordination required in ensuring the type code uniqueness when doing code assignments.

3. Separating Information-Centric and Multi-Hop Forwarding Element in ICN Packet Format Design

In an ICN network, communication is performed using application-level information units. Therefore, the primary function of ICN packet format is to include all the elements that are necessary to describe requests for the data and to describe and secure the data, such as data name, data name constraints, payload, security context, security context constraint, etc. Note that these information-centric elements define the data as requested by consumers and enable communication between ICN consumers and producers in terms of the desired data.

At the same time, there is a separate class of protocol elements that an ICN packet may need to carry in order to aid the information retrieval over multi-hop environment; these elements are not related to the requested information itself. For example, a lifetime element may be needed to avoid a request staying inside the network for too long; a nonce field in the requests can aid the detection of request duplication; QoS labels in requests can enable support for prioritization in retrieving certain information classes; a NACK can be exchanged between adjacent nodes to notify problems in forward requests so that the previous hop routers can promptly try alternatives; and when a packet needs to be fragmented to cross a small MTU hop, certain fields must be added to the packet to facilitate the reassembly at the other end. Note that all of such elements are volatile by nature, they may be needed only in parts of the network, and the perceived need for some of them may change over time as we gain deeper understanding of ICN network operations through trial deployment in coming years.

Given the two separate classes of the elements, the question is how they should be encoded into an ICN packet format. A variety of design proposals exist today. One way of thinking is to clearly separate these two classes of protocol elements. Since the ICN elements are fundamental in information retrieval, they must be implemented in all consumers and producers. At the same time, the need for various forwarding elements may depend on specific forwarding scenarios and environments. This way of thinking leads to the consideration of providing the forwarding elements in a separate shim layer protocol, instead of including them in the base ICN protocol format.

Another way of thinking is to encode all the protocol elements into the based ICN protocol format based on their perceived necessity, and/or based on the processing efficiency consideration, but with no consideration on identifying and separating the two classes of protocol elements.

The basic question here is: what should be considered the narrow waist in an Information-Centric Networking architecture? In principle, the ICN elements are the minimally required ones to enable communication in an ICN network. They are the only ones needed for data flow between directly connected peers or between ICN applications running on the same node. When information retrieval crosses wide areas, additional elements become necessary to assist the information retrieval over multiple hops/networks.

The two ways of encoding information-centric and multi-hop forwarding elements in ICN packet format would define tradeoffs for implementation, usability, and future extensibility. The combined format can be simpler to implement and easier to process, but may lead to inclusion of unnecessary elements. For example, when data packets are cached inside the network and retrieved later by different consumers, most, if not all, forwarding elements may no longer be meaningful. Encoding the two classes of protocol elements gives flexibility of evolving the underlying format for transport function, but also creates a need for standardization, to avoid the danger of multiple incompatible formats being developed for the same multi-hop forwarding function.

In this document we intentionally do not nail down any specific answers to the question of how to combine information-centric and transport-related elements into the packet format. Our goal is to define a framework to clearly articulate the general requirements for the ICN packet format design and describe our understanding about the tradeoffs in the design space. We hope that a set of agreed upon requirements can provide a guideline to packet format design that can serve us many years in the future.

4. ICN Packet Format Success Factors (A Few Lessons from Today's Internet)

TODO

The version number field in the IP header did not help much in incremental introduction of IPv6 deployment.

History of IP address space design: IPv3 used a variable length encoding of IP addresses (IEN28, February 1978), which was changed to fixed length due to the implementers complaints that "it would be too complex to parse the variable length header" and demanded a fixed length header so they did not have to work their way through the header. In retrospect, after we saw the cost of IPv6 deployment, it is no longer clear the choice of fixed length is a good design choice at that time.

BGP stays with BGP4 after adopting TLV encoding.

Technology moves forward fast, this is an important design factor to be considered in ICN packet format design.

Given the diversity of the requirements, a single packet format may not be able to satisfy to the full extent all of the listed requirements. Depending on which requirement is prioritized, the resulting design for the packet format would have different tradeoffs between universality, flexibility, and efficiency.

5. Security Considerations

TODO

6. Informative References

- [1] Clark, D., "The Design Philosophy of the DARPA Internet Protocols", Proceedings of SIGCOMM'88, August 1988.
- [2] "IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", June 2011.
- [3] NSF FIA project, NDN., "<http://www.named-data.net/>", 2010.
- [4] Cheriton, D. and M. Gritter, "TRIAD: A new next-generation Internet architecture", 2000.
- [5] Koponen, T. and et al., "A data-oriented (and beyond) network architecture", ACM SIGCOMM Computer Communication Review, Vol. 37, No. 4, 2007.
- [6] Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggs, N., and R. Braynard, "Networking named content", Proc. of CoNEXT, 2009.
- [7] Zhang, L. and et al., "Named data networking (NDN) project", NDN Tech. Rep. NDN-0001, October 2010.
- [8] Zhang, L., Afanasyev, A., Burke, J., Jacobson, V., claffy, kc., Crowley, P., Papadopoulos, C., Wang, L., and B. Zhang, "Named Data Networking", ACM Computer Communication Reviews, July 2014.
- [9] M. Mosko, , "CCNx Semantics, [draft-mosko-icnrg-ccnxsemantics-00](#)", January 2015.

- [10] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [11] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [12] Stewart, R., "Stream Control Transmission Protocol", [RFC 4960](#), September 2007.
- [13] ITU-T, , "Interface for the Optical Transport Network (OTN)", G.709 Recommendation, February 2012.

Authors' Addresses

Alexander Afanasyev
UCLA
4809 Boelter Hall, UCLA
Los Angeles, CA 90095
US

Email: afanasev@cs.ucla.edu

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

Email: ravi.ravindran@huawei.com

Guoqiang Wang
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
US

Email: gq.wang@huawei.com

Lan Wang
University of Memphis
318 Dunn Hall
Memphis, TN 38152
US

Email: lanwang@memphis.edu

Beichuan Zhang
University of Arizona
Gould-Simpson 723
Tucson, AZ 85721
US

Phone: +1 520 621 4817
Email: bzhang@cs.arizona.edu

Lixia Zhang
UCLA
3713 Boelter Hall, UCLA
Los Angeles, CA 90095
US

Phone: +1 310 825 2695
Email: lixia@cs.ucla.edu

