

ICNRG
Internet-Draft
Intended status: Experimental
Expires: February 27, 2017

S. Mastorakis
UCLA
J. Gibson
I. Moiseenko
R. Droms
D. Oran
Cisco Systems
August 26, 2016

ICN Ping Protocol
draft-mastorakis-icnrg-icnping-00

Abstract

This document presents the design of an ICN Ping protocol. This includes the operations both on the client and the forwarder side.

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 February 27, 2017.

Copyright Notice

Copyright (c) 2016 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
1.1. Requirements Language	2
2. Background on IP-Based Ping Operation	2
3. Ping Functionality Challenges and Opportunities in ICN	3
4. ICN Ping Echo Packet Formats	5
4.1. ICN Ping Echo Request Packet Format	5
4.2. Ping Echo Reply Packet Format	8
5. Forwarder Handling	11
6. Security Considerations	12
7. Acknowledgements	12
8. References	12
8.1. Normative References	12
8.2. Informative References	12
Appendix A. Ping Client Application (Consumer) Operation	13
Authors' Addresses	14

1. Introduction

Determining data plane reachability to a destination and taking coarse performance measurements of round trip time are fundamental facilities for network administration and troubleshooting. In IP, where routing and forwarding are based on IP addresses, ICMP echo and ICMP echo response are the protocol mechanisms used for this purpose, generally exercised through the familiar ping utility. In ICN, where routing and forwarding are based on name prefixes, the ability to determine reachability of names is required.

This document proposes protocol mechanisms for a ping equivalent in ICN networks. A non-normative appendix suggests useful properties for an ICN ping client application, analogous to IP ping, that originates echo requests and process echo replies.

1.1. 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 [RFC2119].

2. Background on IP-Based Ping Operation

In IP-based ping, an IP address is specified, either directly, or via translation of a domain name through DNS. The ping client application sends a number of ICMP Echo Request packets with the

specified IP address as the IP destination address and an IP address from the client's host as the IP source address.

An ICMP Echo Request is forwarded across the network based on its destination IP address. If it eventually reaches the destination, the destination responds by sending back an ICMP Echo Reply packet to the IP source address from the ICMP Echo Request.

If an ICMP Echo Request does not reach the destination or the Echo reply is lost, the ping client times out. Any ICMP error messages, such as "no route to destination", generated by the ICMP Echo Request message are returned to the client and reported.

3. Ping Functionality Challenges and Opportunities in ICN

In ICN protocols (e.g., NDN and CCNx), the communication paradigm is based exclusively on named objects. An Interest is forwarded across the network based on its name. Eventually, it retrieves a content object either from a producer application or some forwarder's Content Store (CS).

IP-based ping was built as an add-on on top of an already existing network architecture. In ICN, we have the opportunity to incorporate diagnostic mechanisms directly in the network layer protocol, and hopefully provide more powerful diagnostic capability than can be realized through the layered ICMP Echo approach.

An ICN network differs from an IP network in at least 4 important ways:

- o IP identifies interfaces to an IP network with a fixed-length number, and delivers IP packets to one or more interfaces. ICN identifies units of data in the network with a variable length name consisting of a list of components.
- o An IP-based network depends on the IP packets having source IP addresses that are used as the destination address for replies. On the other hand, ICN Interests do not have source addresses and they are forwarded based on names, which do not refer to a unique end-point. Data packets follow the reverse path of the Interests based on hop-by-hop state created during Interest forwarding.
- o An IP network supports multi-path, single destination, stateless packet forwarding and delivery via unicast, a limited form of multi-destination selected delivery with anycast, and group-based multi-destination delivery via multicast. In contrast, ICN supports multi-path and multi-destination stateful Interest forwarding and multi-destination data delivery to units of named

data. This single forwarding semantic subsumes the functions of unicast, anycast, and multicast. As a result, consecutive (or retransmitted) ICN Interest messages may be forwarded through an ICN network along different paths, and may be forwarded to different data sources (e.g., end-node applications, in-network storage) holding a copy of the requested unit of data. This can lead to a significant variance in round-trip times, which might not be desirable in the case of a network troubleshooting mechanism like ping.

- o In the case of multiple Interests with the same name arriving at a forwarder, a number of Interests may be aggregated in a common Pending Interest Table (PIT) entry. Depending on the lifetime of a PIT entry, the round-trip time an Interest-Data exchange might significantly vary (e.g., it might be shorter than the full round-trip time to reach the original content producer). To this end, the round-trip time experienced by consumers might also vary.

These differences introduce new challenges, new opportunities and new requirements in the design of an ICN ping protocol. Following this communication model, a ping client should be able to express ping echo requests with some name prefix and receive responses.

Our goals are the following:

- o Test the reachability and the operation of an ICN forwarder.
- o Test the reachability of an application (in the sense of whether Interests for a prefix that it serves can be forwarded to it) and discover the forwarder with local connectivity to (an instance of) the application.
- o Test whether a specific named object is cached in some on-path CS, and, if so, return the corresponding forwarder.
- o Perform some simple network performance measurements.

To this end, a ping name can represent:

- o An administrative name that has been assigned to a forwarder.
- o A name that includes an application's namespace as a prefix.
- o A named object that might reside in some in-network storage.

In order to provide stable and reliable diagnostics, it is desirable that the packet encoding of a ping echo request enables the forwarders to distinguish a ping from a normal Interest, while also

allowing for forwarding behavior to be as similar as possible to that of an Interest packet. In the same way, the encoding of a ping echo reply should allow for forwarder processing similar to that used for data packets.

The ping protocol should also enable relatively stable round-trip time measurements. To this end, it is important to have a mechanism to steer consecutive ping echo requests for the same name towards a common path.

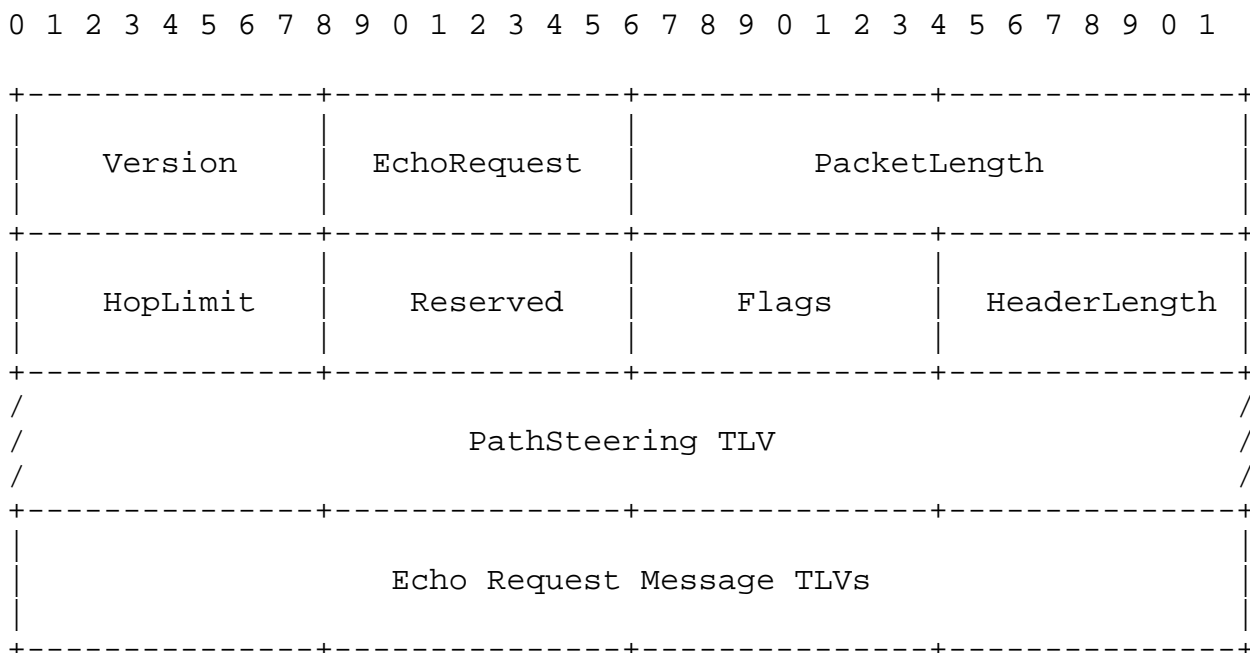
It is also important, in the case of ping echo requests for the same name from different sources, to have a mechanism to avoid aggregating those requests in the PIT. To this end, we need some encoding in the ping echo requests to make each request for a common name unique, and hence avoid PIT aggregation and further enabling the exact matching of a response with a particular ping packet.

4. ICN Ping Echo Packet Formats

Based on the goals mentioned in the previous section, we propose two types of ping packets, an echo request and an echo reply packet type. Both these packets follow the CCNx packet format [CCNMessages], where messages exist within outermost containments (packets).

4.1. ICN Ping Echo Request Packet Format

The format of the ping echo request packet is presented below:

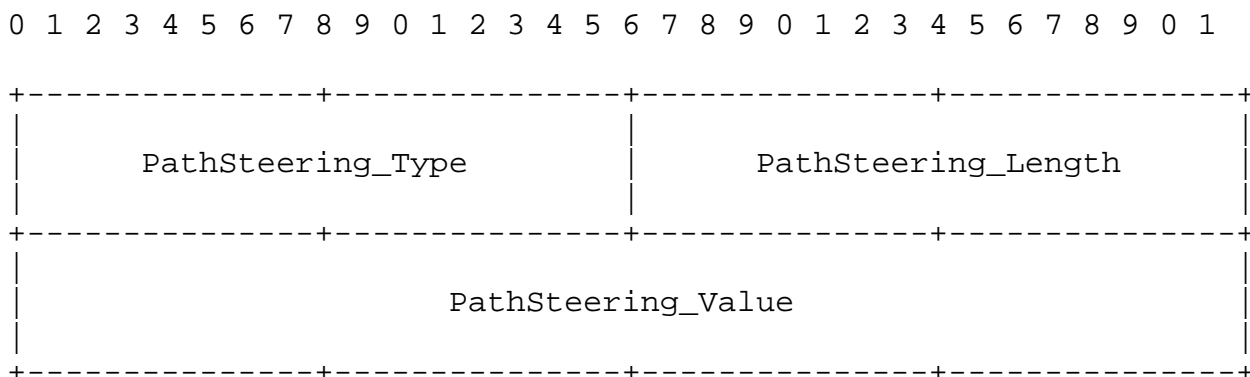


Echo Request Packet Format

The existing packet header fields have similar functionality to the header fields of a CCNx Interest packet. The value of the packet type field is Echo Request. The exact numeric value of this field type is to be determined.

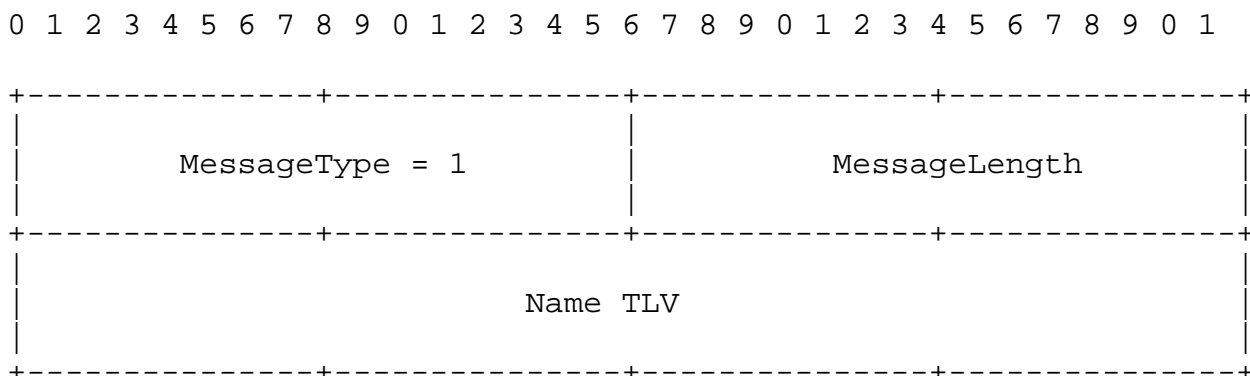
Compared to the typical format of a CCNx packet header [CCNMessages], there is a new optional fixed header TLV added to the packet header:

- o A PathSteering hop-by-hop header TLV, which is constructed hop-by-hop in the echo reply and included in the echo request to steer consecutive echo requests expressed by a ping client towards a common forwarding path. An example of such a scheme is presented in [LIPSIN].



PathSteering TLV


The message of an echo request is presented below:



Echo Request Message Format

The echo request message is of type Interest in order to leverage the Interest forwarding behavior provided by the network. The Name TLV has the structure described in [CCNMessages]. The name consists of the prefix that we would like to ping appended with a nonce typed name component as its last component. The value of this TLV will be a 64-bit nonce. The purpose of the nonce is to avoid Interest aggregation and allow client matching of replies with requests. As described below, the nonce is ignored for CS checking.

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

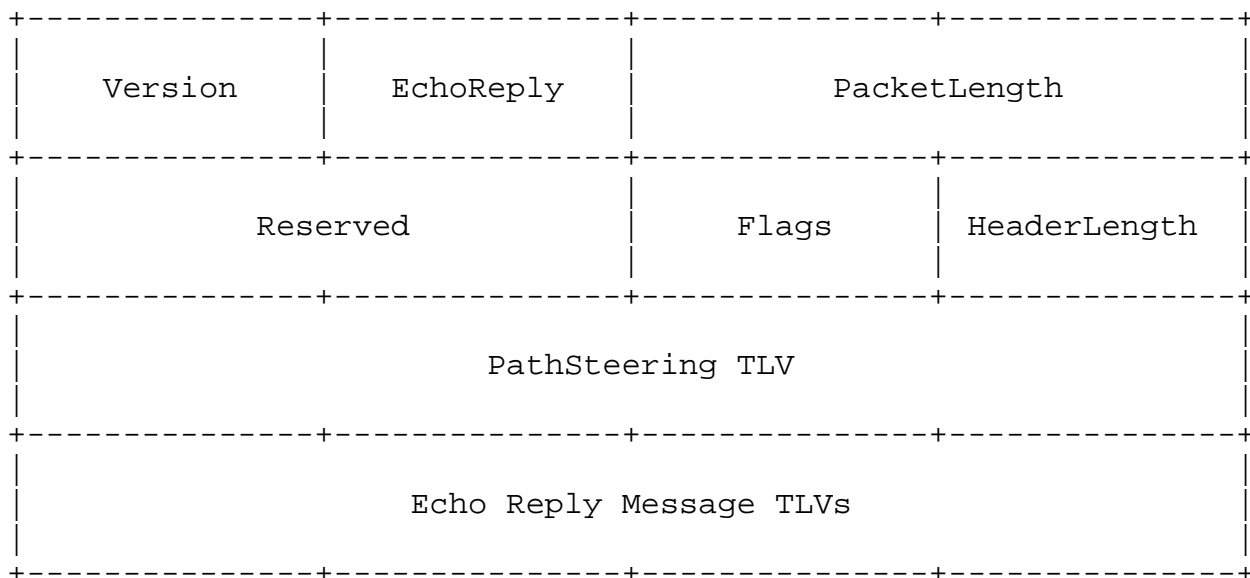


Nonce Typed Name Component TLV

4.2. Ping Echo Reply Packet Format

The format of a ping echo reply packet is presented below:

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

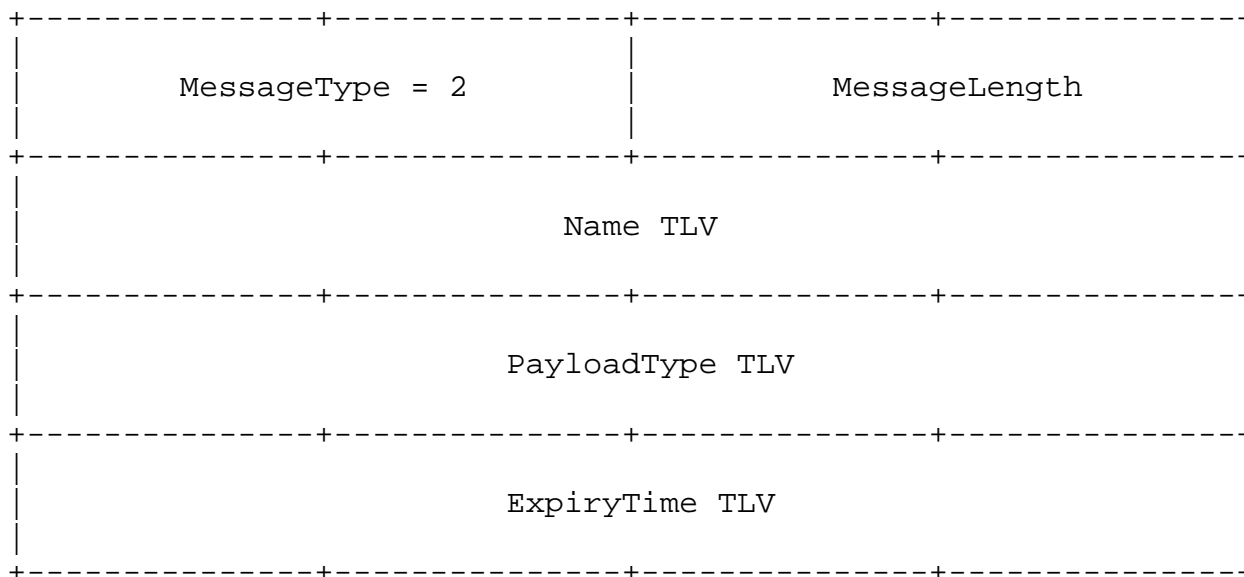


Echo Reply Packet Format

The header of an echo reply consists of the header fields of a CCNx Content Object and a hop-by-hop PathSteering TLV. The value of the packet type field is Echo Reply. The exact numeric value of this field type is to be determined. The PathSteering header TLV is as defined for the echo request packet.

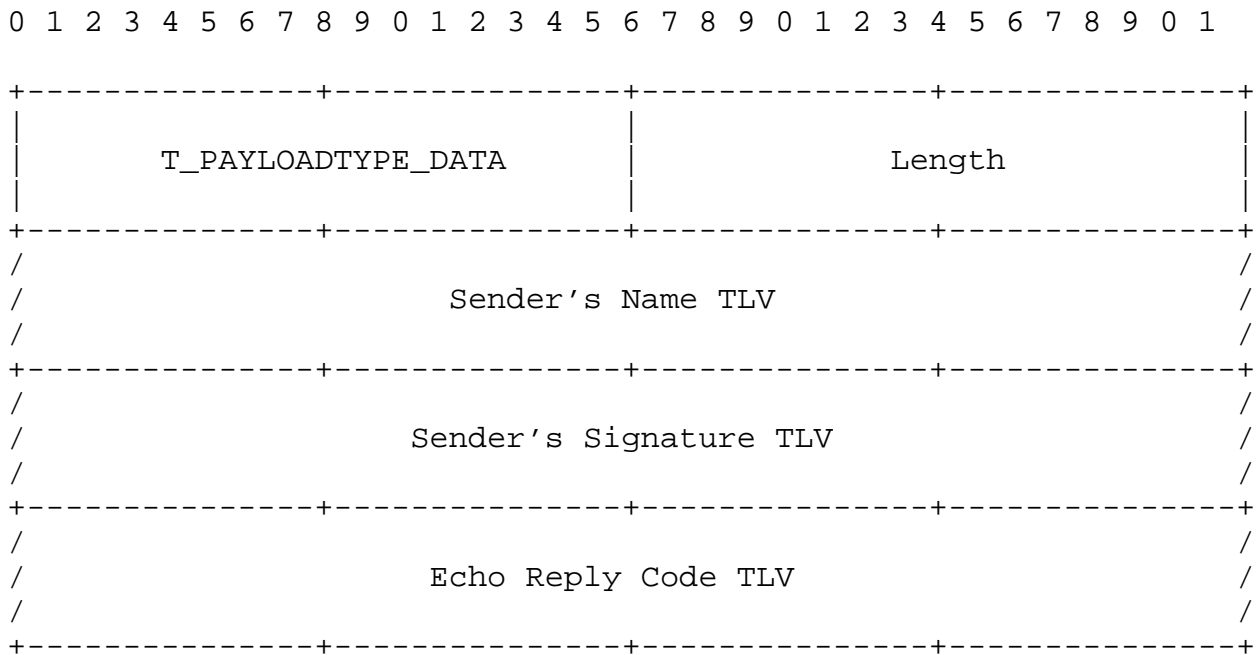
A ping echo reply message is of type Content Object, contains a Name TLV (name of the corresponding echo request), a PayloadType TLV and an ExpiryTime TLV with a value of 0 to indicate that echo replies must not be cached by the network.

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1



Echo Reply Message Format

The PayloadType TLV is presented below. It is of type T_PAYLOADTYPE_DATA, and the data schema consists of 2 TLVs: 1) the name of the sender of this reply (with the same structure as a CCNx Name TLV), 2) the sender's signature of their own name (with the same structure as a CCNx ValidationPayload TLV), 3) a TLV with return codes to indicate what led to the generation of this reply (i.e., existence of a local application, a CS hit or a match with a forwarder's administrative name as specified in Section 5).




Echo Reply Message Format

The goal of including the name of the sender in the echo reply is to enable the user to reach this entity directly to ask for further management/administrative information using generic Interest-Data exchanges after a successful verification of the sender's name.

The structure of the Echo Reply Code TLV is presented below (16-bit value). The potential values are the following:

- o 1: Indicates that the target name matched the administrative name of a forwarder.
- o 2: Indicates that the target name matched a prefix served by an application.
- o 3: Indicates that the target name matched the name of an object in a forwarder's CS.



Echo Reply Code TLV

5. Forwarder Handling

When a forwarder receives an echo request, it will first extract the message's base name (i.e., the request name with the Nonce name component excluded).

In some cases, the forwarder will originate an echo reply, sending the reply downstream through the face on which the echo request was received. An echo reply will include the forwarder's own name and signature, and, the appropriate echo reply code based on the condition that triggered the reply generation. It will also include a path steering TLV, initially a null value (since the echo reply originator does not forward the request and, thus, does not make a path choice).

The forwarder generates an echo reply in the following cases:

- o Assuming that a forwarder has been given one or more administrative names, the echo request base name exactly matches any of the forwarder's administrative name(s).
- o The echo request's base name exactly matches the name of a content-object residing in the forwarder's CS (unless the ping client application has chosen not to receive replies due to CS hits as specified in Appendix A).
- o The echo request base name matches (in a Longest Prefix Match manner) a FIB entry with an outgoing face referring to a local application.

If none of the conditions to reply to the echo request are met, the forwarder will attempt to forward the echo request upstream based on

the path steering value (if present) the results of the FIB LPM lookup and PIT creation (based on the name including the nonce typed name component). If no valid next-hop is found, an InterestReturn is sent downstream (as with a failed attempt to forward an ordinary Interest).

A received echo reply will be matched to an existing PIT entry as usual. On the reverse path, the path steering TLV of an echo reply will be updated by each forwarder to encode its next-hop choice. When included in subsequent echo requests, this path steering TLV will allow the forwarders to steer the requests along the same path.

6. Security Considerations

To avoid reflection attacks, where a compromised forwarder includes in the reply the name of a victim forwarder to redirect the future administrative traffic towards the victim, the forwarder that generates a reply has to sign the name included in the payload. In this way, the client is able to verify that the included name is legitimate and refers to the forwarder that generated the reply. Alternatively, the forwarder can include in the reply payload their routable prefix(es) encoded as a signed NDN Link Object [SNAMP].

7. Acknowledgements

The authors would like to thank Mark Stapp for the fruitful discussion on the objectives of ICN ping protocol.

8. References

8.1. 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>>.

8.2. Informative References

[CCNMessages]

Mosko, M., Solis, I., and C. Wood, "CCNx Messages in TLV Format.", 2016, <<https://tools.ietf.org/html/draft-irtf-icnrg-ccnxmessages-03>>.

[LIPSIN] Jokela, P. and et al, "LIPSIN: line speed publish/subscribe inter-networking, ACM SIGCOMM Computer Communication Review 39.4: 195-206", 2009.

[SNAMP] Afanasyev, A. and et al, "SNAMP: Secure namespace mapping to scale NDN forwarding, IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPs)", 2015.

Appendix A. Ping Client Application (Consumer) Operation

This section is an informative appendix regarding the proposed ping client operation.

The ping client application is responsible for generating echo requests for prefixes provided by users.

When generating a series of echo requests for a specific name, the first echo request will typically not include a PathSteering TLV, since no TLV value is known. After an echo reply containing a PathSteering TLV is received, each subsequent echo request can include the received path steering value in the PathSteering header TLV to drive the requests towards a common path as part of checking the network performance. To discover more paths, a client can omit the path steering TLV in future requests. Moreover, for each new ping echo request, the client has to generate a new nonce and record the time that the request was expressed. It will also set the lifetime of an echo request, which will have semantics similar to the lifetime of an Interest.

Moreover, the client application might like not to receive echo replies due to CS hits. A mechanism to achieve that would be to use a Content Object Hash Restriction TLV with a value of 0 in the payload of an echo request message.

When it receives an echo reply, the client would typically match the reply to a sent request and compute the round-trip time of the request. It should parse the PathSteering value and decode the reply's payload to parse the the sender's name and signature. The client should verify that both the received message and the forwarder's name have been signed by the key of the forwarder, whose name is included in the payload of the reply (by fetching this forwarder's public key and verifying the contained signature). The client can also decode the Echo Reply Code TLV to understand the condition that triggered the generation of the reply.

In the case that an echo reply is not received for a request within a certain time interval (lifetime of the request), the client should time-out and send a new request with a new nonce value up to some maximum number of requests to be sent specified by the user.

Authors' Addresses

Spyridon Mastorakis
UCLA
Los Angeles, CA
US

Email: mastorakis@cs.ucla.edu

Jim Gibson
Cisco Systems
Cambridge, MA
US

Email: gibson@cisco.com

Ilya Moiseenko
Cisco Systems
Cambridge, MA
US

Email: iliamo@mailbox.org

Ralph Droms
Cisco Systems
Cambridge, MA
US

Email: rdroms.ietf@gmail.com

Dave Oran
Cisco Systems
Cambridge, MA
US

Email: daveoran@orandom.net