

Workgroup: Network Working Group
Internet-Draft:
draft-trossen-rtgwg-routing-beyond-
reachability-01

Published: 30 June 2022

Intended Status: Informational

Expires: 1 January 2023

Authors: D. Trossen D. Lou
 Huawei Technologies Huawei Technologies
 S. Jiang
 Huawei Technologies

Continuing to Evolve Internet Routing Beyond 'Mere' Reachability

Abstract

This document discusses the evolution of the Internet routing system beyond mere reachability. We observe, through examples of past development, that such evolution has been taking place to improve on capabilities of the Internet, deal with more complicated network deployments and cater to changing requirements by end users as well as novel and emerging applications.

For achieving a routing system that serves more than a singular reachability purpose, more information is taken into account when performing the purpose-specific functions. Such extra information can be obtained by extending current routing protocols to exchange more information or by carrying that information within packets.

This document is intended to seed discussions of how the observed evolution of the Internet's routing system can continue, what issues may occur when simply continuing the current approach for achieving routing beyond 'mere' reachability and what may be needed to address those issues. Ultimately, however, this document recognizes the positive impact that moving beyond reachability has brought to the Internet and will continue to do so.

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 <https://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 1 January 2023.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Reachability - Original Purpose of the Routing System](#)
- [3. Extension of the Routing System Beyond 'Mere' Reachability](#)
- [4. Issues with Current Approaches](#)
 - [4.1. Limiting Routing Semantics](#)
 - [4.2. Complexity and Efficiency](#)
 - [4.2.1. Repetitive encapsulation](#)
 - [4.2.2. Introducing Path Stretch](#)
 - [4.2.3. Complicating Traffic Engineering](#)
 - [4.3. Security](#)
 - [4.4. Fragility](#)
 - [4.5. Interoperability](#)
- [5. The Driving Need for Evolving Communication Semantics](#)
- [6. Where to Go From Here?](#)
- [7. Security Considerations](#)
- [8. Privacy Considerations](#)
- [9. IANA Considerations](#)
- [10. Acknowledgements](#)
- [11. Contributors](#)
- [12. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

The current routing system was initially designed for a single purpose - reachability. That is, it was built to find paths through the network so as to forward packets to their destinations. The routing system has successfully supported the Internet as it grew

from a very small scale network to a giant system that covers the whole world with billions attached devices and users. This routing system has done a good job for global reachability, however, through the years, many other needs or purposes have arisen in the Internet, such as hostname/address mapping, destination selection, security, privacy, group isolation, various QoS requirements etc. Many of these additional needs or purposes have resulted in the development of extended and distinct systems, such as DNS, patched firewall, DPI, and CDN, etc. These systems have worked well but with costs in terms of quality of experience for the user, particularly with respect to time delay, but also with respect to costs of development, deployment and management throughout (parts of) the Internet.

An alternative approach is the integration of extra capabilities and purposes into the routing system directly. By exchanging necessary additional information or including such information in the packet header, purposes beyond just reachability have found entry into the routing system over the many years of the Internet's development.

This document presents a brief survey of solutions that, when combined, represent a routing system beyond reachability that effectively forms today's Internet. While this survey somewhat relates to that presented in [[I-D.king-irtf-semantic-routing-survey](#)], our focus here is on the identification of the underlying purpose for developing extensions, not on the body of work that represents an approach for doing so (named 'semantic routing' in the above draft). However, [[I-D.king-irtf-semantic-routing-survey](#)] may be useful for more information on the specific extensions.

Some of these extensions are intended to be deployed in limited domains [[RFC8799](#)], while others are intended for use across the Internet. The boundary of limited domains may also be the boundary of purposes and semantics of information defining those purposes. This survey is used to demonstrate the recognized need by those having developed existing solutions for the Internet's routing system to have multiple purposes beyond mere reachability.

Building on the survey and our summary, we recognize that, in many parts, the Internet has already evolved into a 'multi-purpose routing' system. However, we identify issues with the approach that has been taken so far, namely that of purpose-specific extensions. We assert that these issues will increasingly impede the Internet's ability to accommodate future purposes (represented in the form of new use cases), if we simply continue with a 'business as usual' attitude towards developing purpose-specific solutions for them.

Instead, we position this document as the starting point for a discussion on how to evolve the Internet routing system in a

coherent manner that will help us avoid the identified issues outlined in [Section 4](#), while still allowing for integrating evolving the semantics of communication along the lines outlined in [Section 5](#) towards new purposes for Internet routing as they will emerge in the future.

Note: This document does not discuss how routers may use policies, that are coded in, configured at, or signaled to it, to make routing decisions. It does neither pass comment on the advisability or practicality of any of the proposals nor does it define any technical solutions.

2. Reachability - Original Purpose of the Routing System

Network routing protocols were initially designed to enable forwarding of IP packets through the network toward destination addresses. Fundamental to this is the locator semantic of IP addresses, which has been assigned in the context of network topologies. The original routing system was designed on a distributed basis. Each router makes its own decision about the interface/link onto which it forwards a packet. Each decision takes the packet one hop closer to the destination. However, the choices made by distributed devices may not always work well if they are poorly coordinated between the routers, resulting in issues, such as forwarding loops, which may be transitory or permanent. So it is normal to require the use of the same algorithm to decide the forwarding actions at each hop.

A way to avoid routing issues is to select an end-to-end path a priori and consistently execute forwarding on the intermediate routers accordingly. This element of traffic engineering is known as "path steering" [[I-D.ietf-teas-rfc3272bis](#)] and relies on the routing to protocols collect and exchange the reachability within a domain, so that any routers can select an end-to-end path. However, the amount of information needed to support these decisions can become very large (e.g., in large networks, with many possible paths and route metrics), which might impede the scalability both in terms of the storage and the distribution of the information. Although network topology filters are often applied to reduce the storage of the network data and the complexity of the computation algorithm, the path computation accuracy and optimality may be negatively impacted.

The Internet is a very complicated system that is made up of many independently built networks with two types of routing protocols: an interior gateway protocol (IGP) that routes inside a network and an exterior gateway protocol (such as BGP) that routes between networks. For a communication that crosses more than one domain, there could be many possible paths for the given destination. In

principle, the more information that decision-making devices have, the better choices they can make. However, it is often infeasible to have all information of all potential end-to-end paths, particularly for communications through several networks with different ownership. Consequently, the best choices made within each domain may not reach the best overall result. A key challenge here is the tussle between abstraction, needed for scalability, and optimality, which abstraction may impede.

When choosing the best paths or topology structures, the following may need to be considered:

- *The method by which a path, or set of paths, is to be calculated. For example, a path may be selected automatically by the routing protocol or may be imposed (perhaps for traffic engineering reasons) by a central controller or management entity.

- *The criteria used for selecting the best path. For example, classic route preference, or administrative policies such as economic costs, resilience, security, and (if requested) applying geopolitical considerations.

3. Extension of the Routing System Beyond 'Mere' Reachability

In the following, we provide a brief overview of routing extensions with purposes beyond 'mere' reachability. We align our overview with many of the solutions described in [[I-D.king-irtf-semantic-routing-survey](#)] and refer to this draft for more detail, in addition to the example references themselves.

The following [Table 1](#) focusses on three key aspects when considering routing extensions for our discussion in this draft:

- *Purpose: What is the intended purpose of the proposed extension? This aspect may lead to a taxonomy for looking at the capabilities of a multi-purpose routing system.

- *Approach: What is the underlying technical approach to achieve the intended purpose? This aspect may lead to a taxonomy of approaches for achieving desired routing purposes.

- *Examples: What are known examples that have employed the given approach to achieve the given routing purpose? This aspect provides a possibly growing catalogue of explicit examples to study in more detail.

Purpose	Approach	Examples
	Preferential Routing	IS-IS Extensions [RFC5305]

Purpose	Approach	Examples
Path Selection for Traffic Engineering		
	Policy-based Routing	PBR models [RFC1104] Inter-domain policy routing [RFC1479]
	Flow Steering	TBD
	Path Computation	PCE [RFC4655] PCEP [RFC5440] PCEP Extension [RFC8231]
	IRTF	Path-aware Networking RG [PANRGref] Path properties [I-D.irtf-panrg-path-properties] Past efforts evaluation [I-D.irtf-panrg-what-not-to-do]
Path Selection for Multicast	Multicast	IP multicast [RFC1112] IPv6 addressing [RFC4291] MBone [MBONeref] MADCAP [RFC2730] MALLOC [RFC6308] MASC [RFC2909] MZAP [RFC2776] MSDP [RFC3618] SSM [RFC4607]
	Automatic Multicast Tunneling	AMT [RFC7450]
	Path-based Forwarding	BIER [RFC8279] BIER encapsulation [RFC8296] IP-over-ICN [I-D.trossen-icnrg-internet-icn-5glan]
Routing Architectures	Future architectures	[RESEARCHFIAref] [ITUNET2030ref] [SCIONref] [RINAref]
Service Function Chaining	L2/L3 Explicit Header Chaining	SFC [RFC7665] NSH [RFC8300] MPLS encapsulation [RFC8596]
	Name-based Chaining	Name-based SFF [RFC8677]
	Source Routing	Segment Routing [RFC8402]
Application/service-aware Routing	Application-server based	Aalto [RFC7285] APN [I-D.li-apn-framework]
	L3 based	Dyncast use cases [I-D.liu-dyncast-ps-usecases] Dyncast use architecture [I-D.li-dyncast-architecture]
	Network programming	Segment routing [RFC8986]
Anycast Routing	IP Anycast	Architecture considerations [RFC7094] Operation of Anycast [RFC4786]
	Metric-based	

Purpose	Approach	Examples
		Dyncast use cases [I-D.liu-dyncast-ps-usecases] Dyncast architecture [I-D.li-dyncast-architecture] Load-balanced anycast [weightedRef]
Privacy-aware Routing	Crypto routing	CGA [RFC3972] CGA Extension Field [RFC4581]
	Obfuscation	ILNP-based [ILNP_PRIVACY]
Security-enhanced Routing	Routing Architecture	SCION [SCIONref]
Identity Split Routing	Identity/Locator Split	LISP [RFC6830] LISPbis [I-D.ietf-lisp-rfc6830bis] LISPbis [I-D.ietf-lisp-rfc6833bis] HIP[RFC4423] ILNP [RFC6740]
Content Routing	Routing over content names	ICN [ICNref] NDN [NDNref] ICN deployment [RFC8763] HICN [HICNref]
	Routing via indirection points	DNS DONA [DONAref]
Differentiated Routing	QoS Differentiation	DiffServ [RFC2474] IntServ [RFC2210]
	Path differentiation	Segment Routing [RFC8402] SFC [RFC7665]
Extended Routing	EH-based	IPv6 EH [RFC8200]

Table 1: Summary of Routing Extensions

4. Issues with Current Approaches

Developing routing purposes beyond the original 'mere' reachability does come with issues when considering their deployment and use in the Internet; we outline those issues in the following.

We note that those issues are intrinsically linked to the ones stemming from the extension of addressing semantics that may be used to realize the various routing extensions, identified in [[I-D.jia-intarea-internet-addressing-gap-analysis](#)]. We therefore structure our presentation along the same lines.

4.1. Limiting Routing Semantics

Approaches that intend to change the purpose of communication, specifically within the evolution of communication semantics outlined in [Section 5](#) through, e.g., by separating host from network node identification [[RFC7401](#)] or through identification of content

directly [[HICNref](#)], are limited by the reachability purpose of IPv6, as defined by its source and destination address.

This leads to approaches such as [[I-D.trossen-icnrg-internet-icn-5gln](#)] to override addressing semantics, namely replacing the IPv6 source and destination addresses with path information instead, in order to achieve the desired purpose of its routing solution. This, in turn, requires to still carry address information as part of the payload in order to support clients unaware of the routing extension. Furthermore, such approach may lead to 'information leakage' outside the boundaries of the system in which its changed purpose is being realized. Introduction of dedicated gateways to 'translate' from one purpose (new routing) to another (IPv6 routing) is the consequence of this.

But even such approach of 're-writing' packet information towards a new purpose limits the expressible (new) semantic information to the size of the original field, thereby limiting the support of content information in approaches such as [[HICNref](#)] or the size of supported networks in [[I-D.trossen-icnrg-internet-icn-5gln](#)] to the bit size afforded by IPv6 addresses.

4.2. Complexity and Efficiency

Introducing new routing purposes also bring additional complexity. This becomes an issue when new purposes are being introduced in particular parts of the overall Internet, such as the edge of the network, while relying on the existing reachability purpose of the Internet to interconnect those islands over the existing Internet.

This additional complexity therefore often comes with a penalty in the form of efficiency and costs for realizing the novel routing purpose, which in turn may specifically pose an even bigger problem when the solution is introduced at the edge of the network, which is often constrained in resources and therefore costs that can be expensed.

For instance, if the specific new purpose requires compression of packet fields, such as for header compression, additional processing as well as potentially required gateways that restore information towards the Internet may be prohibitive for introducing the desired new routing purpose in this part of the Internet.

Conversely, performance requirements of core networks, in terms of packet processing speed, pose a problem when wanting to accommodate novel routing purposes. Here, not only the possibly additional processing but also the required changes of often HW-based platforms makes adoption of novel routing purposes prohibitive.

4.2.1. Repetitive encapsulation

A routing solution targetting a different purposes often requires encapsulating the relevant information, thereby bloating packet sizes and lowering overall efficiency. This can be seen in routing solutions such as [[I-D.trossen-icnrg-internet-icn-5g1an](#)] (introducing an alternative forwarding solution), [[I-D.ietf-lisp-mn](#)] (handling mobility in LISP), [[RFC8926](#)] and [[RFC7348](#)] (DC networking), [[RFC8986](#)] (traffic engineering) as well as [[TOR](#)] (routing privacy), all of which introduce multiple encapsulations that in turn reduce the forwarding efficiency.

The introduction of dedicated points of encapsulation also introduce complexity and costs at the points of the network where they are required, which may often be at the network edge, while also establishing failure points and therefore increasing the overall fragility of the system; a point we discuss in more detail in [Section 4.4](#).

4.2.2. Introducing Path Stretch

Path stretch is an issue when moving from direct reachability purposes to additional ones, such as dealing with mobility of endpoints, as done in MobileIP [[RFC6275](#)]. In this case, following the typical triangular route affects transmission efficiency as well as overall latency of the communication, instead of communicating directly towards the (new) network location.

Additionally, the realization of novel purposes, such as privacy-compliant routing in systems like TOR [[TOR](#)], often introduce path stretch due to the additional relays being introduced for fulfilling the intended purpose, here the obfuscation of traffic for privacy reasons.

4.2.3. Complicating Traffic Engineering

As outlined in [Section 3](#), many solutions to extend the original reachability purpose of Internet routing aim to introduce or improve on traffic engineering capabilities, e.g., by enabling decisions based on QoS metrics, mobility, chaining and others aspects.

However, realizing each novel purpose as a separate solution in itself likely hampers the goal for which they are developed, namely to improve on traffic engineering, whenever individual solutions are being used in combination. This 'feature interaction' aspect may even prevent combined uses, while at a minimum requiring an understanding if combined uses are possible in the first place or instead incompatible with each other. This is not just an issue that routing purposes may be incompatible at a functional level, e.g.,

through conflicting policies, but may also utilize conflicting realizations for their purposes.

4.3. Security

Security issues, outside the security considerations of the specific design, often arise from the integration of the specific solution into the existing routing system. For instance, HIP [[RFC7401](#)] limits its host identity to 128bit in an effort to be backward compatible, but possibly resulting in weak cryptographical strength. A similar issue can be observed in CGA [[RFC3972](#)], where only 59bits of the 128bit limit may be used for what could be packet signatures not sufficiently robust enough against attacks.

Attempts to introduce privacy purposes into the routing system, e.g., by utilizing ephemeral addresses [[I-D.gont-v6ops-ipv6-addressing-considerations](#)], may in turn pose significant challenges on the routing system through its required renewal rate of addresses.

4.4. Fragility

From the overview of novel routing purposes in [Section 3](#), we can observe that the existence of those additional routing purpose adds many purpose-specific translation/adaptation points, responsible for mapping formats from one meaningful context into another. This in turn creates dependency on this additional functionality to exist for endpoints to communicate with the context of the intended purpose.

While translation/adaptation between purposes and their defining contexts is often not avoidable when going beyond 'mere' reachability, it is the solution-specific nature of those components (required for many if not each extended purpose) that is likely to increase the fragility of the resulting system.

The key problem here is the interaction with other extended purposes that may exist in specific deployments. While needing to operate in the presence of those other purpose-specific components, their design has often not taken into account the specific interaction in question. Given the diversity of extension realizations, utilizing many, almost any packet field, even beyond and entirely different to its intended purpose, conflicting behaviour as well as diverging interpretation of the utilized packet information is clearly an issue. Only careful testing of combinations with possible delineation (of purposes) as well as networks may be required, all of which further increases the costs for utilizing the extended purposes.

4.5. Interoperability

Although routing extensions are developed with their specific purposes in mind, reflected in requirements and behaviours, they are often realized in conjunction with other extensions when it comes to real-world deployments.

This poses an Interoperability challenge, both in terms of backward as well as forward compability. Feature interactions need investigations, often left to operational deployment.

Building extensions on the basis of agreed packet field semantics is one way to achieve the desired interoperability, unless approaches use extensions to packet fields beyond their original intention. As a consequence, translation/adaptation points may be needed to ensure interoperability with other parts of the network, increasing the fragility of the system, as discussed in [Section 4.4](#).

Forward compability aims at ensuring that future extensions will continue to be possible, aiming at an overall extensibility of the system beyond its purpose at the time of developing a specific solution. IPv6 extension headers are one example of enabling future extensions, although not without their own problems in real-world deployments [[SHIMv6ref](#)].

5. The Driving Need for Evolving Communication Semantics

When looking at the evolution of routing beyond reachability, the key question arises on how the purposes of communication, or more concretely the underlying communication semantics, have evolved from the shortest-path routing of packet from sender to a receiver, each of those being originally identified through IP locators and captured as a source and destination address field in packet headers but having evolved through approaches such as those presented in [Section 3](#).

To better understand this evolution, we distinguish communication in networks according to the relationship between senders and receivers and the selection of the paths and endpoints for the delivery of packets, leading us to the following distinct semantics.

- *The Unicast semantic consists of sending a packet from a sender to a single receiver.

- *In Anycast, a packet is sent from the sender to any one of a set of receivers, where the choice of receiver is made by the network.

- *In Multicast, a packet is sent from a sender to all members of a group of receivers.

The identification of endpoints in these semantics may use well-known IP locators for unicast relations or IP multicast groups, while Anycast may use an IP anycast address or a content/service name [[NDNref](#)][[CARDS](#)]. Often, packets also carry higher-layer information, such as ports, to facilitate the endpoint-local handling of received packets.

These relationship semantics can be further constrained through path and endpoint selection semantics:

- *Multicast relations may be defined as (i) by configuration, (ii) dynamically formed through a membership protocol [[RFC3376](#)], (iii) through requests towards the sender [[I-D.trossen-bier-frrm](#)], or (iv) through diffusing towards a sub-group of a larger group, e.g., in Distributed Ledger Technologies (DLTs).

- *In Bestcast, the network applies constraints to determine the best path to the receiver based on the destination address, the state of the network and the compute resources, and information supplied with the packet. Bestcast may also be achieved by extending the anycast address to include multiple virtual unicast representations of the same receiver. The choice of a specific receiver may also determine the network path to reach this receiver. The choice may be made within the network or using a server-based scheduler and a database akin to DNS Resource Records.

- *The Chaincast semantic steers a packet through a specific set of nodes deduced from the value of the destination address, with typical examples being Service Function Chaining [[RFC7665](#)] and Segment Routing Network Programming [[RFC8986](#)].

While we can see many examples of those evolving communication semantics, a crucial question is 'What are the things that are identified by the identifiers?' [[RTGWGinterim](#)]. Behind this question is the observation that 'if you want to put multiple definitions into the same identifier space, then it requires an architecture discussion.'

This interjection is key in understanding the architectural dimension of evolving communication semantics since those evolved meanings are often based on differently identifying the 'ends' of the communication. Information-centric networking (ICN) [[NDNref](#)] is one such example, turning the meaning of an address from being a network location into one where the address represents a piece of information, with the network being tasked to build ephemeral relations between those network components asking for the information and those that may be able to provide it.

The FRRM (forward request return multicast) [[I-D.trossen-bier-frm](#)] semantic for multicast relations is another approach (albeit related to ICN), where the commonality of the forward requests, e.g., in the form of a URL pointing to the same content chunk, identifies the communication relation akin to ICN, while path information (e.g., in the form of BIER forwarding information [[RFC8279](#)]) is used to actually forward the packets from its source to the possible receivers.

Architectures, such as those for ICN and IP, have long lived in parallel, e.g., with ICN deployed in limited domains [[RFC7665](#)] or interconnecting to the Internet through dedicated application-level gateways, while proposals such as [[HICNref](#)] utilize in-address embedding to deploy ICN alongside IP networks.

The architectural question that arises from this is what the overarching architectural principles as well as its resulting frameworks and architectures should look like that would allow not only for rich communication semantics to be implemented but also to emerge over time and continued to be supported without resorting to gateway and in-lay techniques that all come with complexity and fragility issues?

6. Where to Go From Here?

This document outlined the original starting point of the Internet's routing system, namely providing 'mere' reachability, and showed through its survey of existing solutions that have since been developed that Internet routing has, in fact, evolved into a system that serves many purposes beyond its original 'mere' reachability goal.

However, the issues we outlined in [Section 4](#) pose the question on how to move forward in the (future) evolution of Internet routing. We assert that continuing with a 'business as usual' attitude will ultimately compound the identified issues, thereby hampering innovation in novel routing purposes and solutions, and therefore the Internet overall.

As a way forward, we ask the wider RTG WG community to recognize the following cornerstones for an evolution path for Internet routing:

1. Further evolution of the Internet's routing system MUST take a wider architectural approach in order to break with the point solution approach that has led to the identified issues in [Section 4](#).
2. With research and development on routing solutions continuing, as also illustrated in [[I-D.king-irtf-semantic-routing-survey](#)], these works MUST be brought into the process of IETF engagement

and standardization to increase the understanding of what novel trends, works, and possible developments may be just around the corner but also to inform ongoing research and development on paths taken in the IETF.

3. The RTG WG SHOULD play a role in the engagement with research and development since the 'Future of Internet Routing' (FIR) is at the heart of its charter ("The Routing Area working group (RTG WG) is chartered to provide a venue to discuss, evaluate, support and develop proposals for new work in the Routing Area" [[RTG WGref](#)]), a role that goes beyond the "specific small topics that do not fit with an existing working group" [[RTG WGref](#)].

Following on the cornerstones outlined above, we specifically suggest to the RTG WG, aligned with its charter to consider the following actions:

1. Establish suitable efforts within the RTG WG (e.g., as a sub-group) OR
2. Support the establishment of suitable efforts as a standalone FIR WG (or special interest group) OR
3. Support the establishment of suitable efforts within the IRTF, where those efforts directly liaise with the RTG WG through regular updates in its meetings.

7. Security Considerations

[Section 4.3](#) outlines a number of security issues that may occur outside the solution-specific security considerations, such as interactions between protocol behaviours that were previously untested as a combination. With that in mind, security considerations for a wider architectural approach to routing must have the security of the overall routing system as the main goal, not merely the security of a single solution.

8. Privacy Considerations

Protecting user privacy is very important. This extends beyond ensuring that user data cannot be examined in transit, and also requires that a process that inspects the network traffic should not be able to determine which applications or what types of application a user is running.

This makes it critically important to minimize or entirely avoid user and/or application information to be directly used for routing purposes. Instead, applications (or users) should express requirements for traffic delivery in a manner that does not reveal information about the user.

Encryption of user data, which is a common technique to protect user privacy, may obscure information that has previously been used to perform enhanced routing (such as by inspecting or hashing on payload fields), demonstrating that new requirements (here on privacy) may negatively impact previously accepted solutions.

9. IANA Considerations

This draft does not request any IANA action.

10. Acknowledgements

Many thanks go to Daniel King, Mohamed Boucadair for their comments to the text and to Lixia Zhang for raising important questions related to the possible architectural nature of the discussion needed.

11. Contributors

Adrian Farrel
Email: adrian@olddog.co.uk

12. Informative References

- [CARDS] Khandaker, K., Trossen, D., Khalili, R., Despotovic, Z., Hecker, A., and G. Carle, "CARDS: Dealing a New Hand in Reducing Service Request Completion Times", Paper IFIP Networking 2022, 2022.
- [CLNPref] "Protocol for providing the connectionless-mode network service: Protocol specification - Part 1", standard ISO/IEC 8473-1:1998, 1998, <<https://www.iso.org/standard/30931.html>>.
- [CONTENTref] Choi, J., Han, J., and E. Cho, "A survey on content-oriented networking for efficient content delivery", Paper IEEE Communications Magazine, 49(3): 121-127, May 2011., 2011, <<https://ieeexplore.ieee.org/iel5/35/5723785/05723809.pdf>>.
- [DONAref] Koponen, T., Chawla, M., Chun, B.-G., Ermolinskiy, A., Kim, K. H., Shenker, S., and I. Stoica, "A data-oriented (and beyond) network architecture", Paper Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, August 2007, 2007, <<https://dl.acm.org/doi/10.1145/1282380.1282402>>.

[HICNref]

Carofiglio, G., Muscariello, L., Auge, J., Papalini, M., Sardara, M., and A. Compagno, "Enabling ICN in the Internet Protocol: Analysis and Evaluation of the Hybrid-ICN Architecture", Paper Proceedings of the 6th ACM Conference on Information-Centric Networking, 2019., 2019, <https://www.researchgate.net/publication/336344520_Enabling_ICN_in_the_Internet_Protocol_Analysis_and_Evaluation_of_the_Hybrid-ICN_Architecture>.

[I-D.bonica-6man-ext-hdr-update] Bonica, R. and T. Jinmei, "Inserting, Processing And Deleting IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-bonica-6man-ext-hdr-update-07, 24 February 2022, <<https://www.ietf.org/archive/id/draft-bonica-6man-ext-hdr-update-07.txt>>.

[I-D.bryant-arch-fwd-layer-ps] Bryant, S., Chunduri, U., Eckert, T., and A. Clemm, "Forwarding Layer Problem Statement", Work in Progress, Internet-Draft, draft-bryant-arch-fwd-layer-ps-04, 24 January 2022, <<https://www.ietf.org/archive/id/draft-bryant-arch-fwd-layer-ps-04.txt>>.

[I-D.chunduri-lsr-isis-preferred-path-routing]

Chunduri, U., Li, R., White, R., Contreras, L. M., Tantsura, J., and Y. Qu, "Preferred Path Routing (PPR) in IS-IS", Work in Progress, Internet-Draft, draft-chunduri-lsr-isis-preferred-path-routing-07, 12 November 2021, <<https://www.ietf.org/archive/id/draft-chunduri-lsr-isis-preferred-path-routing-07.txt>>.

[I-D.farrel-irtf-introduction-to-semantic-routing] Farrel, A. and D. King, "An Introduction to Semantic Routing", Work in Progress, Internet-Draft, draft-farrel-irtf-introduction-to-semantic-routing-04, 25 April 2022, <<https://www.ietf.org/archive/id/draft-farrel-irtf-introduction-to-semantic-routing-04.txt>>.

[I-D.gont-v6ops-ipv6-addressing-considerations] Gont, F. and G. Gont, "IPv6 Addressing Considerations", Work in Progress, Internet-Draft, draft-gont-v6ops-ipv6-addressing-considerations-02, 1 June 2022, <<https://www.ietf.org/archive/id/draft-gont-v6ops-ipv6-addressing-considerations-02.txt>>.

[I-D.ietf-lisp-mn] Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", Work in Progress, Internet-Draft, draft-ietf-lisp-mn-11, 30 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-mn-11.txt>>.

[I-D.ietf-lisp-rfc6830bis]

Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, "The Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6830bis-38, 7 May 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6830bis-38.txt>>.

[I-D.ietf-lisp-rfc6833bis] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, "Locator/ID Separation Protocol (LISP) Control-Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6833bis-31, 2 May 2022, <<https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6833bis-31.txt>>.

[I-D.ietf-teas-rfc3272bis]

Farrel, A., "Overview and Principles of Internet Traffic Engineering", Work in Progress, Internet-Draft, draft-ietf-teas-rfc3272bis-16, 24 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-rfc3272bis-16.txt>>.

[I-D.ietf-v6ops-ipv6-ehs-packet-drops]

Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. (. Liu, "Operational Implications of IPv6 Packets with Extension Headers", Work in Progress, Internet-Draft, draft-ietf-v6ops-ipv6-ehs-packet-drops-08, 11 June 2021, <<https://www.ietf.org/archive/id/draft-ietf-v6ops-ipv6-ehs-packet-drops-08.txt>>.

[I-D.irtf-panrg-path-properties] Enghardt, T. and C. Krähenbühl, "A Vocabulary of Path Properties", Work in Progress, Internet-Draft, draft-irtf-panrg-path-properties-05, 7 March 2022, <<https://www.ietf.org/archive/id/draft-irtf-panrg-path-properties-05.txt>>.

[I-D.irtf-panrg-what-not-to-do]

Dawkins, S., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", Work in Progress, Internet-Draft, draft-irtf-panrg-what-not-to-do-19, 26 March 2021, <<https://www.ietf.org/archive/id/draft-irtf-panrg-what-not-to-do-19.txt>>.

[I-D.jia-intarea-internet-addressing-gap-analysis]

Jia, Y., Trossen, D., Iannone, L., Mendes, P., Shenoy, N., Toutain, L., Chen, A. Y., and D. Farinacci, "Gap Analysis in Internet Addressing", Work in Progress, Internet-Draft, draft-jia-intarea-internet-addressing-gap-analysis-02, 6 March 2022, <<https://www.ietf.org/>

[archive/id/draft-jia-intarea-internet-addressing-gap-analysis-02.txt](https://www.ietf.org/archive/id/draft-jia-intarea-internet-addressing-gap-analysis-02.txt)>.

[I-D.king-irtf-semantic-routing-survey] King, D. and A. Farrel, "A Survey of Semantic Internet Routing Techniques", Work in Progress, Internet-Draft, draft-king-irtf-semantic-routing-survey-04, 30 May 2022, <<https://www.ietf.org/archive/id/draft-king-irtf-semantic-routing-survey-04.txt>>.

[I-D.li-apn-framework] Li, Z., Peng, S., Voyer, D., Li, C., Liu, P., Cao, C., and G. Mishra, "Application-aware Networking (APN) Framework", Work in Progress, Internet-Draft, draft-li-apn-framework-05, 7 March 2022, <<https://www.ietf.org/archive/id/draft-li-apn-framework-05.txt>>.

[I-D.li-dyncast-architecture] Li, Y., Iannone, L., Trossen, D., Liu, P., and C. Li, "Dynamic-Anycast Architecture", Work in Progress, Internet-Draft, draft-li-dyncast-architecture-03, 7 March 2022, <<https://www.ietf.org/archive/id/draft-li-dyncast-architecture-03.txt>>.

[I-D.liu-dyncast-ps-usecases]
Liu, P., Eardley, P., Trossen, D., Boucadair, M., Contreras, L. M., and C. Li, "Dynamic-Anycast (Dyncast) Use Cases and Problem Statement", Work in Progress, Internet-Draft, draft-liu-dyncast-ps-usecases-03, 7 March 2022, <<https://www.ietf.org/archive/id/draft-liu-dyncast-ps-usecases-03.txt>>.

[I-D.muscariello-intarea-hicn]
Muscariello, L., Carofiglio, G., Augé, J., Papalini, M., and M. Sardara, "Hybrid Information-Centric Networking", Work in Progress, Internet-Draft, draft-muscariello-intarea-hicn-04, 20 May 2020, <<https://www.ietf.org/archive/id/draft-muscariello-intarea-hicn-04.txt>>.

[I-D.trossen-bier-frrm]
Trossen, D., "Realizing Forward Requests Return Multicast Semantic with BIER", Work in Progress, Internet-Draft, draft-trossen-bier-frrm-00, 21 February 2022, <<https://www.ietf.org/archive/id/draft-trossen-bier-frrm-00.txt>>.

[I-D.trossen-icnrg-internet-icn-5glan]
Trossen, D., Robitzsch, S., Reed, M., Al-Naday, M., and J. Riihijarvi, "Internet Services over ICN in 5G LAN Environments", Work in Progress, Internet-Draft, draft-trossen-icnrg-internet-icn-5glan-04, 1 October 2020,

<<https://www.ietf.org/archive/id/draft-trossen-icnrg-internet-icn-5gplan-04.txt>>.

[ICNref] Barbera, D., Xylomenos, G., Ververidis, C., Siris, V., and N. Fotiou, "A Survey of information-centric networking research", Paper IEEE Communications Surveys and Tutorials, vol. 16, Iss. 2, Q2 2014, 2014, <<https://www.scopus.com/record/display.uri?eid=2-s2.0-84901242669>>.

[ILNP_PRIVACY] Bhatti, S., Haywood, G., and R. Yanagida, "End-to-end privacy for identity and location with IP", Paper 2nd Workshop on New Internetworking Protocols, Architecture and Algorithms, 29th IEEE International Conference on Network Protocols, 2021.

[ISISref] "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service", standard ISO/IEC 10589, 2002, <[https://standards.iso.org/ittf/PubliclyAvailableStandards/c030932_ISO_IEC_10589_2002\(E\).zip](https://standards.iso.org/ittf/PubliclyAvailableStandards/c030932_ISO_IEC_10589_2002(E).zip)>.

[ITUNET2030ref] "Network 2030 Architecture Framework", Technical Specification ITU-T Focus Group on Technologies for Network 2030, 2020, <https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Network_2030_Architecture-framework.pdf>.

[MBONeref] Savetz, K., Randall, N., and Y. Lepage, "MBONE: Multicasting Tomorrow's Internet", Book IDG, 1996, <<https://www.savetz.com/mbone/>>.

[NDNref] Zhang, L., Afanasyev, A., and J. Burke, "Named Data Networking", Paper ACM SIGCOMM Computer Communication, Review 44(3): 66-73, 2014, 2014.

[PANRGref] "Path Aware Networking Research Group", RG Path Aware Networking Research Group, <<https://datatracker.ietf.org/rg/panrg/about>>.

[RESEARCHFIaref] Pan, J., Paul, S., and R. Jain, "A Survey of the Research on Future Internet Architectures", Paper IEEE Communications Magazine, vol. 49, no. 7, July 2011, 2014, <<https://ieeexplore.ieee.org/document/5936152>>.

[RFC1069] Callon, R. and H. Braun, "Guidelines for the use of Internet-IP addresses in the ISO Connectionless-Mode

Network Protocol", RFC 1069, DOI 10.17487/RFC1069, February 1989, <<https://www.rfc-editor.org/info/rfc1069>>.

- [RFC1104] Braun, H., "Models of policy based routing", RFC 1104, DOI 10.17487/RFC1104, June 1989, <<https://www.rfc-editor.org/info/rfc1104>>.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [RFC1479] Steenstrup, M., "Inter-Domain Policy Routing Protocol Specification: Version 1", RFC 1479, DOI 10.17487/RFC1479, July 1993, <<https://www.rfc-editor.org/info/rfc1479>>.
- [RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, DOI 10.17487/RFC2210, September 1997, <<https://www.rfc-editor.org/info/rfc2210>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [RFC2730] Hanna, S., Patel, B., and M. Shah, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)", RFC 2730, DOI 10.17487/RFC2730, December 1999, <<https://www.rfc-editor.org/info/rfc2730>>.
- [RFC2776] Handley, M., Thaler, D., and R. Kermode, "Multicast-Scope Zone Announcement Protocol (MZAP)", RFC 2776, DOI 10.17487/RFC2776, February 2000, <<https://www.rfc-editor.org/info/rfc2776>>.
- [RFC2909] Radoslavov, P., Estrin, D., Govindan, R., Handley, M., Kumar, S., and D. Thaler, "The Multicast Address-Set Claim (MASC) Protocol", RFC 2909, DOI 10.17487/RFC2909, September 2000, <<https://www.rfc-editor.org/info/rfc2909>>.
- [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, DOI 10.17487/RFC2992, November 2000, <<https://www.rfc-editor.org/info/rfc2992>>.

[RFC3376]

Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.

[RFC3618]

Fenner, B., Ed. and D. Meyer, Ed., "Multicast Source Discovery Protocol (MSDP)", RFC 3618, DOI 10.17487/RFC3618, October 2003, <<https://www.rfc-editor.org/info/rfc3618>>.

[RFC3972]

Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.

[RFC4291]

Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

[RFC4423]

Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, DOI 10.17487/RFC4423, May 2006, <<https://www.rfc-editor.org/info/rfc4423>>.

[RFC4581]

Bagnulo, M. and J. Arkko, "Cryptographically Generated Addresses (CGA) Extension Field Format", RFC 4581, DOI 10.17487/RFC4581, October 2006, <<https://www.rfc-editor.org/info/rfc4581>>.

[RFC4607]

Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.

[RFC4655]

Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

[RFC4786]

Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, DOI 10.17487/RFC4786, December 2006, <<https://www.rfc-editor.org/info/rfc4786>>.

[RFC5305]

Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

[RFC5440]

Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.

[RFC6275]

Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.

[RFC6308]

Savola, P., "Overview of the Internet Multicast Addressing Architecture", RFC 6308, DOI 10.17487/RFC6308, June 2011, <<https://www.rfc-editor.org/info/rfc6308>>.

[RFC6740]

Atkinson, RJ. and SN. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, DOI 10.17487/RFC6740, November 2012, <<https://www.rfc-editor.org/info/rfc6740>>.

[RFC6830]

Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.

[RFC7094]

McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <<https://www.rfc-editor.org/info/rfc7094>>.

[RFC7285]

Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, DOI 10.17487/RFC7285, September 2014, <<https://www.rfc-editor.org/info/rfc7285>>.

[RFC7348]

Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.

[RFC7401]

Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <<https://www.rfc-editor.org/info/rfc7401>>.

[RFC7450]

Bumgardner, G., "Automatic Multicast Tunneling", RFC 7450, DOI 10.17487/RFC7450, February 2015, <<https://www.rfc-editor.org/info/rfc7450>>.

[RFC7665]

Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/

RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <<https://www.rfc-editor.org/info/rfc7872>>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.

[RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.

[RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8596] Malis, A., Bryant, S., Halpern, J., and W. Henderickx, "MPLS Transport Encapsulation for the Service Function Chaining (SFC) Network Service Header (NSH)", RFC 8596, DOI 10.17487/RFC8596, June 2019, <<https://www.rfc-editor.org/info/rfc8596>>.

[RFC8677] Trossen, D., Purkayastha, D., and A. Rahman, "Name-Based Service Function Forwarder (nSFF) Component within a

Service Function Chaining (SFC) Framework", RFC 8677, DOI 10.17487/RFC8677, November 2019, <<https://www.rfc-editor.org/info/rfc8677>>.

- [RFC8684] Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 8684, DOI 10.17487/RFC8684, March 2020, <<https://www.rfc-editor.org/info/rfc8684>>.
- [RFC8736] Venaas, S. and A. Retana, "PIM Message Type Space Extension and Reserved Bits", RFC 8736, DOI 10.17487/RFC8736, February 2020, <<https://www.rfc-editor.org/info/rfc8736>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8763] Rahman, A., Trossen, D., Kutscher, D., and R. Ravindran, "Deployment Considerations for Information-Centric Networking (ICN)", RFC 8763, DOI 10.17487/RFC8763, April 2020, <<https://www.rfc-editor.org/info/rfc8763>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RINAref] Day, J., "Patterns in Network Architecture: A Return to Fundamentals", Book Prentice Hall, 2008.
- [RTGWGinterim] King (ed.), D., "Minutes interim-2022-rtgwg-01: Tue 08:00", Minutes Minutes RTG WG interim meeting on Semantic Routing at 21.06.2022, 2022, <<https://datatracker.ietf.org/doc/minutes-interim-2022-rtgwg-01-202206210800/>>.
- [RTGWGref] "Routing Working Group Charter", WG Routing Working Group, <<https://datatracker.ietf.org/wg/rtgwg/charter/>>.

[SCIONref]

Chuat, L., Legner, M., Basin, D., Hausherr, D., Hitz, S., Mueller, P., and A. Perrig, "The Complete Guide to SCION. From Design Principles to Formal Verification", Book Springer International Publishing AG, 2022., 2022, <<https://link.springer.com/book/10.1007/978-3-031-05288-0>>.

[SHIMv6ref]

Naderi, H. and B. E. Carpenter, "Putting SHIM6 into practice", Paper 2014 Australasian Telecommunication Networks and Applications Conference (ATNAC), 2014, <<https://dl.acm.org/doi/10.1145/1282380.1282402>>.

[TOR]

"The Tor Project", Website Tor Project, 2022, <<https://www.torproject.org/>>.

[weightedRef]

Lin, C-Y., Lo, J-H., and S-Y. Kuo, "Load-balanced anycast routing", Paper Proceedings of the Tenth International Conference on Parallel and Distributed Systems, 2004., 2004, <https://www.researchgate.net/publication/4083002_Load-balanced_anycast_routing>.

Authors' Addresses

Dirk Trossen
Huawei Technologies
Munich
Germany

Email: dirk.trossen@huawei.com

David Lou
Huawei Technologies
Munich
Germany

Email: zhe.lou@huawei.com

Sheng Jiang
Huawei Technologies
China

Email: jiangsheng@huawei.com