

RTGWG
Internet-Draft
Intended status: Informational
Expires: 27 October 2022

D. King
Lancaster University
A. Farrel
Old Dog Consulting
C. Jacquenet
Orange
25 April 2022

Challenges for the Internet Routing Systems Introduced by Semantic
Routing
draft-king-irtf-challenges-in-routing-08

Abstract

Historically, the meaning of an IP address has been to identify an interface on a network device. Routing protocols were developed based on the assumption that a destination address had this semantic.

Over time, routing decisions have been enhanced to determine paths on which packets could be forwarded according to additional information carried principally within the packet headers, and dependent on policy coded in, configured at, or signaled to the routers.

Many proposals have been made to add semantics to IP packets by placing additional information into existing fields, by adding semantics to IP addresses, or by adding fields to the packets. The intent is always to facilitate routing decisions based on these additional semantics in order to provide differentiated paths to enable forwarding of different packet flows on paths that may be distinct from those derived by shortest path first or path vector routing. We call this approach "Semantic Routing".

This document describes the challenges to the existing routing system that are introduced by Semantic Routing. It then summarizes the opportunities for research into new or modified routing and forwarding approaches that make use of additional semantics.

This document is presented as a study to support further research into clarifying and understanding the issues. It does not pass comment on the advisability or practicality of any of the proposals and does not define any technical solutions.

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 27 October 2022.

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	3
2. Current Challenges to IP Routing	4
3. What is Semantic Routing?	7
3.1. Architectural Considerations	9
4. Challenges for Internet Routing Research	10
4.1. Research Principles	10
4.2. Routing Research Questions to be Addressed	11
5. Security and Privacy Considerations	15
6. IANA Considerations	16
7. Acknowledgements	16
8. Contributors	16
9. Informative References	16
Authors' Addresses	17

1. Introduction

Historically, the meaning of an IP address has been to identify an interface on a network device. Routing protocols were to compute, establish, and maintain paths through networks toward destination prefixes until IP packets eventually reach their destination, and were based on the assumption that a destination address had this semantic. Anycast and multicast addresses were also defined, and those address semantics sometimes required variations to the routing protocols or even encouraged the development of new protocols.

Over time, the mechanisms that enabled routing decisions were enhanced to determine paths on which packets could be forwarded according to additional information carried principally within the packets headers or within 'shim' headers, and dependent on policy coded in, configured at, or signaled to the routers. Perhaps one of the most iconic examples is Equal-Cost Multipath (ECMP) where a router makes a choice about how to forward a packet over a number of parallel links or paths based on the values of a set of fields in the packet header.

Many proposals have been made to add semantics to IP packets by placing additional information into existing fields, by adding semantics to IP addresses, or by adding fields to the packets. The intent is always to facilitate routing decisions based on these additional semantics in order to provide differentiated paths to enable forwarding of different packet flows on paths that may be distinct from those derived by shortest path first or path vector routing. We call this approach "Semantic Routing"
[I-D.farrel-irtf-introduction-to-semantic-routing].

There are many approaches to adding semantics to packet headers: the additional information may be derived from the destination addresses, from other fields in the packet header, or the packet itself. Mechanisms for using the destination address range from assigning an address prefix to have a special purpose and meaning (such as is done for multicast addressing) through allowing the owner of a prefix to use the low-order bits of an address for specific purposes (e.g., to provide an indication of the nature of the service that is associated with these packets). Some proposals suggest variable address lengths, others offer new hierarchical address formats, and some introduce a structure to addresses so that they can carry additional information in a common way. Alternatively, forwarding decisions can be performed based on fields in the packet header (such as the IPv6 Flow Label, or the Traffic Class field), overloading of existing packet fields, or new fields added to the packet headers.

A survey of ways in which routing and forwarding decisions have been made based on additional information carried in packets can be found in [I-D.king-irtf-semantic-routing-survey].

Some Semantic Routing proposals are intended to be deployed in administratively scoped IP domains whose network components (routers, switches, etc.) are operated by a single administrative entity (sometimes referred to as 'limited domains' [RFC8799]), while other proposals are intended for use across the Internet. The impact the proposals have on routing systems may require clean-slate solutions, hybrid solutions, extensions to existing routing protocols, or potentially no changes at all.

This document describes some of the key challenges to the routing system that are already present in today's IP networks. It then briefly outlines the concept of "Semantic Routing" with reference to [I-D.farrel-irtf-introduction-to-semantic-routing] and presents some of the additional challenges to the existing routing system that Semantic Routing may introduce. Finally, this document presents a list of research questions that offer opportunities for future research into new or modified routing protocols and forwarding systems that make use of Semantic Routing.

In this document, the focus is on routing and forwarding at the IP layer. A variety of overlay mechanisms exists to perform service or path routing at higher layers, and those approaches may be based on similar extensions to packet semantics, but that is out of scope for this document. Similarly, it is possible that Semantic Routing can be applied in a number of underlay network technologies, and that, too, is out of scope for this document.

This document is presented as a study to support further research into clarifying and understanding the issues. It does not pass comment on the advisability or practicality of any of the proposals and does not define any technical solutions.

2. Current Challenges to IP Routing

Today's IP routing faces several significant challenges which are a consequence of architectural design decisions and the continued exponential growth in traffic. These challenges include mobility, multihoming, programmable paths, scalability, and security, and were not the focus of the original design of the Internet. Nevertheless, IP networks have, in general, coped well in an incremental manner whenever a new challenge has arisen. The following list is presented to give context to the continuing requirements that routing protocols must meet as new semantics are applied to the routing process.

- * Mobility - Mobility introduces several challenges, including maintaining a relationship between a sender and a receiver in cases where the sender or receiver changes their point of network attachment. The network must always be informed about the mobile node's current location, to allow continuity of services. Mobile users may also consume network resources, while in motion. The mobile user's service instances and attachments will also change due to varying load or latency, e.g., in Multi-access Edge Computing (MEC) environments.
- * Multihoming - Multihomed stations or multihomed networks are connected to the Internet via more than one access circuit or access network and, therefore, may be assigned multiple IP addresses or prefixes from different pools. There are challenges concerning how traffic is forwarded back to the source if the source has originated its traffic using the wrong source address for a particular connection, or if one of the connections to the Internet is degraded.
- * Multi-path - The Internet was initially designed to find the single, "best" path to a destination using a distributed routing algorithm. Current IP network topologies can provide multiple paths to reach a destination, each with different characteristics and with different failure likelihoods. It may be beneficial to send traffic over multiple paths to achieve reliability and enhance throughput, and it may be desirable to select one path or another because of QoS or security considerations for example, or to avoid transiting specific areas of an IP network, based (for example) on the reputation of transit provider for example. However, how packets are forwarded by using the shortest path means that distinguishing these alternate paths and directing traffic to them can be hard. Further, problems concerning scalability, commercial agreements among Service Providers, and the design of BGP make the utilization of multi-path techniques difficult for inter-domain routing. (Note that this discussion is distinct from Equal Cost Multi-path (ECMP) where packets are directed onto several "parallel" paths of identical least cost using a hash algorithm operated on some of the packets' header fields.)
- * Multicast - Delivering the same packet to multiple destinations can place considerable load on a network. Solutions that replicate the packet at the source or at the network edge may obviously cause multiple copies of the packet to flow along the same network links. Solutions that move deterministic replication into the network to make more optimal use of the network resources can be complex to set up and manage since multicast network designs often assume dynamic tree computation where the multicast

distribution tree can be rooted at the source or in the multicast network, thereby leading to specific routing tables whose entries denote the tree structure. More complicated hardware that can replicate packets may also be required within the network. In order that packets can be addressed to a group of destinations and not be forwarded by means of unicast transmission, parts of the addressing space (that is, address prefixes) have been reserved for multicast addressing.

- * Programmable Paths - The ability to decouple IP paths from routing protocols and agreements between Service Providers could allow users and applications to select network paths themselves, based on the required path characteristics. Another option is to let the route computation logic select, establish, and maintain paths on behalf of the user or the application and as a function of their requirements so that Service Providers can participate in the route computation "service". Currently, user and application packets follow the path selected by routing protocols and the way traffic is forwarded through a network is under the control of the Service Provider that operates the said network. The corresponding traffic forwarding policies enforced by the service provider usually comply with the requirements expressed by the user or the application. These requirements may have triggered a dynamic service parameter negotiation cycle that eventually leads to proper (network, CPU, storage) resource allocation.
- * Endpoint Selection - As compute resources and content storage move closer to the edge of the network, there are often multiple points in the network that can satisfy user requests. In order to make the best use of these distributed resources and so as to not overload parts of the network, user traffic needs to be steered to appropriate servers or data centres. In many cases, this function may be achieved in the application layer (such as through DNS [RFC3467]) or in the transport layer (such as using ALTO [RFC5693]). The challenge is to balance higher-layer decisions about which application layer resources to use with information from the lower layers about the availability and load of network resources.
- * Scalability - There are many scaling concerns that pose critical challenges to the Internet. Not least among these challenges is the size of the routing tables that routers in an IP network must maintain. As the number of devices attached to the network grows, so the number of addresses in use also grows, and because of the schemes used to assign address prefixes, the mobility of devices, and the various connectivity options between networks, the routing table sizes also grow, even more so when prefixes are not always amenable to aggregation. This problem is exacerbated by some

services (such as those supported by the IoT where several thousands of objects/sensors may be networked), where, as more devices are added to the network, the size of the routing table may affect the operation of certain routing protocols. It may be noted that scaling issues are also exacerbated by multihoming practices if a host that is multihomed is allocated a different address for each point of attachment.

- * Manageability, Maintainability, and Extensibility - Operational manageability is a key requirement for network technologies: network operators must be able to determine the status of their network and understand the causes of any disruptions or problems. Further, it must be possible to maintain the networks and the technologies running in them without disrupting the services being delivered by the networks. Additionally, the network technologies developed and deployed need to be extensible so that new features can be added and new services supported without the need to invent whole new technologies.
- * Security - Issues of security and privacy have been largely overlooked by the routing systems. However, there is increasing concern that attacks on routing systems can not only be disruptive (for example, causing traffic to be dropped), but may cause traffic to be redirected to inspection points that can breach the security or privacy of the payloads.

Some of the challenges outlined here were previously considered within the IETF by the IAB's "Routing and Addressing Workshop" held in Amsterdam, The Netherlands on October 18-19, 2006 [RFC4984]. Several architectures and protocols have since been developed and worked on within and outside the IETF, and these are examined in [I-D.king-irtf-semantic-routing-survey].

3. What is Semantic Routing?

Semantic Routing is the term applied to routing in an IP network that relies upon additional information to feed the route computation process, to enhance route selection decisions, and to direct the forwarding process. In addition to the routable part of the destination IP address (the prefix), such information may be present in other fields in the packet (chiefly the packet header) and configured or programmed into the routers/forwarders. Semantic Routing includes mechanisms such as "Preferential Routing", "Policy-based Routing", and "Flow steering".

In Semantic Routing, a packet forwarding engine may examine a variety of fields in a packet and match them against forwarding instructions. Those forwarding instructions may be installed by routing protocols,

configured through management protocols or a software defined networking (SDN) controller, or derived by a software component on the router that considers network conditions and traffic loads. The packet fields concerned may be the fields of an IP header, those same fields but with additional semantics, elements of the packet payload, or new fields defined for inclusion in the packet header or as a "shim" between the header and payload. In the case of additional semantics included in existing packet header fields, the approach implies some "overloading" of those fields to include meaning beyond the original definition. In all cases, a well-known definition of the encoding of the additional information is required to enable consistent interpretation within the network.

A more detailed description of Semantic routing can be found in [I-D.farrel-irtf-introduction-to-semantic-routing] and a survey of Semantic Routing proposals and research projects can be found in [I-D.king-irtf-semantic-routing-survey].

Many technical challenges exist for Semantic Routing in IP networks depending on which approach is taken. These challenges include (but are not limited to):

- * The continual growth of routing tables.
- * Convergence times for large networks.
- * Granularity of routing decisions.
- * Address consumption caused by lower address utility rate. The wastage mainly comes from aligning finite allocation for semantic address blocks.
- * Encoding too many semantics into prefixes will require evaluation of which to prioritize.
- * Risk of privacy/information leakage.
- * Lack of visibility of the Semantic Routing information when end-to-end or edge-to-edge encryption is used.
- * Burdening the user, application, or prefix assignment node.
- * Source address spoofing prevention mechanisms are required.
- * Overloading of routing protocols causing stability and scaling problems.

- * Depending on encoding mechanisms, there may be challenges for data planes to scale the processes of finding, reading, and looking up semantic data in order to forward packets at line speed.
- * Backwards compatibility with existing IP networking and routing protocols.
- * Extensibility to support additional functions in the future.
- * Manageability and network diagnostics to be able to determine how the network is functioning and to isolate the causes of any problems.

3.1. Architectural Considerations

Semantic data may be taken into account to integrate with existing routing architectures. An overlay can be built such that Semantic Routing is used to forward traffic between nodes in the overlay, but regular IP is used in the underlay. The application of semantics may also be constrained to within a limited domain. In some cases, such a domain will use IP, but be disconnected from the Internet. In other cases, traffic from within the domain is exchanged with other domains that are connected together across an IP network using tunnels or via application gateways. And in still another case traffic from the domain is forwarded across the Internet to other nodes and this requires backward-compatible routing approaches.

Isolated Domains: Some IP network domains are entirely isolated from the Internet and other IP networks. In these cases, packets cannot "escape" from the isolated domain into external networks and so the Semantic Routing schemes applied within the domain can have no detrimental effects on external domains. Thus, the challenges are limited to enabling the desired function within the domain.

Bridged Domains: In some deployments, it will be desirable to connect together multiple isolated domains to build a larger network. These domains may be connected (or bridged) over an IP network or even over the Internet, possibly using tunnels. An alternative to tunneling is achieved using gateway functionality where packets from a domain are mapped at the domain boundary to produce regular IP packets that are sent across the IP network.

Semantic Prefix Domains: A semantic prefix domain is a portion of the Internet over which a consistent set of semantic-based policies are administered in a coordinated fashion. This is achieved by assigning a routable address prefix (or a set of prefixes) for use with Semantic Routing so that packets may be

forwarded through the regular IP network (or the Internet). Once delivered to the semantic prefix domain, a packet can be subjected to whatever Semantic Routing is enabled in the domain.

Further discussion of architectures for Semantic Routing can be found in [I-D.farrel-irtf-introduction-to-semantic-routing].

4. Challenges for Internet Routing Research

It may not be possible to embrace all emerging scenarios with a single approach or solution. Requirements such as 5G mobility, near-space-networking, and networking for outer-space (inter-planetary networking), may need to be handled using different network technologies. Improving IP network capabilities and capacity to scale, and address a set of growing requirements presents significant research challenges, and will require contributions from the networking research community. Solutions need to be both economically feasible and have the support of the networking equipment vendors as well as the network operators.

4.1. Research Principles

Research into Semantic Routing should be founded on regular scientific research principles [royalsoc]. Given the importance of the Internet today, it is critical that research is targeted, rigorous, and reproducible.

The most valuable research will go beyond an initial hypothesis, a report of the work done, and the results observed. Although that is a required foundation, networking research needs to be independently reproducible so that claims can be verified or falsified. Further, the networks on which the research is carried out need to both reflect the characteristics that are being explicitly tested, and reproduce the variety of real networks that constitute the Internet.

Thus, when conducting experiments and research to address the questions in Section 4.2, attention should be given to how the work is documented and how meaningful the test environment is, with a strong emphasis on making it possible for others to reproduce and validate the work.

4.2. Routing Research Questions to be Addressed

As research into the scenarios and possible uses of Semantic Routing progresses, a number of questions need to be answered. These questions go beyond "Why do we need this function?" and "What could we achieve by carrying additional semantics in an IP address?" The questions are also distinct from issues of how the additional semantics can be encoded within an IP address. All of those issues are, of course, important considerations in the debate about Semantic Routing, but they form only part of the essential groundwork of research into Semantic Routing itself.

This section sets out some of the concerns about how the wider the use of Semantic Routing might impact a routing system. These questions need to be answered in separate research work or folded into the discussion of each Semantic Routing proposal.

1. What is the scope of the Semantic Routing proposal? This question may lead to various answers:

Global: It is intended to apply to all uses of IP.

Backbone: It is intended to apply to IP network connectivity.

Overlay: It is to be used as an overlay network using tunneling over IP or other underlay technologies.

Gateway: The Semantic Routing will be used within a specific domain, and communications with the wider Internet will be handled by IP and probably application gateways.

Domain: The use of the Semantic Routing is strictly limited to within a domain or private network.

Underlying this question is a broader question about the boundaries of the use of IP, and the limit of "the Internet". If a limited domain is used, is it a semantic prefix domain [RFC8799] where a part of the IP address space identifies the domain so that an address is routable to the domain, but the additional semantics are used only within the domain, or is the address used exclusively within the domain so that the external impact of the routability of the address and the additional semantics is not important?

2. What will be the impact on existing routing systems? What would happen if a packet carrying additional semantics was subjected to normal routing operations? How would the existing routing systems react if such a packet escaped (accidentally or

maliciously) from the planned scope of the proposal? For example: how are the semantic parts of an address distinguished from the routable parts (if, indeed, they are separable)?; is there an impact on the size and maintenance of routing tables due to the addition of semantics?; how are cryptographically generated addresses (such as [RFC3972]) made routable and kept simple enough for management?.

3. What path characteristics are needed to describe the desired paths and as input to route computation? Since one of the implications of adding semantics to IP packets is to cause special processing by routers, it is important to understand what behaviors are wanted. Such path characteristics include (but are not limited to):

Quality: Expressed in terms of throughput, latency, jitter, drop precedence, etc.

Resilience: Expressed in terms of survival of network failures and delivery guarantees.

Destination: How is a destination address to be interpreted if it encodes a choice of actual destinations? Can traffic be forwarded over multiple distinct paths if multiple destination addresses are encoded?

Security: What choices of path reduce the vulnerability of the traffic to security or privacy attacks?

In these cases, how do the routers utilize the additional semantics to determine the desired characteristics? Or are such characteristics used to feed the route computation logic, for example, by means of metrics? What additional information about the network do the routing protocols need to gather? What changes to the routing algorithm are needed to deliver packets according to the desired characteristics? How can routes be computed with characteristics that accommodate traffic patterns, requirements, and constraints?

4. Can we solve these routing challenges with existing routing tools and methods? We can break this question into a set of more detailed questions.

- * Is new hardware needed? Existing deployed hardware has certain assumptions about how forwarding is carried out based on IP addresses and routing tables. But hardware is increasingly programmable so that it may be possible to instruct the forwarding components to act on a variety of elements of the packets.
- * Do we need new routing protocols? We might ask some subsidiary questions:
 - Can we make do with existing protocols, possibly by tuning configuration parameters or using them out of the box?
 - Can we make backwards-compatible modifications to existing protocols such that they work equally for today's IP addresses or addresses with extra semantics?
 - Do we need entirely new protocols or radical evolutions of existing protocols in order to enforce advanced Semantic Routing policies?
 - Should we focus on the benefits of routing solutions that are optimized for specific environments (network topologies, technologies, use cases), or should we attempt to generalize to enable wider applicability?
- 5. Do we need new management tools and techniques? How practical is it to debug and operate the routing system? Management of the routing system (especially diagnostic management) is a crucial and often neglected part of the problem space. A critical part of this issue is how packets within the network can be inspected by diagnostic tools (or human operators) and mapped to the routing and forwarding decisions that were made within the network in order to understand the actions made at and by upstream routers.
- 6. What is the impact of Semantic Routing on the security of the routing system?
 - * Does the introduction of Semantic Routing provide a greater attack surface?
 - * Can Semantic Routing provide greater opportunities for security by fine-grain forwarding of flows to be inspected by different security functions?

- * Can Semantic Routing improve security and privacy by obscuring information in the packets, or does the inclusion of additional information risk compromising security and privacy?
 - * To what extent does deployment within a limited domain strengthen security or make it less of a concern?
 - * Does the use of Semantic Routing make it easier or harder to impose censorship, prohibit access to the Internet by specific parties, or block access to certain resources or types of service?
7. What is the scalability impact of Semantic Routing on routing systems? Scalability can be measured as:
- * Routing table size. How many entries need to be maintained in the routing tables by different routers serving different roles in the network? Some approaches to Semantic Routing may be explicitly intended to address this problem.
 - * Forwarding table size. The size of the forwarding table may be less of an issue considering modern hardware, however the more granular the routing/forwarding decisions made in a router, the greater the size of this table. The size of the forwarding table has implications for memory in the forwarding engine, but also for the lookup time for forwarding each packet.
 - * Routing performance. Routing performance may be considered in terms of the volume of data that has to be exchanged both to construct and maintain the routing tables at the participating routers. It may also be measured in terms of how much processing is required to compute new routes when there is a change in the network.
 - * Routing convergence. This is the time that it takes for a routing protocol to discover changes (especially faults) in the network, to distribute the information about any changes to its peers, and to reach a stable state across the network such that packets are forwarded consistently.

For all questions about routing scalability, research that presents figures based on credible example networks is highly desirable. Similar questions may be asked about the amount of forwarding state that has to be maintained in the routers.

8. To what extent can Semantic Routing be applied to multicast transmission schemes:
 - * Can Semantic Routing facilitate the computation and the establishment of (service-inferred) multicast distribution trees?
 - * Can specific semantics be carried in multicast addresses?
9. Is the approach extensible and maintainable? Can new features be added without increasing the complexity and in a backward compatible way? Could the approach be modified to handle evolutions in the rest of the networking infrastructure? Considerations might include the ability to encode additional options or variants within protocol fields, and the ability to add new fields. Such considerations must be actively traded against the processing overhead associated with certain encoding types.
10. What aspects need to be standardized? It is important to understand the necessity of standardization within this research. What degree of interoperability is expected between devices and networks? Is a given domain so constrained (for example, to a single equipment vendor) that standardization would be meaningless? Is the application so narrow (for example, in niche hardware environments) such that interoperability is best handled by agreements among small groups of vendors such as in industry consortia?

5. Security and Privacy Considerations

Research into Semantic Routing must give full consideration to the security and privacy issues that are introduced by these mechanisms. Placing additional information into packet header fields might reveal details of what the packet is for, what function the user is performing, who the user is, etc. Furthermore, in-flight modification of the additional information might not directly change the destination of the packet, but might change how the packet is handled within the network and at the destination.

It should also be considered how packet encryption techniques that are increasingly popular for end-to-end or edge-to-edge security may obscure the semantic information carried in some fields of the packet header or found deeper in the packet. This may render some semantic routing techniques impractical and may dictate other methods of carrying the necessary information to enable Semantic Routing.

6. IANA Considerations

This document makes no requests for IANA action.

7. Acknowledgements

Thanks to Stewart Bryant for useful conversations. Luigi Iannone, Robert Raszuk, Dirk Trossen, Ron Bonica, Marie-Jose Montpetit, Yizhou Li, Toerless Eckert, Tony Li, Joel Halpern, Stephen Farrell, Carsten Bormann, David Hutchison, Jeffery He, Dino Farinacci, Greg Mirsky, and Jeff Haas made helpful suggestions.

This work is partially supported by the European Commission under Horizon 2020 grant agreement number 101015857 Secured autonomic traffic management for a Tera of SDN flows (Teraflow).

8. Contributors

Joanna Dang
Email: dangjuanna@huawei.com

9. Informative References

- [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-03, 22 January 2022, <<https://www.ietf.org/archive/id/draft-farrel-irtf-introduction-to-semantic-routing-03.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-03, 26 November 2021, <<https://www.ietf.org/archive/id/draft-king-irtf-semantic-routing-survey-03.txt>>.
- [RFC3467] Klensin, J., "Role of the Domain Name System (DNS)", RFC 3467, DOI 10.17487/RFC3467, February 2003, <<https://www.rfc-editor.org/info/rfc3467>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.

- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, DOI 10.17487/RFC4984, September 2007, <<https://www.rfc-editor.org/info/rfc4984>>.
- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, DOI 10.17487/RFC5693, October 2009, <<https://www.rfc-editor.org/info/rfc5693>>.
- [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>>.
- [royalsoc] The Royal Society, "Evidence synthesis : Principles", Web page, Principles for good evidence synthesis, 19 September 2018, <<https://royalsociety.org/topics-policy/projects/evidence-synthesis/principles/>>.

Authors' Addresses

Daniel King
Lancaster University
United Kingdom
Email: d.king@lancaster.ac.uk

Adrian Farrel
Old Dog Consulting
United Kingdom
Email: adrian@olddog.co.uk

Christian Jacquenet
Orange
Rennes
France
Email: christian.jacquenet@orange.com