

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

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

An Introduction to Semantic Routing
draft-farrel-irtf-introduction-to-semantic-routing-04

Abstract

Many proposals have been made to add semantics to IP packets by placing additional information in existing fields, by adding semantics to IP addresses themselves, or by adding fields. The intent is to facilitate enhanced routing/forwarding decisions based on these additional semantics to provide differentiated forwarding paths for different packet flows distinct from simple shortest path first routing. The process is defined as Semantic Routing.

This document provides a brief introduction to Semantic Routing.

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. Objectives and Scope	4
3. Approaches to Semantic Routing	6
3.1. Packet and Service Routing	8
4. Semantic Routing Information	8
4.1. Address Space Partitioning	9
4.2. Prefix-based Contextual Address Usage	9
4.3. Semantic Addressing	9
4.4. Flow Marking	10
4.5. Extended Lookup	10
4.6. Semantic Field Overloading	10
4.7. IPv6 Extension Headers	10
4.8. New Extensions	11
5. Architectural Considerations	11
5.1. Isolated Domains	11
5.2. Bridged Domains	12
5.3. Semantic Prefix Domains	12
6. A Brief Discussion of What Constitutes Routing	13
6.1. Application Layer Routing	13
6.2. Higher-Layer Path Selection	13
6.3. Transport Layer Routing	14
6.4. Tunnel-Based Routing	14
6.5. Inter-Domain Routing	14
6.6. Service Function Chaining	15
6.7. Network Layer Traffic Engineering Techniques	15
6.8. Semantic Routing in the Network Layer	16
6.9. Computation In The Network and Semantic Routing	16
7. Security Considerations	18
8. IANA Considerations	18
9. Acknowledgements	18
10. Contributors	18
11. References	18
11.1. Informative References	18
11.2. URL References	23
Authors' Addresses	24

1. Introduction

Historically, the meaning of an IP address has been to identify an interface on a network device or a network to which a host is attached [RFC0814]. Network routing protocols were initially designed to determine paths through a network toward destination addresses so that IP packets with a common destination address converged on that destination. Anycast and multicast addresses were also defined (e.g., Section 2.6.1 of [RFC4291]), and some of these new address semantics necessitated variations to the routing protocols (e.g., [RFC6992]), and in some cases the development of new routing protocols (e.g., Protocol Independent Multicast - Sparse Mode [RFC7761]).

Over time, routing decisions were enhanced to route packets according to additional information carried within the packets and dependent on policy coded in, configured at, or signaled to the routers. Perhaps the most obvious example is Equal-Cost Multipath (ECMP) where a router makes a consistent choice for forwarding packets over a number of parallel links or paths based on the values of a set of fields in the packet header. Another example is Constraint-based Shortest Path First (CSPF) where additional constraints are considered when performing route computation and selection.

Upper-layer applications are placing increasingly sophisticated demands on the network for better quality, more predictability, and increased reliability. Some of these applications are futuristic predictions (for example, haptic augmented reality multiplayer 3D worlds), some are new ideas on the threshold of roll-out (such as holographic conferencing), and many are rapidly developing sectors with established revenue streams (such as multiplayer immersive gaming).

At the same time, lower-layer network technologies are advancing rapidly providing increased bandwidth to the home and to mobile hand-held devices. These advances create an environment that enables the potential of advanced applications being run by very many end-users. This coincides with a massive growth in end-to-end communications that include machines and services, and to introduce routing and addressing behaviors to a particular use case and set of requirements applied within a limited region or domain of the Internet. Examples of these three developments include 5G, predicted wireless evolutions, IoT and vehicular connectivity, space-terrestrial communication, industrial networks, cloud computing, service function chaining and network functions virtualization, digital twins, and data-centric data brokerage platforms.

Despite this plurality of communication scenarios, IP-based addressing and network layer routing have remained focused on identifying locations of communication (i.e., "where") and determining paths between those locations with or without specific constraints (i.e., "how-to-get-there" as per [IEN23]). This has previously depended on higher-layer capabilities (e.g., for name-to-location resolution) to support some of these communication scenarios, but that approach introduces latency and dependencies (e.g., changing locator assignments may depend on the capabilities of the upper-layer capability that are outside the core addressing and routing system). Furthermore, multi-layer lookups and interactions may impact the efficacy of some of the communication scenarios mentioned here, particularly those that employ different routing and addressing approaches beyond just locators.

"Semantic Routing" places the support for advanced routing, forwarding, and location functions directly at the packet routing/forwarding layer, such as through extensions to the identification properties of addresses (so that the address indicates more than just the network location) or through performing routing functions on an extended set of inputs (for example, other fields carried in packet headers). Such an approach should preserve the Internet architecture as it is today while enabling additional routing function.

This document provides a brief introduction to semantic routing and outlines the possible approaches that might be taken. A separate document ([I-D.king-irtf-semantic-routing-survey]) makes a start at a survey of pre-existing work in this area, while [I-D.king-irtf-challenges-in-routing] sets out some of the issues that should be considered when researching, developing, or proposing a semantic routing scheme.

2. Objectives and Scope

As with all advances in Internet protocols, semantic routing may be considered for Internet-wide deployment or may be restricted (possibly only initially) to well-defined and contained networks referred to as "limited domains" (see [RFC8799]). The information used for semantic routing may be opaque within the network (in other words, the additional information is not required to be parsed by the routers and might not even be visible to them), may be transparent (so that routers may see the information, but their processing does not need to be changed to accommodate the information or its encoding), or may be active (so that semantic routing is fully enabled).

When building an end-to-end path across multiple domains, semantic routing may select a path in one domain that is not consistent with the paths selected in other domains in terms of constructing the "best" end-to-end path. That is, the semantic routing decisions within a domain are potentially isolated from knowledge about the other domains.

In any case, concern and consideration must be coexistence with, and backward compatibility to, existing routing and addressing schemes that are widely deployed.

Further understanding of the scope of semantic routing applied to the routing of packets at the network layer may be gained by reading Section 6 to see how various other concepts of routing are out of scope of this work.

A strategic objective of semantic routing, and associated semantic enhancements, is to enable Service Providers to modify the default forwarding behaviour to be based on other information present in the packet and policy configured or dynamically programmed into the routers and devices. This is aimed to cause new and alternative path processing by routers, including:

- * Determinism of quality of delivery in terms of throughput, latency, jitter, and drop precedence.
- * Determinism of resilience in terms of survival of network failures and delivery degradation.
- * Determinism of routing performance in terms of the volume of data that has to be exchanged both to establish and to maintain the routing tables.
- * Deployability in terms of configuration, training, development of new hardware/software, and interaction with the pre-existing network technologies and uses.
- * Efficiency of manageability in terms of:
 1. diagnostic management
 2. management of Service KPIs with/without guarantees
 3. dynamic and controlled instantiation of management information in the packets.

Issues of security and privacy have been largely overlooked within 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 routed via inspection points that can breach the security or privacy of the payloads (e.g., BGP hijack attacks). While semantic routing might offer tools for increasing security and privacy, it is possible that semantic routing and the additional information that may be carried in packets to enable semantic routing may provide vectors for attacks or compromise privacy. This must be examined by any semantic routing proposals. For example, means to control entities that are entitled to access supplied semantic routing information should be considered.

3. Approaches to Semantic Routing

Typically, in an IP-based network packets are forwarded using the least-cost path to the destination IP address. Service Providers may also use techniques to modify the default forwarding behavior based on other information present in the packet and configured or programmed into the routers. These mechanisms, sometimes called semantic routing techniques include "Preferential Routing", "Policy-based Routing", and "Flow Steering".

Examples of existing semantic routing usage in IP-based networks include the following.

- * Using addresses to identify different device types so that their traffic may be handled differently [SEMANTICRTG].
- * Expressing how a packet should be handled, prioritized, or allocated network resources as it is forwarded through the network [TERASTREAMref].
- * Deriving IP addresses from the lower layer identifiers and using addresses depending on the underlying connectivity (for example, [RFC6282]).
- * Building IP addresses from the transport layer identifiers (for example, [RFC7597]).
- * Indicating the application or network function on a destination device or at a specific location, or enable Service Function Chaining (SFC) [RFC7665].
- * Providing semantics specific to mobile networks so that a user or device may move through the network without disruption to their service [CONTENT-RTG-MOBILeref].

- * Enabling optimized multicast traffic distribution by encoding multicast tree and replication instructions within addresses [MULTICAST-SRref].
- * Content-based routing (CBR), forwarding of the packet based on message content rather than the destination addresses [OPENSRRref].
- * Identifying hierarchical connectivity so that routing can be simplified [EIBPref].
- * Providing geographic location information within addresses [GEO-IPref].
- * Using cryptographic algorithms to mask the identity of the source or destination, masking routing tables within the domain, while still enabling packet forwarding across the network [BLIND-FORWARDINGref].

A more comprehensive list of existing implementations and research projects can be found in [I-D.king-irtf-semantic-routing-survey].

Semantic routing, operates to forward packets dependent on information carried in the packets and rules present in the routers. Those rules could be any combination of:

- * Built into the routers
- * Configured network-wide in the routers
- * Configured per-router in a relatively static way
- * Programmed to the routers in a dynamic way, for example, through software defined networking (SDN)
- * Distributed dynamically through the network using routing or signalling protocols

Semantic routing will also require information about network state and capabilities just as existing shortest path first routing systems do. That may require information (such as link delays or other qualitative attributes) to be collected by network nodes and distributed between routers by routing protocols. Alternatively, this information could be collected (centrally) by a set of network controllers and used to derive the rules installed in the routers.

Forwarding by a router is based on a look-up that also considers the semantic routing information carried in the packet (see Section 4) and forwarding instructions programmed into the forwarding element. Some semantic routing proposals may generate the semantic information (e.g., a hash) rather than using information that is directly extracted from the packet. The actions to perform may be derived by the router based on the rules and information that the router has collected, or may be programmed directly from the network controller.

3.1. Packet and Service Routing

Routing is the process of selecting a path for traffic in a network or between or across multiple networks. For example, IP routing uses IP addresses for source and destination identification and is typically used for packet networks, such as the Internet. IP routing assumes that network addresses are structured and facilitates routing entries in a routing table entry to represent a group of IP-capable devices.

While service routing and information-centric networking (ICN) can operate directly on top of layer 2 protocols (for example, [RFC9139]), in the context of this document, we are concerned with the function of service routing and ICN in IP networks. Like any new spanning-layer style protocol, deployment considerations for ICN on the Internet make tunneling through IP a required part of any co-existence or transition. The approach taken in this case, is to create an overlay layer on top of the IP network. Control of the overlay necessitates augmentation of existing routing mechanisms, or entirely new discovery, propagation and resource management protocols and procedures.

By contrast, explicit service-based IP routing [I-D.jiang-service-oriented-ip] abstracts the service actions that the network can provide into a number of classes called Service Action Types (SATs). Each packet is marked with the relevant SAT, and the packets are routed to the next available SAT provider (not the destination IP address). In this approach, a distinct encapsulation is needed and may carry native IP packets as payload, while transition experiments may utilize an overlay on top of IP.

IP Routing and service routing are not the same thing.

4. Semantic Routing Information

The subsections below describe some of the common techniques to enable semantic routing in more detail. The sections are unordered and no meaning should be assigned to how one approach is presented before another. They are not a complete list of possible approaches.

The approaches described here have many advantages and disadvantages. The purpose here is not to determine which approach is best or most appropriate, and so those advantages and disadvantages are not discussed. The reader will inevitably have a preference and see drawbacks.

4.1. Address Space Partitioning

In some cases, an address prefix is assigned a special purpose and meaning. When such an address appears in the packet's address field, a router can know from the prefix that particular routing/forwarding actions are required. An example of this approach is seen in multicast addressing. Another example is the handling of anycast in IPv6 where the nodes to which the address is assigned must be explicitly configured to know that it is an anycast address [RFC4291].

4.2. Prefix-based Contextual Address Usage

The owner of a prefix to use the low-order bits of an address for their own purposes.

The semantics of such an approach might be coordinated between prefix owners, or could be indicated through information that is part of the encoding, and is standardized. An example of such approach is in IPv4/IPv6 Translators [RFC6052].

4.3. Semantic Addressing

Semantic addressing is a term applied to any approach that adds semantics to IP addresses. This includes the mechanisms described in Section 4.1 and Section 4.2. Other semantic addressing proposals suggest variable address lengths, hierarchical addresses, or a structure to addresses so that they can carry additional information in a common way.

In any case, semantic addressing that intends to facilitate routing decisions is based solely on the address and without the need to find and process information carried in other fields within the packets.

Note that not all semantic addressing schemes exist to facilitate routing (for example, content addressing where the interface ID of the address identifies a chunk of the content to be retrieved), but such schemes are naturally out of scope of this document.

4.4. Flow Marking

Flow marking is a way of indicating, in a specific field in the packet header, the treatment that the packet should receive in the network. In IPv4 the six-bit DSCP field is commonly used for this purpose. In IPv6, while the Traffic Class field could be used, it is generally recommended that the Flow Label field should serve this and a more general purpose.

4.5. Extended Lookup

Routers may also examine fields in the packet other than those in the IP header. For example, many router processes may look at the "five-tuple" consisting of:

- * source address
- * destination address
- * next protocol
- * transport protocol source port
- * transport protocol destination port

4.6. Semantic Field Overloading

"Overloading" is a term applied to placing additional semantics on the contents of a field beyond how it is specified. This is relatively hard to do in an IPv6 header because the number of fields is small, and all fields have specific meanings that are needed in all cases. In IPv4 there may be more opportunity to use some fields in very controlled situations to carry additional semantics that can be used for semantic routing.

4.7. IPv6 Extension Headers

IPv6 defines extension headers explicitly for carrying information that may be used by routers along the path. This information can be used to instruct all routers, only the router indicated by the destination address, or by the ultimate destination of the packet.

Extension headers may carry any information to enable semantic routing.

4.8. New Extensions

Another approach is to define a new protocol extension to carry information on which semantic routing can be performed. Such an extension could be in the form of a new extension header (see Section 4.7) or as a new shim encapsulation (e.g., [RFC7665]).

5. Architectural Considerations

Some semantic routing proposals are intended to be deployed in limited domains [RFC8799] (networks) that are IP-based, while other proposals are intended for use across the Internet. The impact that the proposals have on routing systems may require clean-slate solutions, hybrid solutions, extensions to existing routing protocols, or potentially no changes at all.

Semantic data may be applied in several ways to integrate with existing routing architectures. The most obvious is to build an overlay such that IP is used only to route packets between network nodes that utilize the semantics at a higher layer. An overlay may be achieved in a higher protocol layer, or may be performed using tunneling techniques (such as IP-in-IP [RFC1853]) to traverse the areas of the IP network that cannot parse additional semantics thereby joining together those nodes that use the semantic data.

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 Internet (see Section 5.1). In other cases, traffic from within the domain is exchanged with other domains that are connected together across an IP-based network using tunnels or via application gateways (see Section 5.2). And in still another case traffic from the domain is routed across the Internet to other nodes and this requires backward-compatible routing approaches (see Section 5.3).

5.1. Isolated Domains

Some IP network domains are entirely isolated from the Internet and other IP-based networks. In these cases, there is no risk to external networks from any semantic routing schemes carried out within the domain.

Many approaches in isolated domains will utilize environment-specific routing protocols. For example, those suited to constrained environments (for IoT) or mobile environments (for autonomous vehicles). Such routing protocols can be optimized for the exchange of information specific to semantic routing. However, gateways to provide external connectivity are usually deployed in such networks. Appropriate means should be supported in these means to prevent leaking semantic information beyond the boundaries of these domains.

5.2. Bridged Domains

In some deployments, it will be desirable to connect a number of isolated domains to build a larger network. These domains may be connected (or bridged) over an IP network or even over the Internet.

Ideally, the function of the bridged domains should not be impeded by how they are connected, and the operation of the IP network providing the connectivity should not be compromised by the act of carrying traffic between the domains. This can generally be achieved by tunneling the packets between domains using any tunneling technique, and this will not require the IP network to know about the semantic routing used by the domains.

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 to the boundary of the destination domain where they are mapped back into packets for use within that domain.

5.3. Semantic Prefix Domains

A semantic prefix domain [I-D.jiang-semantic-prefix] 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 addressing and routing so that packets may be routed through the regular IP network (or the Internet) using the prefix and without encountering or having to use any semantic addressing. Once delivered to the semantic prefix domain, a packet can be subjected to whatever semantic routing is enabled in the domain.

6. A Brief Discussion of What Constitutes Routing

This section provides an overview of what is considered as "routing" in the scope of this document. There are many functions in the Internet that contain the concept of routing, but not all of them apply to the scope of this document which is concerned with routing packets at the network layer. A more thorough catalogue of approaches to routing and the applications of semantic routing can be found in [I-D.king-irtf-semantic-routing-survey].

6.1. Application Layer Routing

Routing in the application layer concerns the choice of application-level components that are distributed across the network. The choice may be dependent on the services being delivered, knowledge about the locations in the network that can provide the services, knowledge of the network capabilities, and preferences expressed by an application or user. In this sense, the routing choice consists of constructing an "application layer path" and may be performed at the head end or along the path. Packets are carried between components across the underlying network, using normal transport and network layer protocols that may, themselves, involve routing. Thus, application layer routing is concerned with selecting a series of components based on the potential to carry traffic between them, but without concern for how the packets are routed within the network.

Application layer routing may be used in concepts such as Content Distribution Networking (CDN) and computation in the network (COIN) (see Section 6.9).

The ALTO architecture and protocol [RFC7285] is intended to allow the network to answer queries about the availability and characteristics of paths between application-level components to enable choices to be made by providers of function or content about which components to select. This is a server-based approach because it would be impractical to scale the network reporting all available paths to all destinations to every client, or for the network edge to be able to answer queries from their clients.

6.2. Higher-Layer Path Selection

There is another high-level path selection scenario that is more concerned with selecting outbound paths from the source than in determining destinations or next application-layer hops (as described in Section 6.1. For example, consider a mobile phone that is connected to Wi-Fi and 5G. Further, consider that the Wi-Fi network is dual-homed to two different ISPs. This gives an application a choice of three different paths depending on the known (or

advertised) capabilities of the networks.

This type of scenario is being examined by the Path Aware Networking Research Group (PANRG) where, rather than consulting a server to supply the most appropriate path, the source host or application should learn about the potential paths and pick between them.

6.3. Transport Layer Routing

Some transport layer load balancing schemes and proxy-based connection or discovery mechanisms use a mechanism that looks somewhat like routing, but exists in the transport layer. For example, section 2.1.1 of [RFC3135] describes how a transport layer Performance Enhancing Proxy (PEP) may use a concept called TCP spoofing to terminate a TCP connection and initiate a new connection to the next proxy on the transport layer path towards the destination. The IP addresses of the packets are rewritten at the proxies so that the packets can be routed/forwarded to the next proxy, but no change to the underlying routing system is implied, and this is not Semantic Routing.

6.4. Tunnel-Based Routing

Tunnel-based routing schemes, like those in the transport layer (see Section 6.3), are achieved through an overlay. a tunnel-based scheme relies on encapsulating packets so that they can be sent through the normal routing and forwarding network for delivery to an interim node. That node decapsulates the packet and then either continues to forward the contents or encapsulates the contents in another tunnel. Some approaches, such as onion routing in the Tor project (see [ONION]) use a scheme of multiply-nested encapsulation, with each layer being peeled off at the end of a tunnel.

The packets in a tunnel-based approach are routed and forwarded in the packet network as normal packets and so this approach is not Semantic Routing.

6.5. Inter-Domain Routing

A lot of effort has been devoted to consideration of end-to-end paths for IP traffic across multiple autonomous systems (ASes). For example, the BGP Add-Paths feature [RFC7911] allows the advertisement of multiple paths so that a single, "best" path can be determined. These approaches, however, are principally concerned with overall reachability, and then with selecting the path with the fewest transit autonomous systems. They are less capable of selecting an overall least cost path or of considering other traffic engineering constraints in the selection of end-to-end paths. Such path

computation requires the features outlined in Section 6.7 as assembled into an architectural solution in [RFC7926].

Many approaches have been suggested [RFC6115] for improving inter-domain routing performance and scaling using address partitioning schemes including tunneling across domains (see also Section 6.4). However, routing in this inter-domain scenario is about the selection of the next AS along the path, and possibly a choice of the right AS border router (ASBR) to facilitate that route. This choice of ASBRs might be based on additional information carried in the packets so could qualify as Semantic Routing, but packets flowing between these ASBRs are routed and forwarded within the domains as normal packets without the use of Semantic Routing.

6.6. Service Function Chaining

Service Function Chaining (SFC) [RFC7665] is applied at the network layer to steer packet flows through network functions (such as security or load balancing). A chain of services to be delivered (the service function chain) is realized as sequence of service instances (the service function path). Packets are tunneled between the service instances using encapsulation so that the end-to-end payload packet is unchanged. A variety of network layer encapsulation have been considered including the Network Service Header (NSH) [RFC8300], MPLS [RFC8595], and Segment Routing [I-D.li-spring-sr-sfc-control-plane-framework].

The Segment Routing concept of Network Programming [RFC8986], offers a similar approach to SFC, but may be more widely applicable.

The tunneled packets can be freely routed in the network using conventional shortest path techniques or the mechanisms described in Section 6.7 and Section 6.8, thus this approach is not Semantic Routing.

6.7. Network Layer Traffic Engineering Techniques

Techniques for achieving packet-level traffic engineering in the network layer are described in [I-D.ietf-teas-rfc3272bis]. Traffic engineering (TE) is the process of selecting an end-to-end path that considers many attributes of metrics of the links in the network in order to satisfy a set of constraints or requirements imposed by the sender of the traffic. For example, the sender may want to use only secure links, or may know the bandwidth requirements of the flow, or may need at least a specific end-to-end latency, or indeed any combination of this type of constraint.

Routing for TE may be performed in advance of sending the traffic (for example, by computing a path at the sender or by using a tool such as the Path Computation Element (PCE) [RFC4655]). In this case, some form of encapsulation is needed to bind the traffic flow to the selected route: MPLS or Segment Routing may be used.

Alternatively, the network may be tuned through appropriate use of routing protocol metrics, routing algorithms, and statically configured routes, so that packets will be forwarded along traffic engineered paths.

6.8. Semantic Routing in the Network Layer

Semantic routing, as already explained, is about taking routing decisions based on "additional" information carried in packets in order to provide the behavior and network services most suited to the traffic. This approach builds on the techniques described in Section 6.7 but frees up the network to make individual decisions for each packet based on changing network conditions as well as the information in the packets.

A raft of potential solutions have been proposed for carrying the necessary information in the packets, and it is not the purpose of this document to examine them in detail or make suggestions about which is better. The solutions vary from simply using existing fields in the IP header (such as the ToS field), or examining fields below the IP header (such as the transport ports), through "overloading" existing fields in the packet header (such as the destination address), all the way to adding new information in an additional encapsulation as proposed by the Application-aware Networking (APN) effort [I-D.li-apn-framework].

6.9. Computation In The Network and Semantic Routing

The use of semantic enhancements as a key aspect to Semantic Routing (as described in this document) links the development of Semantic Routing solutions to data plane programmability. Novel approaches to semantic routing may inform the evolution of more complex in-network operations, aiding specific Semantic Routing solutions. Further, progress in routing protocols (e.g., on multi-optimality routing [SOBRINHO]) may be seen as a key input into the more general problem within an emerging framework to distribute state needed for in-network computing operations, e.g., through utilizing insights from routing protocols to distribute routing state for more limited routing operations.

As per its charter, the Computation In The Network (COIN) Research Group [COINRG] combines the idea of computing with the programmability of the data plane. Hence, network operations, such as those previously used for routing and forwarding, may be key to the programmability aspects of "computing in the network" within the scope of COIN. Ultimately, as stated in the COIN charter, "The goal is to investigate how to harness and to benefit from this emerging disruption to the Internet architecture to improve network and application performance as well as user experience." From this, we can conclude that data plane programmability and its impact on existing and emerging areas of communication are key to COIN.

The COIN charter further states, "COIN specifically will focus on the evolution necessary for networking to move beyond packet interception as the basis of network operation and into computation." This envisions that data plane programmability is not limited to packet interception, but may evolve towards more complex operations on data flowing across the network. The analysis of use cases and the identification of key areas of study can drive the understanding of what those additional operations may be and how to program them, particularly across several participating network elements and at the endpoints. With this, we can conclude that the areas for applying COIN ideas will ultimately drive the evolution of COIN technologies by identifying emerging requirements and uses for data plane programmability, particularly those beyond simple packet processing, such as packet forwarding and local buffer management.

Given the focus on steering traffic between micro-services instantiated at computational elements within networks and at endpoints, the COIN use cases identify aspects of what is now named Semantic Routing. Thus Semantic Routing is one possible applicability area for COIN.

Conversely, the availability of emerging data plane programmability may enable new capabilities for Semantic Routing. As a distributed problem, Semantic Routing could be enabled by emerging programming frameworks that may be developed within the work of COIN, possibly leading to new ways of orchestrating and deploying distributed routing programs. Thus, the relationship between Semantic Routing and the COIN Research Group can be characterized as a symbiotic process of informing and enabling that may benefit both work areas.

7. Security Considerations

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.

8. IANA Considerations

This document makes no requests for IANA action.

9. Acknowledgements

Thanks to Brian Carpenter, Dave Oran, and Luigi Iannone for helpful discussions and clarifications.

10. Contributors

Mohamed Boucadair
Email: mohamed.boucadair@orange.com

Dirk Trossen
Email: dirk.trossen@huawei.com

11. References

11.1. Informative References

[BLIND-FORWARDINGref]
Simsek, I., "On-Demand Blind Packet Forwarding",
Paper 30th International Telecommunication Networks and
Applications Conference (ITNAC), 2020, 2020,
<<https://www.computer.org/csdl/proceedings-article/itnac/2020/09315187/1qmfFPPggrC>>.

[CONTENT-RTG-MOBILEeref]

Liu, H. and W. He, "Rich Semantic Content-oriented Routing for mobile Ad Hoc Networks", Paper The International Conference on Information Networking (ICOIN2014), 2014, 2014, <<https://ieeexplore.ieee.org/document/6799682>>.

[EIBPref]

Shenoy, N., "Can We Improve Internet Performance? An Expedited Internet Bypass Protocol", Presentation 28th IEEE International Conference on Network Protocols, 2020, <https://icnp20.cs.ucr.edu/Slides/NIPAA/D-3_invited.pptx>.

[GEO-IPref]

Dasu, T., Kanza, Y., and D. Srivastava, "Geotagging IP Packets for Location-Aware Software-Defined Networking in the Presence of Virtual Network Functions", Paper 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL 2017), 2017, <https://about.att.com/ecms/dam/sites/labs_research/content/publications/AI_Geotagging_IP_Packets_for_Location.pdf>.

[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.jiang-semantic-prefix]

Jiang, S., Sun, Q., Farrer, I., Bo, Y., and T. Yang, "Analysis of Semantic Embedded IPv6 Address Schemas", Work in Progress, Internet-Draft, draft-jiang-semantic-prefix-06, 15 July 2013, <<https://www.ietf.org/archive/id/draft-jiang-semantic-prefix-06.txt>>.

[I-D.jiang-service-oriented-ip]

Carpenter, B., Jiang, S., and G. Li, "Service Oriented Internet Protocol", Work in Progress, Internet-Draft, draft-jiang-service-oriented-ip-03, 14 May 2020, <<https://www.ietf.org/archive/id/draft-jiang-service-oriented-ip-03.txt>>.

[I-D.king-irtf-challenges-in-routing]

King, D., Farrel, A., and C. Jacquenet, "Challenges for the Internet Routing Infrastructure Introduced by Semantic Routing", Work in Progress, Internet-Draft, draft-king-irtf-challenges-in-routing-07, 22 January 2022, <<https://www.ietf.org/archive/id/draft-king-irtf-challenges-in-routing-07.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>>.

[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-spring-sr-sfc-control-plane-framework]

Li, C., Sawaf, A. E., Hu, R., and Z. Li, "A Framework for Constructing Service Function Chaining Systems Based on Segment Routing", Work in Progress, Internet-Draft, draft-li-spring-sr-sfc-control-plane-framework-06, 21 April 2022, <<https://www.ietf.org/archive/id/draft-li-spring-sr-sfc-control-plane-framework-06.txt>>.

[IEN23]

Cohen, D., "IEN 23: On Names, Addresses and Routings", Internet Experiment Note IEN 23, Notebook Section 2.3.3.7, 1978, <<https://www.rfc-editor.org/ien/ien23.pdf>>.

[MULTICAST-SRref]

Jia, W. and W. He, "A Scalable Multicast Source Routing Architecture for Data Center Networks", Paper IEEE Journal on Selected Areas in Communications, vol. 32, no. 1, pp. 116-123, January 2014, 2014, <<https://ieeexplore.ieee.org/document/6799682>>.

[OPENSERNref]

Ren, P., Wang, X., Zhao, B., Wu, C., and H. Sun, "OpenSRN: A Software-defined Semantic Routing Network Architecture", Paper IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), Hong Kong, 2015, 2015, <https://www.researchgate.net/publication/308827498_OpenSRN_A_software-defined_semantic_routing_network_architecture>.

[RFC0814] Clark, D., "Name, addresses, ports, and routes", RFC 814, DOI 10.17487/RFC0814, July 1982, <<https://www.rfc-editor.org/info/rfc814>>.

[RFC1853] Simpson, W., "IP in IP Tunneling", RFC 1853, DOI 10.17487/RFC1853, October 1995, <<https://www.rfc-editor.org/info/rfc1853>>.

[RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, DOI 10.17487/RFC3135, June 2001, <<https://www.rfc-editor.org/info/rfc3135>>.

[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>>.

[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>>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.

[RFC6115] Li, T., Ed., "Recommendation for a Routing Architecture", RFC 6115, DOI 10.17487/RFC6115, February 2011, <<https://www.rfc-editor.org/info/rfc6115>>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.

- [RFC6992] Cheng, D., Boucadair, M., and A. Retana, "Routing for IPv4-Embedded IPv6 Packets", RFC 6992, DOI 10.17487/RFC6992, July 2013, <<https://www.rfc-editor.org/info/rfc6992>>.
- [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>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [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>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.
- [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>>.

- [RFC8595] Farrel, A., Bryant, S., and J. Drake, "An MPLS-Based Forwarding Plane for Service Function Chaining", RFC 8595, DOI 10.17487/RFC8595, June 2019, <<https://www.rfc-editor.org/info/rfc8595>>.
- [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>>.
- [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>>.
- [RFC9139] Gündoan, C., Schmidt, T., Wählisch, M., Scherb, C., Marxer, C., and C. Tschudin, "Information-Centric Networking (ICN) Adaptation to Low-Power Wireless Personal Area Networks (LoWPANs)", RFC 9139, DOI 10.17487/RFC9139, November 2021, <<https://www.rfc-editor.org/info/rfc9139>>.
- [SEMANTICRTG] Strassner, J., Sung-Su, K., and J. Won-Ki, "Semantic Routing for Improved Network Management in the Future Internet", Book Chapter Springer, Recent Trends in Wireless and Mobile Networks, 2010, 2010, <https://link.springer.com/chapter/10.1007/978-3-642-14171-3_14>.
- [SOBRINHO] Sobrinho, J. and M. Ferreira, "Routing on Multiple Optimality Criteria", Paper SIGCOMM 2020, 2020, <<https://dl.acm.org/doi/abs/10.1145/3387514.3405864>>.
- [TERASTREAMref] Zaluski, B., Rajtar, B., Habjani, H., Baranek, M., Slibar, N., Petracic, R., and T. Sukser, "Terastream implementation of all IP new architecture", Paper 36th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2013, 2013, <<https://ieeexplore.ieee.org/document/6596297>>.

11.2. URL References

- [COINRG] Internet Engineering Research Group, "Computation in the Network Research Group (COINRG)", 2022, <<https://irtf.org/coinrg>>.

[ONION] The Tor Project, Inc., "The Onion Routing Project :
Anonymity Online", 2022, <<https://torproject.org/>>.

Authors' Addresses

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

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