Authors: D. King   A. Farrel
        Lancaster University   Old Dog Consulting
A Survey of Semantic Internet Routing Techniques

Abstract

   The Internet Protocol (IP) has become the global standard in any
   computer network, independent of the connectivity to the Internet.
   Generally, an IP address is used to identify an interface on a
   network device. Routing protocols are also required and developed
   based on the assumption that a destination address has this semantic
   with routing decisions made on addresses and additional fields in
   the packet headers.

   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. Many proposals have been made to add semantics to IP
   addresses. The intent is usually to facilitate routing decisions
   based solely on the address and without finding and processing
   information carried in other fields within the packets.

   This document is presented as a survey to support the study and
   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

Table of Contents

1. Introduction

   The Internet continues to expand rapidly, and the Internet Protocol
   (IP) has become the global standard in many types of computer
   network independent of whether or what connectivity to the Internet
   it has. At the same time, there are increasingly varied expectations
   of the services and service level objectives that can be required
   from networks. For example, packet-delivery quality expectations
   beyond best effort is a growth area: throughput, latency, error
   recovery, and (absence of) packet or connectivity loss, reordering,
   or jitter. Requirements include relative or absolute guarantees or
   predictable elastic changes under contention on these performance
   factors. This places significant pressure on Service Providers to be
   aware of the type of services being delivered, and to have access to
   sufficient information about how individual packets should be
   treated to meet the user, application and application instance
   requirements.

   IP addresses facilitate the identification of how a device is
   attached to the Internet and how it is distinguished from every
   other device. Addresses are used to direct packets to a destination
   (destination address) and indicate to where the receiver and network
   replies and error messages should be sent (source address). An IP
   address may be assigned to each network interface of a device
   connected to a network that uses IP. Applications use IP addresses

to both identify a host and to indicate the physical or virtual location of the host.

This document presents a brief survey of proposals to extend the semantics of IP addresses by assigning additional meanings to some parts of the address, or by partitioning the address into a set of subfields that give scoped addressing instructions. Some of these proposals are intended to be deployed in limited domains [RFC8799] that are IP-based, while other proposals are intended for use across the Internet. Limited domains may present their own challenges in terms of ensuring the perimeter of the domains, and connecting domains across the Internet.

The impact that some proposals may have on routing systems could require clean-slate solutions, hybrid solutions, extensions to existing routing protocols, or potentially no changes at all. A separate document ([I-D.king-irtf-challenges-in-routing]) describes the challenges to the routing system presented by changes to IP address semantics, and sets out research questions that should be investigated by those proposing new semantic address schemes.

2. Network Path Selection

Two approaches are typically used for network path selection. Firstly, a priori assessment by having the feasible paths and constraints computed in advance. Secondly, real-time computation in response to changing network conditions.

The first approach may be conducted offline and allows for concurrent or global optimization based on constraints and policy. However. as network size and complexity increase, the required computing power may increase exponentially for this type of computation.

The second approach must consider the speed of calculation where complex constraints are applied to the path selection. This processing may delay service setup and the responsiveness to changes (such as failures) in the network. Network topology filters may be applied to reduce the complexity of the network data and the computation algorithm, however, the path computation accuracy and optimality may be negatively affected.

In both approaches, the amount of information that needs to be imported and processed can become very large (e.g., in large networks, with many possible paths and route metrics), which might impede the scalability of either method both in terms of the storage and the distribution of the information.

In the last decade, significant research has been conducted into the architecture of the future Internet (for example, [RESEARCHFIAref])

and [ITUNET2030ref]). During this research, several techniques
emerged, highlighting the benefits of path awareness and path
selection for end-hosts, and multiple path-aware network
architectures have been proposed, including SCION [SCIONref] and
RINA [RINAref], and the work of the Path Aware Networking Research
Group (PANRG) as discussed later in this document.

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 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.

2.1. Path Aware Routing

The current architecture for IP networking is built using a best-
effort philosophy. There are techniques discussed in this document
that attempt better-than-best-effort delivery. The start-point and
end-point of a path are identified using IP addresses, and traffic
is steered along the path that does not necessarily follow the
"shortest path first" route through the network. Furthermore, the
path might not run all the way from a packet's source to its
destination. The assumption is that a packet reaching the end of a
path is forwarded to its destination using best-effort techniques.

Evaluating and building paths that respect requirements beyond the
simple best-effort model is particularly challenging and
computationally heavy since numerous quality-related parameters need
to be considered.

3. What is Semantic Routing?

Networks are often divided into addressing regions for various
administrative or technological reasons. Different routing paradigms
may be applied in each region, and specific "private" semantics may
be applied to the IP addresses within a single region

These address semantics are established using customer types,
customer connections, topological constraints, performance groups,
and security, etc. Service Providers or network operators will apply
local policies to user and application packets as they enter the
network possibly mapping addresses, or encapsulating them with an
additional IP header. In some case, the packet has its source and

destination within a single network and the network operator can apply address semantics policies across the whole network. In other cases (such as general IP-based traffic), a packet will require a path across multiple networks, and each may apply its own set of traffic forwarding policies. In these cases, there is often no consistency or guaranteed performance unless a Service Level Agreement (SLA) is applied to traffic traversing multiple networks.

Semantic routing proposals may apply to addresses in a specific domain, or domain set. In this context, a "limited domain" means that the interpretation of the address, in a semantic routing domain, is only applicable to a well-defined set of network nodes, or specific points in the network. If a packet bearing an address with a modified semantic were to escape from the domain, the special meaning of the address would be lost. Additionally (or alternatively), the meaning of "specific points in the network" may be applied to the source and destination nodes of a packet, while all transit nodes are unaware of the special semantic. However, it could be the case that some key transit nodes are able to access the meaning of the address and so apply special routing or other functions to the packet.

Such proposals include the following:

  *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].

  *Using addresses to identify different device types so that their traffic may be handled differently [SEMANTICRTG].

  *Content-based routing (CBR), forwarding of the packet based on message content rather than the destination addresses [OPENSRNref].

  *Deriving IP addresses from the lower layer identifiers and using addresses depending on the underlying connectivity (for example, [RFC6282].

  *Identifying hierarchical connectivity so that routing can be simplified [EIBPref].

  *Providing geographic location information within addresses [GEO-IPref].

*Indicating the application or network function on a destination device or at a specific location; or enable Service Function Chaining (SFC).

*Expressing how a packet should be handled, prioritized, or allocated network resources as it is forwarded through the network [TERASTREAMref].

*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].

In many cases, it may be argued that existing mechanisms applied on top of the common address semantic defined in [RFC4291] can deliver the correct functionality for these scenarios. That is, packets may be tunneled over IP using several existing encapsulation techniques. Nevertheless, there is pressure to reduce the amount of encapsulation (partly to resist reduction in the maximum transmission unit (MTU) over the network, and partly to achieve a flatter and more transparent network architecture). This leads to investigations into whether the current IP addresses can be "overloaded" (without any negative connotations being attached to that word) by adding semantics to the addresses.

Semantic Routing is the process of routing packets that contain IP addresses with additional semantics, possibly using that information to perform policy-based routing or other enhanced routing functions. Thus, faciliating enhanced routing decisions based on these additional semantics and provide differentiated paths for different packet flows, distinct from simple shortest path first routing. The process of known as Semantic Routing is discussed further in [I-D.farrel-irtf-introduction-to-semantic-routing].

Key use cases exist for semantic routing, typically for specific applications and deplyments, including low earth orbit (LEO) satellite constellations [I-D.lhan-satellite-semantic-addressing].

Based on a variety of use cases, key technical challenges exist for semantic routing: these are discussed further in [I-D.king-irtf-challenges-in-routing].

3.1. Architectural Considerations

Semantics may be applied in multiple 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. There are several uses of this approach, including Service Function Chaining (SFC) (see Section 4.2.3) and Information Centric Networking (ICN) (see Section

[5.4.1](#)). An overlay may be achieved in a higher layer, or may be performed using tunneling techniques (such as IP-in-IP) to traverse the areas of the IP network that cannot parse additional semantics and so join together those nodes that use the semantics.

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. In those cases, the challenges are limited to enabling the desired function within the domain. In other cases, traffic from within the domain is exchanged with other domains that are connected across an IP-based network using tunnels or via application gateways. And in another case traffic from the domain is routed across the Internet to other nodes and this requires backward-compatible routing approaches, tunnels, or gateway functions.

Limited domains [[RFC8799](#)] are a fact of networking life. They are used to safely deploy or test features and functions in a controlled environment so that they cannot contaminate other networks or the Internet in general. Examples of a limited domain in use today include:

 *Internet of Things (IoT) networks such as factory floors or home networks

 *Deterministic Networks (DetNet) that operate in campus networks or private WANs to provide deterministic data paths with bounded latency, low loss, predictable jitter, and high reliability.

 *Content Distribution Networks (CDN) where clusters of servers share content provision but may also need to be interconnected.

 *Physical security may be provided for a site simply by not permitting traffic to enter or leave the site. This may be expanded by connecting multiple sites together using tunnels across the Internet to form a Virtual Private Network (VPN).

Limited domains are also used as a driver for innovation. They provide a safe space to run experiments and deploy new functions such as advances in traffic steering, improvements in security, and new routing protocols. A limited domain is a way to achieve incremental deployability on an isolated island, and this enables innovation that may (or may not) percolate to the whole Internet at a later stage. For example, experiments to increase the programmability of network forwarding functions need to be carried out in networks of similarly capable nodes (to avoid the risks of broken interoperability or forwarding loops), yet these experiments need to use real user data that is flowing between hosts and servers.

Because limited domains don't always operate in isolation they may need to be connected to other domains over the Internet, or other nodes within the wider Internet.

4. Existing Approaches for Routing Based on Additional Semantics

Several IETF-based approaches are available to allow service providers to perform policy-based routing, including identifying and marking IP traffic either by changing the semantic of IP addresses or by adding such a semantic in other fields/namespaces, enabling differentiated handling by transit routers (queuing, dropping, forwarding, etc.). The sections below distinguish between those schemes that perform routing based on information other than IP addresses, those that establish an overlay network in which to apply semantics, and those that add semantics to the addresses. A further separate group of approaches is presented here to cover the concept of group semantics where a single address identifies more than one endpoint.

4.1. Non-Address-Based Routing

Many routing schemes examine the destination address field and other fields in the packet header to make routing decisions. These approaches (sometimes referred to as "policy-based routing") allow packets to follow different paths through the network depending on semantics assigned to these other fields or based on hashing algorithms operating on the values of those fields.

4.1.1. Deep Packet Inspection

Deep Packet Inspection (DPI) may be used by a router to learn the characteristics of packets in order to forward them differently. This involves looking into the packet beyond the top-level network-layer header to identify the payload. Once identified, the traffic type can be used as an input for marking the packets for network handling, or for performing specific policies on the packets.

However, DPI may be expensive both in processing costs and latency. The processing costs means that dedicated infrastructure is necessary to carry out the function, and this may have an associated financial cost. The latency incurred may be too much for use with any delay/jitter sensitive applications. As a result, DPI is difficult for large-scale deployment and its usage is often limited to specific functions at the edge of the network.

Despite this, "shallow DPI" is commonplace in routers today as they examine the five-tuple of source address, destination address, payload protocol, source port, and destination port to perform a hash function for ECMP purposes (a form of policy-based routing).

4.1.2. Differentiated Services

Quality of Service (QoS) based on Differentiated Services [RFC2474] is a widely deployed framework specifying a simple and scalable coarse-grained mechanism for classifying and managing network traffic. However, in a service providers network, DiffServ codepoint (DSCP) values cannot be trusted when they are set by the customer, and may have different meanings as packets are passed between networks.

In real-world scenarios, Service Providers deploy "remarking" points at the edges of their network, re-classifying received packets by rewriting the DSCP field according to local policy using information such as the source/destination address, IP protocol number, transport layer source/destination ports, and possibly applying DPI as described in Section 4.1.1.

The traffic classification process and node-by-node processing leads to increased packet processing overhead and complexity at the edge of the Service Providers network.

4.1.3. IPv6 Extension Headers

[RFC8200] defines the IPv6 header and also a number of extension headers. These extension headers can be used to carry additional information that may be used by transit routers (the hop-by-hop options header) or by the destination identified by the destination address field (the destination options header). In addition, these extension headers could encode additional semantics that might enable routing decisions and determine what functions and operations should be performed on a packet.

[RFC7872] and [I-D.ietf-v6ops-ipv6-ehs-packet-drops] provide some discussions about the operational problems of using IPv6 extension headers, especially in multi-domain environments, while [I-D.bonica-6man-ext-hdr-update] proposes to update RFC 8200 with guidance regarding the processing, insertion and deletion of IPv6 extension headers.

4.2. Semantic Overlays

An overlay network is built on top of an underlay or transport network. Packets are encapsulated with the header for the overlay network to carry the additional information needed to provide the desired function, and then the packets are encapsulated for transport through the underlay network. In this case, no changes are made to the meaning of the IP addresses in the underlay, but the destination address identifies the next hop in the overlay network rather than the ultimate destination of the packet. In this way,

packets can be steered through different overlay nodes where routing
decisions can be made.

4.2.1. Application-Layer Traffic Optimization

Application-Layer Traffic Optimization (ALTO) [RFC7285] is an
architecture and protocol. ALTO defines abstractions and services to
provide simplified network views and network services to guide the
application usage of network resources, including cost.

An ALTO server gathers information about the network and answers
queries from an ALTO client that wants to find a suitable path for
traffic. ALTO responses are typically used to route whole flows (not
individual packets) either to suitable destinations (such as network
functions) or onto paths that have specific qualities.

4.2.2. Multipath TCP

Multipath TCP (MPTCP) [RFC8684] enables the use of TCP in a
multipath network using multiple host addresses. A Multipath TCP
connection provides a bidirectional bytestream between two hosts
communicating like normal TCP and thus does not require any change
to the applications. However, Multipath TCP enables the hosts to use
different paths with different IP addresses to exchange packets
belonging to the MPTCP connection.

MPTCP increases the available bandwidth, and so provides shorter
delays; it increases fault tolerance, by allowing the use of other
routes when one or more routes become unavailable; and it enables
traffic engineering and load balancing.

4.2.3. Service Function Chaining

Service Function Chaining (SFC) [RFC7665] is the process of sending
traffic through an ordered set (a sequence) of abstract service
functions. This may be achieved using an overlay encapsulation such
as the Network Service Header (NSH) [RFC8300] or MPLS [RFC8596] that
rely on tunneling through an underlay without any additional
semantics applied to the IP addresses.

Alternatively, SFC can be performed by adding semantics to the
addresses, for example, as in Section 4.3.3.

4.2.4. Path Computation Element

The Path Computation Element (PCE) [RFC4655] is an architecture and
protocol [RFC5440] that can be used to assist with network path
selection. A PCE is an entity capable of computing paths for a
single or set of services. A PCE might be a network node, network
management station, or dedicated computational platform that is

resource-aware and has the ability to consider multiple constraints for sophisticated path computation. PCE applications compute label switched paths for MPLS and GMPLS traffic engineering, but the PCE has been extended for a variety of additional traffic engineering problems.

4.3. Semantic Routing

In semantic routing, additional information or meaning is placed into the IP address, and this is used to route packets within the network.

4.3.1. Locator/ID Separation Protocol (LISP)

The Locator/ID Separation Protocol (LISP) [RFC6830] was published by the IETF as an Experimental RFC in 2013 and is now being moved to the Standards Track [I-D.ietf-lisp-rfc6830bis] and [I-D.ietf-lisp-rfc6833bis]. LISP separates IP addresses into two numbering spaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). The former, the EIDs, are used to identify communication end-points (as the name states) as well as local routing and forwarding in the edge network. The latter, RLOCs, are used to locate the EIDs in the Internet topology end are usually the address of ASBRs (Autonomous System Border Routers). IP packets addressed with EIDs are encapsulated with RLOCs for routing and forwarding over the Internet.

As end-to-end packet forwarding includes both EIDs and RLOCs an additional control-plane is needed. This control plane provides a mapping system and basic traffic engineering capabilities. Multihoming becomes easier because one EID can be associated to more than one RLOC or even to a local network address prefix.

4.3.2. Identifier-Locator Network Protocol

The Identifier-Locator Network Protocol (ILNP) [RFC6740] is an experimental network protocol designed to separate the two functions of network addresses: identification of network endpoints, topology or location information. Differently from LISP, ILNP encodes both locator and identifier in the IPv6 address format (128 bits). More specifically, the most significant 64 bits of the 128 bits IPv6 address is the locator, while the less significant 64 bits form the identifier. Upon reaching the destination network, a cache is used to find the corresponding node. Furthermore, DNS can be dynamically updated, which is essential for mobility and also for provider-independent addresses. Similar to LISP, multihoming can be set by assigning multiple locators to the same identifier. In addition, identifiers can also be encrypted for privacy reasons. It was

intended that ILNP should be backwards-compatible with existing IP,
and that it should be incrementally deployable.

4.3.3. Segment Routing

Segment Routing (SR) [RFC8402] leverages the source routing
paradigm. A node steers a packet through an ordered list of
instructions, called "segments". A segment can represent any
instruction, topological or service based. A segment can have a
semantic local to an SR node or global within an SR domain. SR
provides a mechanism that allows a flow to be restricted to a
specific topological path, while maintaining per-flow state only at
the ingress node(s) to the SR domain.

In SR for IPv6 networks (SRv6) segment routing functions are used to
achieve a networking objective that goes beyond packet routing, in
order to provide "network programming" [RFC8986]. The network
program is expressed as a list of instructions, which are
represented as 128-bit segments, called Segment Identifiers (SID) -
encoded and presented in the form of an IPv6 address. The first
instruction of the network program is placed in the Destination
Address field of the packet. If the network program requires more
than one instruction, the remaining list of instructions is placed
in the Segment Routing Extension Header (SRH)[RFC8754].

An SRv6 instruction can represent any topological or service-based
instruction. The SRv6 domain is the service provider domain where
SRv6 services are built to transport any kind of customer traffic
including IPv4, IPv6, or frames. SRv6 is the instantiation of
Segment Routing deployed on the IPv6 data plane. Therefore, in order
to support SRv6, the network must first be enabled for IPv6.

The SRH in the IPv6 header is only processed for nodes forwarding
traffic if the destination address identifies the local node. In
this case, the node must take several actions, including reading the
SRH, performing any node-specific actions identified by the
destination address or the next SIDs in the SRH, and re-writing the
IPv6 destination address field using information from the SRH before
forwarding the packet.

4.3.4. Preferred Path Routing

Preferred Path Routing (PPR) [I-D.chunduri-lsr-isis-preferred-path-
routing] is a proposed routing protocol mechanism where alternate
forwarding state is installed for a set of different preferred
paths. Each preferred path is described as an ordered linear list of
nodes, links, and network functions, and the path is identified by a
network-global preferred path identifier. If a packet is marked with
preferred path identifier, it is forwarded according to the

preferred path that has been installed on the router. If a packet is not marked or if the preferred path is not installed on the router, the packet is forwarded using the normal shortest path first algorithm.

In PPR, the preferred path identifier is encoded in an IP address, but the address is only used in an encapsulation of the end-to-end packet. This approach is a hybrid in that it is applying a different meaning to the IP addresses, using that meaning in an encapsulation, but routing the packets through an existing IP network.

4.3.5. Connectionless Network Protocol

The Connectionless Network Protocol (CLNP) [CLNPref] is a network layer encoding that supports variable length, hierarchical addressing. It is widely deployed in many communications networks and is the ITU-T's standardized encoding for packets in the management plane for Synchronous Digital Hierarchy (SDH) networks. For a while, CLNP was considered in competition with IP as the network layer encoding for the Internet, but IP (in conjunction with TCP) won out.

Many of the considerations for semantic addressing can be handled using CLNP, and it is particularly well suited to applications that demand variable-length addresses or that structure addresses hierarchically for routing or geo-political reasons.

Routing for CLNP can be achieved using the IS-IS routing protocol in its full form as documented in [ISISref] rather than its IP-only form [RFC1195]. While this may make it possible to use CLNP alongside IP in some routed networks, it does not integrate the use of IP addresses with additional semantics with the historic use of IP addresses except in "ships that pass in the night" fashion. Alternatively, [RFC1069] explains how to carry regular IP addresses in CLNP.

4.4. Group Semantics

A mayor enhanced addressing semantic in IP is called "group semantics". Here, an IP address identifies more than one individual interface or node. This facilitates the delivery of a packet to any one of a group of destinations, or to all group members.

4.4.1. Multicast

Multicast address semantics support delivery to all members of a group of destinations. This is a controlled variant of broadcasting where packets are delivered to all possible receivers in a particular (static) scope such as a multi-access link. Membership of

a multicast link is dynamically signalled by the group members, and
a group is identified by a specific address.

IP multicast [RFC1112], based on the protocol and service definition
aspects of Steve Deering's PhD, is widely deployed for IPv4. It is
equally adopted and used in IPv6 using the addressing architecture
specified in [RFC4291]. In IP multicast (Any Source Multicast - ASM)
any node can send to the multicast group and have its packets
delivered to all members of the group.

Research deployments in the 1990s (the so called 'MBone' [MBONEref])
indicated that IP multicast gave rise to a number of issues related
to address assignment, implementation, scale, and security. The
problem of allocation and management of IP multicast (group)
addresses led to several proposals, including Multicast Address
Dynamic Client Allocation Protocol (MADCAP) [RFC2730], the Multicast
Address Allocation Architecture (MALLOC) [RFC6308], the Multicast
Address-Set Claim Protocol (MASC) [RFC2909], and the Multicast-Scope
Zone Announcement Protocol (MZAP) [RFC2776], but none was widely
adopted. Attempts to create a complete routing protocol suite for IP
multicast service model within the IETF resulted in the Multicast
Source Discovery Protocol (MSDP) being published as an experimental
RFC [RFC3618].

The popularity of multicast as a concept and the widespread
deployment of commercial IPv4 multicast led to the development of
"Source Specific Multicast" service (SSM) [RFC4607]. In SSM, the
combination of the Source and Group addresses (S,G) of an IP
multicast packet form a so-called SSM channel address, which
identifies group of receivers and implies a single permitted sender.
Receivers subscribe to every SSM channel.

From a service user's perspective, SSM solves the security issue
(only valid sources can send traffic) and the address assignment
issue (all group addresses are relative to the source address). For
the operator, SSM also eliminates the complex operational
requirements of ASM.

4.4.2. Automatic Multicast Tunneling

Automatic Multicast Tunneling (AMT) [RFC7450] is a protocol for
delivering multicast traffic from sources in a multicast-enabled
network to receivers that lack multicast connectivity to the source
network. The protocol uses UDP encapsulation and unicast replication
to provide this functionality as a hybrid solution using both
multicast routing and an overlay approach.

4.4.3. Bit Index Explicit Replication

The IETF standardized or otherwise deployed protocol solutions in
support of ASM and SSM in about 2015 relied all on per-hop, per ASM-
group/per-SSM-channel stateful hop-by-hop forwarding/replication.
Service Provider at that time were starting to removing or reduce
heavy-weight control and per-hop forwarding processing in unicast
caused by MPLS LDP/RSVP-TE driven designs, replacing it with more
lightweight MPLS-SR and later SRv6 forwarding and associated control
planes. But to reduce the cost for multicast service, the only
transit-hop stateless solution available was ingress-replication,
tunnel multicast across unicast, hence trading hop-by-hop state (and
its control and management plane cost) in the network against
traffic overhead and (under congestion) higher latency.

Bit Index Explicit Replication (BIER) [RFC8279] addresses these
problems. BIER does not contain the notion of ASM or SSM groups.
Instead, a sender enumerates the set of receivers to which the
packet is to be delivered. The network routers forward packets and
replicate them onto the shortest paths to the destinations. As the
packets are replicated, so the enumeration of the receivers is
pruned on each copy of the packet.

BIER is able to use existing routing protocols without modification,
but requires enhancements in the forwarding plane to encode, parse,
and act on the set of receivers. The BIER information is carried in
new encapsulations [RFC8296] that is carried hop-by-hop in IP. Thus,
the additional semantic is in an overlay.

5. Overview of Current Routing Research Work

This section presents a limited survey of techniques and projects
that provide mechanisms to facilitate path and forwarding decisions
based on contextual information.

More recently, the proceedings of the June 2021 Semantic Addressing
and Routing for Future Networks (SARNET-21) Workshop was compiled
and published as [I-D.galis-irtf-sarnet21-report]. It captures the
views and positions of the participants as expressed during the
workshop.

5.1. Forwarding

Some research work is engaged in examining the emerging set of new
requirements that exceed the network and transport services of the
current Internet, which only delivers "best effort" service. This
work aims to determine what features can be built on top of existing
solutions by adding additional new components or features. A
starting point for this discussion can be found in [I-D.bryant-arch-
fwd-layer-ps].

Several additional techniques for improving IP-based routing have been proposed, some of these are highlighted below.

5.1.1. Path Aware Networking

The IRTF's Path Aware Networking Research Group [PANRGref] aims to support research in bringing path aware techniques into use in the Internet. This research overlaps with many past and existing IETF and IRTF efforts, including multipath transport protocols, congestion control in multiply-connected environments, traffic engineering, and alternate routing architectures.

[I-D.irtf-panrg-path-properties] offers a vocabulary of path properties. By doing so it gives some clarity of the distinction between path aware routing and semantic routing as considered in this document.

[I-D.irtf-panrg-what-not-to-do] provides a catalog and analysis of past efforts to develop and deploy Path Aware techniques. Most, but not all, of these mechanisms were considered at higher levels, although some apply at the IP routing and forwarding layer.

5.2. Trust and Accountability

5.2.1. Scalability, Control, and Isolation on Next-Generation Networks

The SCION (Scalability, Control, and Isolation on Next-Generation Networks) [SCIONref] inter-domain network architecture has been designed to address security and scalability issues and provides an alternative to current Border Gateway Protocol (BGP) solutions. The SCION proposal combines a globally distributed public key infrastructure, a way to efficiently derive symmetric keys between any network entities, and the forwarding approach of packet-carried forwarding state.

SCION End-hosts fetch viable path segments from the path server infrastructure, and construct the exact forwarding route themselves by combining those path segments. The architecture ensures that a variety of combinations among the path segments are feasible, while cryptographic protections prevent unauthorized combinations or path-segment alteration. The architecture further enables path validation, providing per-packet verifiable guarantees on the path traversed.

5.3. Layering

5.3.1. Recursive InterNetwork Architecture

Recursive Inter Network Architecture (RINA) [RINAref] builds upon the principle that applications communicate through Inter-process

Communication (IPC) facilities. For an application to communicate
through the distributed IPC facility, it only needs to know the name
of the destination application and to use the IPC interface to
request communication.

By leveraging IPC concepts, RINA allows two processes to
communicate, IPC requires certain functions such as locating
processes, determining permission, passing information, scheduling,
and managing memory. Similarly, two applications on different end-
hosts should communicate by utilizing the services of a distributed
IPC facility (DIF). A DIF is an organizing structure, generally
referred to as a "layer".

The scope and functions provided by the different IPC facilities may
vary given the different type of network and performance goals.
Moreover, an IPC layer may recursively request services from other
IPC layers. The idea of recursively using multiple inter-process
communication services creates a multilayer structure repeated until
an IPC facility can fit well for physical technologies, e.g., wired
or wireless networks.

5.4. Naming

5.4.1. Information Naming

Information-Centric Networking (ICN) [ICNref] is an approach to
evolve the Internet infrastructure away from a host-centric
paradigm, based on perpetual connectivity and the end-to-end
principle, to a network architecture in which the focal point is
information (or content or data) that is assigned specific
identifiers.

Several scenarios exist for semantic-based networking, providing
reachability based on Content Routing [CONTENTref] and Name Data
Networking [NDNref]. The technology area of ICN is now reaching
maturity, after many years of research and commercial investigation.
A technical discussion into the deployment and operation of ICNs
continues in the IETF: [RFC8763] provides several important
deployment considerations for facilitating ICN and practical
deployments.

Although ICN is primarily an overlay technology, a more recently
concept, Hybrid-Information-Centric Networking (hICN), has been
introduced [HICNref]. In an hICN environment the ICN aspect is
integrated into the IPv6 architecture, reusing existing IPv6 packet
formats with the intention of maintaining compatibility with
existing and deployed IP network technology without creating
overlays that might require a new packet format or additional

encapsulations. The work is described in [I-D.muscariello-intarea-hicn].

5.4.2. Service Naming

5.4.2.1. Dynamic Anycast

Dyncast (Dynamic anycast) addresses the problem of directing traffic from a client to one service instance among several available, while considering decision metrics beyond shortest path when doing so. Those service instances are therefore possible destinations for a specific service demand. [I-D.liu-dyncast-ps-usecases] outlines several use cases where such traffic steering requirement is desirable and may occur, such as in edge computing scenarios but also in distribution of video content in scenarios like autonomous driving. The draft also outlines problems with existing solutions, most notably latency in changing relations from one service instance to another due to a change in metric, which defines that decision (e.g., load in servers, latency, or a combination of several such metrics).

Key to the proposed dyncast [I-D.li-dyncast-architecture] architecture is to build on the notion of (IP) anycast, while changing the addressing semantic from a locator-based addressing to a service-oriented one. Here, the initial "service demand" packet is being identified through a service identifier as destination address. This identifier is then mapped onto a binding IP (locator-based) address at the ingress of the network, allowing for locator-based routing to be used throughout the network. The ingress-based architecture is designed in such a way that ingress nodes upon arrival of a new service demand can determine which instance (i.e., which binding IP address) to use considering both network- and service-related metrics. Furthermore, these metrics can be distributed among ingress nodes in various ways, including over a routing protocol solution.

The overall forwarding decision is based on the adherence to what is termed "instance affinity", i.e., the need to adhere to a previous routing decision for more than one packet, unlike IP forwarding on locator addresses. This affinity is created, by means of a binding table on the ingress nodes, since often more than one packet is needed for the overall service-level transaction with a specific service instance. For instance, HTTP requests may span more than one routed packet. Also, a service instance may also create ephemeral state, which requires the client to continue communicating with this instance for the duration of this state. While the affinity is entirely defined by the application layer protocol, the network layer takes the affinity marking as input into the decision to renew its routing decision.

5.4.2.2. Prioritycast

A modification to anycast that can be instantiated by additional
engineering in the routing system is called "prioritycast". Instead
of relying on the shortest path forwarding semantic, prioritycast
directs all traffic to the anycast address instance that is
reachable and has the highest priority. This approach only requires
small modifications to routing protocols so that priorities are
advertised along side the addresses.

Prioritycast was originally introduced as a recommended operational
practice for deployments of Bidirectional PIM (Bidir-PIM) [RFC8736]
which requires a single active instance of its Rendezvous Point (RP)
service. The RP is the root of a bidirectional tree and prioritycast
addresses for RP allow fast failover without additional redundancy
protocols beside the routing protocol, which would otherwise be
necessary for such a redundancy service.

5.4.3. Structured Topological Naming

The Internet uses DNS for single-level name resolution, converting
user-level domain names into IP addresses. However, techniques are
being proposed for multiple levels of name resolution; these would
include: application-level and user-level descriptors, service
identifiers, function identifiers and endpoint identifiers, which
may then be mapped to IP addresses. These additional levels of
naming and resolution would allow services and components to
construct the service to be easily identifiable and directly and
persistently named.

5.4.4. Geographical Naming

TBD

5.4.5. Path-based Naming

TBD

5.4.5.1. ICNP

Information-centric networking (ICN) is an approach to evolve the
Internet infrastructure to directly support this use by introducing
uniquely named data as a core Internet principle. Data becomes
independent from location, application, storage, and means of
transportation, enabling in-network caching and replication.

5.4.5.2. Reed

TBD

5.4.6. Content-Based Routing

The OpenSRN [OPENSRNref] project proposed a Content-Based Routing
Scheme (CBR) that uses packet content and header information to
forward traffic conetextually. This proposal uses a novel software
defined networking architecture to provide a semantic routing for
big data network applications.

5.5. Routing

5.5.1. Inter-Domain Routing

5.5.1.1. Expedited Internet Bypass Protocol

The Expedited Internet Bypass Protocol (EIBP) [EIBPref] is a clean
slate approach to routing and forwarding in the Internet using the
Internet infrastructure, but bypassing the Internet Protocol (IP).
The EIBP method may be deployed in current routers and when invoked
for a specific end to end IP hosts or networks, EIBP bypasses the
heavy traffic and security challenges faced at Layer-3. EIBP does
not require routing protocols, instead it abstracts network
structural (physical or logical) information into intelligent
forwarding addresses that are acquired by EIBP routers
automatically.

The Forwarding tables used by EIBP are proportional to the
connectivity (degree) at a routing device making the protocol
scalable. The EIBP routing system does not require network-wide
dissemination. Topology change impacts are local and thus
instabilities on topology changes are minimal. EIBP is a low
configuration protocol, which can be deployed in an AS and extended
to multiple ASes independently. EIBP evaluations were conducted
using GENI testbeds and compared to IP using Open Shortest Path
First and Border Gateway Protocol. Significant performance
improvements in terms of convergence and churn rates highlight the
capabilities of EIBP.

5.5.2. Intra-Domain Routing

5.6. No Changes Needed

It is entirely possible that some forms of modified address semantic
will work perfectly well with existing routing protocols and
mechanisms either across the whole Internet or within limited and
carefully controlled domains. Claims for this sort of functionality
need to be the subject of careful research and analysis as the
existing protocols were developed with a different view of the
meaning of IP addresses, and because routing systems are notoriously
fragile.

5.7. Use Cases

Several documents are availible that discuss the requirements for applications and services that may benefit from Semantic Routing techniques, including:

o [I-D.boucadair-irtf-sdn-and-semantic-routing] This document examines the applicability of SDN techniques to Semantic Routing and provides considerations for the development of Semantic Routing solutions in the context of SDN.

o [I-D.kw-rtgwg-satellite-rtg-add-challanges] This document summerises near-to-mid-term space-networking problems; it outlines the key components, challenges, and requirements for integrating future space-based network infrastructure with existing networks and mechanisms. Furthermore, this document highlights the network control and transport interconnection, and identify the resources and functions required for successful interconnection of space-based and Earth-based Internet infrastructure.

6. Challenges for Internet Routing Research

Improving IP-based semantic network routing capabilities and capacity so that they scale and address a set of growing requirements presents significant research challenges, and will require contributions from the networking research community.

6.1. 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 addressed in the scope of routing. These questions go beyond "Why do we need this function?" and "What could we achieve by carrying this additional semantic 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 part of the essential groundwork of research into semantic routing itself.

The document "Challenges for the Internet Routing Infrastructure Introduced by Changes in Semantic Routing" [I-D.king-irtf-challenges-in-routing] sets out the challenges for the routing system, and how it might be impacted by the use of semantic routing.

7. Security Considerations

This document is a survey of existing work and so introduces no security considerations of itself. However, many of the proposals

referenced either are intended to improve security or have their own security implications. For example:

  *In-network path selection, the criteria used for selecting the best path may include security considerations.

  *Semantic routing, and applied to specific addresses, may be established using security criteria.

  *Physical security may be provided for a site or limited domain simply by not permitting traffic to enter or leave the site. This may be expanded by connecting multiple sites together using tunnels across the Internet to form a Virtual Private Network (VPN) such that the same level of security is shared by all nodes that participate in the VPN provided that the tunnels are themselves secure.

  *There are also additional complexities for security when any form of multicast or anycast is used because of issues of address assignment and the formation of security associations.

8. IANA Considerations

This document makes no requests for IANA action.

9. Acknowledgements

10. Contributors

        Joanna Dang
        Email: dangjuanna@huawei.com

        Dirk Trossen
        Email: dirk.trossen@huawei.com

11. 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>.

[CLEANSLATEref] Feldmann, A., "Internet Clean-Slate Design: What and
            Why?", Paper Annals of telecommunications-annales des
            telecommunications;64(5):271-6, 2009, 2009, <http://
            ccr.sigcomm.org/online/files/p59-feldmannA.pdf>.

[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>.

[CONTENT-RTG-MOBILEref] 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>.

[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>.

[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>.

[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.boucadair-irtf-sdn-and-semantic-routing] Boucadair, M.,
          Trossen, D., and A. Farrel, "Considerations for the use
          of SDN in Semantic Routing Networks", Work in Progress,
          Internet-Draft, draft-boucadair-irtf-sdn-and-semantic-
          routing-00, 2 February 2022, <https://www.ietf.org/
          archive/id/draft-boucadair-irtf-sdn-and-semantic-
          routing-00.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.galis-irtf-sarnet21-report] Galis, A. and D. Lou, "Semantic
          Addressing and Routing for Future Networks (SARNET-21)
          Workshop Report", Work in Progress, Internet-Draft,
          draft-galis-irtf-sarnet21-report-01, 26 July 2021,
          <https://www.ietf.org/archive/id/draft-galis-irtf-
          sarnet21-report-01.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-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.king-irtf-challenges-in-routing] King, D., Farrel, A., and C.
                  Jacquenet, "Challenges for the Internet Routing Systems
                  Introduced by Semantic Routing", Work in Progress,
                  Internet-Draft, draft-king-irtf-challenges-in-routing-08,
                  25 April 2022, <https://www.ietf.org/archive/id/draft-
                  king-irtf-challenges-in-routing-08.txt>.

     [I-D.kw-rtgwg-satellite-rtg-add-challanges] King, D. and N. Wang,
                  "Routing and Addressing Challenges Introduced by New
                  Satellite Constellations", Work in Progress, Internet-
                  Draft, draft-kw-rtgwg-satellite-rtg-add-challanges-00, 7
                  March 2022, <https://www.ietf.org/archive/id/draft-kw-
                  rtgwg-satellite-rtg-add-challanges-00.txt>.

     [I-D.lhan-satellite-semantic-addressing] Han, L., Li, R., Retana,
                  A., Chen, M., and N. Wang, "Satellite Semantic Addressing

for Satellite Constellation", Work in Progress, Internet-
Draft, draft-lhan-satellite-semantic-addressing-01, 6
March 2022, <https://www.ietf.org/archive/id/draft-lhan-
satellite-semantic-addressing-01.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.sarathchandra-coin-appcentres] Trossen, D., Sarathchandra, C.,
and M. Boniface, "In-Network Computing for App-Centric
Micro-Services", Work in Progress, Internet-Draft, draft-
sarathchandra-coin-appcentres-04, 26 January 2021,
<https://www.ietf.org/archive/id/draft-sarathchandra-
coin-appcentres-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>.

[ISISref]   "Intermediate System to Intermediate System intra-domain
routeing 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>.

[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/>.

[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>.

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

[OPENSRNref] 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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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.

[SCIONref] Barbera, D., Chaut, L., Perrig, A., Reischuk, R., and P.
           Szalachowski, "Patterns in Network Architecture: A Return
           to Fundamentals", Paper The ACM, vol. 60, no. 6, June
           2017, 2017, <https://icnp20.cs.ucr.edu/Slides/NIPAA/
           D-3_invited.pptx>.

[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>.

[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>.

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