

Internet Area Working Group
Internet-Draft
Intended status: Informational
Expires: 7 September 2022

Y. Jia
D. Trossen
L. Iannone
Huawei
N. Shenoy
R.I.T.
P. Mendes
Airbus
D. Eastlake 3rd
Futurewei
P. Liu
China Mobile
D. Farinacci
lispers.net
6 March 2022

Challenging Scenarios and Problems in Internet Addressing
draft-jia-intarea-scenarios-problems-addressing-03

Abstract

The Internet Protocol (IP) has been the major technological success in information technology of the last half century. As the Internet becomes pervasive, IP has been replacing communication technology for many domain-specific solutions. However, domains with specific requirements as well as communication behaviors and semantics still exist and represent what [RFC8799] recognizes as "limited domains".

This document describes well-recognized scenarios that showcase possibly different addressing requirements, which are challenging to be accommodated in the IP addressing model. These scenarios highlight issues related to the Internet addressing model and call for starting a discussion on a way to re-think/evolve the addressing model so to better accommodate different domain-specific requirements.

The issues identified in this document are complemented and deepened by a detailed gap analysis in a separate companion document [I-D.jia-intarea-internet-addressing-gap-analysis].

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 7 September 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. Communication Scenarios in Limited Domains	5
2.1. Communication in Constrained Environments	5
2.2. Communication within Dynamically Changing Topologies	7
2.3. Communication among Moving Endpoints	10
2.4. Communication Across Services	13
2.5. Communication Traffic Steering	14
2.6. Communication with built-in security	15
2.7. Communication protecting user privacy	16
2.8. Communication in Alternative Forwarding Architectures	16
3. Desired Network Features	18
4. Issues in Addressing	21
5. Problem Statement	23
6. Security Considerations	25
7. IANA Considerations	25
8. References	25
8.1. Normative References	25
8.2. Informative References	25
Acknowledgments	33
Authors' Addresses	33

1. Introduction

The Internet Protocol (IP), positioned as the unified protocol at the (Internet) network layer, is seen by many as key to the innovation stemming from Internet-based applications and services. Even more so, with the success of TCP/IP protocol stack, IP has been gradually replacing existing domain-specific protocols, evolving into the core protocol of the entire communication eco-system. At its inception, roughly 40 years ago [RFC0791], the Internet addressing system, represented in the form of the IP address and its locator-based (topological) semantics, has brought the notion of a 'common namespace for all communication'. Compared to proprietary technology-specific solutions, such 'common namespace for all communication' advance ensures end-to-end communication from any device connected to the Internet to another.

However, use cases, associated services, node behaviors, and requirements on packet delivery have since been significantly extended, with the Internet technology being developed to accommodate them in the framework of addressing that stood at the beginning of the Internet's development. This evolution is reflected in the concept of "Limited Domains", first introduced in [RFC8799]. It refers to a single physical network, attached to or running in parallel with the Internet, or is defined by a set of users and nodes distributed over a much wider area, but drawn together by a single virtual network over the Internet. Key to a limited domain is that requirements, behaviors, and semantics could be noticeable local and, more importantly, specific to the limited domain. Very often, the realization of a limited domain is defined by specific communication scenario(s) and/or use case(s) that exhibit the domain-specific behaviors and pose the requirements that lead to the establishment of the limited domain. Identifying limited domains may sometime be not obvious because of blurry boundaries depending on the point of view. For instance, from an end user perspective there is no vision at all on limited domains, hence for end users the dichotomy Internet vs limited domains more transparent. In such cases, it is harder to ensure (and detect) that no limited domain specific semantics leak in the Internet or other limited domains.

One key architectural aspect, when communicating within limited domains, is that of addressing and, therefore, the address structure, as well as the semantic that is being used for packet forwarding (e.g., service identification, content location, device type). The topological location centrality of IP is fundamental when reconciling the often differing semantics for 'addressing' that can be found in those limited domains. The result of this fundamental role of the single IP addressing is that limited domains have to adopt specific solutions, e.g., translating/mapping/converting concepts, semantics, and ultimately, domain-specific addressing, into the common IP addressing used across limited domains.

This document advocates flexibility in addressing in order to accommodate limited domain specific semantics, while, if possible, ensuring a single holistic addressing scheme able to reduce, or even entirely remove, the need for aligning the address semantics of different limited domains, such as the current topological location semantic of the Internet. Ultimately, such holistic addressing could be beneficial to those communication scenarios realized within limited domains by improving efficiency, removing of constraints imposed by needing to utilize the limited semantics of IP addressing, and/or in other ways.

In other words, this document revolves around the following question:

"Should interconnected limited domains purely rely on IP addresses and therefore deal with the complexity of translating any semantic mismatch themselves, or should flexibility for supporting those limited domains be a key focus for an evolved Internet addressing?"

To that end, this document describes well-recognized scenarios in limited domains that could benefit from greater flexibility in addressing and overviews the problems encountered throughout these scenarios due to the lack of that flexibility. A detailed gap analysis can be found in {I-D.jia-intarea-internet-addressing-gap-analysis}}, which elaborates on the issues identified in this memo in reference to extensions to Internet addressing that have attempted to address those issues. The purpose of this memo is rather to stimulate discussion on the emerging needs for addressing at large with the possibility to fundamentally re-think the addressing in the Internet beyond the current objectives of IPv6 [RFC8200].

It is important to remark that any change in the addressing, hence at the data plane level, leads to changes and challenges at the control plane level, i.e., routing. The latter is an even harder problem than just addressing and might need more research efforts that are beyond the objective of this document, which focuses solely on the data plane.

2. Communication Scenarios in Limited Domains

The following sub-sections outline a number of scenarios, all of which belong to the concept of "limited domains" [RFC8799]. While the list of scenarios may look long, this document focuses on scenarios with a number of aspects that can be observed in those limited domains, captured in the sub-section titles. For each scenario, possible challenges are highlighted, which are then picked upon in Section 4, when describing more formally the existing shortcomings in current Internet addressing.

2.1. Communication in Constrained Environments

In a number of communication scenarios, such as those encountered in the Internet of Things (IoT), a simple, communication network demanding minimal resources is required, allowing for a group of IoT network devices to form a network of constrained nodes, with the participating network and end nodes requiring as little computational power as possible and having small memory requirements in order to reduce the total cost of ownership of the network. Furthermore, in the context of industrial IoT, real-time requirements and scalability make IP technology not naturally suitable as communication technology ([OCADO]).

In addition to IEEE 802.15.4, i.e., Low-Rate Wireless Personal Area Network [LR-WPAN], several limited domains exist through utilizing link layer technologies such as Bluetooth Low Energy (BLE) [BLE], Digital European Cordless Telecommunications (DECT) - Ultra Low Energy (ULE) [DECT-ULE], Master-Slave/Token-Passing (MS/TP) [BACnet], Near-Field-Communication (NFC) [ECMA-340], and Power Line Communication (PLC) [IEEE_1901.1].

The end-to-end principle (detailed in [RFC2775]) requires IP addresses (e.g., IPv6 [RFC8200]) to be used on such constrained nodes networks, allowing IoT devices using multiple communication technologies to talk on the Internet. Often, devices located at the edge of constrained networks act as gateway devices, usually performing header compression ([RFC4919]). To ensure security and reliability, multiple gateways must be deployed. IoT devices on the network must select one of those gateways for traffic passthrough by the devices on the (limited domain) network.

Given the constraints imposed on the computational and possibly also communication technology, the usage of a single addressing semantic in the form of a 128-bit endpoint identifier, i.e., IPv6 address, may pose a challenge when operating such networks.

Another type of (differently) constrained environment is an aircraft, which encompasses not only passenger communication but also the integration of real-time data exchange to ensure that processes and functions in the cabin are automatically monitored or actuated. The goal for any aircraft network is to be able to send and receive information reliably and seamlessly. From this perspective, the medium with which these packets of information are sent is of little consequence so long as there is a level of determinism to it. However, there is currently no effective method in implementing wireless inter- and intra-communications between all subsystems. The emerging wireless sensor network technology in commercial applications such as smart thermostat systems, and smart washer/dryer units could be transposed onto aircraft and fleet operations. The proposal for having an Wireless Avionics Intra-Communications (WAIC) system promises reduction in the complexity of electrical wiring harness design and fabrication, reduction in wiring weight, increased configuration, and potential monitoring of otherwise inaccessible moving or rotating aircraft parts. Similar to the IoT concept, WAIC systems consist of short-range communications and are a potential candidate for passenger entertainment systems, smoke detectors, engine health monitors, tire pressure monitoring systems, and other kinds of aircraft maintenance systems.

While there are still many obstacles in terms of network security, traffic control, and technical challenges, future WAIC can enable real-time seamless communications between aircraft and between ground teams and aircraft as opposed to the discrete points of data leveraged today in aircraft communications. For that, WAIC infrastructure should also be connected to outside IP based networks in order to access edge/cloud facilities for data storage and mining. However, the restricted capacity (energy, communication) of most aircraft devices (e.g. sensors) and the nature of the transmitted data - periodic transmission of small packets - may pose some challenges for the usage of a single addressing semantic in the form of a 128-bit endpoint identifier, i.e., an IPv6 address. Moreover, most of the aircraft applications and services are focused on the data (e.g. temperature of gas tank on left wing) and not on the topological location of the data source. This means that the current topological location semantic of IP addresses is not beneficial for aircraft applications and services.

Greater flexibility in Internet addressing may avoid complex and energy hungry operations, like header compression and fragmentation, necessary to translate protocol headers from one limited domain to another, while enabling semantics different from locator-based addressing may better support the communication that occurs in those environments.

2.2. Communication within Dynamically Changing Topologies

Communication may occur over networks that exhibit dynamically changing topologies. One such example is that of satellite networks, providing global Internet connections through a combination of inter-satellite and ground station communication. With the convergence of space-based and terrestrial networks, users can experience seamless broadband access, e.g., on cruise ships, flights, and within cars, often complemented by and seamlessly switching between Wi-Fi, cellular, or satellite based networks at any time [WANG19].

The satellite network service provider will plan the transmission path of user traffic based on the network coverage, satellite orbit, route, and link load, providing potentially high-quality Internet connections for users in areas that are not, or hard to be, covered by terrestrial networks. With large scale LEO (Low Earth Orbit) satellites, the involved topologies of the satellite network will be changing constantly while observing a regular flight pattern in relation to other satellites and predictable overflight patterns to ground users [CHEN21].

Although satellite bearer services are capable of transporting IPv4 and IPv6, as well as associated protocols such as IP Multicast, DNS services and routing information, no IP functionality is implemented on-board the spacecraft limiting the capability of leveraging for instance large scale satellite constellations.

One of the major constraints of deploying routing capability on board of a satellite is power consumption. Due to this, space routers may end up being intermittently powered up during a daytime sunlit pass. Another limitation of the first generation of IP routers in space was the lack of capability to remotely manage and upgrade software while in operation.

The limitations faced in early development of IP based satellite communication payloads, showed the need to develop a flexible networking solution that would enable delay tolerant communications in the presence of intermittent connectivity. Further, in order to reduce latency, which is the major impairment of satellite networks, there was a need of a networking solution able to perform in a scenario encompassing mobile devices with the capability of storing data, leading to a significant reduction of latency, which is the major impairment of satellite networks.

Moreover, due to the current IP addressing scheme and its focus on IP unicast addressing with extended deployment of IP multicast and some IP anycast, current deployments do not take advantage of the broadcast nature of satellite networks.

Moreover networking platforms based on a name (data or service) based addressing scheme would bring several potential benefits to satellite networks aiming to tackle their major challenges, including high propagation delay and changing network topology in the case of LEO constellations.

Another example is that of vehicular communication, where services may be accessed across vehicles, such as self-driving cars, for the purpose of collaborative objection recognition (e.g., for collision avoidance), road status conveyance (e.g., for pre-warning of road-ahead conditions), and other purposes. Communication may include Road Side Units (RSU) with the possibility to create ephemeral connections to those RSUs for the purpose of workload offloading, joint computation over multiple (vehicular) inputs, and other purposes [I-D.ietf-lisp-nexagon]. Communication here may exhibit a multi-hop nature, not just involving the vehicle and the RSU over a direct link. Those topologies are naturally changing constantly due to the dynamic nature of the involved communication nodes.

The advent of Flying Ad-hoc NETworks (FANETs) has opened up an opportunity to create new added-value services [CHRIKI19]. Although these networks share common features with vehicular ad hoc networks, they present several unique characteristics such as energy efficiency, mobility degree, the capability of swarming, and the potential large scale of swarm networks. Due to high mobility of FANET nodes, the network topology changes more frequently than in a typical vehicular ad hoc network. From a routing point of view, although ad-hoc reactive and proactive routing approaches can be used, there are other type of routing protocols that have been developed for FANETS, such as hybrid routing protocols and position based routing protocols, aiming to increase efficiency in large scale networks with dynamic topologies.

Both type of protocols challenge the current Internet addressing semantic: in the case of hybrid protocols, two different routing strategies are used inside and outside a network zone. While inside a zone packets are routed to a specific destination IP address, between zones, query packets are routed to a subset of neighbors as determined by a broadcast algorithm. In the case of position based routing protocol, the IP addressing scheme is not used at all, since packets are routed to a different identifier, corresponding to the geographic location of the destination and not its topological location. Hence, what is needed is to consolidate the geo-spatial addressing with that of a locator-based addressing in order to optimize routing policies across the zones.

Moreover most of the application/services deployed in FANETs tend to be agnostic of the topological location of nodes, rather focusing on the location of data or services. This distinction is even more important because in dynamic network such as FANET robust networking solutions may rely on the redundancy of data and services, meaning that they may be found in more than one device in the network. This in turn may bring into play a possible service-centric semantic for addressing the packets that need routing in the dynamic network towards a node providing said service (or content).

In the aforementioned network technologies, there is a significant difference between the high dynamics of the underlying network topologies, compared to the relative static nature of terrestrial network topology, as reported in [HANDLEY]. As a consequence, the notion of a topological network location becomes restrictive in the sense that not only the relation between network nodes and user endpoint may change, but also the relation between the nodes that form the network itself. This may lead to the challenge of maintaining and updating the topological addresses in this constantly changing network topology.

In attempts to utilize entirely different semantics for the addressing itself, geographic-based routing, such as in [CARTISEAN], has been proposed for MANETs (Mobile Ad-hoc NETWORKs) through providing geographic coordinates based addresses to achieve better routing performance, lower overhead, and lower latency [MANET1].

Flexibility in Internet addressing here would allow for accommodating such geographic address semantics into the overall Internet addressing, while also enabling name/content-based addressing, utilizing the redundancy of many network locations providing the possible data.

2.3. Communication among Moving Endpoints

When packet switching was first introduced, back in the 60s/70s, it was intended to replace the rigid circuit switching with a communication infrastructure that was more resilient to failures. As such, the design never really considered communication endpoints as mobile. Even in the pioneering ALOHA [ALOHA] system, despite considering wireless and satellite links, the network was considered static (with the exception of failures and satellites, which fall in what is discussed in Section 2.2). Ever since, a lot of efforts have been devoted to overcome such limitations once it became clear that endpoint mobility will become a main (if not THE main) characteristic of ubiquitous communication systems.

The IETF has for a long time worked on solutions that would allow extending the IP layer with mobility support. Because of the topological semantic of IP addresses, endpoints need to change addresses each time they visit a different network. However, because routing and endpoint identification is also IP address based, this leads to a communication disruption.

To cope with such a situation, sometimes, the transport layer gets involved in mobility solutions, either by introducing explicit in-band signaling to allow for communicating IP address changes (e.g., in SCTP [RFC5061] and MPTCP [RFC6182]), or by introducing some form of connection ID that allows for identifying a communication independently from IP addresses (e.g., the connection ID used in QUIC [RFC9000]).

Concerning network layer only solutions, anchor-based Mobile IP mechanisms have been introduced ([RFC5177], [RFC6626] [RFC5944], [RFC5275]). Mobile IP is based on a relatively complex and heavy mechanism that makes it hard to deploy and it is not very efficient. Furthermore, it is even less suitable than native IP in constrained environments like the ones discussed in Section 2.1.

Alternative approaches to Mobile IP often leverage the introduction of some form of overlay. LISP [I-D.ietf-lisp-introduction], by separating the topological semantic from the identification semantic of IP addresses, is able to cope with endpoint mobility by dynamically mapping endpoint identifiers with routing locators [I-D.ietf-lisp-mn]. This comes at the price of an overlay that needs its own additional control plane [I-D.ietf-lisp-rfc6833bis].

Similarly, the NV03 (Network Virtualization Overlays) Working Group, while focusing on Data Center environments, also explored an overlay-based solution for multi-tenancy purposes, but also resilient to mobility since relocating Virtual Machines (VMs) is common practice.

NVO3 considered for a long time several data planes that implement slightly different flavors of overlays ([RFC8926], [RFC7348], [I-D.ietf-intarea-gue]), but lacks an efficient control plane specifically tailored for DCs.

Alternative mobility architectures have also been proposed in order to cope with endpoint mobility outside the IP layer itself. The Host Identity Protocol (HIP) [RFC7401] introduced a new namespace in order to identify endpoints, namely the Host Identity (HI), while leveraging the IP layer for topological location. On the one hand, such an approach needs to revise the way applications interact with the network layer, by modifying the DNS (now returning an HI instead of an IP address) and applications to use the HIP socket extension. On the other hand, early adopters do not necessarily gain any benefit unless all communicating endpoints upgrade to use HIP. In spite of this, such a solution may work in the context of a limited domain.

Another alternative approach is adopted by Information-Centric Networking (ICN) [RFC7476]. By making content a first class citizen of the communication architecture, the "what" rather than the "where" becomes the real focus of the communication. However, as explained in the next sub-section, ICN can run either over the IP layer or completely replace it, which in turn can be seen as running the Internet and ICN as logically completely separated limited domains.

Unmanned Aircraft Systems (UAS) are examples of moving devices that require a stable mobility management scheme since they consist of a number of Unmanned Aerial Vehicles (UAV) subordinated to a Ground Control Station (GCS) [MAROJEVIC20]. The information produced by the different sensors and electronic devices available at each UAV is collected and processed by a software or hardware data acquisition unit, being transmitted towards the GCS, where it is inspected and/or analyzed. Analogously, control information transmitted from the GCS to the UAV enables the execution of control operations over the aircraft, such as changing the route planning or the direction pointed by a camera.

Although UAVs may have redundant links to maintain communications in long-range missions (e.g., satellite), most of the communications between the GCS and the UAVs take place over wireless data links, e.g., based on a radio line-of-sight technology, Wi-Fi or 3G/4G/5G. While in some scenarios, UAVs will operate always under the range of the same cellular base station, in missions with large range, UAVs will move between different cellular or wireless ground infrastructure, meaning that the UAV needs to upload its topological locator and re-start the ongoing communication sessions. In such cases, most of existing Mobile IP approaches may play a role, as well as approaches to split the UAV identifier and the topological locator, such as HIP.

However, while the industry is given the first steps towards evolved UAS architectures and communication models, the data-centric communication plays an increasing role, where information is named and decoupled from its location, and applications/services operate over these named data rather than on host-to-host communications.

In this context, the Data Distribution Service ([DDS]) has emerged as an industry-oriented open standard that follows this approach. The space and time decoupling allowed by DDS is very relevant in any dynamic and distributed system, since interacting entities are not forced to know each other and are not forced to be simultaneously present to exchange data. Time decoupling can significantly simplify the management of intermittent data-links, in particular for wireless connectivity between UAS, as well as facilitate seamless UAV mobility between GCSs. This model of communication, in turn, questions the locator-based addressing used in IP and instead utilizes a data-centric naming.

In the case of using TCP/IP, mobility of UAVs introduces a significant challenge. Consider the case where a GCS is receiving telemetry information from a specific UAV. Assuming that the UAV moves and changes its point of attachment to the network, it will have to configure a new IP address on its wireless interface. However, this is problematic, as the telemetry information is still being sent by to the previous IP address of the UAV. This simple example illustrates the necessity to deploy mobility management solutions to handle this type of situations.

However, mobility management solutions increase the complexity of the deployment and may impact the performance of data distribution, both in terms of signaling/data overhead and communication path delay. Considering the specific case of multicast data streams, mobility of content producers and consumers is inherently handled by multicast routing protocols, which are able to react to changes of location of mobile nodes by reconstructing the corresponding multicast delivery

trees. Nevertheless, this comes with a cost in terms of signaling and data overhead (data may still flow through branches of a multicast delivery tree where there are no receivers while the routing protocol is still converging).

Another alternative is to perform the mobility management of producers and consumers not at the application layer based on IP multicast trees, but on the network layer based on an Information Centric Network approach, which was already mentioned in this section.

Greater flexibility in addressing may help in dealing with mobility more efficiently, e.g., through an augmented semantic that may fulfil the mobility requirements [RFC7429] in a more efficient way or through moving from a locator- to a content or service-centric semantic for addressing.

2.4. Communication Across Services

As a communication infrastructure spanning many facets of life, the Internet integrates services and resources from various aspects such as remote collaboration, shopping, content production as well as delivery, education, and many more. Accessing those services and resources directly through URIs has been proposed by methods such as those defined in ICN [RFC7476], where providers of services and resources can advertise those through unified identifiers without additional planning of identifiers and locations for underlying data and their replicas. Users can access required services and resources by virtue of using the URI-based identification, with an ephemeral relationship built between user and provider, while the building of such relationship may be constrained with user- as well as service-specific requirements, such as proximity (finding nearest provider), load (finding fastest provider), and others.

While systems like ICN [CCN] provide an alternative to the topological addressing of IP, its deployment requires an overlay (over IP) or native deployment (alongside IP), the latter with dedicated gateways needed for translation. Underlay deployments are also envisioned in [RFC8763], where ICN solutions are being used to facilitate communication between IP addressed network endpoints or URI-based service endpoints, still requiring gateway solutions for interconnection with ICN-based networks as well as IP routing based networks (cf., [ICN5G][ICNIP]).

Although various approaches combining service and location-based addressing have been devised, the key challenge here is to facilitate a "natural", i.e., direct communication, without the need for gateways above the network layer.

Another aspect of communication across services is that of chaining individual services to a larger service. Here, an identifier would be used that serves as a link to next hop destination within the chain of single services, as done in the work on Service Function Chaining (SFC). With this, services are identified at the level of Layer 2/3 ([RFC7665], [RFC8754], [RFC8595]) or at the level of name-based service identifiers like URLs [RFC8677] although the service chain identification is carried as a Network Service header (NSH) [RFC7665], separate to the packet identifiers. The forwarding with the chain of services utilizes individual locator-based IP addressing (for L3 chaining) to communicate the chained operations from one Service Function Forwarder [RFC7665] to another, leading to concerns regarding overhead incurred through the stacking of those chained identifiers in terms of packet overhead and therefore efficiency in handling in the intermediary nodes.

Greater flexibility in addressing may allow for incorporating different information, e.g., service as well as chaining semantics, into the overall Internet addressing.

2.5. Communication Traffic Steering

Steering traffic within a communication scenario may involve at least two aspects, namely (i) limiting certain traffic towards a certain set of communication nodes and (ii) restraining the sending of packets towards a given destination (or a chain of destinations) with metrics that would allow the selection among one or more possible destinations.

One possibility for limiting traffic inside limited domains, towards specific objects, e.g., devices, users, or group of them, is subnet partition with techniques such as VLAN [RFC5517], VxLAN [RFC7348], or more evolved solution like TeraStream [TERASTREAM] realizing such partitioning. Such mechanisms usually involve significant configuration, and even small changes in network and user nodes could result in a repartition and possibly additional configuration efforts. Another key aspect is the complete lack of correlation of the topological address and any likely more semantic-rich identification that could be used to make policy decisions regarding traffic steering. Suitably enriching the semantics of the packet address, either that of the sender or receiver, so that such decision could be made while minimizing the involvement of higher layer mechanisms, is a crucial challenge for improving on network operations and speed of such limited domain traffic.

When making decisions to select one out of a set of possible destinations for a packet, IP anycast semantics can be applied albeit being limited to the locator semantic of the IP address itself.

Recent work in [SFCANYCAST] suggests utilizing the notion of IP anycast address to encode a "service identifier", which is dynamically mapped onto network locations where service instances fulfilling the service request may be located. Scenarios where this capability may be utilized are provided in [SFCANYCAST] and include, but are not limited to, scenarios such as edge-assisted VR/AR, transportation, smart cities, smart homes, smart wearables, and digital twins.

The challenge here lies in the possible encoding of not only the service information itself but the constraint information that helps the selection of the "best" service instance and which is likely a service-specific constraint in relation to the particular scenario. The notion of an address here is a conditional (on those constraints) one where this conditional part is an essential aspect of the forwarding action to be taken. It needs therefore consideration in the definition of what an address is, what is its semantic, and how the address structure ought to look like.

As outlined in the previous sub-section, chaining services are another aspect of steering traffic along a chain of constituent services, where the chain is identified through either a stack of individual identifiers, such as in Segment Routing [RFC8402], or as an identifier that serves as a link to next hop destination within the chain, such as in Service Function Chaining (SFC). The latter can be applied to services identified at the level of Layer 2/3 ([RFC7665], [RFC8754], [RFC8595]) or at the level of name-based service identifiers like URLs [RFC8677]. However, the overhead incurred through the stacking of those chained identifiers is a concern in terms of packet overhead and therefore efficiency in handling in the intermediary nodes.

Flexibility in addressing may enable more semantic rich encoding schemes that may help in steering traffic at hardware level and speed, without complex mechanisms usually resulting in handling packets in the slow path of routers.

2.6. Communication with built-in security

Today, strong security in the Internet is usually implemented as a general network service ([PILA], [RFC6158]). Among the various reasons for such approach is the limited semantic of current IP addresses, which do not allow to natively express security features or trust relationships. Efforts like Cryptographically Generated Addresses (CGA) [RFC3972], provide some security features by embedding a truncated public key in the last 57-bit of IPv6 address, thereby greatly enhancing authentication and security within an IP network via asymmetric cryptography and IPsec [RFC4301]. The

development of the Host Identity Protocol (HIP) [RFC7401] saw the introduction of cryptographic identifiers for the newly introduced Host Identity (HI) to allow for enhanced accountability, and therefore trust. The use of those HIs, however, is limited by the size of IPv6 128bit addresses.

Through a greater flexibility in addressing, any security-related key, certificate, or identifier could instead be included in a suitable address structure without any information loss (i.e., as-is, without any truncation or operation as such), avoiding therefore compromises such as those in HIP. Instead, CGAs could be created using full length certificates, or being able to support larger HIP addresses in a limited domain that uses it. This could significantly help in constructing a trusted and secure communication at the network layer, leading to connections that could be considered as absolute secure (assuming the cryptography involved is secure). Even more, anti-abuse mechanisms and/or DDoS protection mechanisms like the one under discussion in PEARG ([PEARG]) Research Group may leverage a greater flexibility of the overall Internet addressing, if provided, in order to be more effective.

2.7. Communication protecting user privacy

See Comments in Section "Issues".

2.8. Communication in Alternative Forwarding Architectures

The performance of communication networks has long been a focus for optimization due to the immediate impact on cost of ownership for communication service providers. Technologies like MPLS [RFC3031] have been introduced to optimize lower layer communication, e.g., by mapping L3 traffic into aggregated labels of forwarding traffic for the purposes of, e.g., traffic engineering.

Even further, other works have emerged in recent years that have replaced the notion of packets with other concepts for the same purpose of improved traffic engineering and therefore efficiency gains. One such area is that of Software Defined Networks (SDN) [RFC7426], which has highlighted how a majority of Internet traffic is better identified by flows, rather than packets. Based on such observation, alternate forwarding architectures have been devised that are flow-based or path-based. With this approach, all data belonging to the same traffic stream is delivered over the same path, and traffic flows are identified by some connection or path identifier rather than by complete routing information, possibly enabling fast hardware based switching (e.g. [DETNET], [PANRG]).

On the one hand, such a communication model may be more suitable for real-time traffic like in the context of Deterministic Networks ([DETNET]), where indeed a lot of work has focused on how to "identify" packets belonging to the same DETNET flow in order to jointly manage the forwarding within the desired deterministic boundaries.

On the other hand, it may improve the communication efficiency in constrained wireless environments (cf., Section 2.1), by reducing the overhead, hence increasing the number of useful bits per second per Hz.

Also, the delivery of information across similar flows may be combined into a multipoint delivery of a single return flow, e.g., for scenarios of requests for a video chunk from many clients being responded to with a single (multi-destination) flow, as outlined in [BIER-MC] as an example. Another opportunity to improve communication efficiency is being pursued in ongoing IETF/IRTF work to deliver IP- or HTTP-level packets directly over path-based or flow-based transport network solutions, such as in [TROSSEN][BIER-MC][ICNIP][ICN5G] with the capability to bundle unicast forward communication streams flexibly together in return path multipoint relations. Such capability is particularly opportune in scenarios such as chunk-based video retrieval or distributed data storage. However, those solutions currently require gateways to "translate" the flow communication into the packet-level addressing semantic in the peering IP networks. Furthermore, the use of those alternative forwarding mechanisms often require the encapsulation of Internet addressing information, leading to wastage of bandwidth as well as processing resources.

Providing an alternative way of forwarding data has also been the motivation for the efforts created in the European Telecommunication Standards Institute (ETSI), which formed an Industry Specification Group (ISG) named Non-IP Networking (NIN) [ETSI-NIN]. This group sets out to develop and standardize a set of protocols leveraging an alternative forwarding architecture, such as provided by a flow-based switching paradigm. The deployment of such protocols may be seen to form limited domains, still leaving the need to interoperate with the (packet-based forwarding) Internet; a situation possibly enabled through a greater flexibility of the addressing used across Internet-based and alternative limited domains alike.

As an alternative to IP routing, EIBP (Extended Internet Bypass Protocol) [EIBP] offers a communications model that can work with IP in parallel and entirely transparent and independent to any operation at network layer. For this, EIBP proposes the use of physical and/or virtual structures in networks and among networks to auto assign

routable addresses that capture the relative position of routers in a network or networks in a connected set of networks, which can be used to route the packets between end domains. EIBP operates at Layer 2.5 and provides encapsulation (at source domain), routing, and de-encapsulation (at destination domain) for packets. EIBP can forward any type of packets between domains. A resolver to map the domain ID to EIBP's edge router addresses is required. When queried for a specific domain, the resolver will return the corresponding edge router structured addresses.

EIBP decouples routing operations from end domain operations, offering to serve any domain, without point solutions to specific domains. EIBP also decouples routing IDs or addresses from end device/domain addresses. This allows for accommodation of new and upcoming domains. A domain can extend EIBP's structured addresses into the domain, by joining as a nested domain under one or more edge routers, or by extending the edge router's structure addresses to its devices.

A greater flexibility in addressing semantics may reduce the aforementioned wastage by accommodating Internet addressing in the light of such alternative forwarding architectures, instead enabling the direct use of the alternative forwarding information.

3. Desired Network Features

From the previous subsection, we recognize that Internet technologies are used across a number of scenarios, each of which brings their own (vertical) view on needed capabilities in order to work in a satisfactory manner to those involved.

In the following, we complement those vertical-specific insights with answers to the question of network features that end users (in the form of individuals or organizations alike) desire from the networked system at large. Answers to this question look at the network more from a horizontal perspective, i.e. not with a specific usage in mind beyond communication within and across networks. The text here summarizes the discussion that took place on the INT Area mailing list after IETF112 on this issue. For some of those identified features, we can already identify how innovations on addressing may impact the realization of a particular feature.

We then combine the insights from both scenario-specific and wider horizontal views for the identification of issues when realizing the specific capability of addressing, presented in Section 4.

1. Always-On: The world is getting more and more connected, leading to being connected to the Internet, anywhere, by any technology (e.g., cable, fiber, or radio), even simultaneously, "all the time", and, most importantly, automatically (without any switch turning). However, when defining "all the time" there is a clear and important difference to be made between availability and reliability vs "desired usage". In other words, "always on" can be seen as a desirable perception at the end user level or as a characteristic of the underlying system. From an end user perspective, clearly the former is of importance, not necessarily leading to an "always on" system notion but instead "always-app-available", merely requiring the needed availability and reliability to realize the perception of being "always on" (e.g., for earthquake alerts), possibly complemented by app-specific methods to realize the "always on" perception (e.g., using local caching rather than communication over the network).
2. Transparency: Being agnostic with respect to local domains network protocols (Bluetooth, ZigBee, Thread, Airdrop, Airplay, or any others) is key to provide an easy and straightforward method for contacting people and devices without any knowledge of network issues, particularly those specific to network-specific solutions. While having a flexible addressing model that accommodates a wide range of use cases is important, the centrality of the IP protocol remains key as a mean to provide global connectivity.
3. Multi-homing: Seamless multi-homing capability for the host is key to best use the connectivity options that may be available to an end user, e.g., for increasing resilience in cases of failures of one available option. Protocols like LISP, SHIM6, QUIC, MPTCP, SCTP (to cite a few) have been successful at providing this capability in an incremental way, but too much of that capability is realized within the application, making it hard to leverage across all applications. While today each transport protocol has its own way to perform multi-address discovery, the network layer should provide the multi-homing feature (e.g., SHIM6 can be used to discover all addresses on both ends), and then leave the address selection to the transport. With that, multi-address discovery remains a network feature exposed to the upper layers. This may also mean to update the Socket API (which may be actually the first thing to do), which does not necessarily mean to expose more network details to the applications but instead be more address agnostic yet more expressive.

4. **Mobility:** A lot of work has been put in MobileIP ([RFC5944],[RFC6275]) to provide seamless and lossless communications for moving nodes (vehicle, satellites). However, it has never been widely deployed for several reasons, like complexity of the protocol and the fact that the problem has often been tackled at higher layers, with applications resilient to address changes. However, similar to multi-homing, solving the problem at higher layers means that each and every transport protocol and application have their own way to deal with mobility, leading to similar observations as those for the previous multi-homing aspect.
5. **Security and Privacy:** The COVID-19 pandemic has boosted end users' desire to be protected and protect their privacy. The balance among privacy, security, and accountability is not simple to achieve. There exist different views on what those properties should be, however the network should provide the means to provide what is felt as the best trade-off for the specific use case.
6. **Performance:** While certainly desirable, "performance" is a complex issue that depends on the objectives of those building for but also paying for performance. Examples are (i) speed (shorter paths/direct communications), (ii) bandwidth (10petabit/s for a link), (iii) efficiency (less overlays/encapsulations), (iv) high efficacy or sustainability (avoid waste). From an addressing perspective, length/format/semantics that may adapt to the specific use case (e.g. use short addresses for low power IoT, or, where needed, longer for addresses embedding certificates for strong authentication, authorization and accountability) may contribute to the performance aspects that end users desire, such as reducing waste through not needed encapsulation or needed conversion at network boundaries.
7. **Availability, Reliability, Predictability:** These three properties are important to enable wide-range of services and applications according to the desired usage (cf. point 1).
8. **Do not do harm:** Access to the Internet is considered a human right [RFC8280]. Access to and expression through it should align with this core principle. This issue transcends through a variety of previously discussed 'features' that are desired, such as privacy, security but also availability and reliability. However, lifting the feature of network access onto a basic rights level also brings in the aspect of "do not do harm" through the use of the Internet with respect to wider societal objectives. Similar to other industries, such as electricity or cars, preventing harm usually requires an interplay of

commercial, technological, and regulatory efforts, such as the enforcement of seat belt wearing to reduce accident death. As a first step, the potential harmfulness of a novel method must be recognized and weighted against the benefits of its introduction and use. One increasingly important consideration in the technology domain is "sustainability" of resource usage for an end user's consumption of and participation in Internet services. As an example, Distributed Ledger Technologies (DLT) are seen as an important tool for a variety of applications, including Internet decentralization ([DINRG]). However, the non-linear increase in energy consumption means that extending proof-of-work systems to the entire population of the planet would not only be impractical but also possibly highly wasteful, not just at the level of computational but also communication resource usage [DLT-draft]. This poses the question on how novel methods for addressing may improve on sustainability of such technologies, particularly if adopted more widely.

9. Maximum Transmission Unit (MTU): One long standing issue in the Internet is related to the MTU and how to discover the path MTU in order to avoid fragmentation ([I-D.ietf-6man-mtu-option], [I-D.templin-6man-aero]). While it makes sense to always leverage as much performance from local systems as possible, this should come without sacrificing the ability to communicate with all systems. Having a solid solution to solve the issue would make the overall interconnection of systems more robust.

4. Issues in Addressing

The desired properties outlined in the previous section have implications that go beyond addressing and need to be tackled from a larger architectural point of view. Such a discussion is left as future action, limiting the present document at discussing only the addressing viewpoint and identifying shortcomings perceived from this perspective.

There are a number of issues that we can identify from the communication scenarios in Section 2 and the network features generally desire from the network, presented in Section 3. We do not claim to be exhaustive in our list:

1. Limiting Alternative Address Semantics: Several communication scenarios pursue the use of alternative semantics of what constitute an 'address' of a packet traversing the Internet, which may fall foul of the defined network interface semantic of IP addresses.

2. **Hampering Security:** Aligning with the semantic and length limitations of IP addressing may hamper the security objectives of any new semantic, possibly leading to detrimental effects and possible other workarounds (at the risk of introducing fragility rather than security).
1. **Hampering Privacy:**
 - * Easy individual identification
 - * Flow linkability
 - * App/Activity profiling
2. **Complicating Traffic Engineering:** Utilizing a plethora of non-address inputs into the traffic steering decision in real networks complicates traffic engineering in that it makes the development of suitable policies more complex, while also leading to possible contention between methods being used.
3. **Hampering Efficiency:** Extending IP addressing through point-wise solutions also hampers efficiency, e.g., through needed re-encapsulation (therefore increasing the header processing overhead as well as header-to-payload ratio), through introducing path stretch, or through requiring compression techniques to reduce the header proportion of large addresses when operating in constrained environments.
4. **Fragility:** The introduction of point solutions, each of which comes with possibly own usages of address or packet fields, together with extension-specific operations, increases the overall fragility of the resulting system, caused, for instance, through contention between feature extensions that were neither foreseen in the design nor tested during the implementation phase.
5. **Extensibility:** Accommodating new requirements through ever new extensions as an extensibility approach to addressing compounds aspects discussed before, i.e., fragility, efficiency etc. It complicates engineering due to the clearly missing boundaries against which contentions with other extensions could be managed. It complicates standardization since extension-based extensibility requires independent, and often lengthy, standardization processes. And ultimately, deployments are complicated due to backward compatibility testing required for any new extension being integrated into the deployed system.

The table below shows how the above identified issues do arise somehow in our outlined communication scenarios in Section 2. This overview will be deepened in more details in the gap analysis document [I-D.jia-intarea-internet-addressing-gap-analysis].

	Issue 1	Issue 2	Issue 3	Issue 4	Issue 5	Issue 6
Constrained Environments				x	x	x
Dynamically Changing Topologies	x		x	x	x	x
Moving Endpoints	x		x	x	x	x
Across Services	x		x	x	x	x
Traffic Steering	x		x	x	x	x
Built-in Security	x	x		x	x	x
Alternative Forwarding Architectures	x			x		x

Table 1: Issues Involved in Challenging Scenarios

5. Problem Statement

This document identifies a number of scenarios as well as general features end users would want from the network, positioning the existing Internet addressing structure itself as a potential hindrance in solving key problems for Internet service provisioning. Such problems include supporting new, e.g., service-oriented, scenarios more efficiently, with improved security and efficient traffic engineering, as well as large scale mobility. We can observe that those new forms of communication are particularly driven by the conceptual framework of limited domains, realizing the requirements of stakeholders for an optimized communication in those limited domains, while still utilizing the Internet for interconnection as

well as for access to the wealth of existing Internet services.

This co-existence of optimized LD-level as well as Internet communication creates a tussle between those requirements on addressing stemming from those limited domains and those coming from the Internet in the form of agreed IPv6 semantics. This tussle directly refers back to our introductory question on flexibility in addressing (or leaving the problem to limited domain solutions to deal with). It is also captured in the discussion on where new features are being introduced, i.e. at the edge or core of the Internet.

But more importantly, the question on 'what is an address anyway' (derived from what features we may want from the network) should not be guided by the answers that the Internet can give us today, e.g., being a mere ephemeral token for accessing PoP-based services (as indicated in related arch-d mailing list discussions), but instead what features could be enabled by a particular view of what an address is. However, that is not to 'second guess' the market and its possible evolution, but to outline clear features from which to derive clear principles for a design.

For this, it is important to recognize that skewing the technical capabilities of any feature, let alone addressing, to the current economic situation of the Internet bears the danger of locking down innovation capabilities as an outcome of those technical limitations being introduced. Instead, addressing must align with enabling the model of permissionless but compatible innovation that the IETF has been promoting, ultimately enabling the serendipity of new applications that has led to many of those applications we can see in the Internet today.

At this stage, this document does not provide a definite answer nor does it propose or promote specific solutions to the problems here portrayed. Instead, this document aims at stimulating discussion on the emerging needs for addressing, with the possibility to fundamentally re-think the addressing in the Internet beyond the current objectives of IPv6, in order to provide the flexibility to suitably support the many new forms of communication that will emerge. Addressing can be rather flexible and can be of any form that applications may need. There is no limitation on the address to preclude any future applications.

To complement the problem statement in this document, the companion gap analysis document [I-D.jia-intarea-internet-addressing-gap-analysis] deepens the issues identified in Section 4 along key properties of today's Internet addressing.

6. Security Considerations

The present memo does not introduce any new technology and/or mechanism and as such does not introduce any security threat to the TCP/IP protocol suite.

Nevertheless, it is worth to observe whether or not greater flexibility of addressing (as suggested in previous sections) would allow to introduce fully featured security in endpoint identification, potentially able to eradicate the spoofing problem, as one example. Furthermore, it may be used to include application gateways' certificates in order to provide more efficiency, e.g., using web certificates also in the addressing of web services. While increasing security, privacy protection may also be improved.

7. IANA Considerations

This document does not include an IANA request.

8. References

8.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

8.2. Informative References

- [ALOHA] Kuo, F., "The ALOHA System", ACM SIGCOMM Computer Communication Review Vol. 25, pp. 41-44, DOI 10.1145/205447.205451, January 1995, <<https://doi.org/10.1145/205447.205451>>.
- [BACnet] "BACnet-A Data Communication Protocol for Building Automation and Control Networks", ANSI/ASHRAE Standard 135-2016, January 2016, <https://www.techstreet.com/ashrae/standards/ashrae-135-2016?product_id=1918140>.
- [BIER-MC] Trossen, D., Rahman, A., Wang, C., and T. Eckert, "Applicability of BIER Multicast Overlay for Adaptive Streaming Services", Work in Progress, Internet-Draft, draft-ietf-bier-multicast-http-response-06, 10 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-bier-multicast-http-response-06.txt>>.

- [BLE] "Bluetooth Specification", Bluetooth SIG Working Groups, n.d., <<https://www.bluetooth.com/specifications>>.
- [CARTISEAN] Hughes, L., Shumon, K., and Y. Zhang, "Cartesian Ad Hoc Routing Protocols", Ad-Hoc, Mobile, and Wireless Networks pp. 287-292, DOI 10.1007/978-3-540-39611-6_27, 2003, <https://doi.org/10.1007/978-3-540-39611-6_27>.
- [CCN] Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggs, N., and R. Braynard, "Networking named content", Proceedings of the 5th international conference on Emerging networking experiments and technologies - CoNEXT '09, DOI 10.1145/1658939.1658941, 2009, <<https://doi.org/10.1145/1658939.1658941>>.
- [CHEN21] Chen, Y., Li, H., Liu, J., Wu, Q., and Z. Lai, "GAMS: An IP Address Management Mechanism in Satellite Mega-constellation Networks", 2021 International Wireless Communications and Mobile Computing (IWCMC), DOI 10.1109/iwcmc51323.2021.9498722, June 2021, <<https://doi.org/10.1109/iwcmc51323.2021.9498722>>.
- [CHRIKI19] Chriki, A., Touati, H., Snoussi, H., and F. Kamoun, "FANET: Communication, mobility models and security issues", Computer Networks Vol. 163, pp. 106877, DOI 10.1016/j.comnet.2019.106877, November 2019, <<https://doi.org/10.1016/j.comnet.2019.106877>>.
- [DDS] AL-Madani, B., Elkhider, S., and S. El-Ferik, "DDS-Based Containment Control of Multiple UAV Systems", Applied Sciences Vol. 10, pp. 4572, DOI 10.3390/app10134572, July 2020, <<https://doi.org/10.3390/app10134572>>.
- [DECT-ULE] "Digital Enhanced Cordless Telecommunications (DECT); Common Interface (CI); Part 1: Overview", ETSI European Standard, EN 300 175-1, V2.6.1, May 2009, <https://www.etsi.org/deliver/etsi_en/300100_300199/30017501/02.06.01_60/en_30017501v020601p.pdf>.
- [DETNET] "Deterministic Networking (DetNet)", n.d., <<https://datatracker.ietf.org/wg/detnet/about/>>.
- [DINRG] "Decentralized Internet Infrastructure - DINRG", n.d., <<https://datatracker.ietf.org/rg/dinrg/about/>>.

[DLT-draft]

Trossen, D., Guzman, D., Bride, M. M., and X. Fan, "Impact of DLTs on Provider Networks", Work in Progress, Internet-Draft, draft-trossen-rtgwg-impact-of-dlts-01, 2 March 2022, <<https://www.ietf.org/archive/id/draft-trossen-rtgwg-impact-of-dlts-01.txt>>.

[ECMA-340] EECMA-340, "Near Field Communication - Interface and Protocol (NFCIP-1) 3rd Ed.", June 2013.

[EIBP] Shenoy, S Chandraiah, P Willis, N., "A Structured Approach to Routing in the Internet", June 2021, <First Intl Workshop on Semantic Addressing and Routing for Future Networks>.

[ETSI-NIN] ETSI - European Telecommunication Standards Institute, "Non-IP Networking - NIN", n.d., <<https://www.etsi.org/technologies/non-ip-networking>>.

[HANDLEY] Handley, M., "Delay is Not an Option: Low Latency Routing in Space", Proceedings of the 17th ACM Workshop on Hot Topics in Networks, DOI 10.1145/3286062.3286075, November 2018, <<https://doi.org/10.1145/3286062.3286075>>.

[I-D.ietf-6man-mtu-option]

Hinden, R. M. and G. Fairhurst, "IPv6 Minimum Path MTU Hop-by-Hop Option", Work in Progress, Internet-Draft, draft-ietf-6man-mtu-option-13, 28 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-6man-mtu-option-13.txt>>.

[I-D.ietf-intarea-gue]

Herbert, T., Yong, L., and O. Zia, "Generic UDP Encapsulation", Work in Progress, Internet-Draft, draft-ietf-intarea-gue-09, 26 October 2019, <<https://www.ietf.org/archive/id/draft-ietf-intarea-gue-09.txt>>.

[I-D.ietf-lisp-introduction]

Cabellos, A. and D. S. (Ed.), "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-introduction-15, 20 September 2021, <<https://www.ietf.org/archive/id/draft-ietf-lisp-introduction-15.txt>>.

[I-D.ietf-lisp-mn]

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

[I-D.ietf-lisp-nexagon]

Barkai, S., Fernandez-Ruiz, B., Tamir, R., Rodriguez-Natal, A., Maino, F., Cabellos-Aparicio, A., and D. Farinacci, "Network-Hexagons: H3-LISP GeoState & Mobility Network", Work in Progress, Internet-Draft, draft-ietf-lisp-nexagon-19, 14 September 2021, <<https://www.ietf.org/archive/id/draft-ietf-lisp-nexagon-19.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-30, 18 November 2020, <<https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6833bis-30.txt>>.

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

Jia, Y., Trossen, D., Iannone, L., Shenoy, N., and P. Mendes, "Gap Analysis in Internet Addressing", Work in Progress, Internet-Draft, draft-jia-intarea-internet-addressing-gap-analysis-01, 23 October 2021, <<https://www.ietf.org/archive/id/draft-jia-intarea-internet-addressing-gap-analysis-01.txt>>.

[I-D.templin-6man-aero]

Templin, F. L., "Automatic Extended Route Optimization (AERO)", Work in Progress, Internet-Draft, draft-templin-6man-aero-39, 22 February 2022, <<https://www.ietf.org/archive/id/draft-templin-6man-aero-39.txt>>.

[ICN5G]

Ravindran, R., Suthar, P., Trossen, D., Wang, C., and G. White, "Enabling ICN in 3GPP's 5G NextGen Core Architecture", Work in Progress, Internet-Draft, draft-irtf-icnrg-5gc-icn-04, 10 January 2021, <<https://www.ietf.org/archive/id/draft-irtf-icnrg-5gc-icn-04.txt>>.

- [ICNIP] Trossen, D., Robitzsch, S., Reed, M., Al-Naday, M., and J. Riihijarvi, "Internet Services over ICN in 5G LAN Environments", Work in Progress, Internet-Draft, draft-trossen-icnrg-internet-icn-5gln-04, 1 October 2020, <<https://www.ietf.org/archive/id/draft-trossen-icnrg-internet-icn-5gln-04.txt>>.
- [IEEE_1901.1] "Standard for Medium Frequency (less than 15 MHz) Power Line Communications for Smart Grid Applications", IEEE 1901.1 IEEE-SA Standards Board, May 2018, <<https://ieeexplore.ieee.org/document/8360785>>.
- [LR-WPAN] "IEEE 802.15.4 - IEEE Standard for Low-Rate Wireless Networks", IEEE 802.15 WPAN Task Group 4, May 2020, <https://standards.ieee.org/standard/802_15_4-2020.html>.
- [MANET1] Abdallah, A., Abdallah, E., Bsoul, M., and A. Ootom, "Randomized geographic-based routing with nearly guaranteed delivery for three-dimensional ad hoc network", International Journal of Distributed Sensor Networks Vol. 12, pp. 155014771667125, DOI 10.1177/1550147716671255, October 2016, <<https://doi.org/10.1177/1550147716671255>>.
- [MAROJEVIC20] Marojevic, V., Guvenc, I., Dutta, R., Sichitiu, M., and B. Floyd, "Advanced Wireless for Unmanned Aerial Systems: 5G Standardization, Research Challenges, and AERPAAW Architecture", IEEE Vehicular Technology Magazine Vol. 15, pp. 22-30, DOI 10.1109/mvt.2020.2979494, June 2020, <<https://doi.org/10.1109/mvt.2020.2979494>>.
- [OCADO] "Ocado Technologys robot warehouse a Hive of IoT innovation", n.d., <<https://techmonitor.ai/tech-leaders/ocado-technology-robot-hive-innovation>>.
- [PANRG] "Path Aware Networking Research Group - PANRG", n.d., <<https://datatracker.ietf.org/rg/panrg/about/>>.
- [PEARG] "Privacy Enhancements and Assessments Research Group - PEARG", n.d., <<https://irtf.org/pearg>>.
- [PILA] Krahenbuhl, C., Legner, M., Bitterli, S., and A. Perrig, "Pervasive Internet-Wide Low-Latency Authentication", 2021 International Conference on Computer Communications and Networks (ICCCN), DOI 10.1109/icccn52240.2021.9522235, July 2021, <<https://doi.org/10.1109/icccn52240.2021.9522235>>.

- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <<https://www.rfc-editor.org/info/rfc2775>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC5061] Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M. Kozuka, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", RFC 5061, DOI 10.17487/RFC5061, September 2007, <<https://www.rfc-editor.org/info/rfc5061>>.
- [RFC5177] Leung, K., Dommety, G., Narayanan, V., and A. Petrescu, "Network Mobility (NEMO) Extensions for Mobile IPv4", RFC 5177, DOI 10.17487/RFC5177, April 2008, <<https://www.rfc-editor.org/info/rfc5177>>.
- [RFC5275] Turner, S., "CMS Symmetric Key Management and Distribution", RFC 5275, DOI 10.17487/RFC5275, June 2008, <<https://www.rfc-editor.org/info/rfc5275>>.
- [RFC5517] HomChaudhuri, S. and M. Foschiano, "Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment", RFC 5517, DOI 10.17487/RFC5517, February 2010, <<https://www.rfc-editor.org/info/rfc5517>>.
- [RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, DOI 10.17487/RFC5944, November 2010, <<https://www.rfc-editor.org/info/rfc5944>>.

- [RFC6158] DeKok, A., Ed. and G. Weber, "RADIUS Design Guidelines", BCP 158, RFC 6158, DOI 10.17487/RFC6158, March 2011, <<https://www.rfc-editor.org/info/rfc6158>>.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, DOI 10.17487/RFC6182, March 2011, <<https://www.rfc-editor.org/info/rfc6182>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC6626] Tsirtsis, G., Park, V., Narayanan, V., and K. Leung, "Dynamic Prefix Allocation for Network Mobility for Mobile IPv4 (NEMOv4)", RFC 6626, DOI 10.17487/RFC6626, May 2012, <<https://www.rfc-editor.org/info/rfc6626>>.
- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.
- [RFC7401] Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <<https://www.rfc-editor.org/info/rfc7401>>.
- [RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S., Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, DOI 10.17487/RFC7426, January 2015, <<https://www.rfc-editor.org/info/rfc7426>>.
- [RFC7429] Liu, D., Ed., Zuniga, JC., Ed., Seite, P., Chan, H., and CJ. Bernardos, "Distributed Mobility Management: Current Practices and Gap Analysis", RFC 7429, DOI 10.17487/RFC7429, January 2015, <<https://www.rfc-editor.org/info/rfc7429>>.
- [RFC7476] Pentikousis, K., Ed., Ohlman, B., Corujo, D., Boggia, G., Tyson, G., Davies, E., Molinaro, A., and S. Eum, "Information-Centric Networking: Baseline Scenarios", RFC 7476, DOI 10.17487/RFC7476, March 2015, <<https://www.rfc-editor.org/info/rfc7476>>.

- [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>>.
- [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>>.
- [RFC8280] ten Oever, N. and C. Cath, "Research into Human Rights Protocol Considerations", RFC 8280, DOI 10.17487/RFC8280, October 2017, <<https://www.rfc-editor.org/info/rfc8280>>.
- [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>>.
- [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>>.
- [RFC8677] Trossen, D., Purkayastha, D., and A. Rahman, "Name-Based Service Function Forwarder (nSFF) Component within a Service Function Chaining (SFC) Framework", RFC 8677, DOI 10.17487/RFC8677, November 2019, <<https://www.rfc-editor.org/info/rfc8677>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8763] Rahman, A., Trossen, D., Kutscher, D., and R. Ravindran, "Deployment Considerations for Information-Centric Networking (ICN)", RFC 8763, DOI 10.17487/RFC8763, April 2020, <<https://www.rfc-editor.org/info/rfc8763>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.

- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.
- [SFCANYCAST] Wion, A., Bouet, M., Iannone, L., and V. Conan, "Distributed Function Chaining with Anycast Routing", Proceedings of the 2019 ACM Symposium on SDN Research, DOI 10.1145/3314148.3314355, April 2019, <<https://doi.org/10.1145/3314148.3314355>>.
- [TERASTREAM] "Deutsche Telekom tests TeraStream, the network of the future, in Croatia", n.d., <<https://www.telekom.com/en/media/media-information/archive/deutsche-telekom-tests-terastream-the-network-of-the-future-in-croatia-358444>>.
- [TROSSEN] Trossen, D., Sarela, M., and K. Sollins, "Arguments for an information-centric internetworking architecture", ACM SIGCOMM Computer Communication Review Vol. 40, pp. 26-33, DOI 10.1145/1764873.1764878, April 2010, <<https://doi.org/10.1145/1764873.1764878>>.
- [WANG19] Wang, P., Zhang, J., Zhang, X., Yan, Z., Evans, B., and W. Wang, "Convergence of Satellite and Terrestrial Networks: A Comprehensive Survey", IEEE Access Vol. 8, pp. 5550-5588, DOI 10.1109/access.2019.2963223, 2020, <<https://doi.org/10.1109/access.2019.2963223>>.

Acknowledgments

Thanks to all the people that shared insightful comments both privately to the authors as well as on various mailing list, especially on the INTArea Mailing List. Also thanks for the interesting discussions to Stewart Bryant, Ron Bonica, Toerless Eckert, Brian E. Carpenter, Kiran Makhijani, Fred Templin.

Authors' Addresses

Yihao Jia
Huawei Technologies Co., Ltd
156 Beiqing Rd.
Beijing
100095
P.R. China
Email: jiayihao@huawei.com

Dirk Trossen
Huawei Technologies Duesseldorf GmbH
Riesstr. 25C
80992 Munich
Germany
Email: dirk.trossen@huawei.com

Luigi Iannone
Huawei Technologies France S.A.S.U.
18, Quai du Point du Jour
92100 Boulogne-Billancourt
France
Email: luigi.iannone@huawei.com

Nirmala Shenoy
Rochester Institute of Technology
New-York, 14623
United States of America
Email: nxsvks@rit.edu

Paulo Mendes
Airbus
Willy-Messerschmitt Strasse 1
81663 Munich
Germany
Email: paulo.mendes@airbus.com

Donald E. Eastlake 3rd
Futurewei Technologies
2386 Panoramic Circle
Apopka, FL, 32703
United States of America
Email: d3e3e3@gmail.com

Peng Liu
China Mobile
32 Xuanwumen West Ave
Xicheng, Beijing
100053
P.R. China
Email: liupengygy@chinamobile.com

Dino Farinacci
lispers.net
United States of America
Email: farinacci@gmail.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 8 September 2022

S. Peng
Z. Li
Huawei Technologies
G. Mishra
Verizon Inc.
7 March 2022

APN Scope and Gap Analysis
draft-peng-apn-scope-gap-analysis-04

Abstract

The APN work in IETF is focused on developing a framework and set of mechanisms to derive, convey and use an attribute allowing the implementation of fine-grain user group-level and application group-level requirements in the network layer. APN aims to apply various policies in different nodes along a network path onto a traffic flow altogether, for example, at the headend to steer into corresponding path, at the midpoint to collect corresponding performance measurement data, and at the service function to execute particular policies. Currently there is still no way to efficiently realize this composite network service provisioning along the path. This document further clarifies the scope of the APN work and describes the solution gap analysis.

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 8 September 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	2
2. Requirements Language	3
3. Terminologies	3
4. APN Framework and Scope	3
5. Example Use Case and Existing Issues	4
6. Basic Solution and Benefits	5
7. Solution Gap Analysis	7
7.1. IPv6/MPLS Flow Label	7
7.2. SFC ServiceID	7
7.3. IOAM Flow ID	8
7.4. Binding SID	9
7.5. FlowSpec Label	9
7.6. Group Policy ID	9
7.7. Detnet Flow Identification	9
7.8. Network Slicing Resource ID	10
7.9. Service Path ID	10
7.10. Summary	10
8. IANA Considerations	11
9. Acknowledgements	11
10. Informative References	11
Authors' Addresses	15

1. Introduction

Application-aware Networking (APN) is introduced in [I-D.li-apn-framework] and [I-D.li-apn-problem-statement-usecases]. APN conveys an attribute along with data packets into network and makes the network aware about data flow requirements at different granularity levels.

Such an attribute is acquired, constructed in a structured value, and then encapsulated in the packet. Such structured value is treated as an opaque object in the network to which the network operator applies policies in various nodes/service functions along the path and provides corresponding services.

This structured attribute can be encapsulated in various data planes adopted within a Network Operator controlled limited domain, e.g. MPLS, VXLAN, SR/SRv6 and other tunnel technologies, which waits to be further specified.

With APN, it becomes possible to apply various policies in different nodes along a network path onto a traffic flow altogether in a more efficient way, e.g., at the headend to steer into corresponding path, at the midpoint to collect corresponding performance measurement data, and at the service function to execute particular policies. Currently there is still no way to realize this composite network service provisioning along the path very efficiently. It may be possible to stack those various policies in a list of TLVs at the headend. However, this approach would introduce great complexities and impose big challenges on the hardware processing and forwarding.

The example use-case presented in this draft further expands on the rationale for such an attribute and how it can be derived and used in that specific context.

This document further clarifies the scope of the APN work and describes the solution gap analysis.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 RFC 2119 [RFC2119] RFC 8174 [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminologies

APN: Application-aware Networking

CPE: Customer Premises Equipment

DPI: Deep Packet Inspection

OS: Operating System

4. APN Framework and Scope

The APN framework is introduced in [I-D.li-apn-framework], as shown in the Figure 1.




Figure 1. APN Framework and Scope

APN is only applied to an edge-to-edge tunnel encapsulation within a limited trusted domain. It means that the source and destination addresses of the packet are the endpoints of the tunnel (i.e. the domain edges), and nothing about the payload source and destination can be deduced, which substantially reduces the privacy concerns. Typically, an APN domain is defined as a Network Operator controlled limited domain (see Figure 1), in which MPLS, VXLAN, SR/SRv6 and other tunnel technologies are adopted to provide network services.

With APN, the attribute is acquired based on the existing information in the packet header (i.e. source and destination addresses, incoming L2 (or) MPLS encapsulation, incoming physical/virtual port information, the other fields of the 5-tuple if they are not encrypted) at the edge devices of the APN domain, added to the data packets along with the tunnel encapsulation, and delivered to the network, wherein, according to this attribute, corresponding network services are provisioned. When the packets leave the APN domain, the attribute is removed together with the tunnel encapsulation header.

5. Example Use Case and Existing Issues

To be more specific and more concrete, here we use SD-WAN as an example use case to further expand on the rationale for such attribute and how it can be derived and used in that specific context.

In the case of SD-WAN, an enterprise obtains WAN services from an SD-WAN provider so that its employees have access to the applications in the Cloud, and then the SD-WAN provider may buy WAN lines from a Network Operator. The enterprise may know what applications will use the SD-WAN services, but it will only provide the 5 tuples (i.e. source IP address, source port, destination IP address, destination port, transport protocol) of those applications to the SD-WAN

provider. So, the SD-WAN provider does not know what applications it is serving, and will only provide 5 tuples to the Network Operator and the service performance requirements for steering their customer's traffic. In this way, the Network Operator does not know anything else about the traffic except the 5 tuples and requirements. Nowadays, SD-WAN is usually using 5-tuple to steer the traffic into corresponding WAN lines across the Network Operator's network [SD-WAN].

However, there are two main issues in the current SD-WAN deployments.

1) It is complicated to resolve the 5 tuples. Even worse, as the traffic is encrypted, it becomes impossible to obtain any transport layer information. Moreover, in the IPv6 data plane, with the extension headers being added before the upper layer, in some implementations it becomes very difficult and even impossible to obtain transport layer information because that information is located deep in the packet. So, there is no 5 tuples anymore, and maybe only 2 tuples are available.

2) Currently there is still no way to apply various policies in different nodes along the network path onto a traffic flow altogether, that is, at the headend to steer into corresponding path, at the midpoint to collect corresponding performance measurement data, and at the service function to execute particular policies. It may be possible to stack those various policies in a list of TLVs at the headend. However, this approach would introduce great complexities and impose big challenges on the hardware processing and forwarding.

6. Basic Solution and Benefits

With APN, at the edge node, i.e. CPE, of the SD-WAN (see Figure 2), the 5-tuple, plus information related to user or application group-level requirements is constructed into a structured value, called APN attribute. This attribute is only meaningful for the network operators to apply various policies in different nodes/service functions, which can be enforced from the Controllers.

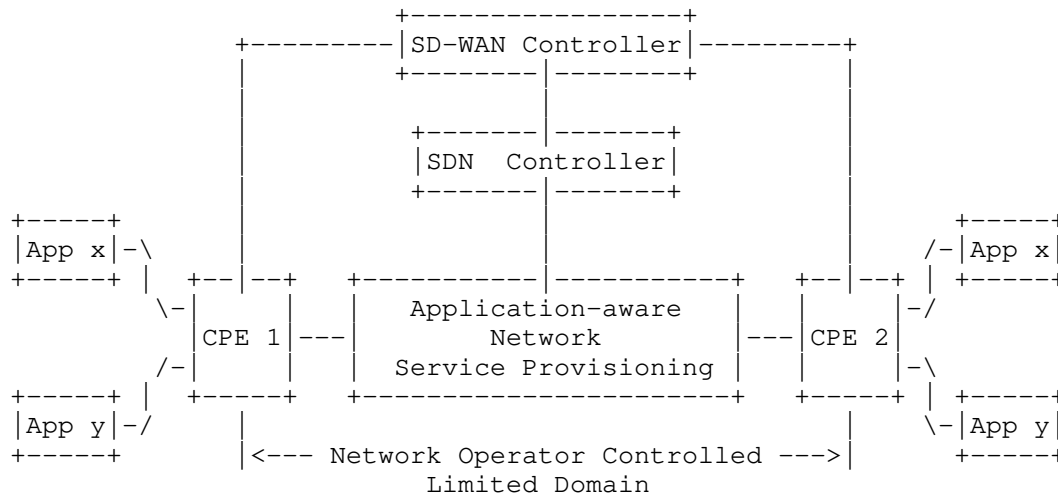


Figure 2. SD-WAN using the APN Framework

With such an attribute in the network, we can easily solve the two issues above-mentioned. For example, when the packet is sent from the CPE1 and the attribute is added along with the tunnel encapsulation, then it is not necessary to resolve the 5-tuple and perform the deep inspection in every node along the path. This attribute is encapsulated in the network layer and can be easily read by the routers and service functions. If the tunnel is based on the IPv6 data plane, for example, such an attribute can be encapsulated in an option of IPv6 hop-by-hop options header.

Since this attribute is taken as an object to the network, the network operators will simply place the policies in the nodes/service functions where this indicated traffic will go through, and the corresponding node/service function will just apply policies for this object. This can be easily done by utilizing this attribute, which is not possible with any current existing mechanism.

Such attribute will also bring other benefits, for example,

- * Improve the forwarding performance since it will only use 1 field in the IP layer instead of resolving 5 tuples, which will also improve the scalability.
- * Very flexible policy enforcement in various nodes and service functions along the network path.

Furthermore, with such attribute, more new services could be enabled, for example,

- * Even more fine-granularity performance measurement could be achieved and the granularity to be monitored and visualized can be controllable, which is able to relieve the processing pressure on the controller when it is facing the massive monitoring data.
- * The policy execution on the service function can be based only on this value and not based on 5-tuple, which can eliminate the need of deep packet inspection.
- * The underlay performance guarantee could be achieved for SD-WAN overlay services, such as explicit traffic engineering path satisfying SLA and selective visualized accurate performance measurement.

7. Solution Gap Analysis

There are already some solutions specified in IETF, which use identifier to perform traffic steering and service provisioning. However, the existing solutions are specific to a particular scenario or data plane. None of them is the same as APN and able to achieve the same effects.

7.1. IPv6/MPLS Flow Label

[RFC6437] specifies the IPv6 flow label which enables the IPv6 flow classification. However, the IPv6 flow label is mainly used for Equal Cost Multipath Routing (ECMP) and Link Aggregation [RFC6438].

Similarly, [RFC6391] describes a method of adding an additional Label Stack Entry (LSE) at the bottom of the stack in order to facilitate the load balancing of the flows within a pseudowire (PW) over the available ECMPs. A similar design for general MPLS use has also been proposed in [RFC6790] using the concept of Entropy Label.

7.2. SFC ServiceID

Subscriber Identifier and Performance Policy Identifier are specified in [RFC8979]. These identifiers are carried only in the Network Service Header (NSH) [RFC8300] Context Header, as shown in Figure 3, while the APN attribute can be carried in various data plane encapsulations.

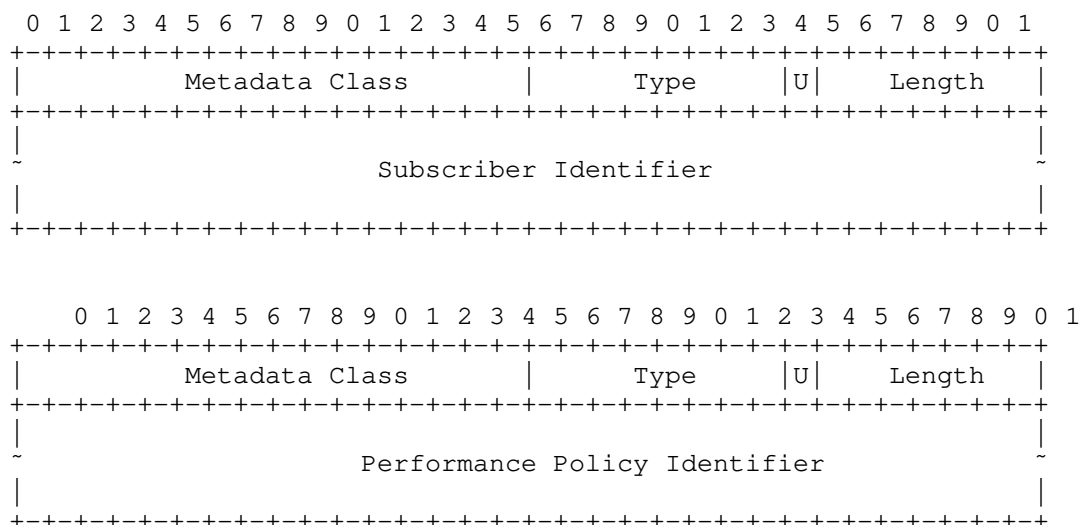


Figure 3. Subscriber Identifier and Performance Policy Identifier

In this draft [RFC8979], the Subscriber Identifier carries an opaque local identifier that is assigned to a subscriber by a network operator, and the Performance Policy Identifier represents an opaque value pointing to specific performance policy to be enforced. In this way, in order to apply various policies in different nodes along the network path onto a traffic flow altogether, e.g., at the headend to steer into corresponding path, at the midpoint to collect corresponding performance measurement data, and at the service function to execute particular policies, those various policies would have to be stacked in a list of TLVs at the headend, introducing great complexities and big challenges on the hardware processing and forwarding.

The APN attribute is treated as an opaque object in the network, to which the network operator applies policies in various nodes/service functions along the path and provide corresponding services.

7.3. IOAM Flow ID

A 32-bit Flow ID is specified in [I-D.ietf-ippm-ioam-direct-export], which is used to correlate the exported data of the same flow from multiple nodes and from multiple packets, while the APN attribute can serve more various purposes.

7.4. Binding SID

The Binding SID (BSID) [RFC8402] is bound to an SR Policy, instantiation of which may involve a list of SIDs. Any packets received with an active segment equal to BSID are steered onto the bound SR Policy. A BSID may be either a local or a global SID. While the APN attribute is not bound to SR only, and it can be carried in various data plane encapsulations.

7.5. FlowSpec Label

The flow specification (FlowSpec) [RFC5575] is actually an n-tuple consisting of several matching criteria that can be applied to IP traffic, which include elements such as source and destination address prefixes, IP protocol, and transport protocol port numbers. In BGP VPN/MPLS networks, BGP FlowSpec can be extended to identify and change (push/swap/pop) the label(s) for traffic that matches a particular FlowSpec rule in [I-D.ietf-idr-flowspec-mpls-match] and [I-D.ietf-idr-bgp-flowspec-label]. In [I-D.liang-idr-bgp-flowspec-route], BGP is used to distribute the FlowSpec rule bound with label(s). While the APN attribute is not bound to MPLS only, and it can be carried in various data plane encapsulations.

7.6. Group Policy ID

The capabilities of the VXLAN-GPE protocol can be extended by defining next protocol "shim" headers that are used to implement new data plane functions. For example, Group Policy ID is carried in the Group-Based Policy (GBP) Shim header [I-D.lemon-vxlan-lisp-gpe-gbp]. GENEVE has similar ability as VXLAN-GPE to carry metadata.

7.7. Detnet Flow Identification

Identification and Specification of DetNet Flows is specified in [RFC9016]. DetNet MPLS flows can be identified and specified by the SLabel and the FLabelStack. The IP 6-tuple is used for DetNet IP flow identification, which consists of SourceIpAddress, DestinationIpAddress, Dscp, Protocol, SourcePort, and DestinationPort. IPv6FlowLabel and IPsecSpi are additional attributes that can be used for DetNet flow identification in addition to the 6-tuple. Therefore, the Detnet IP Flow ID is logical and there is no such Flow ID carried for Detnet, but only the 6-tuple is directly used to identify the Detnet flows.

Only one exceptional case, in [I-D.ietf-spring-sr-redundancy-protection], the 32-bit flow identification (FID) identifies one specific Detnet flow of

redundancy protection. This FID is usually allocated from centralized controller to the SR ingress node or redundancy node in SR network.

7.8. Network Slicing Resource ID

In [I-D.dong-6man-enhanced-vpn-vtn-id], VTN Resource ID is a 4-octet identifier which uniquely identifies the set of network resources allocated to a VTN. For network slicing, the ID is used to indicate the network resources to be allocated to the network slices and it is not bound to any traffic flow.

APN is for traffic steering, while network slicing is about resource partition [I-D.ietf-teas-rfc3272bis].

7.9. Service Path ID

In [RFC8300], Service Path Identifier (SPI) uniquely identifies a Service Function Path (SFP). Participating nodes MUST use this identifier for SFP selection. The initial Classifier MUST set the appropriate SPI for a given classification result. For SFC, the ID is used to indicate a SF path and it is not bound to any traffic flow.

7.10. Summary

The comparison of the identifiers for the typical network services (incl. iOAM, Detnet, Network Slicing (NS), and Service Function Chaining (SFC)) is shown in the following Table from different aspects (incl. ID, Identification Object, Source (for generating the ID), Configuration (Conf.) node, and Size).

	ID	Identification Object	Source	Conf. node	Size
APN	APN ID	The flow that needs fine-granular services	5-tuple Layer 2	Controller	32bits 128b
iOAM	Flow ID	The flow that needs performance monitoring	-	Controller Ingress	32bits
Detnet	Flow ID (6-tuple)	The flow that needs Detnet services	-	Controller	-
Detnet	Flow ID	The redundant protection flow	-	Detnet Controller	32bits
NS	Resource ID	The network resources that are allocated to network slices	-	Controller	32bits
SFC	SPI	The SF Path	-	Controller	24bits
SFC	Performance Policy ID	The performance policy	-	Controller	-

Table 1. Comparison of the Identifiers

As driven by ever-emerging new 5G services, fine-granularity service provisioning becomes urgent. The existing solutions are either specific to a particular scenario or data plane. While APN aims to define a generalized attribute used for fine-granularity service provisioning, and can be carried in various data plane encapsulations.

8. IANA Considerations

There are no IANA considerations in this document.

9. Acknowledgements

The authors would like to acknowledge Martin Vigoureux, Alvaro Retana, Barry Leiba, Stefano Previdi, Adrian Farrel, and Daniel King for their valuable review and comments.

10. Informative References

[I-D.brockners-ippm-ioam-vxlan-gpe]

Brockners, F., Bhandari, S., Govindan, V. P., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "VXLAN-GPE Encapsulation for In-situ OAM Data", Work in Progress, Internet-Draft, draft-brockners-ippm-ioam-vxlan-gpe-03, 4 November 2019, <<https://www.ietf.org/archive/id/draft-brockners-ippm-ioam-vxlan-gpe-03.txt>>.

[I-D.dong-6man-enhanced-vpn-vtn-id]

Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network (VTN) Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-dong-6man-enhanced-vpn-vtn-id-06, 24 October 2021, <<https://www.ietf.org/archive/id/draft-dong-6man-enhanced-vpn-vtn-id-06.txt>>.

[I-D.ietf-idr-bgp-flowspec-label]

Liang, Q., Hares, S., You, J., Raszuk, R., and D. Ma, "Carrying Label Information for BGP FlowSpec", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-flowspec-label-01, 6 December 2016, <<https://www.ietf.org/archive/id/draft-ietf-idr-bgp-flowspec-label-01.txt>>.

[I-D.ietf-idr-flowspec-mpls-match]

Yong, L., Hares, S., Liang, Q., and J. You, "BGP Flow Specification Filter for MPLS Label", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-mpls-match-01, 6 December 2016, <<https://www.ietf.org/archive/id/draft-ietf-idr-flowspec-mpls-match-01.txt>>.

[I-D.ietf-ippm-ioam-direct-export]

Song, H., Gafni, B., Zhou, T., Li, Z., Brockners, F., Bhandari, S., Sivakolundu, R., and T. Mizrahi, "In-situ OAM Direct Exporting", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-direct-export-07, 13 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-direct-export-07.txt>>.

[I-D.ietf-sfc-serviceid-header]

Sarikaya, B., Hugo, D. V., and M. Boucadair, "Subscriber and Performance Policy Identifier Context Headers in the Network Service Header (NSH)", Work in Progress, Internet-Draft, draft-ietf-sfc-serviceid-header-14, 11 December 2020, <<https://www.ietf.org/archive/id/draft-ietf-sfc-serviceid-header-14.txt>>.

- [I-D.ietf-spring-sr-redundancy-protection]
Geng, X., Chen, M., Yang, F., Garvia, P. C., and G. Mishra, "SRv6 for Redundancy Protection", Work in Progress, Internet-Draft, draft-ietf-spring-sr-redundancy-protection-01, 15 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-redundancy-protection-01.txt>>.
- [I-D.ietf-teas-rfc3272bis]
Farrel, A., "Overview and Principles of Internet Traffic Engineering", Work in Progress, Internet-Draft, draft-ietf-teas-rfc3272bis-15, 24 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-rfc3272bis-15.txt>>.
- [I-D.lemon-vxlan-lisp-gpe-gbp]
Lemon, J., Maino, F., Smith, M., and A. Isaac, "Group Policy Encoding with VXLAN-GPE and LISP-GPE", Work in Progress, Internet-Draft, draft-lemon-vxlan-lisp-gpe-gbp-02, 30 April 2019, <<https://www.ietf.org/archive/id/draft-lemon-vxlan-lisp-gpe-gbp-02.txt>>.
- [I-D.li-6man-app-aware-ipv6-network]
Li, Z., Peng, S., Li, C., Xie, C., Voyer, D., Li, X., Liu, P., Cao, C., and K. Ebisawa, "Application-aware IPv6 Networking (APN6) Encapsulation", Work in Progress, Internet-Draft, draft-li-6man-app-aware-ipv6-network-03, 22 February 2021, <<https://www.ietf.org/archive/id/draft-li-6man-app-aware-ipv6-network-03.txt>>.
- [I-D.li-apn-framework]
Li, Z., Peng, S., Voyer, D., Li, C., Liu, P., Cao, C., Mishra, G., Ebisawa, K., Previdi, S., and J. N. Guichard, "Application-aware Networking (APN) Framework", Work in Progress, Internet-Draft, draft-li-apn-framework-04, 25 October 2021, <<https://www.ietf.org/archive/id/draft-li-apn-framework-04.txt>>.
- [I-D.li-apn-problem-statement-usecases]
Li, Z., Peng, S., Voyer, D., Xie, C., Liu, P., Qin, Z., Mishra, G., Ebisawa, K., Previdi, S., and J. N. Guichard, "Problem Statement and Use Cases of Application-aware Networking (APN)", Work in Progress, Internet-Draft, draft-li-apn-problem-statement-usecases-05, 20 December 2021, <<https://www.ietf.org/archive/id/draft-li-apn-problem-statement-usecases-05.txt>>.

- [I-D.liang-idr-bgp-flowspec-route]
Liang, Q. and J. You, "BGP FlowSpec based Multi-dimensional Route Distribution", Work in Progress, Internet-Draft, draft-liang-idr-bgp-flowspec-route-00, 20 October 2014, <<https://www.ietf.org/archive/id/draft-liang-idr-bgp-flowspec-route-00.txt>>.
- [I-D.peng-apn-security-privacy-consideration]
Peng, S., Li, Z., Voyer, D., Li, C., Liu, P., and C. Cao, "APN Security and Privacy Considerations", Work in Progress, Internet-Draft, draft-peng-apn-security-privacy-consideration-02, 16 June 2021, <<https://www.ietf.org/archive/id/draft-peng-apn-security-privacy-consideration-02.txt>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC6391] Bryant, S., Ed., Filsfils, C., Drafz, U., Kompella, V., Regan, J., and S. Amante, "Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network", RFC 6391, DOI 10.17487/RFC6391, November 2011, <<https://www.rfc-editor.org/info/rfc6391>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [RFC8979] Sarikaya, B., von Hugo, D., and M. Boucadair, "Subscriber and Performance Policy Identifier Context Headers in the Network Service Header (NSH)", RFC 8979, DOI 10.17487/RFC8979, February 2021, <<https://www.rfc-editor.org/info/rfc8979>>.
- [RFC9016] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 2021, <<https://www.rfc-editor.org/info/rfc9016>>.
- [SD-WAN] MEF 70.1 Draft (R1), available at <https://www.mef.net/wp-content/uploads/2020/08/MEF-70-1-Draft-R1.pdf>/, "SD-WAN Service Attributes and Service Framework", August 2020.

Authors' Addresses

Shuping Peng
Huawei Technologies
Beijing
China
Email: pengshuping@huawei.com

Zhenbin Li
Huawei Technologies
Beijing
China
Email: lizhenbin@huawei.com

Gyan Mishra
Verizon Inc.
United States of America
Email: gyan.s.mishra@verizon.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 19 May 2021

L. Colitti
Google
T. Pauly
Apple Inc.
15 November 2020

Per-Application Networking Considerations
draft-per-app-networking-considerations-00

Abstract

This document describes considerations for and implications of using application identifiers as a method of differentiating traffic on networks. Specifically, it discusses privacy considerations, possible mitigations, and considerations for user experience and API design.

Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at <https://github.com/tfpauly/per-app-networking-considerations>.

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 19 May 2021.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Conventions and Definitions	3
2. Requesting differential treatment	3
3. Open Internet implications	4
4. Privacy implications	4
5. Mitigating implications via traffic categories	5
6. User experience considerations	6
7. API considerations	6
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Acknowledgments	7
Authors' Addresses	7

1. Introduction

There are a number of use cases where network operators, or applications, might desire for application traffic to be treated differently by the network. Some examples are:

- * Network-specific services. Applications might want to access local resources on a network that does not otherwise provide Internet access (for example, an entertainment system on an airplane).
- * Per-application private networks. Certain applications, such as enterprise applications, might want to connect directly to the enterprise network in a secure fashion without using a device-wide VPN.
- * Mobile network services. In mobile networks, applications like voice over LTE, IMS and RCS often use a different virtual network than general Internet traffic.

- * Applications with specific performance requirements. Certain applications would benefit from particular scheduling or QoS policies - for example applications requiring low latency such as voice might be scheduled and queued differently from latency-insensitive traffic.
- * Local breakout. In a mobile networks, applications might want to access resources through a different network interface (e.g., one that uses IPv6 addresses that are local to a specific area, and do not have a wide mobility).
- * Zero-rating traffic. As allowed by regulators, certain classes of traffic (e.g., messaging or streaming video) might be exempt from metering on networks that are otherwise metered.

In existing networks, this is sometimes implemented by the network using deep packet inspection (e.g., flow tracking coupled with inspection of the SNI handshake). This is complex, implicates public policy concerns, and generally conflicts with the recommendations in [RFC7258]. The move towards encrypted protocols such as [RFC8484] and [I-D.ietf-tls-esni] will make this more difficult for some operators. Thus, if an application is to receive different treatment, the host or the application itself should be involved in requesting specific network treatment. This document explores the implications.

In this document, the term "application" refers to an application as understood by the user of the device.

1.1. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Requesting differential treatment

There are already mechanisms for applications to request and obtain particular treatment by the network, or to communicate application identity to the network in order to obtain particular treatment. These include:

- * Diffserv
- * APN6

- * Network tokens
- * Slicing in 3GPP 5G networks
- * Explicit application selection of a given Provisioning Domain (PvD) [RFC8801]

3. Open Internet implications

In certain regulatory regions, networks that provide general Internet access may not be permitted to discriminate between traffic sent to or from different lawful applications or websites, or such discrimination may be prohibited if commercially based. In a situation where the network operator has influence on the implementation of the user host (e.g., mobile networks where the handset is sold by the carrier), the device may be able to implement network policies directly, and thus may be impacted by neutrality considerations.

Neutrality concerns can be addressed by providing user control over assignment of particular applications to the particular network resources available to that user. Further, network neutrality implications may be reduced or avoided in some jurisdictions if the differential treatment occurs between different classes of traffic with different network requirements (e.g., bandwidth-intensive traffic vs. low-latency traffic) as opposed to between different applications with similar network requirements, and thus, by ensuring that the mechanism used to communicate requests to the network only specifies traffic classes and not individual applications.

4. Privacy implications

IETF guidance to avoid pervasive monitoring [RFC7258] is for network protocols to expose as little information as possible. Some of the proposed technologies for application signalling rely on the application exposing its identity to the network so that the network can then implement appropriate policies. This may provide the network with much more information than is needed to implement the desired behaviour. Information about which users are using specific applications, or visiting certain destinations, and when, can be highly privacy-sensitive.

Note that application identity can be exposed to the network even in the absence of explicit signalling. For example, if the host were to implement a network-set policy that requires that traffic from application X be sent on a different network path than all other traffic, the identity of application X would be exposed to the network as soon as it sends traffic.

Privacy concerns may also be reduced or avoided if the mechanism to request a different class of service only specifies the class of service (e.g., "low latency" or "streaming video") instead of the application originating the traffic.

In a situation where the network operator has influence over the implementation of the user host, the operator can still impose policies on what requests are possible - for example, the operator might choose to limit access to specialized services such as carrier messaging only to carrier applications. It is possible for such policies to preserve privacy if the policies specify general categories of traffic as opposed to specifying applications.

5. Mitigating implications via traffic categories

Many of these implications can be mitigated if the mechanism does not request different treatment of a service for a particular application, but instead specifies a general category of traffic, especially one that is defined based on traffic properties rather than commercial agreements.

Categories of traffic need to be sufficiently broad to not identify individual applications, and should be general enough that details about a user cannot be inferred merely by use of the category.

Consider the example a network that wants to provide differentiated service for a role-playing game application that can take advantage of a low-latency path. Several levels of categories could be defined. The following list shows some examples, in order of decreasing specificity:

1. Role-playing game
2. Game
3. Real-time/low-latency

The first category would not be an appropriate choice due to the privacy implications of identifying what kind of game a user plays. The second category is preferable, but the third is best since it defines a way to manage the network traffic without identifying anything about the content of the application.

Some use cases for traffic differentiation might need other kinds of categories. For example, operators might wish to zero-rate applications using categories based on payment tiers and rate-limiting.

6. User experience considerations

Privacy and neutrality concerns can be mitigated if the host's user is informed that particular applications are seeking or designated for particular treatment and consents to it. In order for consent to be meaningful, the user should be presented with a message that they understand. It may be difficult to balance the goal of providing complete and accurate information with the goal of ensuring that the user understands the implications.

7. API considerations

It is desirable to provide an API layer that is not tied to specific network technologies (e.g., URSP, VPN, etc.). Having applications select a specific Provisioning Domain (PvD) could provide a useful layer of abstraction, as described in [I-D.ietf-taps-interface].

Any API should not involve revealing an application or user identity to the network via metadata without network authentication. Instead, the API should allow a given setting to be conditional on the identity of the network. For example, an application should express "use the zero-rated service for my app when on a particular carrier network", instead of blindly saying "this is my application identifier".

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

- [I-D.ietf-taps-interface] Trammell, B., Welzl, M., Enghardt, T., Fairhurst, G., Kuehlewind, M., Perkins, C., Tiesel, P., Wood, C., and T. Pauly, "An Abstract Application Layer Interface to Transport Services", Work in Progress, Internet-Draft, draft-ietf-taps-interface-10, 2 November 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-taps-interface-10.txt>>.

[I-D.ietf-tls-esni]

Rescorla, E., Oku, K., Sullivan, N., and C. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-08, 16 October 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-tls-esni-08.txt>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

[RFC8801] Pfister, P., Vyncke, É., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", RFC 8801, DOI 10.17487/RFC8801, July 2020, <<https://www.rfc-editor.org/info/rfc8801>>.

Acknowledgments

Thanks to Adi Masputra and Elliot Briggs for their inputs to this discussion.

Authors' Addresses

Lorenzo Colitti
Google
Shibuya 3-21-3,
Japan

Email: lorenzo@google.com

Tommy Pauly
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America

Email: tpauly@apple.com

LPWAN
Internet-Draft
Updates: 5172 (if approved)
Intended status: Standards Track
Expires: 23 October 2021

P. Thubert, Ed.
Cisco Systems
21 April 2021

SCHC over PPP
draft-thubert-intarea-schc-over-ppp-03

Abstract

This document extends RFC 5172 to signal the use of SCHC as the compression method between a pair of nodes over PPP. Combined with RFC 2516, this enables the use of SCHC over Ethernet and Wi-Fi.

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 23 October 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. BCP 14	3
3. Extending RFC 5172	3
4. Profiling SCHC for high speed links	4
4.1. Mapping the SCHC Architecture	4
4.2. SCHC Parameters	5
4.2.1. Resulting Packet Format	6
4.3. Security Considerations	8
5. IANA Considerations	8
6. Acknowledgments	9
7. Normative References	9
8. Informative References	9
Author's Address	10

1. Introduction

The Point-to-Point Protocol (PPP) [RFC5172] provides a standard method of encapsulating network-layer protocol information over serial (point-to-point and bus) links. "A Method for Transmitting PPP Over Ethernet (PPPoE)" [RFC2516] transports PPP over Ethernet between a pair of nodes. It is compatible with a translating bridge to Wi-Fi, and therefore enables PPP over Wi-Fi as well.

PPP also proposes an extensible Link Control Protocol and a family of Network Control Protocols (NCPs) for establishing and configuring different network-layer protocols. "IP Version 6 over PPP" [RFC5072] specifies the IPv6 Control Protocol (IPV6CP), which is an NCP for a PPP link, and allows for the negotiation of desirable parameters for an IPv6 interface over PPP. "Negotiation for IPv6 Datagram Compression Using IPv6 Control Protocol" [RFC5172] defines the IPv6 datagram compression option that can be negotiated by a node on the link through the IPV6CP.

PPP is not commonly used in Low-Power Wide Area Networks (LPWAN) but the extreme compression techniques that are defined for use in LPWAN may be applicable to more traditional links where PPP applies.

The "Static Context Header Compression (SCHC) and fragmentation for LPWAN, application to UDP/IPv6" [SCHC] is a new technology that can provide an extreme compression performance but requires a same state to be provisioned on both ends before it can be operated.

The "SCHC Architecture" [I-D.pelov-lpwan-architecture] enables a peer to peer SCHC operation in addition to the classical device to network LPWAN paradigm, e.g., over a PPP connection. To enable SCHC over PPP and therefore Ethernet and Wi-Fi, this specification extends [RFC5172] to signal SCHC as an additional compression method for use over PPP.

An example use case for SCHC over PPP over Ethernet (SCHCoPPPoE) is to apply SCHC to periodic flows and maintain them at a protocol-independent size and rate. The constant size may be too small for a particular flow or protocol. The SCHC fragmentation can then be used to transport a protocol data unit (PDU) as N compressed SCHC fragments, in which case the effective PDU rate is the TSN frame rate divided by N.

This can be useful to streamline the frames and simplifies the scheduling of Deterministic Networking [DetNet] and Operational Technology (OT) control flows over IEEE Std 802.1 Time-Sensitive Networking (TSN) [IEEE802.1TSNTG] or one of the RAW Technologies [RAW Technologies].

2. BCP 14

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Extending RFC 5172

With this specification, a PPP session defines a virtual link where a SCHC context is established with a particular set of Rules, which is indicated at the set up of the PPP session as follows:

[RFC5172] defines an IPV6CP option called the IPv6-Compression-Protocol Configuration option with a type of 2. The option contains an IPv6-Compression-Protocol field value that indicates a compression protocol and an optional data field as shown in Figure 1:




Figure 1: The IPv6-Compression-Protocol Configuration Option

This specification indicates a new IPv6-Compression-Protocol field value for [SCHC] (see Section 5), and enables to transport a Uniform Resource Identifier (URI) [RFC3986] of the set of rules in the optional data. The default format for the set of rules is YANG using the "Data Model for SCHC" [SCHC_DATA_MODEL] encoded in JSON as specified in [RFC7951]. The size of the URL is computed based on the Length of the option as Length-4. If the encoding is asymmetrical, the initiator of the session is considered downstream, playing the role of the device in an LPWAN network.

4. Profiling SCHC for high speed links

Appendix D of [SCHC] specifies the profile information that technology specifications such as this must provide. The following section address this requirement.

4.1. Mapping the SCHC Architecture

This specification leverages SCHC between an end point that is an IP Host and possibly a serial DTE (Data Terminal Equipment), and another that is an IP Node (either another IP Host or a Router) and possibly a serial DCE (Data Control Equipment), or a more modern physical or emulated endpoint, e.g., Ethernet devices that exchange IP packets over PPPoE.

Both endpoints MUST support the function of SCHC Compressor/Decompressor (C/D) as shown in Figure 2.




Figure 2: Typical Deployment

The SCHC Fragmenter/Reassembler (F/R) is generally not needed, because the maximum transmission unit (MTU) is expected to be large enough and SCHC only reduces the frame size vs. native IP. But it may be used to obtain a small protocol-independant frame size for the compressed packets, possibly way smaller than MTU.

A context may be generated for a particular upper layer application, such as a control loop using an industrial automation protocol, to protect the particular flow with a DetNet service. The context can be asymmetric, e.g., when connecting a primary and a secondary endpoints, a client and a server, or a programmable logic controller with a sensor or an actuator.

4.2. SCHC Parameters

Compared to typical LPWANs, most serial links and emulations such as PPPoE are very fast and most of the constraints can be alleviated. For this reason, the SCHC profile for PPP is defined as follows:

RuleID numbering scheme: The RuleID for a compression rule is expressed as 2 bytes. The first (leftmost) 2 bits of that RuleID MUST be set to 0 This leaves 14 bits to index the rule. A SCHC compressed packet is always in the form:




Figure 3: SCHC Compressed Packet

This specification only supports the No-ACK Mode of SCHC fragmentation as specified in section 8.4.1 of [SCHC]. The SCHC Fragment Header is 2 bytes long.

The RuleID for a fragmentation rule is expressed as 4 bits. The bits MUST all set to 1 for a fragmentation rule in No-ACK Mode. The DTag field is 11 bits long (T=11) and the FCN field is one bit (N=1), which is set to 1 on the last fragment as illustrated in Appendix B of [SCHC] and to 0 otherwise. There is no W field (M=0).




Figure 4: SCHC Fragment

The No-ACK mode has been designed under the assumption that data unit out-of-sequence delivery does not occur between the entity performing fragmentation and the entity performing reassembly and a DetNet PREOF function might be needed to reorder the fragments.

Maximum packet size: MAX_PACKET_SIZE is aligned to the PPP Link MTU.

Padding: The Compression Residue MUST be aligned to the L2 word.

For Ethernet, the L2 word is one byte, so padding is needed up to the next byte boundary. If a compression rule produces a residue that is not byte aligned, then it is implicitly terminated with a statement that indicates padding till the next byte boundary. The padding bit is 0.

4.2.1. Resulting Packet Format

In the case of PPPoE, the sequence of compression and encapsulation is as follows:

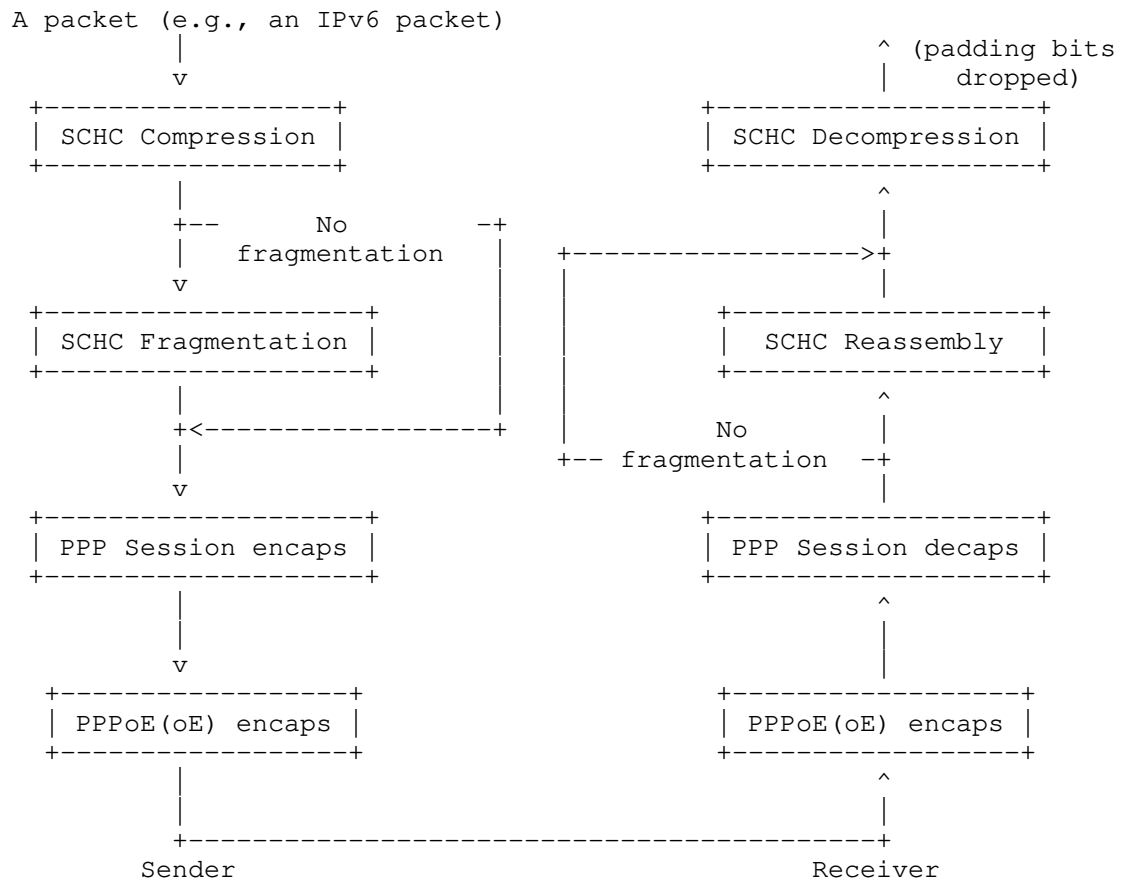


Figure 5: Stack Operation (no fragment)

In the case of PPPoE, a frame that transports an IPv6 packet compressed with SCHC with no fragmentation shows as follows:




Figure 6: SCHC over PPP over Ethernet Format

4.3. Security Considerations

This draft enables to use the SCHC compression and fragmentation over PPP and therefore Ethernet and Wi-Fi with PPPoE. It inherits the possible threats against SCHC listed in the "Security considerations" section of [SCHC].

5. IANA Considerations

This document requests the allocation of a new value in the registry "IPv6-Compression-Protocol Types" for "SCHC". A suggested value is proposed in Table 1:

Value	Description	Reference
4	Static Context Header Compression (SCHC)	This document

Table 1: IP Header Compression Configuration Option Suboption Types

6. Acknowledgments

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, DOI 10.17487/RFC2516, February 1999, <<https://www.rfc-editor.org/info/rfc2516>>.
- [RFC5072] Varada, S., Ed., Haskins, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, DOI 10.17487/RFC5072, September 2007, <<https://www.rfc-editor.org/info/rfc5072>>.
- [RFC5172] Varada, S., Ed., "Negotiation for IPv6 Datagram Compression Using IPv6 Control Protocol", RFC 5172, DOI 10.17487/RFC5172, March 2008, <<https://www.rfc-editor.org/info/rfc5172>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [SCHC] Minaburo, A., Toutain, L., Gomez, C., Barthel, D., and JC. Zúñiga, "SCHC: Generic Framework for Static Context Header Compression and Fragmentation", RFC 8724, DOI 10.17487/RFC8724, April 2020, <<https://www.rfc-editor.org/info/rfc8724>>.

8. Informative References

[I-D.pelov-lpwan-architecture]

Pelov, A., Thubert, P., and A. Minaburo, "Static Context Header Compression (SCHC) Architecture", Work in Progress, Internet-Draft, draft-pelov-lpwan-architecture-00, 19 January 2021, <<https://tools.ietf.org/html/draft-pelov-lpwan-architecture-00>>.

[SCHC_DATA_MODEL]

Minaburo, A. and L. Toutain, "Data Model for Static Context Header Compression (SCHC)", Work in Progress, Internet-Draft, draft-ietf-lpwan-schc-yang-data-model-03, 10 July 2020, <<https://tools.ietf.org/html/draft-ietf-lpwan-schc-yang-data-model-03>>.

[RAW Technologies]

Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-thubert-raw-technologies-05, 18 May 2020, <<https://tools.ietf.org/html/draft-thubert-raw-technologies-05>>.

[RFC7951] Lhotka, L., "JSON Encoding of Data Modeled with YANG", RFC 7951, DOI 10.17487/RFC7951, August 2016, <<https://www.rfc-editor.org/info/rfc7951>>.

[DetNet] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

[IEEE802.1TSNTG]

IEEE, "Time-Sensitive Networking (TSN) Task Group", <<https://1.ieee802.org/tsn/>>.

Author's Address

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allée des Ormes - BP1200
06254 Mougins - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com