

---

Stream: Independent Submission  
RFC: [8799](#)  
Category: Informational  
Published: July 2020  
ISSN: 2070-1721  
Authors: B. Carpenter      B. Liu  
          *Univ. of Auckland*    *Huawei Technologies*

# RFC 8799

## Limited Domains and Internet Protocols

---

### Abstract

There is a noticeable trend towards network behaviors and semantics that are specific to a particular set of requirements applied within a limited region of the Internet. Policies, default parameters, the options supported, the style of network management, and security requirements may vary between such limited regions. This document reviews examples of such limited domains (also known as controlled environments), notes emerging solutions, and includes a related taxonomy. It then briefly discusses the standardization of protocols for limited domains. Finally, it shows the need for a precise definition of "limited domain membership" and for mechanisms to allow nodes to join a domain securely and to find other members, including boundary nodes.

This document is the product of the research of the authors. It has been produced through discussions and consultation within the IETF but is not the product of IETF consensus.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8799>.

### 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.

## Table of Contents

1. Introduction
2. Failure Modes in Today's Internet
3. Examples of Limited Domain Requirements
4. Examples of Limited Domain Solutions
5. The Scope of Protocols in Limited Domains
6. Functional Requirements of Limited Domains
7. Security Considerations
8. IANA Considerations
9. Informative References

### Appendix A. Taxonomy of Limited Domains

- A.1. Domain as a Whole
- A.2. Individual Nodes
- A.3. Domain Boundary
- A.4. Topology
- A.5. Technology
- A.6. Connection to the Internet
- A.7. Security, Trust, and Privacy Model
- A.8. Operations
- A.9. Making Use of This Taxonomy

### Acknowledgements

### Contributors

### Authors' Addresses

## 1. Introduction

As the Internet continues to grow and diversify, with a realistic prospect of tens of billions of nodes being connected directly and indirectly, there is a noticeable trend towards network-specific and local requirements, behaviors, and semantics. The word "local" should be understood in a special sense, however. In some cases, it may refer to geographical and physical locality -- all the nodes in a single building, on a single campus, or in a given vehicle. In other cases, it may refer to a defined set of users or nodes distributed over a much wider area, but drawn together by a single virtual network over the Internet, or a single physical network running in parallel with the Internet. We expand on these possibilities below. To capture the topic, this document refers to such networks as "limited domains". Of course, a similar situation may arise for a network that is completely disconnected from the Internet, but that is not our direct concern here. However, it should not be forgotten that interoperability is needed even within a disconnected network.

Some people have concerns about splintering of the Internet along political or linguistic boundaries by mechanisms that block the free flow of information. That is not the topic of this document, which does not discuss filtering mechanisms (see [\[RFC7754\]](#)) and does not apply to protocols that are designed for use across the whole Internet. It is only concerned with domains that have specific technical requirements.

The word "domain" in this document does not refer to naming domains in the DNS, although in some cases, a limited domain might incidentally be congruent with a DNS domain. In particular, with a "split horizon" DNS configuration [\[RFC6950\]](#), the split might be at the edge of a limited domain. A recent proposal for defining definite perimeters within the DNS namespace [\[DNS-PERIMETER\]](#) might also be considered to be a limited domain mechanism.

Another term that has been used in some contexts is "controlled environment". For example, [\[RFC8085\]](#) uses this to delimit the operational scope within which a particular tunnel encapsulation might be used. A specific example is GRE-in-UDP encapsulation [\[RFC8086\]](#), which explicitly states that "The controlled environment has less restrictive requirements than the general Internet." For example, non-congestion-controlled traffic might be acceptable within the controlled environment. The same phrase has been used to delimit the useful scope of quality-of-service protocols [\[RFC6398\]](#). It is not necessarily the case that protocols will fail to operate outside the controlled environment, but rather that they might not operate optimally. In this document, we assume that "limited domain" and "controlled environment" mean the same thing in practice. The term "managed network" has been used in a similar way, e.g., [\[RFC6947\]](#). In the context of secure multicast, a "group domain of interpretation" is defined by [\[RFC6407\]](#).

Yet more definitions of types of domains are to be found in the routing area, such as [\[RFC4397\]](#), [\[RFC4427\]](#), and [\[RFC4655\]](#). We conclude that the notion of a limited domain is very widespread in many aspects of Internet technology.

The requirements of limited domains will depend on the deployment scenario. Policies, default parameters, and the options supported may vary. Also, the style of network management may vary between a completely unmanaged network, one with fully autonomic management, one with traditional central management, and mixtures of the above. Finally, the requirements and solutions for security and privacy may vary.

This document analyzes and discusses some of the consequences of this trend and how it may impact the idea of universal interoperability in the Internet. First, we list examples of limited domain scenarios and of technical solutions for limited domains, with the main focus being the Internet layer of the protocol stack. An appendix provides a taxonomy of the features to be found in limited domains. With this background, we discuss the resulting challenge to the idea that all Internet standards must be universal in scope and applicability. To the contrary, we assert that some protocols, although needing to be standardized and interoperable, also need to be specifically limited in their applicability. This implies that the concepts of a limited domain, and of its membership, need to be formalized and supported by secure mechanisms. While this document does not propose a design for such mechanisms, it does outline some functional requirements.

This document is the product of the research of the authors. It has been produced through discussions and consultation within the IETF but is not the product of IETF consensus.

## 2. Failure Modes in Today's Internet

Today, the Internet does not have a well-defined concept of limited domains. One result of this is that certain protocols and features fail on certain paths. Earlier analyses of this topic have focused either on the loss of transparency of the Internet [[RFC2775](#)] [[RFC4924](#)] or on the middleboxes responsible for that loss [[RFC3234](#)] [[RFC7663](#)] [[RFC8517](#)]. Unfortunately, the problems persist both in application protocols and even in very fundamental mechanisms. For example, the Internet is not transparent to IPv6 extension headers [[RFC7872](#)], and Path MTU Discovery has been unreliable for many years [[RFC2923](#)] [[RFC4821](#)]. IP fragmentation is also unreliable [[FRAG-FRAGILE](#)], and problems in TCP MSS negotiation have been reported [[IPV6-USE-MINMTU](#)].

On the security side, the widespread insertion of firewalls at domain boundaries that are perceived by humans but unknown to protocols results in arbitrary failure modes as far as the application layer is concerned. There are operational recommendations and practices that effectively guarantee arbitrary failures in realistic scenarios [[IPV6-EXT-HEADERS](#)].

Domain boundaries that are defined administratively (e.g., by address filtering rules in routers) are prone to leakage caused by human error, especially if the limited domain traffic appears otherwise normal to the boundary routers. In this case, the network operator needs to take active steps to protect the boundary. This form of leakage is much less likely if nodes must be explicitly configured to handle a given limited-domain protocol, for example, by installing a specific protocol handler.

Investigations of the unreliability of IP fragmentation [[FRAG-FRAGILE](#)] and the filtering of IPv6 extension headers [[RFC7872](#)] strongly suggest that at least for some protocol elements, transparency is a lost cause and middleboxes are here to stay. In the following two sections, we show that some application environments require protocol features that cannot, or should not, cross the whole Internet.

### 3. Examples of Limited Domain Requirements

This section describes various examples where limited domain requirements can easily be identified, either based on an application scenario or on a technical imperative. It is, of course, not a complete list, and it is presented in an arbitrary order, loosely from smaller to bigger.

1. A home network. It will be mainly unmanaged, constructed by a non-specialist. It must work with devices "out of the box" as shipped by their manufacturers and must create adequate security by default. Remote access may be required. The requirements and applicable principles are summarized in [[RFC7368](#)].
2. A small office network. This is sometimes very similar to a home network, if whoever is in charge has little or no specialist knowledge, but may have differing security and privacy requirements. In other cases, it may be professionally constructed using recommended products and configurations but operate unmanaged. Remote access may be required.
3. A vehicle network. This will be designed by the vehicle manufacturer but may include devices added by the vehicle's owner or operator. Parts of the network will have demanding performance and reliability requirements with implications for human safety. Remote access may be required to certain functions but absolutely forbidden for others. Communication with other vehicles, roadside infrastructure, and external data sources will be required. See [[IPWAVE-NETWORKING](#)] for a survey of use cases.
4. Supervisory Control And Data Acquisition (SCADA) networks and other hard real-time networks. These will exhibit specific technical requirements, including tough real-time performance targets. See, for example, [[RFC8578](#)] for numerous use cases. An example is a building services network. This will be designed specifically for a particular building but using standard components. Additional devices may need to be added at any time. Parts of the network may have demanding reliability requirements with implications for human safety. Remote access may be required to certain functions but absolutely forbidden for others. An extreme example is a network used for virtual reality or augmented reality applications where the latency requirements are very stringent.
5. Sensor networks. The two preceding cases will all include sensors, but some networks may be specifically limited to sensors and the collection and processing of sensor data. They may be in remote or technically challenging locations and installed by non-specialists.
6. Internet-of-Things (IoT) networks. While this term is very flexible and covers many innovative types of networks, including ad hoc networks that are formed spontaneously and some applications of 5G technology, it seems reasonable to expect that IoT edge networks will have special requirements and protocols that are useful only within a specific domain, and that these protocols cannot, and for security reasons should not, run over the Internet as a whole.

7. Constrained Networks. An important subclass of IoT networks consists of constrained networks [RFC7228] in which the nodes are limited in power consumption and communications bandwidth and are therefore limited to using very frugal protocols.
8. Delay-tolerant networks. These may consist of domains that are relatively isolated and constrained in power (e.g., deep space networks) and are connected only intermittently to the outside, with a very long latency on such connections [RFC4838]. Clearly, the protocol requirements and possibilities are very specialized in such networks.
9. "Traditional" enterprise and campus networks, which may be spread over many kilometers and over multiple separate sites, with multiple connections to the Internet. Interestingly, the IETF appears never to have analyzed this long-established class of networks in a general way, except in connection with IPv6 deployment (e.g., [RFC7381]).
10. Unsuitable standards. A situation that can arise in an enterprise network is that the Internet-wide solution for a particular requirement may either fail locally or be much more complicated than is necessary. An example is that the complexity induced by a mechanism such as Interactive Connectivity Establishment (ICE) [RFC8445] is not justified within such a network. Furthermore, ICE cannot be used in some cases because candidate addresses are not known before a call is established, so a different local solution is essential [RFC6947].
11. Managed wide-area networks run by service providers for enterprise services such as Layer 2 (Ethernet, etc.) point-to-point pseudowires, multipoint Layer 2 Ethernet VPNs using Virtual Private LAN Service (VPLS) or Ethernet VPN (EVPN), and Layer 3 IP VPNs. These are generally characterized by service-level agreements for availability, packet loss, and possibly multicast service. These are different from the previous case in that they mostly run over MPLS infrastructures, and the requirements for these services are well defined by the IETF.
12. Data centers and hosting centers, or distributed services acting as such centers. These will have high performance, security, and privacy requirements and will typically include large numbers of independent "tenant" networks overlaid on shared infrastructure.
13. Content Delivery Networks (CDNs), comprising distributed data centers and the paths between them, spanning thousands of kilometers, with numerous connections to the Internet.
14. Massive Web Service Provider Networks. This is a small class of networks with well-known trademarked names, combining aspects of distributed enterprise networks, data centers, and CDNs. They have their own international networks bypassing the generic carriers. Like CDNs, they have numerous connections to the Internet, typically offering a tailored service in each economy.

Three other aspects, while not tied to specific network types, also strongly depend on the concept of limited domains:

1. Many of the above types of networks may be extended throughout the Internet by a variety of virtual private network (VPN) techniques. Therefore, we argue that limited domains may overlap each other in an arbitrary fashion by use of virtualization techniques. As noted above in the discussion of controlled environments, specific tunneling and encapsulation techniques may be tailored for use within a given domain.
2. Intent-Based Networking. In this concept, a network domain is configured and managed in accordance with an abstract policy known as "Intent" to ensure that the network performs as

required [IBN-CONCEPTS]. Whatever technologies are used to support this will be applied within the domain boundary, even if the services supported in the domain are globally accessible.

3. Network Slicing. A network slice is a form of virtual network that consists of a managed set of resources carved off from a larger network [ENHANCED-VPN]. This is expected to be significant in 5G deployments [USER-PLANE-PROTOCOL]. Whatever technologies are used to support slicing will require a clear definition of the boundary of a given slice within a larger domain.

While it is clearly desirable to use common solutions, and therefore common standards, wherever possible, it is increasingly difficult to do so while satisfying the widely varying requirements outlined above. However, there is a tendency when new protocols and protocol extensions are proposed to always ask the question "How will this work across the open Internet?" This document suggests that this is not always the best question. There are protocols and extensions that are not intended to work across the open Internet. On the contrary, their requirements and semantics are specifically limited (in the sense defined above).

A common argument is that if a protocol is intended for limited use, the chances are very high that it will in fact be used (or misused) in other scenarios including the so-called open Internet. This is undoubtedly true and means that limited use is not an excuse for bad design or poor security. In fact, a limited use requirement potentially adds complexity to both the protocol and its security design, as discussed later.

Nevertheless, because of the diversity of limited domains with specific requirements that is now emerging, specific standards (and ad hoc standards) will probably emerge for different types of domains. There will be attempts to capture each market sector, but the market will demand standardized solutions within each sector. In addition, operational choices will be made that can in fact only work within a limited domain. The history of RSVP [RFC2205] illustrates that a standard defined as if it could work over the open Internet might not in fact do so. In general, we can no longer assume that a protocol designed according to classical Internet guidelines will in fact work reliably across the network as a whole. However, the "open Internet" must remain as the universal method of interconnection. Reconciling these two aspects is a major challenge.

## 4. Examples of Limited Domain Solutions

This section lists various examples of specific limited domain solutions that have been proposed or defined. It intentionally does not include Layer 2 technology solutions, which by definition apply to limited domains. It is worth noting, however, that with recent developments such as Transparent Interconnection of Lots of Links (TRILL) [RFC6325] or Shortest Path Bridging [SPB], Layer 2 domains may become very large.

1. Differentiated Services. This mechanism [RFC2474] allows a network to assign locally significant values to the 6-bit Differentiated Services Code Point field in any IP packet. Although there are some recommended code point values for specific per-hop queue management behaviors, these are specifically intended to be domain-specific code points with traffic being classified, conditioned, and mapped or re-marked at domain boundaries

- (unless there is an inter-domain agreement that makes mapping or re-marking unnecessary).
2. Integrated Services. Although it is not intrinsic in the design of RSVP [RFC2205], it is clear from many years' experience that Integrated Services can only be deployed successfully within a limited domain that is configured with adequate equipment and resources.
  3. Network function virtualization. As described in [RFC8568], this general concept is an open research topic in which virtual network functions are orchestrated as part of a distributed system. Inevitably, such orchestration applies to an administrative domain of some kind, even though cross-domain orchestration is also a research area.
  4. Service Function Chaining (SFC). This technique [RFC7665] assumes that services within a network are constructed as sequences of individual service functions within a specific SFC-enabled domain such as a 5G domain. As that RFC states: "Specific features may need to be enforced at the boundaries of an SFC-enabled domain, for example to avoid leaking SFC information". A Network Service Header (NSH) [RFC8300] is used to encapsulate packets flowing through the service function chain: "The intended scope of the NSH is for use within a single provider's operational domain."
  5. Firewall and Service Tickets (FAST). Such tickets would accompany a packet to claim the right to traverse a network or request a specific network service [FAST]. They would only be meaningful within a particular domain.
  6. Data Center Network Virtualization Overlays. A common requirement in data centers that host many tenants (clients) is to provide each one with a secure private network, all running over the same physical infrastructure. [RFC8151] describes various use cases for this, and specifications are under development. These include use cases in which the tenant network is physically split over several data centers, but which must appear to the user as a single secure domain.
  7. Segment Routing. This is a technique that "steers a packet through an ordered list of instructions, called segments" [RFC8402]. The semantics of these instructions are explicitly local to a segment routing domain or even to a single node. Technically, these segments or instructions are represented as an MPLS label or an IPv6 address, which clearly adds a semantic interpretation to them within the domain.
  8. Autonomic Networking. As explained in [REF-MODEL], an autonomic network is also a security domain within which an autonomic control plane [ACP] is used by autonomic service agents. These agents manage technical objectives, which may be locally defined, subject to domain-wide policy. Thus, the domain boundary is important for both security and protocol purposes.
  9. Homenet. As shown in [RFC7368], a home networking domain has specific protocol needs that differ from those in an enterprise network or the Internet as a whole. These include the Home Network Control Protocol (HNCP) [RFC7788] and a naming and discovery solution [HOMENET-NAMING].
  10. Creative uses of IPv6 features. As IPv6 enters more general use, engineers notice that it has much more flexibility than IPv4. Innovative suggestions have been made for:
    - The flow label, e.g., [RFC6294].



- Extension headers, e.g., for segment routing [[RFC8754](#)] or Operations, Administration, and Maintenance (OAM) marking [[IPV6-ALT-MARK](#)].
- Meaningful address bits, e.g., [[EMBEDDED-SEMANTICS](#)]. Also, segment routing uses IPv6 addresses as segment identifiers with specific local meanings [[RFC8402](#)].
- If segment routing is used for network programming [[SRV6-NETWORK](#)], IPv6 extension headers can support rather complex local functionality.

The case of the extension header is particularly interesting, since its existence has been a major "selling point" for IPv6, but new extension headers are notorious for being virtually impossible to deploy across the whole Internet [[RFC7045](#)] [[RFC7872](#)]. It is worth noting that extension header filtering is considered an important security issue [[IPV6-EXT-HEADERS](#)]. There is considerable appetite among vendors or operators to have flexibility in defining extension headers for use in limited or specialized domains, e.g., [[IPV6-SRH](#)], [[BIGIP](#)], and [[APP-AWARE](#)]. Locally significant hop-by-hop options are also envisaged, that would be understood by routers inside a domain but not elsewhere, e.g., [[IN-SITU-OAM](#)].

11. Deterministic Networking (DetNet). The Deterministic Networking Architecture [[RFC8655](#)] and encapsulation [[DETNET-DATA-PLANE](#)] aim to support flows with extremely low data loss rates and bounded latency but only within a part of the network that is "DetNet aware". Thus, as for Differentiated Services above, the concept of a domain is fundamental.
12. Provisioning Domains (PvDs). An architecture for Multiple Provisioning Domains has been defined [[RFC7556](#)] to allow hosts attached to multiple networks to learn explicit details about the services provided by each of those networks.
13. Address Scopes. For completeness, we mention that, particularly in IPv6, some addresses have explicitly limited scope. In particular, link-local addresses are limited to a single physical link [[RFC4291](#)], and Unique Local Addresses [[RFC4193](#)] are limited to a somewhat loosely defined local site scope. Previously, site-local addresses were defined, but they were obsoleted precisely because of "the fuzzy nature of the site concept" [[RFC3879](#)]. Multicast addresses also have explicit scoping [[RFC4291](#)].
14. As an application-layer example, consider streaming services such as IPTV infrastructures that rely on standard protocols, but for which access is not globally available.

All of these suggestions are only viable within a specified domain. Nevertheless, all of them are clearly intended for multivendor implementation on thousands or millions of network domains, so interoperable standardization would be beneficial. This argument might seem irrelevant to private or proprietary implementations, but these have a strong tendency to become de facto standards if they succeed, so the arguments of this document still apply.

## 5. The Scope of Protocols in Limited Domains

One consequence of the deployment of limited domains in the Internet is that some protocols will be designed, extended, or configured so that they only work correctly between end systems in such domains. This is to some extent encouraged by some existing standards and by the assignment of code points for local or experimental use. In any case, it cannot be prevented. Also, by endorsing efforts such as Service Function Chaining, Segment Routing, and Deterministic

Networking, the IETF is in effect encouraging such deployments. Furthermore, it seems inevitable, if the Internet of Things becomes reality, that millions of edge networks containing completely novel types of nodes will be connected to the Internet; each one of these edge networks will be a limited domain.

It is therefore appropriate to discuss whether protocols or protocol extensions should sometimes be standardized to interoperate only within a limited-domain boundary. Such protocols would not be required to interoperate across the Internet as a whole. Various scenarios could then arise if there are multiple domains using the limited-domain protocol in question:

- A. If a domain is split into two parts connected over the Internet directly at the IP layer (i.e., with no tunnel encapsulating the packets), a limited-domain protocol could be operated between those two parts regardless of its special nature, as long as it respects standard IP formats and is not arbitrarily blocked by firewalls. A simple example is any protocol using a port number assigned to a specific non-IETF protocol.

Such a protocol could reasonably be described as an "inter-domain" protocol because the Internet is transparent to it, even if it is meaningless except in the two limited domains. This is, of course, nothing new in the Internet architecture.

- B. If a limited-domain protocol does not respect standard IP formats (for example, if it includes a non-standard IPv6 extension header), it could not be operated between two domains connected over the Internet directly at the IP layer.

Such a protocol could reasonably be described as an "intra-domain" protocol, and the Internet is opaque to it.

- C. If a limited-domain protocol is clearly specified to be invalid outside its domain of origin, neither scenario A nor B applies. The only solution would be a single virtual domain. For example, an encapsulating tunnel between two domains could be used to create the virtual domain. Also, nodes at the domain boundary must drop all packets using the limited-domain protocol.

- D. If a limited-domain protocol has domain-specific variants, such that implementations in different domains could not interoperate if those domains were unified by some mechanism as in scenario C, the protocol is not interoperable in the normal sense. If two domains using it were merged, the protocol might fail unpredictably. A simple example is any protocol using a port number assigned for experimental use. Related issues are discussed in [\[RFC5704\]](#), including the complex example of Transport MPLS.

To provide a widespread example, consider Differentiated Services [\[RFC2474\]](#). A packet containing any value whatsoever in the 6 bits of the Differentiated Services Code Point (DSCP) is well formed and falls into scenario A. However, because the semantics of DSCP values are locally significant, the packet also falls into scenario D. In fact, Differentiated Services are only interoperable across domain boundaries if there is a corresponding agreement between the operators; otherwise, a specific gateway function is required for meaningful interoperability. Much more detailed discussion is found in [\[RFC2474\]](#) and [\[RFC8100\]](#).

To provide a provocative example, consider the proposal in [IPV6-SRH] that the restrictions in [RFC8200] should be relaxed to allow IPv6 extension headers to be inserted on the fly in IPv6 packets. If this is done in such a way that the affected packets can never leave the specific limited domain in which they were modified, scenario C applies. If the semantic content of the inserted headers is locally defined, scenario D also applies. In neither case is the Internet outside the limited domain disturbed. However, inside the domain, nodes must understand the variant protocol. Unless it is standardized as a formal version, with all the complexity that implies [RFC6709], the nodes must all be non-standard to the extent of understanding the variant protocol. For the example of IPv6 header insertion, that means non-compliance with [RFC8200] within the domain, even if the inserted headers are themselves fully compliant. Apart from the issue of formal compliance, such deviations from documented standard behavior might lead to significant debugging issues. The possible practical impact of the header insertion example is explored in [IN-FLIGHT-IPV6].

The FAST proposal mentioned in Section 4, Paragraph 2, Item 5 is also an interesting case study. The semantics of FAST tickets [FAST] have limited scope. However, they are designed in a way that, in principle, allows them to traverse the open Internet, as standardized IPv6 hop-by-hop options or even as a proposed form of IPv4 extension header [IPV4-EXT-HEADERS]. Whether such options can be used reliably across the open Internet remains unclear [IPV6-EXT-HEADERS].

We conclude that it is reasonable to explicitly define limited-domain protocols, either as standards or as proprietary mechanisms, as long as they describe which of the above scenarios apply and they clarify how the domain is defined. As long as all relevant standards are respected outside the domain boundary, a well-specified limited-domain protocol need not damage the rest of the Internet. However, as described in the next section, mechanisms are needed to support domain membership operations.

Note that this conclusion is not a recommendation to abandon the normal goal that a standardized protocol should be global in scope and able to interoperate across the open Internet. It is simply a recognition that this will not always be the case.

## 6. Functional Requirements of Limited Domains

Noting that limited-domain protocols have been defined in the past, and that others will undoubtedly be defined in the future, it is useful to consider how a protocol can be made aware of the domain within which it operates and how the domain boundary nodes can be identified. As the taxonomy in Appendix A shows, there are numerous aspects to a domain. However, we can identify some generally required features and functions that would apply partially or completely to many cases.

Today, where limited domains exist, they are essentially created by careful configuration of boundary routers and firewalls. If a domain is characterized by one or more address prefixes, address assignment to hosts must also be carefully managed. This is an error-prone method, and a combination of configuration errors and default routing can lead to unwanted traffic escaping

the domain. Our basic assumption is therefore that it should be possible for domains to be created and managed automatically, with minimal human configuration. We now discuss requirements for automating domain creation and management.

First, if we drew a topology map, any given domain -- virtual or physical -- will have a well-defined boundary between "inside" and "outside". However, that boundary in itself has no technical meaning. What matters in reality is whether a node is a member of the domain and whether it is at the boundary between the domain and the rest of the Internet. Thus, the boundary in itself does not need to be identified, but boundary nodes face both inwards and outwards. Inside the domain, a sending node needs to know whether it is sending to an inside or outside destination, and a receiving node needs to know whether a packet originated inside or outside. Also, a boundary node needs to know which of its interfaces are inward facing or outward facing. It is irrelevant whether the interfaces involved are physical or virtual.

To underline that domain boundaries need to be identifiable, consider the statement from the Deterministic Networking Problem Statement [[RFC8557](#)] that "there is still a lack of clarity regarding the limits of a domain where a deterministic path can be set up". This remark can certainly be generalized.

With this perspective, we can list some general functional requirements. An underlying assumption here is that domain membership operations should be cryptographically secured; a domain without such security cannot be reliably protected from attack.

1. **Domain Identity.** A domain must have a unique and verifiable identifier; effectively, this should be a public key for the domain. Without this, there is no way to secure domain operations and domain membership. The holder of the corresponding private key becomes the trust anchor for the domain.
2. **Nesting.** It must be possible for domains to be nested (see, for example, the network-slicing example mentioned above).
3. **Overlapping.** It must be possible for nodes and links to be in more than one domain (see, for example, the case of PvDs mentioned above).
4. **Node Eligibility.** It must be possible for a node to determine which domain(s) it can potentially join and on which interface(s).
5. **Secure Enrollment.** A node must be able to enroll in a given domain via secure node identification and to acquire relevant security credentials (authorization) for operations within the domain. If a node has multiple physical or virtual interfaces, individual enrollment for each interface may be required.
6. **Withdrawal.** A node must be able to cancel enrollment in a given domain.
7. **Dynamic Membership.** Optionally, a node should be able to temporarily leave or rejoin a domain (i.e., enrollment is persistent but membership is intermittent).
8. **Role, implying authorization to perform a certain set of actions.** A node must have a verifiable role. In the simplest case, the role choices are "interior node" and "boundary node". In a boundary node, individual interfaces may have different roles, e.g., "inward facing" and "outward facing".

9. Peer Verification. A node must be able to verify whether another node is a member of the domain.
10. Role Verification. A node should be able to learn the verified role of another node. In particular, it should be possible for a node to find boundary nodes (interfacing to the Internet).
11. Domain Data. In a domain with management requirements, it must be possible for a node to acquire domain policy and/or domain configuration data. This would include, for example, filtering policy to ensure that inappropriate packets do not leave the domain.

These requirements could form the basis for further analysis and solution design.

Another aspect is whether individual packets within a limited domain need to carry any sort of indicator that they belong to that domain or whether this information will be implicit in the IP addresses of the packet. A related question is whether individual packets need cryptographic authentication. This topic is for further study.

## 7. Security Considerations

As noted above, a protocol intended for limited use may well be inadvertently used on the open Internet, so limited use is not an excuse for poor security. In fact, a limited use requirement potentially adds complexity to the security design.

Often, the boundary of a limited domain will also act as a security boundary. In particular, it will serve as a trust boundary and as a boundary of authority for defining capabilities. For example, segment routing [RFC8402] explicitly uses the concept of a "trusted domain" in this way. Within the boundary, limited-domain protocols or protocol features will be useful, but they will in many cases be meaningless or harmful if they enter or leave the domain.

The boundary also serves to provide confidentiality and privacy for operational parameters that the operator does not wish to reveal. Note that this is distinct from privacy protection for individual users within the domain.

The security model for a limited-scope protocol must allow for the boundary and in particular for a trust model that changes at the boundary. Typically, credentials will need to be signed by a domain-specific authority.

## 8. IANA Considerations

This document has no IANA actions.

## 9. Informative References

- [ACP] Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", Work in Progress, Internet-Draft, draft-ietf-anima-autonomic-control-plane-27, 2 July 2020, <<https://tools.ietf.org/html/draft-ietf-anima-autonomic-control-plane-27>>.

- [APP-AWARE]** Li, Z., Peng, S., Li, C., Xie, C., Voyer, D., Li, X., Liu, P., Liu, C., and K. Ebisawa, "Application-aware IPv6 Networking (APN6) Encapsulation", Work in Progress, Internet-Draft, draft-li-6man-app-aware-ipv6-network-02, 2 July 2020, <<https://tools.ietf.org/html/draft-li-6man-app-aware-ipv6-network-02>>.
- [BIGIP]** Li, R., "HUAWEI - Big IP Initiative", 2018, <<https://www.iaria.org/announcements/HuaweiBigIP.pdf>>.
- [DETNET-DATA-PLANE]** Varga, B., Farkas, J., Berger, L., Malis, A., and S. Bryant, "DetNet Data Plane Framework", Work in Progress, Internet-Draft, draft-ietf-detnet-data-plane-framework-06, 6 May 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-data-plane-framework-06>>.
- [DNS-PERIMETER]** Crocker, D. and T. Adams, "DNS Perimeter Overlay", Work in Progress, Internet-Draft, draft-dcrocker-dns-perimeter-01, 11 June 2019, <<https://tools.ietf.org/html/draft-dcrocker-dns-perimeter-01>>.
- [EMBEDDED-SEMANTICS]** Jiang, S., Qiong, Q., Farrer, I., Bo, Y., and T. Yang, "Analysis of Semantic Embedded IPv6 Address Schemas", Work in Progress, Internet-Draft, draft-jiang-semantic-prefix-06, 15 July 2013, <<https://tools.ietf.org/html/draft-jiang-semantic-prefix-06>>.
- [ENHANCED-VPN]** Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Service", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-06, 13 July 2020, <<https://tools.ietf.org/html/draft-ietf-teas-enhanced-vpn-06>>.
- [FAST]** Herbert, T., "Firewall and Service Tickets", Work in Progress, Internet-Draft, draft-herbert-fast-04, 10 April 2019, <<https://tools.ietf.org/html/draft-herbert-fast-04>>.
- [FRAG-FRAGILE]** Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", Work in Progress, Internet-Draft, draft-ietf-intarea-frag-fragile-17, 30 September 2019, <<https://tools.ietf.org/html/draft-ietf-intarea-frag-fragile-17>>.
- [HOMENET-NAMING]** Lemon, T., Migault, D., and S. Cheshire, "Homenet Naming and Service Discovery Architecture", Work in Progress, Internet-Draft, draft-ietf-homenet-simple-naming-03, 23 October 2018, <<https://tools.ietf.org/html/draft-ietf-homenet-simple-naming-03>>.
- [IBN-CONCEPTS]** Clemm, A., Ciavaglia, L., Granville, L., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", Work in Progress, Internet-Draft, draft-irtf-nmrg-ibn-concepts-definitions-01, 9 March 2020, <<https://tools.ietf.org/html/draft-irtf-nmrg-ibn-concepts-definitions-01>>.
- [IN-FLIGHT-IPV6]** Smith, M., Kottapalli, N., Bonica, R., Gont, F., and T. Herbert, "In-Flight IPv6 Extension Header Insertion Considered Harmful", Work in Progress, Internet-Draft, draft-smith-6man-in-flight-eh-insertion-harmful-02, 30 May 2020, <<https://tools.ietf.org/html/draft-smith-6man-in-flight-eh-insertion-harmful-02>>.

- [IN-SITU-OAM]** Bhandari, S., Brockners, F., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., Spiegel, M., Krishnan, S., and R. Asati, "In-situ OAM IPv6 Options", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-ipv6-options-02, 13 July 2020, <<https://tools.ietf.org/html/draft-ietf-ippm-ioam-ipv6-options-02>>.
- [IPV4-EXT-HEADERS]** Herbert, T., "IPv4 Extension Headers and Flow Label", Work in Progress, Internet-Draft, draft-herbert-ipv4-eh-01, 2 May 2019, <<https://tools.ietf.org/html/draft-herbert-ipv4-eh-01>>.
- [IPV6-ALT-MARK]** Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", Work in Progress, Internet-Draft, draft-ietf-6man-ipv6-alt-mark-01, 22 June 2020, <<https://tools.ietf.org/html/draft-ietf-6man-ipv6-alt-mark-01>>.
- [IPV6-EXT-HEADERS]** Gont, F. and W. LIU, "Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-ietf-opsec-ipv6-eh-filtering-06, 2 July 2018, <<https://tools.ietf.org/html/draft-ietf-opsec-ipv6-eh-filtering-06>>.
- [IPV6-SRH]** Voyer, D., Filsfils, C., Dukes, D., Matsushima, S., Leddy, J., Li, Z., and J. Guichard, "Deployments With Insertion of IPv6 Segment Routing Headers", Work in Progress, Internet-Draft, draft-voyer-6man-extension-header-insertion-09, 19 May 2020, <<https://tools.ietf.org/html/draft-voyer-6man-extension-header-insertion-09>>.
- [IPV6-USE-MINMTU]** Andrews, M., "TCP Fails To Respect IPV6\_USE\_MIN\_MTU", Work in Progress, Internet-Draft, draft-andrews-tcp-and-ipv6-use-minmtu-04, 18 October 2015, <<https://tools.ietf.org/html/draft-andrews-tcp-and-ipv6-use-minmtu-04>>.
- [IPWAVE-NETWORKING]** Jeong, J., "IPv6 Wireless Access in Vehicular Environments (IPWAVE): Problem Statement and Use Cases", Work in Progress, Internet-Draft, draft-ietf-ipwave-vehicular-networking-16, 7 July 2020, <<https://tools.ietf.org/html/draft-ietf-ipwave-vehicular-networking-16>>.
- [REF-MODEL]** Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., and J. Nobre, "A Reference Model for Autonomic Networking", Work in Progress, Internet-Draft, draft-ietf-anima-reference-model-10, 22 November 2018, <<https://tools.ietf.org/html/draft-ietf-anima-reference-model-10>>.
- [RFC2205]** Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<https://www.rfc-editor.org/info/rfc2205>>.
- [RFC2474]** Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.

- 
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <<https://www.rfc-editor.org/info/rfc2775>>.
- [RFC2923] Lahey, K., "TCP Problems with Path MTU Discovery", RFC 2923, DOI 10.17487/RFC2923, September 2000, <<https://www.rfc-editor.org/info/rfc2923>>.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, DOI 10.17487/RFC3234, February 2002, <<https://www.rfc-editor.org/info/rfc3234>>.
- [RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, DOI 10.17487/RFC3879, September 2004, <<https://www.rfc-editor.org/info/rfc3879>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4397] Bryskin, I. and A. Farrel, "A Lexicography for the Interpretation of Generalized Multiprotocol Label Switching (GMPLS) Terminology within the Context of the ITU-T's Automatically Switched Optical Network (ASON) Architecture", RFC 4397, DOI 10.17487/RFC4397, February 2006, <<https://www.rfc-editor.org/info/rfc4397>>.
- [RFC4427] Mannie, E., Ed. and D. Papadimitriou, Ed., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, DOI 10.17487/RFC4427, March 2006, <<https://www.rfc-editor.org/info/rfc4427>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", RFC 4838, DOI 10.17487/RFC4838, April 2007, <<https://www.rfc-editor.org/info/rfc4838>>.
- [RFC4924] Aboba, B., Ed. and E. Davies, "Reflections on Internet Transparency", RFC 4924, DOI 10.17487/RFC4924, July 2007, <<https://www.rfc-editor.org/info/rfc4924>>.
- [RFC5704] Bryant, S., Ed., Morrow, M., Ed., and IAB, "Uncoordinated Protocol Development Considered Harmful", RFC 5704, DOI 10.17487/RFC5704, November 2009, <<https://www.rfc-editor.org/info/rfc5704>>.
- [RFC6294] Hu, Q. and B. Carpenter, "Survey of Proposed Use Cases for the IPv6 Flow Label", RFC 6294, DOI 10.17487/RFC6294, June 2011, <<https://www.rfc-editor.org/info/rfc6294>>.



- 
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, DOI 10.17487/RFC6325, July 2011, <<https://www.rfc-editor.org/info/rfc6325>>.
- [RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <<https://www.rfc-editor.org/info/rfc6398>>.
- [RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, DOI 10.17487/RFC6407, October 2011, <<https://www.rfc-editor.org/info/rfc6407>>.
- [RFC6709] Carpenter, B., Aboba, B., Ed., and S. Cheshire, "Design Considerations for Protocol Extensions", RFC 6709, DOI 10.17487/RFC6709, September 2012, <<https://www.rfc-editor.org/info/rfc6709>>.
- [RFC6947] Boucadair, M., Kaplan, H., Gilman, R., and S. Veikkolainen, "The Session Description Protocol (SDP) Alternate Connectivity (ALTC) Attribute", RFC 6947, DOI 10.17487/RFC6947, May 2013, <<https://www.rfc-editor.org/info/rfc6947>>.
- [RFC6950] Peterson, J., Kolkman, O., Tschofenig, H., and B. Aboba, "Architectural Considerations on Application Features in the DNS", RFC 6950, DOI 10.17487/RFC6950, October 2013, <<https://www.rfc-editor.org/info/rfc6950>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <<https://www.rfc-editor.org/info/rfc7045>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <<https://www.rfc-editor.org/info/rfc7368>>.
- [RFC7381] Chittimaneni, K., Chown, T., Howard, L., Kuarsingh, V., Pouffary, Y., and E. Vyncke, "Enterprise IPv6 Deployment Guidelines", RFC 7381, DOI 10.17487/RFC7381, October 2014, <<https://www.rfc-editor.org/info/rfc7381>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.
- [RFC7663] Trammell, B., Ed. and M. Kuehlewind, Ed., "Report from the IAB Workshop on Stack Evolution in a Middlebox Internet (SEMI)", RFC 7663, DOI 10.17487/RFC7663, October 2015, <<https://www.rfc-editor.org/info/rfc7663>>.
- [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>>.
-

- 
- [RFC7754] Barnes, R., Cooper, A., Kolkman, O., Thaler, D., and E. Nordmark, "Technical Considerations for Internet Service Blocking and Filtering", RFC 7754, DOI 10.17487/RFC7754, March 2016, <<https://www.rfc-editor.org/info/rfc7754>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <<https://www.rfc-editor.org/info/rfc7872>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [RFC8086] Yong, L., Ed., Crabbe, E., Xu, X., and T. Herbert, "GRE-in-UDP Encapsulation", RFC 8086, DOI 10.17487/RFC8086, March 2017, <<https://www.rfc-editor.org/info/rfc8086>>.
- [RFC8100] Geib, R., Ed. and D. Black, "Diffserv-Interconnection Classes and Practice", RFC 8100, DOI 10.17487/RFC8100, March 2017, <<https://www.rfc-editor.org/info/rfc8100>>.
- [RFC8151] Yong, L., Dunbar, L., Toy, M., Isaac, A., and V. Manral, "Use Cases for Data Center Network Virtualization Overlay Networks", RFC 8151, DOI 10.17487/RFC8151, May 2017, <<https://www.rfc-editor.org/info/rfc8151>>.
- [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>>.
- [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>>.
- [RFC8445] Keranen, A., Holmberg, C., and J. Rosenberg, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal", RFC 8445, DOI 10.17487/RFC8445, July 2018, <<https://www.rfc-editor.org/info/rfc8445>>.
- [RFC8517] Dolson, D., Ed., Snellman, J., Boucadair, M., Ed., and C. Jacquenet, "An Inventory of Transport-Centric Functions Provided by Middleboxes: An Operator Perspective", RFC 8517, DOI 10.17487/RFC8517, February 2019, <<https://www.rfc-editor.org/info/rfc8517>>.

- 
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8568] Bernardos, CJ., Rahman, A., Zuniga, JC., Contreras, LM., Aranda, P., and P. Lynch, "Network Virtualization Research Challenges", RFC 8568, DOI 10.17487/RFC8568, April 2019, <<https://www.rfc-editor.org/info/rfc8568>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [RFC8655] 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>>.
- [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>>.
- [SPB] "IEEE Standard for Local and metropolitan area networks - Bridges and Bridged Networks", DOI 10.1109/IEEESTD.2018.8403927, IEEE 802.1Q-2018, July 2018, <<https://ieeexplore.ieee.org/document/8403927>>.
- [SRV6-NETWORK] Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-network-programming-16, 27 June 2020, <<https://tools.ietf.org/html/draft-ietf-spring-srv6-network-programming-16>>.
- [USER-PLANE-PROTOCOL] Homma, S., Miyasaka, T., Matsushima, S., and D. Voyer, "User Plane Protocol and Architectural Analysis on 3GPP 5G System", Work in Progress, Internet-Draft, draft-ietf-dmm-5g-uplane-analysis-03, 3 November 2019, <<https://tools.ietf.org/html/draft-ietf-dmm-5g-uplane-analysis-03>>.

## Appendix A. Taxonomy of Limited Domains

This appendix develops a taxonomy for describing limited domains. Several major aspects are considered in this taxonomy:

- The domain as a whole
- The individual nodes
- The domain boundary
- The domain's topology
- The domain's technology
- How the domain connects to the Internet
- The security, trust, and privacy model
- Operations

The following sub-sections analyze each of these aspects.

## **A.1. Domain as a Whole**

- Why does the domain exist? (e.g., human choice, administrative policy, orchestration requirements, technical requirements such as operational partitioning for scaling reasons)
- If there are special requirements, are they at Layer 2, Layer 3, or an upper layer?
- Where does the domain lie on the spectrum between completely managed by humans and completely autonomic?
- If managed, what style of management applies? (Manual configuration, automated configuration, orchestration?)
- Is there a policy model? (Intent, configuration policies?)
- Does the domain provide controlled or paid service or open access?

## **A.2. Individual Nodes**

- Is a domain member a complete node or only one interface of a node?
- Are nodes permanent members of a given domain, or are join and leave operations possible?
- Are nodes physical or virtual devices?
- Are virtual nodes general purpose or limited to specific functions, applications, or users?
- Are nodes constrained (by battery, etc.)?
- Are devices installed "out of the box" or pre-configured?

## **A.3. Domain Boundary**

- How is the domain boundary identified or defined?
- Is the domain boundary fixed or dynamic?
- Are boundary nodes special, or can any node be at the boundary?

## **A.4. Topology**

- Is the domain a subset of a Layer 2 or 3 connectivity domain?
- Does the domain overlap other domains? (In other words, is a node allowed to be a member of multiple domains?)
- Does the domain match physical topology, or does it have a virtual (overlay) topology?
- Is the domain in a single building, vehicle, or campus? Or is it distributed?
- If distributed, are the interconnections private or over the Internet?
- In IP addressing terms, is the domain Link local, Site local, or Global?
- Does the scope of IP unicast or multicast addresses map to the domain boundary?

## A.5. Technology

- What routing protocol(s) or different forwarding mechanisms (MPLS or other non-IP mechanism) are used?
- In an overlay domain, what overlay technique is used (L2VPN, L3VPN, etc.)?
- Are there specific QoS requirements?
- Link latency - Normal or long latency links?
- Mobility - Are nodes mobile? Is the whole network mobile?
- Which specific technologies, such as those in [Section 4](#), are applicable?

## A.6. Connection to the Internet

- Is the Internet connection permanent or intermittent? (Never connected is out of scope.)
- What traffic is blocked, in and out?
- What traffic is allowed, in and out?
- What traffic is transformed, in and out?
- Is secure and privileged remote access needed?
- Does the domain allow unprivileged remote sessions?

## A.7. Security, Trust, and Privacy Model

- Must domain members be authorized?
- Are all nodes in the domain at the same trust level?
- Is traffic authenticated?
- Is traffic encrypted?
- What is hidden from the outside?

## A.8. Operations

- Safety level - Does the domain have a critical (human) safety role?
- Reliability requirement - Normal or 99.999%?
- Environment - Hazardous conditions?
- Installation - Are specialists needed?
- Service visits - Easy, difficult, or impossible?
- Software/firmware updates - Possible or impossible?

## A.9. Making Use of This Taxonomy

This taxonomy could be used to design or analyze a specific type of limited domain. For the present document, it is intended only to form a background to the scope of protocols used in limited domains and the mechanisms required to securely define domain membership and properties.

## Acknowledgements

Useful comments were received from Amelia Andersdotter, Edward Birrane, David Black, Ron Bonica, Mohamed Boucadair, Tim Chown, Darren Dukes, Donald Eastlake, Adrian Farrel, Tom Herbert, Ben Kaduk, John Klensin, Mirja Kuehlewind, Warren Kumari, Andy Malis, Michael Richardson, Mark Smith, Rick Taylor, Niels ten Oever, and others.

## Contributors

### Sheng Jiang

Huawei Technologies  
Q14, Huawei Campus  
No. 156 Beiqing Road  
Hai-Dian District, Beijing  
100095  
China  
Email: [jiangsheng@huawei.com](mailto:jiangsheng@huawei.com)

## Authors' Addresses

### Brian Carpenter

The University of Auckland  
School of Computer Science  
University of Auckland  
PB 92019  
Auckland 1142  
New Zealand  
Email: [brian.e.carpenter@gmail.com](mailto:brian.e.carpenter@gmail.com)

**Bing Liu**

Huawei Technologies

Q14, Huawei Campus

No. 156 Beiqing Road

Hai-Dian District, Beijing

100095

China

Email: [leo.liubing@huawei.com](mailto:leo.liubing@huawei.com)