                   Limited Domains and Internet Protocols
                      draft-carpenter-limited-domains-04

Abstract

   There is a noticeable trend towards network requirements, behaviours
   and semantics that are specific to a limited region of the Internet
   and a particular set of requirements. Policies, default parameters,
   the options supported, the style of network management and security
   requirements may vary. This document reviews examples of such
   limited domains and emerging solutions, and develops a related
   taxonomy. It then briefly discusses the standardization of protocols
   for limited domains. Finally, it shows the needs 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.

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 April 17, 2019.

Table of Contents

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 local requirements,
   behaviours 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 partially in parallel with the
   Internet. We expand on these possibilities below. To capture the
   topic, this document refers to such networks as "limited domains".

Some people have concerns about splintering of the Internet along
political or linguistic boundaries by mechanisms that block the free
flow of information across the network. That is not the topic of
this document, which does not discuss filtering mechanisms 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.

The requirements of limited domains will be different in different
scenarios. 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 documents analyses and discusses some of the consequences of
this trend, and how it impacts the idea of universal interoperability
in the Internet. Firstly 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. Then we
develop a taxonomy of the features to be found in limited domains.
With this background, we discuss the resulting challenge to the
notion that all Internet standards must be universal in scope and
applicability. To the contrary, we assert that some protocols need
to be specifically limited in their applicability. This implies that
the concepts of a limited domain, and of its membership, need to be
formalised and supported by secure mechanisms. While this document
does not propose a design for such mechanisms, it does outline some
resulting functional requirements.

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. Previously, this has been analysed in terms
of transparency [RFC2775], [RFC4924] or of intrusive middleboxes
[RFC3234], [RFC7663], [I-D.dolson-plus-middlebox-benefits].
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

[I-D.bonica-intarea-frag-fragile], and problems in TCP MSS
negotiation have been reported [I-D.andrews-tcp-and-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.

This situation is not acceptable, so it seems that a new approach is
needed.

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 unmanaged, constructed by a non-
        specialist, and will possibly include wiring errors such as
        physical loops. 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 summarised 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
        [I-D.ietf-ipwave-vehicular-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 [I-D.ietf-detnet-use-cases] 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.

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 network, including
ad hoc networks that are formed spontaneously, 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.   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 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 specialised
in such networks.

9.   "Traditional" enterprise and campus networks, which may be
spread over many kilometres and over multiple separate sites,
with multiple connections to the Internet. Interestingly, the
IETF appears never to have analysed this long-established class
of networks in a general way, except in connection with IPv6
deployment (e.g. [RFC7381]).

10. Data centres and hosting centres, or distributed services acting
as such centres. These will have high performance, security and
privacy requirements and will typically include large numbers of
independent "tenant" networks overlaid on shared infrastructure.

11. Content Delivery Networks (CDNs), comprising distributed data
     centres and the paths between them, spanning thousands of
     kilometres, with numerous connections to the Internet.

12. Massive Web Service Provider Networks. This is a small class of
     networks with well known trademarked names, combining aspects of
     distributed enterprise networks, data centres 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. 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 [I-D.moulchan-nmrg-network-intent-concepts]. Whatever
    technologies are used to support this, they will be applied
    within the domain boundary.

2. Many of the above types of network may be extended throughout the
    Internet by a variety of virtual private network (VPN)
    techniques. Therefore we may argue that limited domains may
    overlap each other in an arbitrary fashion by use of
    virtualization techniques.

3. Network Slicing. A network slice is a virtual network that
    consists of a managed set of resources carved off from a larger
    network [I-D.geng-netslices-architecture]. Whatever technologies
    are used to support slicing, they will require a clear definition
    of the boundary of a given slice.

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 right 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 environments with
specific requirements that is now emerging, specific standards will
necessarily emerge. There will be attempts to capture each market
sector, but the market will demand standardised limited solutions.
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.

   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 codepoint values for
        specific per-hop queue management behaviours, these are
        specifically intended to be domain-specific codepoints with
        traffic being classified, conditioned and re-marked at domain
        boundaries (unless there is an inter-domain agreement that makes
        re-marking unnecessary).

   2.   Network function virtualisation. As described in
        [I-D.irtf-nfvrg-gaps-network-virtualization], 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.

   3.   Service Function Chaining (SFC). This technique [RFC7665]
        assumes that services within a network are constructed as
        sequences of individual functions within a specific SFC-enabled
        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."

   4.   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 [I-D.herbert-fast]". They
would only be valid within a particular domain.

5.   Data Centre Network Virtualization Overlays. A common
     requirement in data centres 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 centres, but which
     must appear to the user as a single secure domain.

6.   Segment Routing. This is a technique which "steers a packet
     through an ordered list of instructions, called segments"
     [I-D.ietf-spring-segment-routing]. 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.

7.   Autonomic Networking. As explained in
     [I-D.ietf-anima-reference-model], an autonomic network is also a
     security domain within which an autonomic control plane
     [I-D.ietf-anima-autonomic-control-plane] is used by service
     agents. These service 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.

8.   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 [I-D.ietf-homenet-simple-naming].

9.   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],
        [I-D.fioccola-v6ops-ipv6-alt-mark].

     * Extension headers, e.g. for segment routing
        [I-D.ietf-6man-segment-routing-header].

     * Meaningful address bits, e.g. [I-D.jiang-semantic-prefix].
        Also, segment routing uses IPv6 addresses as segment

identifiers with specific local meanings
[I-D.ietf-spring-segment-routing].

All of these suggestions are only viable within a specified
domain. The case of the extension header is particularly
interesting, since its existence has been a major "selling
point" for IPv6, but it is notorious that new extension headers
are virtually impossible to deploy across the whole Internet
[RFC7045], [RFC7872]. It is worth noting that extension header
filtering is considered as an important security issue
[I-D.ietf-opsec-ipv6-eh-filtering]. There is considerable
appetite among vendors or operators to have flexibility in
defining extension headers for use in limited or specialised
domains, e.g. [I-D.voyer-6man-extension-header-insertion] and
[BIGIP].

10. Deterministic Networking (DetNet). The Deterministic Networking
     Architecture [I-D.ietf-detnet-architecture] and encapsulation
     [I-D.ietf-detnet-dp-sol] 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.

11. 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.

5. Taxonomy of Limited Domains

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

   o The domain as a whole.

   o The individual nodes.

   o The domain boundary.

   o The domain's topology.

   o The domain's technology.

   o How the domain connects to the Internet.

   o The security, trust and privacy model.

   o Operations.

The following sub-sections analyse each of these aspects.

5.1. The Domain as a Whole

   o Why does the domain exist? (e.g., human choice, administrative
      policy, orchestration requirements, technical requirements)

   o If there are special requirements, are they at Layer 2, Layer 3 or
      an upper layer?

   o Is the domain managed by humans or fully autonomic?

   o If managed, what style of management applies? (Manual
      configuration, automated configuration, orchestration?)

   o Is there a policy model? (Intent, configuration policies?)

   o Does the domain provide controlled or paid service or open access?

5.2. Individual Nodes

   o Is a domain member a complete node, or only one interface of a
      node?

   o Are nodes permanent members of a given domain, or are join and
      leave operations possible?

   o Are nodes physical or virtual devices?

   o Are virtual nodes general-purpose, or limited to specific
      functions, applications or users?

   o Are nodes constrained (by battery etc)?

   o Are devices installed "out of the box" or pre-configured?

5.3. The Domain Boundary

   o How is the domain boundary identified or defined?

   o Is the domain boundary fixed or dynamic?

   o Are boundary nodes special? Or can any node be at the boundary?

5.4. Topology

   o Is the domain a subset of a layer 2 or 3 connectivity domain?

   o In IP addressing terms, is the domain Link-local, Site-local, or
      Global?

   o Does the domain overlap other domains? (In other words, a node
      may or may not be allowed to be a member of multiple domains.)

   o Does the domain match physical topology, or does it have a virtual
      (overlay) topology?

   o Is the domain in a single building, vehicle or campus? Or is it
      distributed?

   o If distributed, are the interconnections private or over the
      Internet?

   o In IP addressing terms, is the domain Link-local, Site-local, or
      Global?

5.5. Technology

   o In routing terms, what routing protocol(s) are used, or even
      different forwarding mechanisms (MPLS or other non-IP mechanism)?

   o In an overlay domain, what overlay technique is used (L2VPN,
      L3VPN,...)?

   o Are there specific QoS requirements?

   o Link latency - normal or long latency links?

   o Mobility - are nodes mobile? Is the whole network mobile?

   o Which specific technologies, such as those in Section 4, are
      applicable?

5.6. Connection to the Internet

   o Is the Internet connection permanent or intermittent? (Never
      connected is out of scope.)

   o What traffic is blocked, in and out?

   o What traffic is allowed, in and out?

   o What traffic is transformed, in and out?

   o Is secure and privileged remote access needed?

   o Does the domain allow unprivileged remote sessions?

5.7. Security, Trust and Privacy Model

   o Must domain members be authorized?

   o Are all nodes in the domain at the same trust level?

   o Is traffic authenticated?

   o Is traffic encrypted?

   o What is hidden from the outside?

5.8. Operations

   o Safety level - does the domain have a critical (human) safety
      role?

   o Reliability requirement - normal or 99.999% ?

   o Environment - hazardous conditions?

   o Installation - are specialists needed?

   o Service visits - easy, difficult, impossible?

   o Software/firmware updates - possible or impossible?

5.9. Making use of this taxonomy

   This taxonomy could be used to design or analyse a specific type of
   limited domain. For the present document, it is intended to form a
   background to the following two sections, concerning the scope of
   protocols used in limited domains, and mechanisms reuqired to
   securely define domain membership and properties.

6. 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 node 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 standardised to interoperate only
within a Limited Domain Boundary. Such protocols would not be
required to interoperate across the Internet as a whole. Several
possibly overlapping scenarios could then arise:

   A. If a limited 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 parts of the limited domain. 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 parts of a domain
   split 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 two domains need to be unified as a single virtual
   domain. For example, an encapsulating tunnel between the parts of
   the split domain could be used. 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, 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. Such a protocol usually also falls into
scenario C.

To provide an existing example, consider Differentiated Services
[RFC2474]. A packet containing any value whatever in the 6 bits of
the Differentiated Service 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 to be
found in [RFC2474] and [RFC8100].

To provide a provocative example, consider the proposal in
[I-D.voyer-6man-extension-header-insertion] 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 disturbed.

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 is not harmful to the Internet. However, as
described in the next section, mechanisms are needed to support
domain membership operations.

7. Functional Requirements of Limited Domains

As the preceding taxonomy shows, there are very numerous aspects to a
domain, so the common features are not immediately obvious. It would
be possible, but tedious, to apply the taxonomy to each of the domain
types described in Section 3. However, we can deduce some generally
required features and functions without doing so.

Firstly, if we drew a topology map, any 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. However, a sending node needs to know whether it is
sending to an inside or outside destination; a receiving node needs

to know whether a packet originated inside or outside; and 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.

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. Node Eligibility. It must be possible for a node to determine which domain(s) it can potentially join, and on which interface(s).

3. Secure Enrolment. A node must be able to enrol in a given domain via secure node identfication and to acquire relevant security credentials (authorization) for operations within the domain. If a node has multiple physical or virtual interfaces, they may require to be individually enrolled.

4. Withdrawal. A node must be able to cancel enrolment in a given domain.

5. Dynamic Membership. Optionally, a node should be able temporarily leave or rejoin a domain (i.e. enrolment is persistent but membership is intermittent).

6. Role, implying authorization to perform a certain set of actions. A node must have a verifiable role. In the simplest case, the choices of role are "interior node" and "boundary node". In a boundary node, individual interfaces may have different roles, e.g. "inward facing" and "outward facing".

7. Verify Peer. A node must be able to verify whether another node is a member of the domain.

8. Verify Role. A node must be able to learn the verified role of another node. In particular, it must be possible for a node to find boundary nodes (interfacing to the Internet).

9. 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.

8. Security Considerations

Clearly, the boundary of a limited domain will almost always also act
as a security boundary. In particular, it will serve as a trust
boundary, and as a boundary of authority for defining capabilities.
Within the boundary, limited-domain protocols or protocol features
will be useful, but they will be meaningless if they enter or leave
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.

9. IANA Considerations

This document makes no request of the IANA.

10. Contributors

Sheng Jiang made important contributions to this document.

11. Acknowledgements

Useful comments were received from Amelia Andersdotter, Edward
Birrane, Ron Bonica, Tim Chown, Darren Dukes, Tom Herbert, John
Klensin, Michael Richardson, Rick Taylor, Niels ten Oever, and other
members of the ANIMA and INTAREA WGs.

12. Informative References

[BIGIP]   Li, R., "HUAWEI - Big IP Initiative.", 2018,
            <https://www.iaria.org/announcements/HuaweiBigIP.pdf>.

[I-D.andrews-tcp-and-ipv6-use-minmtu]
            Andrews, M., "TCP Fails To Respect IPV6_USE_MIN_MTU",
            draft-andrews-tcp-and-ipv6-use-minmtu-04 (work in
            progress), October 2015.

[I-D.bonica-intarea-frag-fragile]
          Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O.,
          and F. Gont, "IP Fragmentation Considered Fragile", draft-
          bonica-intarea-frag-fragile-03 (work in progress), July
          2018.

[I-D.dolson-plus-middlebox-benefits]
          Dolson, D., Snellman, J., Boucadair, M., and C. Jacquenet,
          "Beneficial Functions of Middleboxes", draft-dolson-plus-
          middlebox-benefits-03 (work in progress), March 2017.

[I-D.fioccola-v6ops-ipv6-alt-mark]
          Fioccola, G., Velde, G., Cociglio, M., and P. Muley, "IPv6
          Performance Measurement with Alternate Marking Method",
          draft-fioccola-v6ops-ipv6-alt-mark-01 (work in progress),
          June 2018.

[I-D.geng-netslices-architecture]
          67, 4., Dong, J., Bryant, S., kiran.makhijani@huawei.com,
          k., Galis, A., Foy, X., and S. Kuklinski, "Network Slicing
          Architecture", draft-geng-netslices-architecture-02 (work
          in progress), July 2017.

[I-D.herbert-fast]
          Herbert, T., "Firewall and Service Tickets", draft-
          herbert-fast-03 (work in progress), September 2018.

[I-D.ietf-6man-segment-routing-header]
          Filsfils, C., Previdi, S., Leddy, J., Matsushima, S., and
          d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header
          (SRH)", draft-ietf-6man-segment-routing-header-14 (work in
          progress), June 2018.

[I-D.ietf-anima-autonomic-control-plane]
          Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic
          Control Plane (ACP)", draft-ietf-anima-autonomic-control-
          plane-18 (work in progress), August 2018.

[I-D.ietf-anima-reference-model]
          Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L.,
          and J. Nobre, "A Reference Model for Autonomic
          Networking", draft-ietf-anima-reference-model-08 (work in
          progress), October 2018.

[I-D.ietf-detnet-architecture]
          Finn, N., Thubert, P., Varga, B., and J. Farkas,
          "Deterministic Networking Architecture", draft-ietf-
          detnet-architecture-08 (work in progress), September 2018.

   [I-D.ietf-detnet-dp-sol]
             Korhonen, J., Andersson, L., Jiang, Y., Finn, N., Varga,
             B., Farkas, J., Bernardos, C., Mizrahi, T., and L. Berger,
             "DetNet Data Plane Encapsulation", draft-ietf-detnet-dp-
             sol-04 (work in progress), March 2018.

   [I-D.ietf-detnet-use-cases]
             Grossman, E., "Deterministic Networking Use Cases", draft-
             ietf-detnet-use-cases-19 (work in progress), October 2018.

   [I-D.ietf-homenet-simple-naming]
             Lemon, T., Migault, D., and S. Cheshire, "Simple Homenet
             Naming and Service Discovery Architecture", draft-ietf-
             homenet-simple-naming-02 (work in progress), July 2018.

   [I-D.ietf-ipwave-vehicular-networking]
             Jeong, J., "IP Wireless Access in Vehicular Environments
             (IPWAVE): Problem Statement and Use Cases", draft-ietf-
             ipwave-vehicular-networking-04 (work in progress), July
             2018.

   [I-D.ietf-opsec-ipv6-eh-filtering]
             Gont, F. and W. LIU, "Recommendations on the Filtering of
             IPv6 Packets Containing IPv6 Extension Headers", draft-
             ietf-opsec-ipv6-eh-filtering-06 (work in progress), July
             2018.

   [I-D.ietf-spring-segment-routing]
             Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B.,
             Litkowski, S., and R. Shakir, "Segment Routing
             Architecture", draft-ietf-spring-segment-routing-15 (work
             in progress), January 2018.

   [I-D.irtf-nfvrg-gaps-network-virtualization]
             Bernardos, C., Rahman, A., Zuniga, J., Contreras, L.,
             Aranda, P., and P. Lynch, "Network Virtualization Research
             Challenges", draft-irtf-nfvrg-gaps-network-
             virtualization-10 (work in progress), September 2018.

   [I-D.jiang-semantic-prefix]
             Jiang, S., Qiong, Q., Farrer, I., Bo, Y., and T. Yang,
             "Analysis of Semantic Embedded IPv6 Address Schemas",
             draft-jiang-semantic-prefix-06 (work in progress), July
             2013.

   [I-D.moulchan-nmrg-network-intent-concepts]
              Sivakumar, K. and M. Chandramouli, "Concepts of Network
              Intent", draft-moulchan-nmrg-network-intent-concepts-00
              (work in progress), October 2017.

   [I-D.voyer-6man-extension-header-insertion]
              daniel.voyer@bell.ca, d., Leddy, J., Filsfils, C., Dukes,
              D., Previdi, S., and S. Matsushima, "Insertion of IPv6
              Segment Routing Headers in a Controlled Domain", draft-
              voyer-6man-extension-header-insertion-04 (work in
              progress), June 2018.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

Appendix A. Change log [RFC Editor: Please remove]

   draft-carpenter-limited-domains-00, 2018-06-11:

   Initial version

   draft-carpenter-limited-domains-01, 2018-07-01:

   Minor terminology clarifications

   draft-carpenter-limited-domains-02, 2018-08-03:

   Additions following IETF102 discussions

   Updated authorship/contributors

   draft-carpenter-limited-domains-03, 2018-09-12:

   First draft of taxonomy

   Editorial improvements

   draft-carpenter-limited-domains-04, 2018-10-14:

   Reorganized section 3

   Newly written sections 6 and 7

   Editorial improvements

Authors' Addresses

    Brian Carpenter
    Department of Computer Science
    University of Auckland
    PB 92019
    Auckland   1142
    New Zealand

    Email: brian.e.carpenter@gmail.com

    Bing Liu
    Huawei Technologies
    Q14, Huawei Campus
    No.156 Beiqing Road
    Hai-Dian District, Beijing   100095
    P.R. China

    Email: leo.liubing@huawei.com