Workgroup: Network Working Group
Internet-Draft:
draft-bormann-t2trg-interconnect-declared-00
Published: 12 July 2021
Intended Status: Informational
Expires: 13 January 2022
A C. Bormann
uUniversität Bremen TZI
t
h
o
r
s
:

Interconnecting Limited Domains Based on Declared Communication Requirements

Abstract

"Limited Domains" are parts of an internet that may have notable differences or are just convenient to separate from the general Internet and can be delimited from that and from other Limited Domains by a defined boundary (the "border").

This memo focuses on the case where the nodes inside the Limited Domain want to interact with nodes on (or reachable via) the general Internet, but need some assistance at the border that is cognizant about the specific properties of the nodes in the Limited Domain.

Self-Descriptions can provide the information needed for this assistance.

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 13 January 2022.

Table of Contents

1. Introduction

[RFC8799] introduces the concept of "Limited Domains", i.e., parts of an internet that may have notable differences or are just convenient to separate from the general Internet and can be delimited from that and from other Limited Domains by a defined boundary (the "border").

Limited Domains are not a new concept, but they recently have gained significant attention as a way to accelerate innovation without always having to wait for the whole Internet to accept a new feature.

Some Limited Domains can be directly connected to or interconnected via the Internet -- rules they use internally simply lose their force outside the Limited Domain. Some require stripping off some structures or translating some fields on the border to the Internet. Some can only be interconnected by running tunnels on top of the Internet.

This memo focuses on the case where the nodes inside the Limited Domain want to interact with nodes on (or reachable via) the general Internet, but need some assistance at the border that is cognizant about the specific properties of the nodes in the Limited Domain.

6LoWPAN header compression [RFC6282] actually is such an example, which can be considered a very small Limited Domain -- initially just the link and adaptation layer between two LoWPAN nodes, which themselves otherwise feel like standard Internet nodes. (6LoWPAN neighbor discovery [RFC6775] already extends the Limited Domain out to the border router (6LBR), but let's focus on header compression itself for now.) Extending the Limited Domain to more than two nodes may allow the nodes inside the Limited Domain to make use of the knowledge that all of them share some common procedures, such as

using the [RFC8138] routing header (6LoRH); it is then the job of the border router (6LBR) to decapsulate this form into packets that can be used in the global Internet and to appropriately encapsulate global Internet packets on the way in. (Virtual Reassembly Buffers (VRBs, [RFC8930]) simulate a subnet-size Limited Domain based on [RFC6282]'s hop-by-hop ones.)

This memo uses examples from the area of the Internet of Things (IoT), both because the author is most familiar with it and because a concept of self-descriptions has already been developed for this area, which provide new opportunities for organizing Limited Domains (Section 3). (To do: add more examples from outside the IoT core.)

1.1. Terminology

Limited Domain: An area in a network that is separate from others by notable internal differences and/or by a strong administrative demarcation. Examples are found in Section 4 of [RFC8799], however this document is not limiting itself to those or to the definition in [RFC8799]. In contrast to some other usage, the nodes in a Limited Domain are expected to normally form a connected graph, possibly by employing tunnels between them. However, not all nodes in a Limited Domain always need to be aware of their situation or implement all the internal differences.

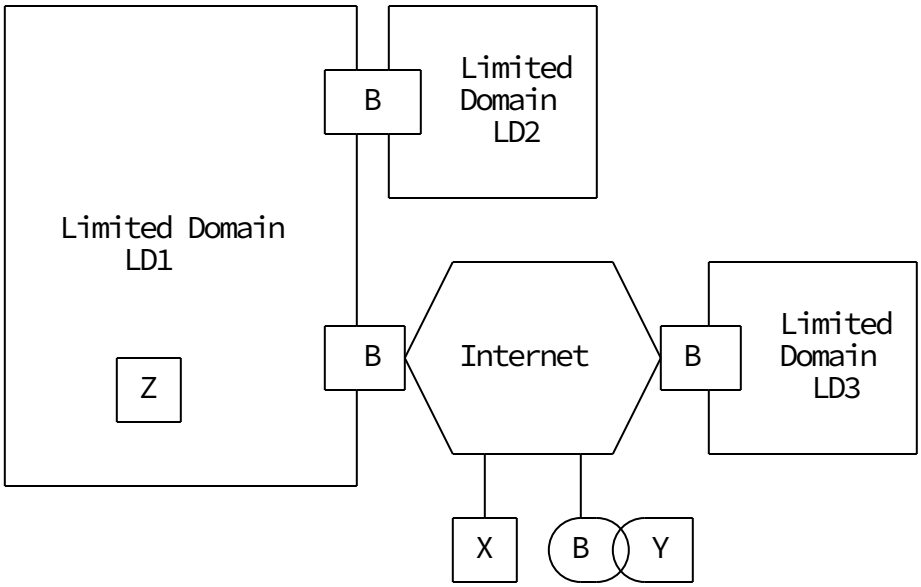

Figure 1: Illustration for terms

Figure 1 illustrates the following additional terms:

Border: qualifier for network elements (B) (as in "border router" etc.) that are situated on the boundary between a Limited Domain and a different one and/or the general Internet.

Internally interoperable Limited Domains: Limited Domains that can accommodate nodes that can operate as if they were in the Internet (Z).

Globally interoperable Limited Domains:
Limited Domains that can interoperate with nodes in the global Internet (X).

Externally interoperable Limited Domains: Limited Domains that can interoperate via the Internet (LD1), but possibly limited to interoperation with other Limited Domains (LD3) or with specially equipped Internet nodes (Limited Domains of size 1, Y logically containing a B).

Internet, Global Internet, General Internet: (TBD, clarify usage here)

2. Desirable Communication

All the examples so far presume an environment where it is desirable that any node can communicate with any other node. This is certainly a guiding principle for quickly improving the value of a network: Leaving open the potential to communicate maximizes the potential network effect [METCALFE]. However, firewalls are then widely used to suppress some of these potential communication paths [FIREWALL]. MUD [RFC8520] was designed to aid routers and switches in setting up limited connectivity to this end.

In the MUD architecture, we first build a system that fundamentally supports unlimited connectivity from everyone to everyone and then restrict it based on self-descriptions of the nodes. An alternative approach would be an architecture that does not provide any connectivity unless and until that is authorized by declared communication requirements that have in turn been authorized by some management entity.

2.1. Example: Addressing Desirable Peers Only

There may be other reasons for pursuing an architecture that limits itself to desirable communication only. A trivial, but maybe not overly useful example would be to number all the addresses of correspondent hosts allowed by the self-descriptions and replace the IP addresses in the packets by these numbers.

If this limitation becomes part of the architecture, protocols used inside the Limited Domains could completely get rid of IP addresses and use just the correspondent numbers (possibly packaged in something that looks like an IP address, but is limited in its variability to just encapsulating that number).

The border router would become a NAT, but one that is acting based on extensive, precomputable information about the communication requirements inside the Limited Domain, instead of learning and potentially losing dynamic state that becomes a single point of failure.

Again, this is a trivial example, but it should be sufficient as a motivation for having a look at employing knowledge about the nodes and their communication requirements in a Limited Domain for interconnecting this with the Internet (and thus possibly to like-minded (Limited Domains on the other side of an Internet path).

3. Self-Descriptions

Note that not all of the information that may be needed as a description of Limited Domain nodes can come from MUD-like class definitions. Limited Domain nodes are instantiations of these classes, where the instantiations will differ between each other in details. Different Limited Domain nodes may also be assigned a different purpose in life, causing a need to further parameterize the self-description.

Alongside a discussion of an interconnection architecture that can make use of self-descriptions would therefore need to be a discussion on how to structure self-description classes with this purpose in mind and how to parameterize these and derive description instances.

For the Internet of Things (IoT), additional self-description techniques have been defined that can provide information for Limited Domain network elements. A fully instance-oriented description of an IoT device is provided by a W3C WoT (Web of Things) Thing Description [TD]. W3C WoT is presently in the process of adding a class-based description technique to TDs, the Thing Model (previously called Thing Description Template, TDT) [TD-WD]. The communication patterns offered by the device are detailed in *Protocol Bindings*, which can contain URIs combined with protocol-specific vocabulary ([TD-PB], currently defined for HTTP, CoAP, MQTT). An experimental extension to TDs that enables deriving the configuration of Time-Sensitive Networking (TSN) networks from the self-description is described in [TD-OPC-UA].

An IoT-oriented description technique that, unlike TD, is class-based from the outset is the Semantic Definition Format (SDF) for Data and Interactions of Things [I-D.ietf-asdf-sdf]. A concept similar to WoT Protocol Bindings is not defined yet, but a combination of MUD and SDF descriptions could provide a basic description of a device situated in a Limited Domain.

The Constrained RESTful Environments (CoRE) architecture also provides instance-oriented self-descriptions in the form of the CoRE Link Format [RFC6690], an instance of which is provided by each CoAP server under /.well-known/core. Link-format information, as well as self-describing information in the newer CoRAL format [I-D.ietf-core-coral], can be stored in the CoRE Resource Directory [I-D.ietf-core-resource-directory].

All these potential sources of (self-)description only provide meager information about purpose-in-life, i.e., beyond the intrinsic properties of the device. Obtaining a full description of the communication requirements of a node (including its desirable correspondence nodes) will therefore require additional input, beyond the class-based self-descriptions of the devices.

4. Directions of Work

The above discussion leads us to the following interrelated areas for further exploration:

1. Extending the self-description mechanisms to provide more information that may be useful in a Limited Domain.

2. Merging the self-description information with other configuration/management information (such as purpose-in-life) that may be available for the Limited Domain.

3. Defining Limited Domain architectures that can benefit from information made available by (1) and (2), including defining the operation of network elements and nodes inside the Limited Domain.

4. Defining border network element functionality that makes such a Limited Domain a Globally Interoperable Limited Domain.

5. Defining border network element functionality that makes such a Limited Domain an Externally Interoperable Limited Domain.

6. Discovery between Limited Domains, between Limited Domain nodes (Rendezvous problem); establishment of communications (cf. [RFC8445]).

7. Defining appropriate security workflows and the supporting security mechanisms for items 1 to 6.

8. Addressing operational considerations for items 1 to 7.

9. Addressing privacy considerations for items 1 to 8.

5. IANA Considerations

This document contains no requests to IANA.

6. Security Considerations

The security considerations of [RFC8799] apply.

Item 7 of Section 4 raises the need for security work, one example of which might be:

Self-descriptions of nodes in many cases need to undergo an authorization process before they can be used as the basis of network configuration. The authorization process sketched by [RFC8520] may be too simplistic, in particular the simplified number of stakeholders assumed. The present document is not providing answers in this space, but needs to raise the issue.

7. Informative References

[FIREWALL] Bellovin, S. and W. Cheswick, "Network firewalls", IEEE Communications Magazine Vol. 32, pp. 50-57, DOI 10.1109/35.312843, September 1994, <https://doi.org/10.1109/35.312843>.

[I-D.ietf-asdf-sdf] Koster, M. and C. Bormann, "Semantic Definition Format (SDF) for Data and Interactions of Things", Work in Progress, Internet-Draft, draft-ietf-asdf-sdf-07, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-asdf-sdf-07.txt>.

[I-D.ietf-core-coral] Hartke, K., "The Constrained RESTful Application Language (CoRAL)", Work in Progress, Internet-Draft, draft-ietf-core-coral-03, 9 March 2020, <https://www.ietf.org/archive/id/draft-ietf-core-coral-03.txt>.

[I-D.ietf-core-resource-directory] Amsüss, C., Shelby, Z., Koster, M., Bormann, C., and P. V. D. Stok, "CoRE Resource Directory", Work in Progress, Internet-Draft, draft-ietf-core-resource-directory-28, 7 March 2021, <https://www.ietf.org/archive/id/draft-ietf-core-resource-directory-28.txt>.

[METCALFE] Metcalfe, B., "Metcalfe's Law after 40 Years of Ethernet", Computer Vol. 46, pp. 26-31, DOI 10.1109/mc.2013.374, December 2013, <https://doi.org/10.1109/mc.2013.374>.

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

[RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <https://www.rfc-editor.org/info/rfc6690>.

[RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <https://www.rfc-editor.org/info/rfc6775>.

[RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <https://www.rfc-editor.org/info/rfc8138>.

[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>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/

RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

[RFC8930] Watteyne, T., Ed., Thubert, P., Ed., and C. Bormann, "On Forwarding 6LoWPAN Fragments over a Multi-Hop IPv6 Network", RFC 8930, DOI 10.17487/RFC8930, November 2020, <https://www.rfc-editor.org/info/rfc8930>.

[TD] "Web of Things (WoT) Thing Description", (Link errors corrected 23 June 2020), W3C Recommendation, April 2020, <https://www.w3.org/TR/wot-thing-description/>.

[TD-OPC-UA] Sciullo, L., Bhattacharjee, S., and M. Kovatsch, "Bringing deterministic industrial networking to the W3C web of things with TSN and OPC UA", Proceedings of the 10th International Conference on the Internet of Things, DOI 10.1145/3410992.3410997, October 2020, <https://doi.org/10.1145/3410992.3410997>.

[TD-PB] "Web of Things (WoT) Binding Templates", W3C Working Group Note, January 2020, <https://www.w3.org/TR/wot-binding-templates/>.

[TD-WD] "Web of Things (WoT) Thing Description 1.1", W3C Editor's Draft, May 2021, <https://w3c.github.io/wot-thing-description/>.

Acknowledgments

Adrian Farrel provided substantive comments as well as the basis for Figure 1.

Author's Address

Carsten Bormann
Universität Bremen TZI
Postfach 330440
D-28359 Bremen
Germany

Phone: +49-421-218-63921
Email: cabo@tzi.org