LWIG Working Group C. Bormann
Internet-Draft Universitaet Bremen TZI
Intended status: Informational M. Ersue
Expires: September 29, 2013 Nokia Siemens Networks
March 28, 2013

Terminology for Constrained Node Networks
draft-ietf-lwig-terminology-02

Abstract

The Internet Protocol Suite is increasingly used on small devices with severe constraints, creating constrained node networks. This document provides a number of basic terms that have turned out to be useful in the standardization work for constrained environments.

Table of Contents

1. Introduction

Small devices with limited CPU, memory, and power resources, so called constrained devices (aka. sensor, smart object, or smart device) can constitute a network, becoming "constrained nodes" in that network. Such a network may itself exhibit constraints, e.g. with unreliable or lossy channels, limited and unpredictable bandwidth, and a highly dynamic topology.

Constrained devices might be in charge of gathering information in diverse settings including natural ecosystems, buildings, and factories and send the information to one or more server stations. Constrained devices may work under severe resource constraints such as limited battery and computing power, little memory and insufficient wireless bandwidth, and communication capabilities. Other entities on the network, e.g., a base station or controlling server, might have more computational and communication resources and can support the interaction between the constrained devices and applications in more traditional networks.

Today diverse sizes of constrained devices with different resources and capabilities are becoming connected. Mobile personal gadgets, building-automation devices, cellular phones, Machine-to-machine (M2M) devices, etc. benefit from interacting with other "things" in the near or somewhere in the Internet. With this, the Internet of Things (IoT) becomes a reality, built up out of uniquely identifiable

and addressable objects (things). And over the next decade, this could grow to large numbers [fifty-billion] of Internet-connected constrained devices, greatly increasing the Internet's size and scope.

The present document provides a number of basic terms that have turned out to be useful in the standardization work for constrained environments. The intention is not to exhaustingly cover the field, but to make sure a few core terms are used consistently between different groups cooperating in this space.

2. Terminology

The main focus of this field of work appears to be _scaling_:

o Scaling up Internet technologies to a large number [fifty-billion] of inexpensive nodes, while

o scaling down the characteristics of each of these nodes and of the networks being built out of them, to make this scaling up econmically and physically viable.

The need for scaling down the characteristics of nodes leads to _constrained nodes_.

2.1. Constrained Nodes

The term "constrained node" is best defined on not meeting certain widely held expectations:

Constrained Node: A node where some of the characteristics that are otherwise pretty much taken for granted for Internet nodes in 2013 are not attainable, often due to cost constraints and/or physical constraints on characteristics such as size, weight, and available power.

While this is less than satisfying as a rigorous definition, it is grounded in the state of the art and clearly sets apart constrained nodes from server systems, desktop or laptop computers, powerful mobile devices such as smartphones etc.

(An alternative name, when the properties as a network node are not in focus, is "constrained device".)

There are multiple facets to the constraints on nodes, often applying in combination, e.g.:

o constraints on the maximum code complexity (ROM/Flash);

o constraints on the size of state and buffers (RAM);

o constraints on the available power.

Section 3 defines a small number of interesting classes ("class-N" for N=0,1,2) of constrained nodes focusing on relevant combinations of the first two constraints. With respect to available power, [RFC6606] distinguishes "power-affluent" nodes (mains-powered or regularly recharged) from "power-constrained nodes" that draw their power from primary batteries or using energy harvesting.

The use of constrained nodes in networks often also leads to constraints on the networks themselves. However, there may also be constraints on networks that are largely independent from those of the nodes. We therefore distinguish _constrained networks_ and _constrained node networks_.

2.2. Constrained Networks

We define "constrained network" in a similar way:

Constrained Network: A network where some of the characteristics pretty much taken for granted for Internet link layers in 2013 are not attainable.

Again, there may be several reasons for this:

o cost constraints on the network

o constraints of the nodes (for constrained node networks)

o physical constraints (e.g., power constraints, media constraints such as underwater operation, limited spectrum for very high density).

Constraints may include:

o low achievable bit rate

o high packet loss, packet loss (delivery rate) variability

o severe penalties for using larger packets (e.g., high packet loss due to link layer fragmentation)

o lack of (or severe constraints on) advanced services such as IP multicast

2.2.1. Challenged Networks

A constrained network is not necessarily a _challenged_ network [FALL]:

Challenged Network: A network that has serious trouble maintaining what an application would today expect of the end-to-end IP model, e.g., by:

- not being able to offer end-to-end IP connectivity at all;
- exhibiting serious interruptions in end-to-end IP connectivity;
- exhibiting delay well beyond the MSL defined by TCP.

All challenged networks are constrained networks in some sense, but not all constrained networks are challenged networks. There is no well-defined boundary between the two, though. Delay-Tolerant Networking (DTN) has been designed to cope with challenged networks [RFC4838].

2.3. Constrained Node Networks

Constrained Node Network: A network whose characteristics are influenced by being composed of a significant portion of constrained nodes.

A constrained node network always is a constrained network because of the network constraints stemming from the node constraints, but may also have other constraints that already make it a constrained network.

2.3.1. LLN ("low-power lossy network")

A related term that has been used recently is "low-power lossy network" (LLN). The ROLL working group currently is struggling with its definition [I-D.ietf-roll-terminology]:

> LLN: Low power and Lossy networks (LLNs) are typically composed of many embedded devices with limited power, memory, and processing resources interconnected by a variety of links, such as IEEE 802.15.4 or Low Power WiFi. There is a wide scope of application areas for LLNs, including industrial monitoring, building automation (HVAC, lighting, access control, fire), connected home, healthcare, environmental monitoring, urban sensor networks, energy management, assets tracking and refrigeration.. [sic]

It is not clear that "LLN" is much more specific than "interesting" or "the network characteristics that RPL has been designed for". LLNs do appear to have significant loss at the physical layer, with significant variability of the delivery rate, and some short-term unreliability, coupled with some medium term stability that makes it worthwhile to construct medium-term stable directed acyclic graphs for routing and do measurements on the edges such as ETX [RFC6551]. Actual "low power" does not seem to be required for an LLN [I-D.hui-vasseur-roll-rpl-deployment], and the positions on scaling of LLNs appear to vary widely [I-D.clausen-lln-rpl-experiences].

Also, LLNs seem to be composed of constrained nodes; otherwise operation modes such as RPL's "non-storing mode" would not be sensible. So an LLN seems to be a constrained node network with certain constraints on the network as well.

2.3.2. LoWPAN, 6LoWPAN

One interesting class of a constrained network often used as a constrained node network is the "LoWPAN" [RFC4919], a term inspired from the name of the IEEE 802.15.4 working group (low-rate wireless personal area networks (LR-WPANs)). The expansion of that acronym, "Low-Power Wireless Personal Area Network" contains a hard to justify "Personal" that is due to IEEE politics more than due to an orientation of LoWPANs around a single person. Actually, LoWPANs have been suggested for urban monitoring, control of large buildings, and industrial control applications, so the "Personal" can only be considered a vestige. Maybe the term is best read as "Low-Power Wireless Area Networks" (LoWPANs) [WEI]. Originally focused on IEEE 802.15.4, "LoWPAN" (or when used for IPv6, "6LoWPAN") is now also being used for networks built from similarly constrained link layer technologies [I-D.ietf-6lowpan-btle] [I-D.mariager-6lowpan-v6over-dect-ule].

3. Classes of Constrained Devices

Despite the overwhelming variety of Internet-connected devices that can be envisioned, it may be worthwhile to have some succinct terminology for different classes of constrained devices. In this document, the following class designations may be used as rough indications of device capabilities:

Name	data size (e.g., RAM)	code size (e.g., Flash)
Class 0, C0	<< 10 KiB	<< 100 KiB
Class 1, C1	~ 10 KiB	~ 100 KiB
Class 2, C2	~ 50 KiB	~ 250 KiB

Table 1: Classes of Constrained Devices

As of the writing of this document, these characteristics correspond to distinguishable sets of commercially available chips and design cores for constrained devices. While it is expected that the boundaries of these classes will move over time, Moore's law tends to be less effective in the embedded space than in personal computing devices: Gains made available by increases in transistor count and density are more likely to be invested in reductions of cost and power requirements than into continual increases in computing power.

Class 0 devices are very constrained sensor-like motes. Most likely they will not be able to communicate directly with the Internet in a secure manner. Class 0 devices will participate in Internet communications with the help of larger devices acting as proxies, gateways or servers. Class 0 devices generally cannot be secured or managed comprehensively in the traditional sense. They will be most likely preconfigured and if ever will be reconfigured rarely with a very small data set. For management purposes, they could answer keepalive signals and send on/off or basic health indications.

Class 1 devices cannot easily talk to other Internet nodes employing a full protocol stack such as using HTTP, TLS and related security protocols and XML-based data representations. However, they have enough power to use a protocol stack specifically designed for constrained nodes (e.g., CoAP over UDP) and participate in meaningful conversations without the help of a gateway node. In particular, they can provide support for the security functions required on a large network. Therefore, they can be integrated as fully developed peers into an IP network, but they need to be parsimonious with state

memory, code space, and often power expenditure for protocol and application usage.

Class 2 can already support mostly the same protocol stacks as used on notebooks or servers. However, even these devices can benefit from lightweight and energy-efficient protocols and from consuming less bandwidth. Furthermore, using fewer resources for networking leaves more resources available to applications. Thus, using the protocol stacks defined for very constrained devices also on Class 2 devices might reduce development costs and increase the interoperability.

Constrained devices with capabilities significantly beyond Class 2 devices exist. They are less demanding from a standards development point of view as they can largely use existing protocols unchanged. The present document therefore does not make any attempt to define classes beyond Class 2. These devices can still be constrained by a limited energy supply.

With respect to examining the capabilities of constrained nodes, particularly for Class 1 devices, it is important to understand what type of applications they are able to run and which protocol mechanisms would be most suitable. Because of memory and other limitations, each specific Class 1 device might be able to support only a few selected functions needed for its intended operation. In other words, the set of functions that can actually be supported is not static per device type: devices with similar constraints might choose to support different functions. Even though Class 2 devices have some more functionality available and may be able to provide a more complete set of functions, they still need to be assessed for the type of applications they will be running and the protocol functions they would need. To be able to derive any requirements, the use cases and the involvement of the devices in the application and the operational scenario need to be analyzed. Use cases may combine constrained devices of multiple classes as well as more traditional Internet nodes.

4. Power Terminology

Devices not only differ in their computing capabilities, but also in available electrical power and/or energy. While it is harder to find recognizable clusters in this space, it is still useful to introduce some common terminology.

4.1. Scaling Properties

The power and/or energy available to a device may vastly differ, from kilowatts to microwatts, from essentially unlimited to hundreds of microjoules.

Instead of defining classes or clusters, we propose simply stating one or both of the following quantities in SI units:

o Ps: Sustainable average power available for the device over the time it is functioning (in W).

o Et: Total electrical energy available before the energy source is exhausted (in J).

The value of Et may need to be interpreted in conjunction with an indication over which period of time the value is given; see the next subsection.

4.2. Energy Limitation Classes

As discussed above, some devices are limited in available energy as opposed to (or in addition to) being limited in available power. Where no relevant limitations exist with respect to energy, the device is classified as E3. The energy limitation may be in total energy available in the usable lifetime of the device (e.g. a device with a non-replaceable primary battery, which is discarded when this battery is exhausted), classified as E2. Where the relevant limitation is for a specific period, this is classified as E1, e.g. a limited amount of energy available for the night with a solar-powered device, or for the period between recharges with a device that is manually connected to a charger. Finally, there may be a limited amount of energy available for a specific event, e.g. for a button press in an energy harvesting light switch; this is classified as E0. Note that many E1 devices in a sense also are E2, as the rechargeable battery has a limited number of useful recharging cycles.

In summary, we distinguish:

o E0: Event energy-limited

- o E1: Period energy-limited

- o E2: Lifetime energy-limited

- o E3: No direct quantitative limitations to available energy

4.3. Power Usage Strategies

Especially when wireless transmission media is used, the radio often consumes a big portion of the total energy consumed by the device. Design parameters such as desired range and the spectrum available and bitrate aimed for influence the power consumed during transmission and reception; the duration of transmission and reception (including potential reception) influence the total energy consumption.

Based on the type of the energy source (e.g., battery or mains power) and how often device needs to communicate, it may use different kinds of strategies for power usage and network attachment.

The general strategies for power usage can be described as follows:

Always-on: This strategy is most applicable if there is no reason for extreme measures for power saving. The device can stay on in the usual manner all the time. It may be useful to employ power-friendly hardware or limit the number of wireless transmissions, CPU speeds, and other aspects for general power saving and cooling needs, but the device can be connected to the network all the time.

Always-off: Under this strategy, the device sleeps such long periods at a time that once it wakes up, it makes sense for it to not pretend that it has been connected to the network during sleep: The device re-attaches to the network as it is woken up. The main optimization goal is to minimize the effort during such re-attachment process and any resulting application communications.

If the device sleeps for long periods of time, for infrequent communication the relative increase in energy expenditure during reattachment may be acceptable.

Low-power: This strategy is most applicable to devices that need to operate on a very small amount of power, but still need to be able to communicate on a relatively frequent basis. This implies that extremely low power solutions needs to be used for the hardware, chosen link layer mechanisms, and so on. Typically, given the small amount of time between transmissions, despite their sleep state these devices retain some form of network attachment to the

network. Techniques used for minimizing power usage for the network communications include minimizing any work from re-establishing communications after waking up, tuning the frequency of communications, and other parameters appropriately.

In summary, we distinguish the power usage strategies:

- S0: Always-off
- S1: Low-power
- S2: Always-on

5. Security Considerations

This draft introduces common terminology that does not raise any new security issue.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

Ari Keranen, Dominique Barthel, and Peter van der Stok provided useful comments.

Peter van der Stok insisted that we should have power terminology, hence Section 4. The text for Section 4.3 is mostly lifted from [I-D.arkko-lwig-cellular] and has been adapted for this document.

8. Informative References

[FALL] Fall, K., "A Delay-Tolerant Network Architecture for Challenged Internets", SIGCOMM 2003, 2003.

[I-D.arkko-lwig-cellular]
Arkko, J., Eriksson, A., and A. Keraenen, "Building Power-Efficient CoAP Devices for Cellular Networks", draft-arkko-lwig-cellular-00 (work in progress), February 2013.

[I-D.clausen-lln-rpl-experiences]
Clausen, T., Verdiere, A., Yi, J., Herberg, U., and Y. Igarashi, "Observations of RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-clausen-lln-rpl-experiences-06 (work in progress), February 2013.

[I-D.hui-vasseur-roll-rpl-deployment]

Vasseur, J., Hui, J., Dasgupta, S., and G. Yoon, "RPL deployment experience in large scale networks", draft-hui-vasseur-roll-rpl-deployment-01 (work in progress), July 2012.

[I-D.ietf-6lowpan-btle]
Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "Transmission of IPv6 Packets over BLUETOOTH Low Energy", draft-ietf-6lowpan-btle-12 (work in progress), February 2013.

[I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-12 (work in progress), March 2013.

[I-D.mariager-6lowpan-v6over-dect-ule]
Mariager, P. and J. Petersen, "Transmission of IPv6 Packets over DECT Ultra Low Energy", draft-mariager-6lowpan-v6over-dect-ule-02 (work in progress), May 2012.

[RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", RFC 4838, April 2007.

[RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, August 2007.

[RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, March 2012.

[RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, May 2012.

[WEI] Shelby, Z. and C. Bormann, "6LoWPAN: the Wireless Embedded Internet", ISBN 9780470747995, 2009.

[fifty-billion]
Ericsson, -., "More Than 50 Billion Connected Devices", Ericsson White Paper 284 23-3149 Uen, February 2011, <http://www.ericsson.com/res/docs/whitepapers/wp-50-billions.pdf>.

Authors' Addresses

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

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

Mehmet Ersue
Nokia Siemens Networks

Email: mehmet.ersue@nsn.com