CoRE Working Group A. Rahman, Ed.
Internet-Draft InterDigital Communications, LLC
Intended status: Informational E. Dijk, Ed.
Expires: September 10, 2012 Philips Research
March 9, 2012

Group Communication for CoAP
draft-ietf-core-groupcomm-01

Abstract

This is a working document intended to develop draft text for the CoAP protocol specification in the area of group communication. A solution based on IP multicast is proposed and detailed. Also, guidance is provided for deployment in various constrained network topologies.

Status of this Memo

Copyright Notice

Table of Contents

1. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The following are definitions of specific terminology used in this draft.

Group Communication: A source node sends a message to more than one destination node, where all destinations are identified to belong to a specific group. The set of source nodes and/or the set of destination nodes may consist of an arbitrary mix of constrained and non-constrained nodes.

Multicast: Sending a message to multiple receiving nodes simultaneously. Typically, this is done as part of a group communication process. There are various options to implement multicast including layer 2 (Media Access Control) or layer 3 (IP) mechanisms.

IP Multicast: A specific multicast solution based on the use of IP multicast addresses as defined in "IANA Guidelines for IPv4 Multicast Address Assignments" [RFC5771] and "IP Version 6 Addressing Architecture" [RFC4291].

Low power and Lossy Network (LLN): LLNs are made up of constrained devices. These devices may be interconnected by a variety of links, such as IEEE 802.15.4, Bluetooth, WiFi, wired or low power power-line communication links.

2. Introduction

2.1. Background

The Constrained Application Protocol (CoAP) is an application protocol (analogous to HTTP) for resource constrained devices operating in an IP network [I-D.ietf-core-coap]. Constrained devices can be large in number, but are often highly correlated to each other (e.g. by type or location). For example, all the light switches in a building may belong to one group and all the thermostats belong to another group. All the smart meters in the same region can belong to a group as well. Groups may be composed by function; for example, the group "all lights in building one" may consist of the groups "all lights on floor one of building one", "all lights on floor two of building one", etc. Groups may be preconfigured or dynamically formed. If information needs to be sent to or received from a group

of devices, group communication mechanisms can improve efficiency and latency of communication and reduce bandwidth requirements for a given application.

2.2. Problem Statement and Scope

In this draft, we address the issues related to group communication in detail, with requirements, use cases, proposed solutions and analysis of their impact to the CoAP protocol and to implementations. We assume that all, or a substantial part of, CoAP devices participating in group communication are constrained devices (e.g. Low Power and Lossy Network (LLN) devices). The guiding principle is to apply wherever possible existing IETF protocols to achieve group communication functionality. In many cases the contribution of this document lies in explaining how existing mechanisms may be used to together fulfill CoAP group communication needs for specific use cases.

2.3. Potential Solutions for Group Communication

The classic concept of group communications is that of a single source distributing content to multiple recipients that are all part of a group, as shown in the example sequence diagram in Figure 1. Also shown there is the pre-requisite step of forming the group before content can be distributed to it. The source may be either a member or non-member of the group.

Group communication solutions have evolved from "bottom" to "top", i.e., from the network layer (IP multicast) to application layer group communication, also referred to as application layer multicast. A study published in 2005 [Lao05] identified new solutions in the "middle" (referred to as overlay multicast) that utilize an infrastructure based on proxies.

Each of these classes of solutions may be compared [Lao05] using metrics such as link stress and level of host complexity [Banerjee01]. The results show for a realistic internet topology that IP Multicast is the most resource-efficient, with the downside being that it requires the most effort to deploy in the infrastructure. IP Multicast is the solution recommended by this draft with detailed analysis and guidance for this choice being provided in the following sections.

```
                                             Group
          Node 1           Node 2            Coordinator       Node 3
            |                 |               |                   |
            |   REQUEST       |               |                   |
            |  (Join Group X) |               |                   |
            |-----------------|------------- >|                   |
            |   RESPONSE      |               |                   |
            |< ---------------|---------------|                   |
            |                 |               |                   |
            |                 |    REQUEST    |                   |
            |                 |(Join Group X) |                   |
            |                 |------------- >|                   |
            |                 |   RESPONSE    |                   |
            |                 |< -------------|                   |
            |                 |               |     REQUEST       |
            |                 |               |    (Send to       |
            |                 |               |     Group X )     |
            |                 |               |< -----------------|
            |                 |               |                   |
            |                 |          Map to                   |
            |                 |      Group X addresses            |
            |                 |               |                   |
            |   REQUEST (to multicast addr)   |                   |
            |< ---------------|< -------------|                   |
            |                 |               |                   |
            |     (optional) RESPONSE         |                   |
            |                 |------------- >|                   |
            |-----------------|-------------->|                   |
            |                 |               |       RESPONSE    |
            |                 |               |----------------- >|
            |                 |               |                   |
```

Figure 1: Example Group Communication Concept

2.4. CoAP-Observe Alternative to Group Communication

The CoAP Observation extension [I-D.ietf-core-observe] can be used as a simple (but very limited) alternative for group communication. A group in this case consists of a CoAP server hosting a specific resource, plus all CoAP clients observing that resource. The server is in that case the only group member that can send a group message. It does this by modifying the state of a resource under observation and subsequently notifying its observers of the change. Serial unicast is used for sending the notifications. This approach can be a simple alternative for networks where IP multicast is not available or too expensive.

The CoAP-Observe approach is unreliable in the sense that, even though Confirmable CoAP messages may be used, there are no guarantees that an update will be received. For example, a client may believe it is observing a resource while in reality the server rebooted and lost its listener state.

3. Use Cases and Requirements

3.1. Use Cases

The use of CoAP group communication is shown in the context of several use cases. The following use cases are identified at this point:

- o Lighting Control: synchronous operation of a group of 6LoWPAN IPv6-connected lights

- o Discovery: discovering CoAP devices and the Resource and Services they offer

- o Parameter Update: updating parameters/settings simultaneously in a large group of devices in a building/campus control ([I-D.vanderstok-core-bc]) application

In a future version of this document, more use cases should be added and described in more detail.

3.2. Requirements

Requirements that a CoAP group communication solution should fulfill can be found in existing documents ([RFC5867], [I-D.ietf-6lowpan-routing-requirements], [I-D.vanderstok-core-bc], and [I-D.shelby-core-coap-req]). Below, a set of high-level requirements is listed that a group communication solution should ideally fulfill. In practice, all these requirements can never be satisfied at once in an LLN context. Furthermore, different use cases will have different needs i.e. an elaboration of a subset of below requirements.

3.2.1. Background

The requirements for CoAP are documented in [I-D.shelby-core-coap-req]. In this draft, we focus and expand discussions on the requirements pertaining to CoAP "group communication" and "multicast" support as stated in [I-D.shelby-core-coap-req]:

REQ 9: CoAP will support a non-reliable IP multicast message to be sent to a group of Devices to manipulate a resource on all the Devices simultaneously. The use of multicast to query and advertise descriptions must be supported, along with the support of unicast responses.

Currently, the CoAP protocol [I-D.ietf-core-coap] supports unreliable IP multicast using UDP. It defines the unreliable multicast operation as follows in Section 4.5:

"CoAP supports sending messages to multicast destination addresses. Such multicast messages MUST be Non-Confirmable. Some mechanisms for avoiding congestion from multicast requests are being considered in [I-D.eggert-core-congestion-control]."

Additional requirements were introduced in [I-D.vanderstok-core-bc] driven by quality of experience issues in commercial lighting; the need for large numbers of devices to respond with near simultaneity to a command (multicast PUT), and for that command to be received reliably (reliable multicast).

3.2.2. General Requirements

A CoAP group communication solution should (ideally) meet the following general requirements:

GEN-REQ 1: Optional Reliability: the application can select between unreliable group communication and reliable group communication.

GEN-REQ 2: Efficiency: delivers messages more efficiently than a "serial unicast" solution. Provides a balance between group data traffic and control overhead.

GEN-REQ 3: Low latency: deliver a message as quickly as possible.

GEN-REQ 4: Synchrony: allows near-simultaneous modification of a resource on all devices in a target group, providing a perceived effect of synchrony or simultaneity. For example a specified time span D such that a message is delivered to all destinations in a time interval [t,t+D].

GEN-REQ 5: Ordering: message ordering may be required for reliable group communication use cases.

GEN-REQ 6: Security: see Section 6 for security requirements for group communication.

GEN-REQ 7: Flexibility: support for one or many source(s), both dense and sparse networks, for high or low listener density, small or large number of groups, and multi-group membership.

GEN-REQ 8: Robust group management: functionality to join groups, leave groups, view group membership, and persistent group membership in failure or sleeping node situations.

GEN-REQ 9: Network layer independence: a solution is independent from specific unicast and/or IP multicast routing protocols.

GEN-REQ 10: Minimal specification overhead: a group communication solution should preferably re-use existing/established (IETF) protocols that are suitable for LLN deployments, instead of defining new protocols from scratch.

GEN-REQ 11: Minimal implementation overhead: e.g. a solution allows to re-use existing (software) components that are already present on constrained nodes such as (typical) 6LoWPAN/CoAP nodes.

GEN-REQ 12: Mixed backbone/LLN topology support: a solution should work within a single LLN, and in combined LLN/backbone network topologies, including multi-LLN topologies. Both the senders and receivers of CoAP group messages may be attached to different network links or be part of different LLNs, possibly with routers or switches in between group members. In addition, different routing protocols may operate on the LLN and backbone networks. Preferably a solution also works with existing, common backbone IP infrastructure (e.g. switches or routers).

GEN-REQ 13: CoAP Proxying support: a CoAP proxy can handle distribution of a message to a group on behalf of a (constrained) CoAP client.

GEN-REQ 14: Suitable for operation on LLNs with constrained nodes.

3.2.3. Security Requirements

Security for group communications at the IP level has been studied extensively in the IETF MSEC (Multicast Security) WG, and to a lesser extent in the IRTF SAMRG (Scalable Adaptive Multicast Research Group). In particular, [RFC3740], [RFC5374] and [RFC4046] are very instructive. A set of requirements for securing group communications in CoAP were derived from a study of these previous investigations as well as understanding of CoAP specific needs. These are listed below.

A CoAP group communication solution should (ideally) meet the following security requirements:

SEC-REQ 1: Group communications data encryption: Important CoAP group communications shall be encrypted (using a group key) to preserve confidentiality. It shall also be possible to send CoAP group communications in the clear (i.e. unencrypted) for low value data.

SEC-REQ 2: Group communications source data authentication: Important CoAP group communications shall be authenticated by verifying the source of the data (i.e. that it was generated by a given and trusted group member). It shall also be possible to send unauthenticated CoAP group communications for low value data.

SEC-REQ 3: Group communications limited data authentication: Less important CoAP group communications shall be authenticated by simply verifying that it originated from one of the group members (i.e. without explicitly identifying the source node). This is a weaker requirement (but simpler to implement) than REQ2. It shall also be possible to send unauthenticated CoAP group communications for low value data.

SEC-REQ 4: Group key management: There shall be a secure mechanism to manage the cryptographic keys (e.g. generation and distribution) belonging to the group; the state (e.g. current membership) associated with the keys; and other security parameters.

SEC-REQ 5: Use of Multicast IPSec: The CoAP protocol [I-D.ietf-core-coap] allows IPSec to be used as one option to secure CoAP. If IPSec is used as a way to security CoAP communications, then multicast IPSec [RFC5374] should be used for securing CoAP group

communications.

SEC-REQ 6: Independence from underlying routing security: CoAP group communication security shall not be tied to the security of underlying routing and distribution protocols such as PIM [RFC4601] and RPL [I-D.ietf-roll-rpl]. Insecure or inappropriate routing (including IP multicast routing) may cause loss of data to CoAP but will not affect the authenticity or secrecy of CoAP group communications.

SEC-REQ 7: Interaction with HTTPS: The security scheme for CoAP group communications shall account for the fact that it may need to interact with HTTPS (Hypertext Transfer Protocol Secure) when a transaction involves a node in the general Internet (non-constrained network) communicating via a HTTP-CoAP proxy.

4. IP Multicast Solution

4.1. Introduction

IP Multicast protocols have been evolving for decades, resulting in proposed standards such as Protocol Independent Multicast - Sparse Mode (PIM-SM) [RFC4601]. Yet, due to various technical and marketing reasons, IP Multicast is not widely deployed on the general Internet. However, IP Multicast is popular in specific deployments such as in enterprise networks (e.g. for video conferencing or general IP multicast PC applications within a single LAN broadcast domain) and carrier IPTV deployments. The packet economy and minimal host complexity of IP multicast make it attractive for group communication in constrained environments. IP multicast is the recommended solution for CoAP group communications.

4.2. Multicast Listener Discovery (MLD) & Multicast Router Discovery (MRD)

In order to extend the scope of IP multicast beyond link-local, an IP multicast routing protocol has to be active in routers on an LLN. To achieve efficient multicast routing (i.e. avoid always flooding multicast IP packets), routers have to learn which hosts need to receive packets addressed to specific IP multicast destinations.

The Multicast Listener Discovery (MLD) protocol [RFC3810] (or its IPv4 pendant IGMP) is today the method of choice used by an (IP multicast enabled) router to discover the presence of multicast listeners on directly attached links, and to discover which multicast

addresses are of interest to those listening nodes. MLD was specifically designed to cope with fairly dynamic situations in which multicast listeners may join and leave at any time.

IGMP/MLD Snooping is a technique implemented in some corporate LAN routing/switching devices. An MLD snooping switch listens to MLD State Change Report messages from MLD listeners on attached links. Based on this, the switch learns on what LAN segments there is interest for what IP multicast traffic. If the switch receives at some point an IP multicast packet, it uses the stored information to decide onto which LAN segment(s) to send the packet. This improves network efficiency compared to the regular behavior of forwarding every incoming multicast packet onto all LAN segments. An MLD snooping switch may also send out MLD Query messages (which is normally done by an MLD Router) if no MLD router is present.

The Multicast Router Discovery (MRD) protocol [RFC4286] defines a way to discover multicast routers, for the purpose of using this information by IGMP/MLD snooping devices.

[I-D.ietf-multimob-igmp-mld-tuning] discusses optimal tuning of the parameters of MLD for routers for mobile and wireless networks. These guidelines may be useful when implementing MLD in LLNs.

4.3. Group URIs and IP Multicast Addresses

An approach to map group authorities onto IP multicast addresses using DNS was proposed in [I-D.vanderstok-core-bc]. Based on this, examples of group URI naming (and scoping) for a building control application are shown below. Group URIs MUST follow the URI syntax defined in [RFC3986].

```
URI authority                 Targeted group
all.bldg6.example.com         "all nodes in building 6"
all.west.bldg6.example.com    "all nodes in west wing, building 6"
all.floor1.west.bldg6.examp... "all nodes in floor 1, west wing,
                                building 6"
all.bu036.floor1.west.bldg6... "all nodes in office bu036, floor1,
                                west wing, building 6"
```

The authority portion of the URI is used to identify a node (or group) and the resulting DNS name is bound to a unicast or multicast IP address. Each example group URI shown above can be mapped to a unique multicast IP address. This may be a site-local or global address allocated according to [RFC3956], [RFC3306] or [RFC3307].

4.4. Group Discovery and Member Discovery

CoAP defines a resource discovery capability but, in the absence of a standardized group communication infrastructure, it is limited to link-local scope IP multicast; examples may be found in [I-D.ietf-core-link-format]. A service discovery capability is required to extend discovery to other subnets and scale beyond a certain point, as originally proposed in [I-D.vanderstok-core-bc]. Discovery includes both discovering groups (e.g. find a group to join or send a multicast message to) and discovering members of a group (e.g. to address selected group members by unicast).

4.4.1. DNS-SD

DNS-based Service Discovery [I-D.cheshire-dnsext-dns-sd] defines a conventional way to configure DNS PTR, SRV, and TXT records to enable enumeration of services, such as services offered by CoAP nodes, or enumeration of all CoAP nodes, within specified subdomains. A service is specified by a name of the form <Instance>.<ServiceType>.<Domain>, where the service type for CoAP nodes is _coap._udp and the domain is a DNS domain name that identifies a group as in the examples above. For each CoAP end-point in a group, a PTR record with the name _coap._udp and/or a PTR record with the name _coap._udp.<Domain> is defined and it points to an SRV record having the <Instance>.<ServiceType>.<Domain> name.

All CoAP nodes in a given subdomain may be enumerated by sending a query for PTR records named _coap._udp to the authoritative DNS server for that zone. A list of SRV records is returned. Each SRV record contains the port and host name (AAAA record) of a CoAP node. The IP address of the node is obtained by resolving the host name. DNS-SD also specifies an optional TXT record, having the same name as the SRV record, which can contain "key=value" attributes. This can be used to store information about the device, e.g. schema=DALI, type=switch, group=lighting.bldg6, etc.

Another feature of DNS-SD is the ability to specify service subtypes using PTR records. For example, one could represent all the CoAP groups in a subdomain by PTR records with the name _group._sub._coap._udp or alternatively _group._sub._coap._udp.<Domain>.

4.4.2. CoRE Resource Directory

CoRE Resource Directory [I-D.shelby-core-resource-directory] defines the concept of a Resource Directory (RD) server where CoAP servers can register their resources offered and CoAP clients can discover these resources by querying the RD server. RD syntax can be mapped

to DNS-SD syntax and vice versa [I-D.lynn-core-discovery-mapping], such that the above approach can be reused for group discovery and group member discovery.

Specifically, the Domain (d) parameter can be set to the group URI by an end-point registering to the RD. If an end-point wants to join multiple groups, it has to repeat the registration process for each group it wants to join.

4.5. Group Resource Manipulation

Group communications shall only be used for idempotent messages (i.e. CoAP GET, PUT, DELETE). Group communications shall NOT be used for non-idempotent messages (i.e. CoAP POST). The CoAP messages that are sent via group communications shall only be of the Non-Confirmable type. A response may be sent back to the group message (.e.g "Response 2.01" to a group GET request). This will typically require a CoAP proxy in the message processing path to process the multiple responses. See also Section 5.2.

Ideally, all nodes in a given group (defined by its multicast IP address) must receive the same request with high probability. This will not be the case if there is diversity in the authority port (i.e. a diversity of dynamic port addresses across the group) or if the targeted resource is located at different paths on different nodes. Extending the definition of group membership to include port and path discovery is not desirable.

Therefore, some measures must be present to ensure uniformity in port number and resource name/location within a group.

A first solution in this respect is to couple groups to service descriptions in DNS (using DNS-SD as in Section 4.4 and [I-D.vanderstok-core-bc]). A service description for a multicast group may have a TXT record in DNS defining a schema X (e.g. "schema=DALI"), which defines by service standard X (e.g. "DALI") which resources a node supporting X MUST have. Therefore a multicast source can safely refer to all resources with corresponding operations as prescribed by standard X. For port numbers (which can be found using DNS-SD also) the same holds. Alternatively, only the default CoAP port may be used in all CoAP multicast requests.

A second solution is to impose the following restrictions, e.g. for groups not found using, or advertised in, DNS-SD:

- All CoAP multicast requests MUST be sent to the well-known CoAP port.

- o All CoAP multicast requests SHOULD operate on /.well-known/core URIs

4.6. Congestion Control

CoAP requests may be multicast, resulting a multitude of replies from different nodes, potentially causing congestion. [I-D.eggert-core-congestion-control] suggests to conservatively control sending multicast requests.

CoAP already addresses the congestion problem to some extent by requiring all multicast CoAP requests to be Non-Confirmable. In CoAP a MAX_RETRANSMIT value set by default to 4 is used for retransmission of Confirmable messages, but since CoAP multicast messages are Non-Confirmable their effective retransmission value is 0. However, as responses to multicast requests SHOULD be sent ([I-D.ietf-core-coap]), using CoAP multicast still may lead to congestion issues.

Various means can be implemented to prevent congestion. For a multicast request that leads to the sending of a response by a server, CoAP currently recommends a required random delay, within a specified TIMEOUT period, before the server can send the response. In order to cope with the different requirements for TIMEOUT imposed by different use cases and network topologies, one recommended approach is to define a CoAP Option via which a CoAP client can indicate a preference for TIMEOUT for a specific response. This Option proposal will be done in a separate draft.

4.7. CoAP Multicast and HTTP Unicast Interworking

Within the constrained network, CoAP runs over UDP for which IP multicast is supported. In a non-constrained network (i.e. general Internet), HTTP over TCP is used for which IP multicast is not supported. Therefore a CoAP/HTTP Proxy node that supports group communication needs to have functionalities to support interworking of unicast and multicast. One possible way of operation of the Proxy is illustrated in Figure 2. Note that this topic is covered in more detail in [I-D.castellani-core-http-mapping].

```
           CoAP             CoAP           CoAP/HTTP            HTTP
          Node 1           Node 2           Proxy           Node 3
            |                 |               |                   |
            |   REQUEST       |               |                   |
            |  (Join Group X) |               |                   |
            |-----------------|------------- >|                   |
            |   RESPONSE      |               |                   |
            |< ---------------|---------------|                   |
            |                 |               |                   |
            |                 |    REQUEST    |                   |
            |                 | (Join Group X)|                   |
            |                 |------------- >|                   |
            |                 |   RESPONSE    |                   |
            |                 |< -------------|                   |
            |                 |               |                   |
            |                 |               |                   |
            |                 |               |  HTTP REQUEST     |
            |                 |               |    (URI to        |
            |                 |               |   unicast addr)   |
            |                 |               |< -----------------|
            |                 |               |                   |
            |                 |          Map URI                  |
            |                 |     to Group X multicast address  |
            |                 |               |                   |
            |   REQUEST (to multicast addr)   |                   |
            |< ---------------|< -------------|                   |
            |                 |               |                   |
            |                 |               |                   |
            |     (optional) RESPONSE         |                   |
            |                 |------------- >|                   |
            |-----------------|-------------->|                   |
            |                 |               |  HTTP RESPONSE    |
            |                 |               |----------------- >|
            |                 |               |                   |
```

Figure 2: CoAP Multicast and HTTP Unicast Interworking

Note that Figure 2 illustrates the case of IP multicast as the underlying group communications mechanism.

A key point in Figure 2 is that the incoming HTTP Request (from node 3) will carry a URI (with the HTTP scheme) that resolves in the general Internet to the proxy node. At the proxy node, the URI will then possibly be mapped (as detailed in [I-D.castellani-core-http-mapping]) and again resolved (with the CoAP scheme) to an IP multicast destination. This may be accomplished, for example, by using DNS-SD (Section 4.4). The proxy node will then

IP multicast the CoAP Request (corresponding to the received HTTP Request) to the appropriate nodes (i.e. nodes 1 and 2).

In terms of the HTTP Response, Figure 2 illustrates that it will be generated by the proxy node based on aggregated responses of the CoAP nodes and sent back to the client in the general Internet that sent the HTTP Request (i.e. node 1). In [I-D.castellani-core-http-mapping] the HTTP Response that the Proxy may use to aggregate multiple CoAP responses is described in more detail. So in terms of overall operation, the CoAP proxy can be considered to be a "non-transparent" proxy according to [RFC2616]. Specifically, [RFC2616] states that a "non-transparent proxy is a proxy that modifies the request or response in order to provide some added service to the user agent, such as group annotation services, media type transformation, protocol reduction or anonymity filtering."

An alternative to the above is using a Forward Proxy. In this case, the CoAP request URI could be carried in the HTTP Request Line (as defined in [I-D.ietf-core-coap] Section 8) in a HTTP request sent to the IP address of the Proxy.

5. Deployment Guidelines

5.1. Overview

We recommend to use IP multicast as outlined in Section 4 as the base solution for CoAP Group Communication, provided that the use case and network characteristics allow this. It has the advantage that it re-uses the IP multicast suite of protocols and can operate even if group members are distributed over both constrained and non-constrained network segments. Still, this approach may require specifying or implementing additional IP Multicast functionality in an LLN, in a backbone network, or in both - this will be evaluated in more detail in this section.

5.2. Example Lighting Use Case

We first present an example use case to illustrate the overall steps in an IP Multicast based CoAP Group Communication solution. We assume the following network configuration for this example (see Figure 3):

1) A large room (Room-A) with three lights (Light-1, Light-2, Light-3) controlled by a Light Switch. The devices are organized into two 6LoWPAN subnets.

2) Light-1 and the Light Switch are connected to a router (Rtr-1) which is also a CoAP Proxy and a 6LoWPAN Border Router (6LBR).

3) Light-2 and the Light-3 are connected to another router (Rtr-2) which is also a CoAP Proxy and a 6LBR.

4) The routers are connected to a an IPv6 network backbone which is also multicast enabled. In the general case, this means the network backbone and 6LBRs support a PIM based multicast routing protocol, and MLD for forming groups. In a limited case, if the network backbone is one link, then the routers only have to support MLD-snooping for the example use case to work.

```
                                                                 Network
                                                                Backbone
                                                                       |
      ################################################                 |
      #                                       Room-A #                 |
      #         **********************               #                 |
      #       **     LoWPAN-1         **             #                 |
      #     *                            *           #                 |
      #    *     +----------+             *          #                 |
      #   *      |  Light   |-------+      *         #                 |
      #  *       |  Switch  |       |       *        #                 |
      #  *       +----------+  +---------+  *        #                 |
      #  *                     |  Rtr-1  |-----------------------------|
      #  *                     +---------+  *        #                 |
      #  *       +----------+        |      *        #                 |
      #   *      |  Light-1 |--------+     *         #                 |
      #    *     +----------+             *          #                 |
      #     *                            *           #                 |
      #       **                      **             #                 |
      #         **********************               #                 |
      #                                              #                 |
      #                                              #                 |
      #        **********************                #                 |
      #       **     LoWPAN-2         **             #                 |
      #     *                            *           #                 |
      #    *     +----------+             *          #                 |
      #   *      |  Light-2 |-------+      *         #                 |
      #  *       |          |       |       *        #                 |
      #  *       +----------+  +---------+  *        #                 |
      #  *                     |  Rtr-2  |-----------------------------|
      #  *                     +---------+  *        #                 |
      #  *       +----------+        |      *        #                 |
      #   *      |  Light-3 |--------+     *         #                 |
      #    *     +----------+             *          #                 |
      #     *                            *           #                 |
      #       **                      **             #                 |
      #         **********************               #                 |
      #                                              #                 |
     #################################################                 |
                                                                       |
                                           +--------+                  |
                                           |  DNS   |------------------|
                                           | Server |
                                           +--------+
```

Figure 3: Network Topology of a Large Room (Room-A)

The corresponding protocol flow for an IP Multicast based CoAP Group Communication solution for the network shown in Figure 3 is shown in sequence in Figure 4, Figure 5, and Figure 6. We assume the following steps occur before the illustrated flow:

1) Startup phase: 6LoWPANs are formed. IPv6 addresses assigned to all devices. The CoAP network is formed.

2) Commissioning phase (by applications): The IP multicast address of the group (Room-A-Lights) has been set in all the Lights. The URI of the group (Room-A-Lights) has been set in the Light Switch.

The indicated MLD Report messages are link-local multicast. In each LoWPAN, it is assumed that a multicast routing protocol in 6LRs will propagate the Join information over multiple hops to the 6LBR.

```
                                    Light      Rtr-1     Rtr-2   Network
   Light-1   Light-2    Light-3     Switch    (CoAP     (CoAP   Backbone
    |          |          |          |         Proxy)    Proxy)       |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    | MLD Report: Join    |          |          |          |          |
    | Group (Room-A-Lights)          |          |          |          |
    |------------------------------------------>|          |          |
    |          |          |          |          |MLD Report: Join     |
    |          |          |          |          |Group (Room-A-Lights)|
    |          |          |          |          |-------------------->|
    |          |          |          |          |          |          |
    |          | MLD Report: Join    |          |          |          |
    |          | Group (Room-A-Lights)          |          |          |
    |          |------------------------------------------>|          |
    |          |          |          |          |          |          |
    |          |          | MLD Report: Join    |          |          |
    |          |          | Group (Room-A-Lights)          |          |
    |          |          |------------------------------->|          |
    |          |          |          |          |          |          |
    |          |          |          |          |MLD Report: Join     |
    |          |          |          |          |Group (Room-A-Lights)|
    |          |          |          |          |          |--------->|
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
```

Figure 4: Joining Groups in a Large Room

```
                                    Light      Rtr-1     Rtr-2   Network
   Light-1   Light-2    Light-3     Switch    (CoAP     (CoAP   Backbone
    |          |          |          |         Proxy)    Proxy)       |
    |          |          |          |          |          |          |
    |          |          ***********************          |          |
    |          |          *   User flips on     *          |          |
    |          |          *   light switch to   *          |          |
    |          |          *   turn on all the   *          |          |
    |          |          *   lights in Room A  *          |          |
    |          |          ***********************          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          | COAP NON (PUT       |          |          |
    |          |          |          (Proxy-URI |          |          |
    |          |          |          (URI for Room-A-Lights))         |
    |          |          |          turn on lights)       |          |
    |          |          |          |--------->|          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          |          |     Request DNS resolution of  |
    |          |          |          |     URI for Room-A-Lights      |
    |          |          |          |          |-------------------->|
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          |          |     DNS returns: AAAA          |
    |          |          |          |     Group (Room-A-Lights)      |
    |          |          |          |     IPv6 multicast address     |
    |          |          |          |          |<--------------------|
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          |          | COAP NON (Put                  |
    |          |          |          |          (URI Path)            |
    |          |          |          |          turn on lights)       |
    |          |          |          |  * Destination IP Address =    |
    |          |          |          |       IP multicast address     |
    |          |          |          |       for Group (Room-A-Lights)|
    |          |          |          |  * Originating IP Address =    |
    |          |          |          |        RTR-1                   |
    |          |          |          |          |-------------------->|
    |<------------------------------------------|          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |<---------|
    |          |<---------|<-------------------------------|          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
```

Figure 5: Sending Multicast Message in a Large Room

```
                                      Light      Rtr-1     Rtr-2   Network
   Light-1   Light-2    Light-3     Switch    (CoAP     (CoAP   Backbone
     |          |          |          |         Proxy)    Proxy)        |
     |          |          |          |          |          |          |
    ***********************           |          |          |          |
    *   Lights in Room-A  *           |          |          |          |
    *   turn on (nearly   *           |          |          |          |
    *   simultaneously)   *           |          |          |          |
    ***********************           |          |          |          |
     |          |          |          |          |          |          |
     |          |          |          |          |          |          |
     |     COAP NON (Response         |          |          |          |
     |                (Success))      |          |          |          |
     |------------------------------------------>|          |          |
     |          |          |          |          |          |          |
     |          |          |          |          |          |          |
     |          COAP NON (Response    |          |          |          |
     |                    (Success))  |          |          |          |
     |          |------------------------------->|          |          |
     |          |          |          |          |          |          |
     |          |          |          |          |          |          |
     |          |        COAP NON (Response      |          |          |
     |          |                  (Success))    |          |          |
     |          |          |-------------------->|          |          |
     |          |          |          |          |          |          |
     |          |          |          ******************************   |
     |          |          |          *  Rtr-1 as CoAP Proxy       *   |
     |          |          |          *  processes all reponses    *   |
     |          |          |          *  to multicast message      *   |
     |          |          |          *  and formulates one        *   |
     |          |          |          *  consolidated response     *   |
     |          |          |          *  to originator             *   |
     |          |          |          ******************************   |
     |          |          |          |          |          |          |
     |          |          |      COAP NON (Response                   |
     |          |          |           (Success))                      |
     |          |          |          |<---------|          |          |
     |          |          |          |          |          |          |
```

         Figure 6: Sending Response to Mullticast in a Large Room

5.3. Implementation in Target Network Topologies

This section looks in more detail how an IP Multicast based solution can be deployed onto the various network topologies that we consider important for group communication use cases. Note that the chosen solution of IP Multicast for CoAP group communication works mostly

independently from the underlying network topology and its specific IP multicast implementation.

Starting from the simplest case of a single LLN topology, we move to more complex topologies involving a backbone network or multiple LLNs. With "backbone" we refer here typically to a corporate LAN or VLAN, which constitutes a single broadcast domain by design. It could also be an in-home network. A multi-link backbone is also possible, if there is proper IP multicast routing or forwarding configured between these links. (The term 6LoWPAN Border Router or "6LBR" is used here for a border router, though our evaluation is not necessarily restricted to 6LoWPAN networks.)

5.3.1. Single LLN Topology

The simplest topology is a single LLN, where all the IP multicast source(s) and destinations are constrained nodes within this same LLN. Possible implementations of IP multicast routing and group administration for this topology are listed below.

5.3.1.1. Mesh-Under Multicast Routing

The LLN may be set up in either a mesh-under or a route-over configuration. In the former case, the mesh routing protocol should take care of routing IP multicast messages throughout the LLN.

Because conceptually all nodes in the LLN are attached to a single link, there is in principle no need for nodes to announce their interest in multicast IP addresses via MLD (see Section 4.2). A multicast message to a specific IP destination, which is delivered to all 6LoWPAN nodes by the mesh routing algorithm, is accepted by the IP network layer of that node only if it is listening on that specific multicast IP address and port.

5.3.1.2. RPL Multicast Routing

The RPL routing protocol for LLNs provides support for routing to multicast IP destinations (Section 12 of [I-D.ietf-roll-rpl]). Like regular unicast destinations, multicast destinations are advertised by nodes using RPL DAO messages. This functionality requires "Storing mode with multicast support" (Mode Of Operation, MOP is 3) in the RPL network.

Once all RPL routing tables in the network are populated, any RPL node can send packets to an IP multicast destination. The RPL protocol performs distribution of multicast packet both upward towards the DODAG root and downwards into the DODAG.

The text in Section 12 of the RPL specification clearly implies that IP multicast packets are distributed using link-layer unicast transmissions, looking at the use of the word "copied" in this section. Specifically in 6LoWPAN networks, this behavior conflicts with the requirement that IP multicast packets MUST be carried as link-layer 802.15.4 broadcast frames [RFC4944].

Assuming that link-layer unicast is indeed meant, this approach seems efficient only in a balanced, sparse tree network topology, or in situations where the fraction of nodes listening to a specific multicast IP address is low, or in duty cycled LLNs where link-layer broadcast is a very expensive operation.

5.3.1.3. RPL Routers with Non-RPL Hosts

Now we consider the case that hosts exist in a RPL network that are not RPL-aware themselves, but rely on RPL routers for their IP connectivity beyond link-local scope. Note that the current RPL specification [I-D.ietf-roll-rpl] leaves this case for future specification (see Section 16.4). Non-RPL hosts cannot advertise their IP multicast groups of interest via RPL DAO messages as defined above. Therefore in that case MLD could be used for such advertisements (State Change Report messages), with all or a subset of RPL routers acting in the role of MLD Routers as defined in [RFC3810]. However, as the MLD protocol is not designed specifically for LLNs it may be a burden for the constrained RPL router nodes to run the full MLD protocol. Alternatives are therefore proposed in Section 5.4.1.

5.3.1.4. Trickle Multicast Forwarding

Trickle Multicast Forwarding [I-D.ietf-roll-trickle-mcast] is an IP multicast routing protocol suitable for LLNs, that uses the Trickle algorithm as a basis. It is a simple protocol in the sense that no topology maintenance is required. It can deal especially well with situations where the node density is a-priori unknown.

Nodes from anywhere in the LLN can be the multicast source, and nodes anywhere in the LLN can be multicast destinations.

Using Trickle Multicast Forwarding it is not required for IP multicast destinations (listeners) to announce their interest in a specific multicast IP address, e.g. by means of MLD. Instead, all multicast IP packets regardless of IP destination address are stored and forwarded by all routers. Because forwarding is always done by multicast, both hosts and routers will be able to receive all multicast IP packets. Routers that receive multicast packets they are not interested in, will only buffer these for a limited time

until retransmission can be stopped as specified by the protocol. Hosts that receive multicast packets they are not interested in, will discard multicast packets that are not of interest. Above properties seem to make Trickle especially efficient for cases where the multicast listener density is high and the number of distinct multicast groups relatively low.

5.3.1.5. Other Route-Over Methods

Other known IP multicast routing methods may be used, for example flooding or other to be defined methods suitable for LLNs. An important design consideration here is whether multicast listeners need to advertise their interest in specific multicast addresses, or not. If they do, MLD is a possible option but also protocol-specific means (as in RPL) is an option. See Section 5.4.1 for more efficient substitutes for MLD targeted towards a LLN context.

5.3.2. Single LLN with Backbone Topology

A LLN may be connected via a Border Router (e.g. 6LBR) to a backbone network, on which IP multicast listeners and/or sources may be present. This section analyzes cases in which IP multicast traffic needs to flow from/to the backbone, to/from the LLN.

5.3.2.1. Mesh-Under Multicast Routing

Because in a mesh routing network conceptually all nodes in the LLN are attached to a single link, a multicast IP packet originating in the LLN is typically delivered by the mesh routing algorithm to the 6LBR as well, although there is no guaranteed delivery. The 6LBR may be configured to accept all IP multicast traffic from the LLN and then may forward such packets onto its backbone link. Alternatively, the 6LBR may act in an MLD Router or MLD Snooper role on its backbone link and decide whether to forward a multicast packet or not based on information learned from previous MLD Reports received on its backbone link.

Conversely, multicast packets originating on the backbone network will reach the 6LBR if either the backbone is a single link (LAN/VLAN) or IPv6 multicast routing is enabled on the backbone. Then, the 6LBR could simply forward all IP multicast traffic from the backbone onto the LLN. However, in practice this situation may lead to overload of the LLN caused by unnecessary multicast traffic. Therefore the 6LBR SHOULD only forward traffic that one or more nodes in the LLN have expressed interest in, effectively filtering inbound LLN multicast traffic.

To realize this "filter", nodes on the LLN may use MLD to announce

their interest in specific multicast IP addresses to the 6LBR. One option is for the 6LBR to act in an MLD Router role on its LLN interface. However, this may be too much of a "burden" for constrained nodes. Light-weight alternatives for MLD are discussed in Section 5.4.1.

5.3.2.2. RPL Multicast Routing

For RPL routing within the 6LoWPAN, we first consider the case of an IP multicast source on the backbone network with one or more IP multicast listeners on the RPL LLN. Typically, the 6LBR would be the root of a DODAG so that the 6LBR can easily forward the IP multicast packet received on its backbone interface to the right RPL nodes in the LLN down along this DODAG (based on previously DAO-advertized destinations).

Second, a multicast source may be in the RPL LLN and listeners may be both on the LLN and on the backbone. For this case RPL defines that the multicast packet will propagate both up and down the DODAG, eventually reaching the DODAG root (typically a 6LBR) from which the packet can be routed onto the backbone in a manner specified in the previous section.

5.3.2.3. RPL Routers with Non-RPL Hosts

For the case that a RPL LLN contains non-RPL hosts, the solutions from the previous section can be used if in addition RPL routers implement MLD or "MLD like" functionality similar to as described in Section 5.3.1.3.

5.3.2.4. Trickle Multicast Forwarding

First, we consider the case of an IP multicast source node on the LLN (where all 6LRs support Trickle Multicast Forwarding) and IP multicast listeners that may be on the LLN and on the backbone. As Trickle will eventually deliver multicast packets also to a 6LBR, which acts as a Trickle Multicast router as well, the 6LBR can then forward onto the backbone in the ways described earlier in Section 5.3.2.1.

Second, for the case of an IP multicast source on the backbone and multicast listeners on both backbone and/or LLN, the 6LBR needs to forward multicast traffic from the backbone onto the LLN. Here, the aforementioned problem (Section 5.3.2.1) of potentially overloading the LLN with unwanted backbone IP multicast traffic appears again.

A possible solution to this is (again) to let multicast listeners advertise their interest using MLD as described in Section 5.3.2.1 or

to use an MLD alternative suitable for LLNs as described in Section 5.4.1. However, following this approach requires possibly an extension to Trickle Multicast Forwarding: the protocol should ensure that MLD-advertised information is somehow communicated to the 6LBR, possibly over multiple hops. MLD itself supports link-local communication only.

5.3.2.5. Other Route-Over Methods

For other multicast routing methods used on the LLN, there are similar considerations to the ones in sections above: the strong need to filter IP multicast traffic coming into the LLN, the need for reporting multicast listener interest (e.g. with MLD or a to-be-defined MLD alternative) by constrained (6LoWPAN) nodes, and the need for LLN-internal routing as identified in the previous section such that the MLD communicated information can reach the 6LBR to be used there in multicast traffic filtering decisions.

5.3.3. Multiple LLNs with Backbone Topology

Now the case of a single backbone network with two or more LLNs attached to it via 6LBRs is considered. For this case all the considerations and solutions of the previous section can be applied.

For the specific case that a source on a backbone network has to send to a very large number of destination located on many LLNs, the use of IGMP/MLD Proxying [RFC4605] with a leaf IGMP/MLD Proxy located in each 6LBR may be useful. This method only is defined for a tree topology backbone network with the IP multicast source at the root of the tree.

5.3.4. LLN(s) with Multiple 6LBRs

[TBD: an LLN with multiple 6LBRs may require some additional consideration. Any need to synchronize mutually on multicast listener information?]

5.3.5. Conclusions

For all network topologies that were evaluated, CoAP group communication can be in principle supported with IP Multicast, making use of existing protocols. For the case of Trickle Multicast Forwarding, it appears that an addition to the protocol is required such that information about multicast listeners can be distributed towards the 6LBR. Opportunities were identified for an "MLD-like" or "MLD-lightweight" protocol specifically suitable for LLNs, which should inter-work with regular MLD on the backbone network. Such MLD variants are further analyzed in Section 5.4.1.

5.4. Implementation Considerations

In this section various implementation aspects are considered such as required protocol implementations, additional functionality of the 6LBR and backbone network equipment.

5.4.1. MLD Implementation on LLNs

In previous sections, it was mentioned that the MLDv2 protocol [RFC3810] may be too costly for use in a LLN. MLD relies on periodic link-local multicast operations to maintain state. Also it is optimized to fairly dynamic situations where multicast listeners may come and go over time. Such dynamic situations are less frequently found in typical LLN use cases such as building control, where multicast group membership can remain constant over longer periods of time (e.g. months) after commissioning.

Hence, a viable strategy is to implement a subset of MLD functionality in 6LoWPAN nodes which is just enough for the required functionality. A first option is that 6LoWPAN Routers, like MLD Snoopers, passively listen to MLD State Change Report messages and handle the learned ("snooped") IP multicast destinations in the way defined by the multicast routing protocol they are running (e.g. for RPL, Routers advertise these destinations using DAO messages).

A second option is to use MLD as-is but adapt the recommended parameter values such that operation on a LLN becomes more efficient.

A third option is to standardize a new protocol, taking a subset of MLD functionality into a "MLD for 6LoWPAN" protocol to support constrained nodes optimally.

A fourth option is now presented, which seems attractive in that it minimizes standardization, implementation and network communication overhead all at the same time. This option is to specify a new Multicast Listener Option (MLO) as an addition to the 6LoWPAN-ND [I-D.ietf-6lowpan-nd] protocol communication that is anyway ongoing between a 6LoWPAN host and router(s). This MLO is preferably designed to be maximally similar to the Address Registration Option (ARO), which minimizes the need for additional program code on constrained nodes. With an MLO, instead of registering a unicast IP address, a host "registers" its interest in a multicast IP address. Unlike ARO, multiple MLO can be used in the same ND packet. A registration period is also defined just like in the ARO. MLO allows a host to persistently register as a listener to IP multicast traffic and to avoid the overhead of periodic multicast communication which is required for full MLD.

[TBD: consider what aspects are needed/not needed for CoAP/LLN applications. Will MLDv1 suffice? What to do with options like 'source specific' and include/exclude. Source-specific can also be dealt with at the destination host by filtering? Do we need limits on number of records per packet? Do we need a higher MLD reliability setting - see the parameters in the MLD RFC]

5.4.2. 6LBR Implementation

To support mixed backbone/LLN scenarios in CoAP group communication, it is RECOMMENDED that a 6LowPAN Border Router (6LBR) will act in an MLD Router role on the backbone link. If this is not possible then the 6LBR SHOULD be configured to act as an MLD Multicast Address Listener and/or MLD Snooper on the backbone link.

5.4.3. Backbone IP Multicast Infrastructure

For corporate/professional applications, most routing and switching equipment that is currently on the market is IPv6 capable. For that reason backbone infrastructure operating IPv4 only is considered out of scope in this document, at least for the backbone network segment(s) where IP multicast destinations are present. What is still in scope is for example an IPv4-only HTTP client that wants to send a group communication message via a HTTP-CoAP proxy as considered in [I-D.castellani-core-http-mapping].

The availability of, and requirements for, IP multicast support may depend on the specific installation use case. For example, the following cases may be relevant for new IP based building control installations:

1. System deployed on existing IP (Ethernet/WiFi/...) infrastructure, shared with existing IP devices (PCs)

2. Newly designed and deployed IP (Ethernet/WiFi/...) infrastructure, to be shared with other IP devices (PCs)

3. Newly designed and deployed IP (Ethernet/WiFi/...) infrastructure, exclusively used for building control.

Besides physical separation the building control backbone can be separated from regular (PC) infrastructure by using a different VLAN. A typical corporate installation will have many LAN switches and/or routing switches, which pass through IP multicast traffic but on the other hand do not support acting in the Router role of MLD/IGMP. Perhaps for case 2) and 3) above it is acceptable to add a MLD/IGMP capable router somewhere in the network, while for case 1) this may not be the case.

[TBD: consider the influence of WiFi based backbone networks. What if 6LBRs are at the same time also WiFi routers? What if 6LBRs have an Ethernet connection to legacy WiFI routers? Check if equivalent with Ethernet backbone.]

6. Security Considerations

TBD

7. IANA Considerations

This document makes no request of IANA.

8. Conclusions

IP multicast as outlined in Section 4 is recommended to be adopted as the base solution for CoAP Group Communication for situations where the use case and network characteristics allow use of IP multicast. This approach requires no standards changes to the IP multicast suite of protocols and it provides interoperability with IP multicast group communication on unconstrained backbone networks.

The proposals for group communication described in this draft should be considered for incorporation into the overall CoAP protocol specification.

9. Acknowledgements

Thanks to Peter Bigot, Carsten Bormann, Anders Brandt, Angelo Castellani, Guang Lu, Salvatore Loreto, Kerry Lynn, Dale Seed, Zach Shelby, Peter van der Stok, and Juan Carlos Zuniga for their helpful comments and discussions that have helped shape this document.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, August 2002.

[RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

[RFC3740] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC 3740, March 2004.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC3956] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[RFC4046] Baugher, M., Canetti, R., Dondeti, L., and F. Lindholm, "Multicast Security (MSEC) Group Key Management Architecture", RFC 4046, April 2005.

[RFC4286] Haberman, B. and J. Martin, "Multicast Router Discovery", RFC 4286, December 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.

[RFC5374] Weis, B., Gross, G., and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", RFC 5374, November 2008.

[RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for

IPv4 Multicast Address Assignments", BCP 51, RFC 5771, March 2010.

[RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.

[I-D.ietf-core-coap] Frank, B., Bormann, C., Hartke, K., and Z. Shelby, "Constrained Application Protocol (CoAP)", draft-ietf-core-coap-08 (work in progress), October 2011.

10.2. Informative References

[I-D.cheshire-dnsext-dns-sd] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", draft-cheshire-dnsext-dns-sd-11 (work in progress), December 2011.

[I-D.eggert-core-congestion-control] Eggert, L., "Congestion Control for the Constrained Application Protocol (CoAP)", draft-eggert-core-congestion-control-01 (work in progress), January 2011.

[I-D.ietf-6lowpan-routing-requirements] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for 6LoWPAN Routing", draft-ietf-6lowpan-routing-requirements-10 (work in progress), November 2011.

[I-D.ietf-6lowpan-hc] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams in Low Power and Lossy Networks (6LoWPAN)", draft-ietf-6lowpan-hc-15 (work in progress), February 2011.

[I-D.ietf-6lowpan-nd] Shelby, Z., Chakrabarti, S., and E. Nordmark, "Neighbor Discovery Optimization for Low Power and Lossy Networks (6LoWPAN)", draft-ietf-6lowpan-nd-18 (work in progress), October 2011.

[I-D.ietf-core-link-format] Shelby, Z., "CoRE Link Format", draft-ietf-core-link-format-11 (work in progress), January 2012.

[I-D.ietf-core-observe]
Hartke, K., "Observing Resources in CoAP", draft-ietf-core-observe-04 (work in progress), February 2012.

[I-D.shelby-core-coap-req]
Shelby, Z., Stuber, M., Sturek, D., Frank, B., and R. Kelsey, "CoAP Requirements and Features", draft-shelby-core-coap-req-02 (work in progress), October 2010.

[I-D.shelby-core-resource-directory]
Krco, S. and Z. Shelby, "CoRE Resource Directory", draft-shelby-core-resource-directory-02 (work in progress), October 2011.

[I-D.vanderstok-core-bc]
Stok, P. and K. Lynn, "CoAP Utilization for Building Control", draft-vanderstok-core-bc-05 (work in progress), October 2011.

[I-D.lynn-core-discovery-mapping]
Lynn, K. and Z. Shelby, "CoRE Link-Format to DNS-Based Service Discovery Mapping", draft-lynn-core-discovery-mapping-01 (work in progress), July 2011.

[I-D.castellani-core-http-mapping]
Castellani, A., Loreto, S., Rahman, A., Fossati, T., and E. Dijk, "Best practices for HTTP-CoAP mapping implementation", draft-castellani-core-http-mapping-02 (work in progress), October 2011.

[I-D.ietf-roll-rpl]
Brandt, A., Vasseur, J., Hui, J., Pister, K., Thubert, P., Levis, P., Struik, R., Kelsey, R., Clausen, T., and T. Winter, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-19 (work in progress), March 2011.

[I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Forwarding Using Trickle", draft-ietf-roll-trickle-mcast-00 (work in progress), April 2011.

[I-D.ietf-multimob-igmp-mld-tuning]
Asaeda, H., Liu, H., and Q. Wu, "Tuning the Behavior of IGMP and MLD for Routers in Mobile and Wireless Networks",

draft-ietf-multimob-igmp-mld-tuning-05 (work in progress), March 2012.

[Lao05] Lao, L., Cui, J., Gerla, M., and D. Maggiorini, "A Comparative Study of Multicast Protocols: Top, Bottom, or In the Middle?", 2005, <http://www.cs.ucla.edu/NRL/hpi/AggMC/papers/comparison_gi_2005.pdf>.

[Banerjee01] Banerjee, B. and B. Bhattacharjee, "A Comparative Study of Application Layer Multicast Protocols", 2001, <http://wmedia.grnet.gr/P2PBackground/a-comparative-study-ofALM.pdf>.

Authors' Addresses

Akbar Rahman (editor)
InterDigital Communications, LLC

Email: Akbar.Rahman@InterDigital.com

Esko Dijk (editor)
Philips Research

Email: esko.dijk@philips.com