CoRE Working Group A. Rahman, Ed.
Internet-Draft InterDigital Communications, LLC
Intended status: Informational E. Dijk, Ed.
Expires: January 11, 2013 Philips Research
July 10, 2012

Group Communication for CoAP
draft-ietf-core-groupcomm-02

Abstract

CoAP is a RESTful transfer protocol for constrained devices. It is anticipated that constrained devices will often naturally operate in groups (e.g. in a building automation scenario all lights in a given room may need to be switched on/off as a group). This document defines how the CoAP protocol should be used in a group communication context. An approach for using CoAP on top of IP multicast is detailed for both constrained and un-constrained networks. Also, various use causes and corresponding protocol flows are provided to illustrate important concepts. Finally, guidance is provided for deployment in various network topologies.

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].

This document assumes readers are familiar with the terms and concepts that are used in [I-D.ietf-core-coap]. In addition, this document defines the following terminology:

Group Communication
A source node sends a single message which is delivered to multiple destination nodes, where all destinations are identified to belong to a specific group. The source node may or may not be part of the group. The underlying mechanism for group communication is assumed to be multicast based. The network where the group communication takes place can be either a constrained or a regular (un-constrained) network

Multicast
Sending a message to multiple destination nodes simultaneously. 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)
Low power and Lossy Network (LLN): A type of constrained network where the devices are interconnected by a variety of low power, lossy 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 may belong to another group. 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. HTTP does not support any equivalent functionality to CoAP group communication.

2.2. Scope

In this draft, we address the issues related to CoAP group communication in detail, with use cases, recommended approaches and analysis of the impact to the CoAP protocol and to implementations. 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 destination recipients that are all part of a group. Before content can be distributed, there is a separate process to form the group. 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 layer 2 (Media Access Control broadcast/multicast) and layer 3 (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 adopted by this draft for CoAP group communication.

3. IP Multicast Based Group Communication

3.1. Introduction

IP Multicast routing 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 routing is not widely deployed on the general Internet. However, IP Multicast is very popular in specific deployments such as in enterprise networks (e.g. for video conferencing), smart home networks (e.g. UPnP/mDNS) and carrier IPTV deployments. The packet economy and minimal host complexity of IP multicast make it attractive for group communication in constrained environments. Therefore IP multicast is the recommended underlying mechanism for CoAP group communications, and the approach assumed in this document.

To achieve IP multicast beyond a subnet, an IP multicast routing protocol needs to be active on routers. The RPL protocol [RFC6550] for example is able to route multicast traffic in constrained LLNs. While PIM-SM [RFC4601] is often used for multicast routing in un-constrained networks.

IP multicast can also be run in a Link-Local (LL) scope. This means that there is no routing involved and the IP multicast message is only sent and received in the local subnet.

3.2. Group URIs and IP Multicast Addresses

A group of CoAP nodes can be addressed using its IP multicast addresses or a group URI ([I-D.vanderstok-core-dna]) which can be mapped to a site-local or global multicast IP address via DNS resolution. A CoAP node can become a group member by listening for CoAP messages on the corresponding IP multicast address. Group URIs MUST follow the URI syntax [RFC3986]. Examples of hierarchical group naming (and scoping) for a building control application are shown below.

```
  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"
```

Reverse mapping (from IP multicast address to group authority) is supported using the reverse DNS resolution technique ([I-D.vanderstok-core-dna]).

3.3. 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). These topics are elaborated in more detail in [I-D.vanderstok-core-dna] including examples for using DNS-SD and CoRE Resource Directory.

3.3.1. 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>.

3.3.2. 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.

3.4. 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 be Non-Confirmable. A unicast response MAY be sent back to answer the group request (e.g. response "2.05 Content" to a group GET request) taking into account the security and congestion control rules defined in [I-D.ietf-core-coap].

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 solution is to impose the following restrictions:

o All CoAP multicast requests MUST be sent either to the default CoAP port (i.e. default Uri-Port as defined in [I-D.ietf-core-coap]), or to a port number obtained via a service discovery lookup operation being a valid CoAP port for the targeted multicast group.

o All CoAP multicast requests SHOULD operate only on URIs (links) which were retreived either from a "/.well-known/core" lookup on at least one group member node, or from equivalent service discovery lookup.

3.5. Congestion Control

Multicast CoAP requests may result in a multitude of replies from different nodes, potentially causing congestion. Therefore sending multicast requests should be conservatively controlled.

CoAP reduces multicast-specific congestion risks through the following measures:

- o A server MAY choose not to respond to a multicast request if there's nothing useful to respond (e.g. error or empty response).

- o A server SHOULD limit the support for multicast requests to specific resources where multicast operation is required.

- o A multicast request MUST be Non-Confirmable.

- o A server does not respond immediately to a multicast request, but SHOULD first wait for a time that is randomly picked within a predetermined time interval called the Leisure.

- o A server SHOULD NOT accept multicast requests that can not be authenticated.

Additional guidelines to reduce congestion risks are:

- o A server in an LLN should only support multicast GET for resources that are small i.e. where the payload of the response fits into a single link-layer frame.

- o A server can minimize the payload length in response to a multicast GET on "/.well-known/core" by using hierarchy in arranging link descriptions for the response. An example of this is given in Section 5 of [I-D.ietf-core-link-format].

- o Preferably IP multicast with link-local scope should be used, rather than global or site-local.

- o The Hop Limit field in the IPv6 packet should be chosen as low as possible (if the CoAP/IP stack allows setting of this value. TBD - discuss whether this guideline is relevant/realistic in CoAP context)

3.6. CoAP Multicast and HTTP Unicast Interworking

CoAP supports operation over UDP multicast, while HTTP does not. For use cases where it is required that CoAP group communication is initiated from an HTTP end-point, it would be advantageous if the

HTTP-CoAP Proxy supports mapping of HTTP unicast to CoAP group communication based on IP multicast. One possible way of operation of such HTTP-CoAP Proxy is illustrated in Figure 1. Note that this topic is covered in more detail in [I-D.castellani-core-advanced-http-mapping].

```
        CoAP    Mcast    CoAP    Mcast    HTTP-CoAP          HTTP
       Node 1   Rtr1    Node 2    Rtr2    Proxy            Node 3
         |       |         |       |       |                   |
         |MLD REQUEST      |       |       |                   |
         |(Join Group X)   |       |       |                   |
         |--LL-->|         |       |       |                   |
         |       |         |MLD REQUEST    |                   |
         |       |         |(Join Group X) |                   |
         |       |         |--LL-->|       |                   |
         |       |         |       |       |  HTTP REQUEST     |
         |       |         |       |       |    (URI to        |
         |       |         |       |       |   unicast addr)   |
         |       |         |       |       |< -----------------|
         |       |         |       |       |                   |
         |       |         |   Resolve HTTP Request-Line URI   |
         |       |         |   to Group X multicast address    |
         |       |         |       |       |                   |
         | CoAP REQUEST (to multicast addr)|                   |
         |< ------<---------< ------<------|                   |
         |       |         |       |       |                   |
         |                 |               |                   |
         |     (optional) CoAP RESPONSE(s) |                   |
         |                 |------------- >|                   |
         |-----------------|-------------->|                   |
         |                 |               |  HTTP RESPONSE    |
         |                 |               |----------------- >|
         |                 |               |                   |
```

Figure 1: CoAP Multicast and HTTP Unicast Interworking

Note that Figure 1 illustrates the case of IP multicast as the underlying group communications mechanism. MLD denotes the Multicast Listener Discovery protocol ([RFC3810], Appendix A) and LL denotes a Link-Local multicast.

A key point in Figure 1 is that the incoming HTTP Request (from node 3) will carry a Host request-header field that resolves in the general Internet to the proxy node. At the proxy node, this hostname and/or the Request-Line 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 address. This may be accomplished, for example, by using DNS or DNS-SD (Section 3.3). 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 1 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-advanced-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 is 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.

4. Use Cases and Corresponding Protocol Flows

4.1. Introduction

The use of CoAP group communication is shown in the context of the following use cases and corresponding protocol flows:

- Discovery of Resource Directory: discovering the local CoAP RD which contains links (URIs) to resources stored on other servers [I-D.ietf-core-link-format].

- Lighting Control: synchronous operation of a group of 6LoWPAN [RFC4944] IPv6-connected lights

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

- Firmware Update: efficiently updating firmware simultaneously in a large group of devices in a building/campus control ([I-D.vanderstok-core-bc]) application --- TBD suggests a multicast extension of core-block.

- Group Status Report: requesting status information or event reports from a group of devices in a building/campus control application --- TBD, may require reliable group communication to be feasible.

4.2. Network Configuration

We assume the following network configuration for all the use cases as shown in Figure 2:

- 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.

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

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

- The routers are connected to 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 (Appendix A) for the following use cases to work.

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

Figure 2: Network Topology of a Large Room (Room-A)

4.3. Discovery of Resource Directory

The protocol flow for discovery of a RD for the given network (of Figure 2) is shown in Figure 3:

- o The fixture for Light-2 is installed and powered on for the first time.

- o Light-2 will then search for the local RD (RD-2) by sending out a GET request (for the "/.well-known/core" resource) via a LL IP multicast message. In this case, the group is assumed to include all nodes in the subnet.

- o This LL IP multicast message will then go to each node in subnet-2. However, only Rtr-2 (RD-2) will respond because the GET is qualified by the query string "?rt=core-rd".

- o Note that the flow is shown only for Light-2 for clarity. Similar flows will happen for Light-1, Light-3 and the Light Switch when they are first powered on.

The RD may also be discovered by other means such as by assuming a default location (e.g. on a 6LBR), using DHCP, etc. However, these approaches do not invoke CoAP group communication.

For other discovery use cases such as discovering local CoAP servers, services or resources group communication can be used in a similar fashion as in the above use case.

```
                                     Light      Rtr-1     Rtr-2   Network
   Light-1   Light-2    Light-3     Switch    (RD-1)    (RD-2)  Backbone
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    **********************************          |          |          |
    *   Light-2 is installed         *          |          |          |
    *   and powers on for first time *          |          |          |
    **********************************          |          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          | COAP NON (GET                             |          |
    |          |          /.well-known/core?rt=core-rd)   |          |
    |          |--------LL------------------------------->|          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
    |          | COAP NON (Response                        |          |
    |          |          2.05 Content                     |          |
    |          |        </rd>; rt="core-rd; ins="Primary")|          |
    |          |<-----------------------------------------|          |
    |          |          |          |          |          |          |
    |          |          |          |          |          |          |
```

Figure 3: Resource Directory Discovery via Multicast Message

4.4. Lighting Control

The protocol flow for a building automation lighting control scenario for the network (Figure 2) is shown in sequence in Figure 4, Figure 5, and Figure 6. We assume the following steps occur before the illustrated flow:

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

o 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.

o 3) The indicated MLD Report messages are link-local multicast. In each LoWPAN, it is assumed that a multicast routing protocol in 6LRs will then propagate the Join information contained in the MLD Report 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 Lighting Groups

```
                                       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           |
       |          |          |          Payload=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               |
       |          |          |          |         Payload=turn on lights)|
       |          |          |          |    Destination IP Address =    |
       |          |          |          |       IP multicast address     |
       |          |          |          |       for Group (Room-A-Lights)|
       |          |          |          |    Originating IP Address =    |
       |          |          |          |        RTR-1                   |
       |          |          |          |          |-------------------->|
       |<------------------------------------------|          |          |
       |          |          |          |          |          |          |
       |          |          |          |          |          |<---------|
       |          |<---------|<-------------------------------|          |
       |          |          |          |          |          |          |
       |          |          |          |          |          |          |
```

Figure 5: Sending Lighting Control Multicast Message

```
                                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 responses   *   |
    |          |          |          *  to multicast message      *   |
    |          |          |          *  and formulates one        *   |
    |          |          |          *  consolidated response     *   |
    |          |          |          *  to originator             *   |
    |          |          |          ******************************   |
    |          |          |          |          |          |          |
    |          |          |      COAP NON (Response        |          |
    |          |          |                Success)        |          |
    |          |          |          |<---------|          |          |
    |          |          |          |          |          |          |
```

Figure 6: Sending Lighting Control Response to Multicast Message

NOTE: In the last step of Figure 6, instead of a single consolidated response the CoAP Proxy Rtr-1 could also return multiple individual CoAP responses, similar to the case that a CoAP client sends a CoAP multicast request directly. The format of a consolidated response is currently not defined in [I-D.ietf-core-coap].

5. Deployment Guidelines

This section provides some guidelines how an IP Multicast based solution for CoAP group communication can be deployed in various network configurations.

5.1. Target Network Topologies

CoAP group communication can be deployed in various network topologies. First, the target network may be a regular IP network, or a LLN such as e.g. a 6LoWPAN network, or consist of mixed constrained/unconstrained network segments. Second, it may be a single subnet only or multi-subnet; e.g. multiple 6LoWPAN networks joined by a single backbone LAN. Third, a wireless network segment may have all nodes reachable in a single IP hop, or it may require multiple IP hops for some pairs of nodes to reach eachother.

Each topology may pose different requirements on the configuration of routers and protocol(s), in order to enable efficient CoAP group communication.

5.2. Multicast Routing

If a network (segment) requires multiple IP hops to reach certain nodes, a multicast routing protocol is required to propagate multicast UDP packets to these nodes. Examples of routing protocols specifically for LLNs, able to route multicast, are RPL (Section 12 of [RFC6550]) and Trickle Multicast Forwarding [I-D.ietf-roll-trickle-mcast].

5.3. Use of the Multicast Listener Discovery (MLD) protocol

CoAP nodes that are IP hosts (not routers) are unaware of the specific multicast routing protocol being used. When such a host needs to join a specific (CoAP) multicast group, it usually requires a way to signal to the multicast routers which multicast traffic it wants to receive. For efficient multicast routing (i.e. avoid always flooding multicast IP packets), routers must know which hosts need to receive packets addressed to specific IP multicast destinations.

The Multicast Listener Discovery (MLD) protocol ([RFC3810], Appendix A) is the standard IPv6 method to achieve this. [RFC6636] discusses tuning of MLD for mobile and wireless networks. These guidelines may be useful when implementing MLD in LLNs.

Alternatively, to avoid the addition of MLD in LLN deployments, all nodes can be configured as multicast routers.

5.4. 6LoWPAN-Specific Guidelines

To support multi-LoWPAN scenarios for 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 (Appendix A) on the backbone link.

To avoid that backbone IP multicast traffic needlessly congests 6LoWPAN network segments, it is RECOMMENDED that a filtering means is implemented to block IP multicast traffic from 6LoWPAN segments where none of the 6LoWPAN nodes listen to this traffic. Possible means are:

- o Filtering in 6LBRs based on information from the routing protocol. This allows a 6LBR to only forward multicast traffic onto the LoWPAN, for which it is known that there exists at least one listener on the LoWPAN.

- o Filtering in 6LBRs based on MLD reports. Similar as previous but based directly on MLD reports from 6LoWPAN nodes. This only works in a single-IP-hop 6LoWPAN network such as a mesh-under routing network.

- o Filtering in 6LBRs based on settings. Filtering tables with blacklists/whitelists can be configured in the 6LBR by system administration for all 6LBRs or configured on a per-6LBR basis.

- o Filtering in router(s) that provide access to 6LoWPAN network segments. For example, in an access router/bridge that connects a regular intranet LAN to a building control IPv6 backbone. This backbone connects multiple 6LoWPAN segments.

6. Security Considerations

TBD

7. IANA Considerations

A request is made to IANA for reserving a range of IP addresses for "CoAP group communication" for:

- o IPv4 link-local scope multicast.

- o IPv6 link-local scope multicast.

- o IPv4 general multicast.
- o IPv6 general multicast.

8. Conclusions

IP multicast as outlined in Section 3 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 un-constrained backbone networks.

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.

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

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

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

[RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 5771, March 2010.

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.

[RFC6636] Asaeda, H., Liu, H., and Q. Wu, "Tuning the Behavior of the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) for Routers in Mobile and Wireless Networks", RFC 6636, May 2012.

[I-D.ietf-core-coap] Shelby, Z., Hartke, K., Bormann, C., and B. Frank, "Constrained Application Protocol (CoAP)", draft-ietf-core-coap-10 (work in progress), June 2012.

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.ietf-core-link-format] Shelby, Z., "CoRE Link Format", draft-ietf-core-link-format-14 (work in progress), June 2012.

[I-D.ietf-core-observe] Hartke, K., "Observing Resources in CoAP", draft-ietf-core-observe-05 (work in progress), March 2012.

[I-D.shelby-core-resource-directory] Shelby, Z. and S. Krco, "CoRE Resource Directory", draft-shelby-core-resource-directory-03 (work in progress), May 2012.

[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.vanderstok-core-dna]
Stok, P., Lynn, K., and A. Brandt, "CoRE Discovery, Naming, and Addressing", draft-vanderstok-core-dna-01 (work in progress), March 2012.

[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-05 (work in progress), July 2012.

[I-D.castellani-core-advanced-http-mapping]
Castellani, A., Loreto, S., Rahman, A., Fossati, T., and E. Dijk, "Best Practices for HTTP-CoAP Mapping Implementation", draft-castellani-core-advanced-http-mapping-00 (work in progress), July 2012.

[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.

[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>.

Appendix A. Multicast Listener Discovery (MLD)

In order to extend the scope of IP multicast beyond link-local scope, 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 a device in MLD Router role) if no MLD Router is present.

[RFC6636] 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.

Appendix B. 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 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.

Appendix C. Change Log

Changes from ietf-01 to ietf-02:

- Rewrote congestion control section based on latest CoAP text including Leisure concept (#188)

- Updated the CoAP/HTTP interworking section and example use case with more details and use of MLD for multicast group joining

- Key use cases added (#185)

- References to [I-D.vanderstok-core-dna] and [I-D.castellani-core-advanced-http-mapping] added

- Moved background sections on "MLD" and "CoAP-Observe" to Appendices

- Removed requirements section (and moved it to draft-dijk-core-groupcomm-misc)

- Added details for IANA request for group communication multicast addresses

- Clarified text to distinguish between "link local" and general multicast cases

- Moved lengthy background section 5 to draft-dijk-core-groupcomm-misc and replaced with a summary

- Various editorial updates for improved readibility

- Changelog added

Changes from ietf-00 to ietf-01:

- Moved CoAP-observe solution section to section 2

- Editorial changes

- Moved security requirements into requirements section

- Changed multicast POST to PUT in example use case

- Added CoAP responses in example use case

Changes from rahman-07 to ietf-00:

- Editorial changes
- Use cases section added
- CoRE Resource Directory section added
- Removed section 3.3.5. IP Multicast Transmission Methods
- Removed section 3.4 Overlay Multicast
- Removed section 3.5 CoAP Application Layer Group Management
- Clarified section 4.3.1.3 RPL Routers with Non-RPL Hosts case
- References added and some normative/informative status changes

Authors' Addresses

Akbar Rahman (editor)
InterDigital Communications, LLC

Email: Akbar.Rahman@InterDigital.com

Esko Dijk (editor)
Philips Research

Email: esko.dijk@philips.com