

CoRE Working Group
Internet-Draft
Obsoletes: [7390](#) (if approved)
Updates: [7252](#), [7641](#) (if approved)
Intended status: Standards Track
Expires: August 26, 2021

E. Dijk
IoTconsultancy.nl
C. Wang
InterDigital
M. Tiloca
RISE AB
February 22, 2021

**Group Communication for the Constrained Application Protocol (CoAP)
draft-ietf-core-groupcomm-bis-03**

Abstract

This document specifies the use of the Constrained Application Protocol (CoAP) for group communication, using UDP/IP multicast as the underlying data transport. Both unsecured and secured CoAP group communication are specified. Security is achieved by use of the Group Object Security for Constrained RESTful Environments (Group OSCORE) protocol. The target application area of this specification is any group communication use cases that involve resource-constrained devices or networks. This document replaces [RFC7390](#), while it updates [RFC7252](#) and [RFC7641](#).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 26, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Scope	4
1.2.	Terminology	5
2.	Group Definition and Group Configuration	5
2.1.	Group Definition	5
2.1.1.	CoAP Group	5
2.1.2.	Application Group	6
2.1.3.	Security Group	6
2.1.4.	Relations Between Group Types	6
2.2.	Group Configuration	9
2.2.1.	Group Naming	9
2.2.2.	Group Creation and Membership	10
2.2.3.	Group Discovery	11
2.2.4.	Group Maintenance	12
3.	CoAP Usage in Group Communication	13
3.1.	Request/Response Model	13
3.2.	Caching	16
3.2.1.	Freshness Model	17
3.2.2.	Validation Model	17
3.3.	Port and URI Path Selection	21
3.4.	Proxy Operation	22
3.4.1.	Forward-Proxies	22
3.4.2.	Reverse-Proxies	24
3.4.3.	Caching	25
3.5.	Congestion Control	30
3.6.	Observing Resources	31
3.7.	Block-Wise Transfer	33
3.8.	Transport	34
3.8.1.	UDP/IPv6 Multicast Transport	34
3.8.2.	UDP/IPv4 Multicast Transport	34
3.8.3.	6LoWPAN	34
3.9.	Interworking with Other Protocols	35
3.9.1.	MLD/MLDv2/IGMP/IGMPv3	35
3.9.2.	RPL	36
3.9.3.	MPL	36
4.	Unsecured Group Communication	37
5.	Secured Group Communication using Group OSCORE	37
5.1.	Secure Group Maintenance	39

5.2.	Caching of Responses at Proxies	40
5.2.1.	Using Deterministic Requests to Achieve Cachability .	40
5.2.2.	Validation of Responses	41
6.	Security Considerations	41
6.1.	CoAP NoSec Mode	41
6.2.	Group OSCORE	42
6.2.1.	Group Key Management	42
6.2.2.	Source Authentication	43
6.2.3.	Countering Attacks	43
6.3.	Replay of Non-Confirmable Messages	45
6.4.	Use of CoAP No-Response Option	45
6.5.	6LoWPAN	46
6.6.	Wi-Fi	46
6.7.	Monitoring	47
6.7.1.	General Monitoring	47
6.7.2.	Pervasive Monitoring	47
7.	IANA Considerations	48
7.1.	CoAP Option Numbers Registry	48
8.	References	48
8.1.	Normative References	48
8.2.	Informative References	51
Appendix A.	Use Cases	53
A.1.	Discovery	53
A.1.1.	Distributed Device Discovery	53
A.1.2.	Distributed Service Discovery	54
A.1.3.	Directory Discovery	54
A.2.	Operational Phase	55
A.2.1.	Actuator Group Control	55
A.2.2.	Device Group Status Request	55
A.2.3.	Network-wide Query	55
A.2.4.	Network-wide / Group Notification	56
A.3.	Software Update	56
Appendix B.	Document Updates	56
B.1.	Version -02 to -03	56
B.2.	Version -01 to -02	57
B.3.	Version -00 to -01	57
Acknowledgments	57
Authors' Addresses	58

1. Introduction

This document specifies group communication using the Constrained Application Protocol (CoAP) [[RFC7252](#)] together with UDP/IP multicast. CoAP is a RESTful communication protocol that is used in resource-constrained nodes, and in resource-constrained networks where packet sizes should be small. This area of use is summarized as Constrained RESTful Environments (CoRE).

One-to-many group communication can be achieved in CoAP, by a client using UDP/IP multicast data transport to send multicast CoAP request messages. In response, each server in the addressed group sends a response message back to the client over UDP/IP unicast. Notable CoAP implementations supporting group communication include the framework "Eclipse Californium" 2.0.x [[Californium](#)] from the Eclipse Foundation and the "Implementation of CoAP Server & Client in Go" [[Go-OCF](#)] from the Open Connectivity Foundation (OCF).

Both unsecured and secured CoAP group communication over UDP/IP multicast are specified in this document. Security is achieved by using Group Object Security for Constrained RESTful Environments (Group OSCORE) [[I-D.ietf-core-oscore-groupcomm](#)], which in turn builds on Object Security for Constrained Restful Environments (OSCORE) [[RFC8613](#)]. This method provides end-to-end application-layer security protection of CoAP messages, by using CBOR Object Signing and Encryption (COSE) [[I-D.ietf-cose-rfc8152bis-struct](#)][[I-D.ietf-cose-rfc8152bis-algs](#)].

All guidelines in [[RFC7390](#)] are updated by this document, which replaces and obsoletes [[RFC7390](#)]. Furthermore, this document updates [[RFC7252](#)], by specifying: a group request/response model; cachability of responses to group requests at proxies; a response validation model for responses to group requests; and the use of Group OSCORE [[I-D.ietf-core-oscore-groupcomm](#)] to achieve security for CoAP group communication. Finally, this document also updates [[RFC7641](#)], by defining the multicast usage of the CoAP Observe Option for both the GET and FETCH methods.

All sections in the body of this document are normative, while appendices are informative. For additional background about use cases for CoAP group communication in resource-constrained devices and networks, see [Appendix A](#).

[1.1.1](#). Scope

For group communication, only solutions that use CoAP over UDP/IP multicast are in the scope of this document. There are alternative methods to achieve group communication using CoAP, for example Publish-Subscribe [[I-D.ietf-core-coap-pubsub](#)] which uses a central broker server that CoAP clients access via unicast communication. These methods may be usable for the same or similar use cases as are targeted in this document.

Furthermore, this document defines Group OSCORE [[I-D.ietf-core-oscore-groupcomm](#)] as the default group communication security solution for CoAP. Security solutions for group communication and configuration other than Group OSCORE are not in

scope. General principles for secure group configuration are in scope.

1.2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

This specification requires readers to be familiar with CoAP terminology [[RFC7252](#)]. Terminology related to group communication is defined in [Section 2.1](#).

Furthermore, "Security material" refers to any security keys, counters or parameters stored in a device that are required to participate in secure group communication with other devices.

2. Group Definition and Group Configuration

In the following, different group types are first defined in [Section 2.1](#). Then, Group configuration, including group creation and maintenance by an application, user or commissioning entity is considered in [Section 2.2](#).

2.1. Group Definition

Three types of groups and their mutual relations are defined in this section: CoAP group, application group, and security group.

2.1.1. CoAP Group

A CoAP group is defined as a set of CoAP endpoints, where each endpoint is configured to receive CoAP multicast messages that are sent to the group's associated IP multicast address and UDP port. An endpoint may be a member of multiple CoAP groups by subscribing to multiple IP multicast groups and/or listening on multiple UDP ports. Group membership(s) of an endpoint may dynamically change over time. A device sending a CoAP multicast message to a CoAP group is not necessarily itself a member of this CoAP group: it is a member only if it also has a CoAP endpoint listening on the group's associated IP multicast address and UDP port. A CoAP group can be encoded within a Group URI. This is defined as a CoAP URI that has the "coap" scheme and includes in the authority part either an IP multicast address or a group hostname (e.g., a Group Fully Qualified Domain Name (FQDN)) that can be resolved to an IP multicast address. A Group URI also

contains an optional UDP port number in the authority part. Group URIs follow the regular CoAP URI syntax (see [Section 6 of \[RFC7252\]](#)).

2.1.2. Application Group

Besides CoAP groups, that have relevance at the level of IP networks and CoAP endpoints, there are also application groups. An application group is a set of CoAP server endpoints that share a common set of CoAP resources. An endpoint may be a member of multiple application groups. An application group has relevance at the application level - for example an application group could denote all lights in an office room or all sensors in a hallway. A client endpoint that sends a group communication message to an application group is not necessarily itself a member of this application group. There can be a one-to-one or a one-to-many relation between a CoAP group and application group(s). An application group identifier is optionally encoded explicitly in the CoAP request, for example as a name in the URI path. If not explicitly encoded, the application group is implicitly derived by the receiver, based on information in the CoAP request. See [Section 2.2.1](#) for more details on identifying the application group.

2.1.3. Security Group

For secure group communication, a security group is required. A security group is a group of endpoints that each store group security material, such that they can mutually exchange secured messages and verify secured messages. So, a client endpoint needs to be a member of a security group in order to send a valid secured group communication message to this group. An endpoint may be a member of multiple security groups. There can be a one-to-one or a one-to-many relation between security groups and CoAP groups. Also, there can be a one-to-one or a one-to-many relation between security groups and application groups. A special security group named "NoSec" identifies group communication without any security at the transport layer nor at the CoAP layer.

2.1.4. Relations Between Group Types

Using the above group type definitions, a CoAP group communication message sent by an endpoint can be represented as a tuple that contains one instance of each group type:

(application group, CoAP group, security group)

A special note is appropriate about the possible relation between security groups and application groups.

On one hand, multiple application groups may use the same security group. Thus, the same group security material is used to protect the messages targeting any of those application groups. This has the benefit that typically less storage, configuration and updating are required for security material. In this case, a CoAP endpoint is supposed to know the exact application group to refer to for each message that is sent or received, based on, e.g., the used server port number, the targeted resource, or the content and structure of the message payload.

On the other hand, a single application group may use multiple security groups. Thus, different messages targeting the resources of the application group can be protected with different security material. This can be convenient, for example, if the security groups differ with respect to the cryptographic algorithms and related parameters they use. In this case, a CoAP client can join just one of the security groups, based on what it supports and prefers, while a CoAP server in the application group would rather have to join all of them.

Beyond this particular case, applications should be careful in associating a same application group to multiple security groups. In particular, it is NOT RECOMMENDED to use different security groups to reflect different access policies for resources in a same application group. That is, being a member of a security group actually grants access only to exchange secured messages and enables authentication of group members, while access control (authorization) to use resources in the application group belongs to a separate security domain. It has to be separately enforced by leveraging the resource properties or through dedicated access control credentials assessed by separate means.

Figure 1 summarizes the relations between the different types of groups described above in UML class diagram notation. The class attributes in square brackets are optionally defined.

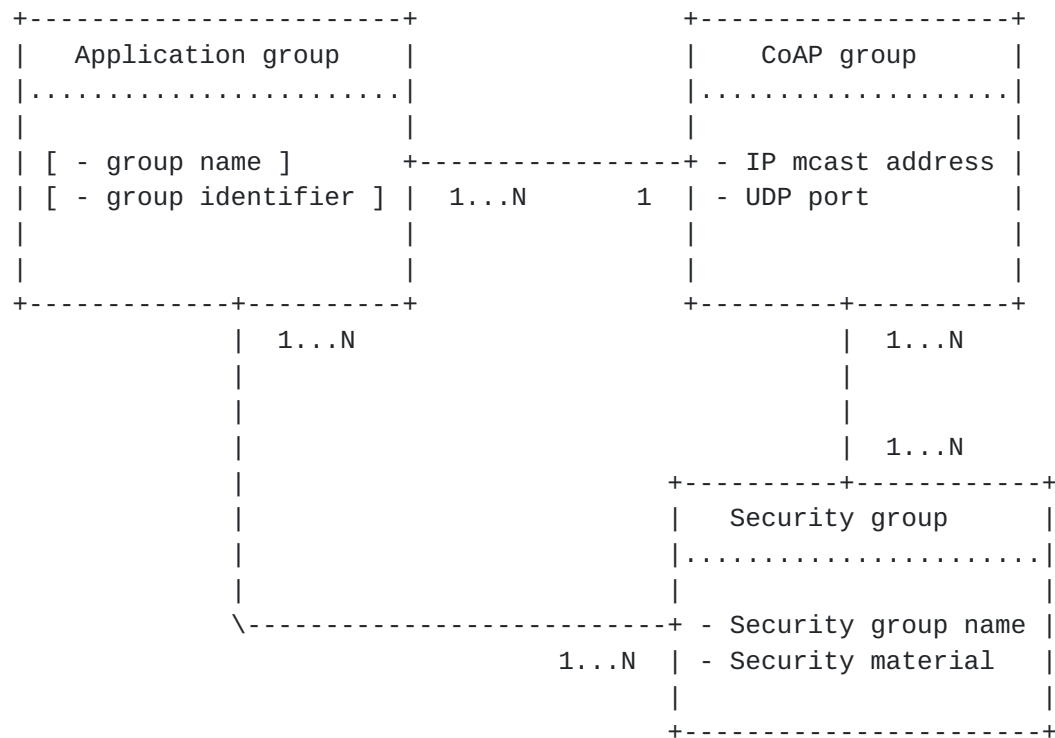


Figure 1: Relations Among Different Group Types

Figure 2 provides a deployment example of the relations between the different types of groups. It shows six CoAP servers (Srv1-Srv6) and their respective resources hosted (/resX). There are three application groups (1, 2, 3) and two security groups (1, 2). Security Group 1 is used by both Application Group 1 and 2. Three clients (Cli1, Cli2, Cli3) are configured with security material for Security Group 1. Two clients (Cli2, Cli4) are configured with security material for Security Group 2. All the shown application groups use the same CoAP group (not shown in the figure), i.e. one specific multicast IP address and UDP port on which all the shown resources are hosted for each server.

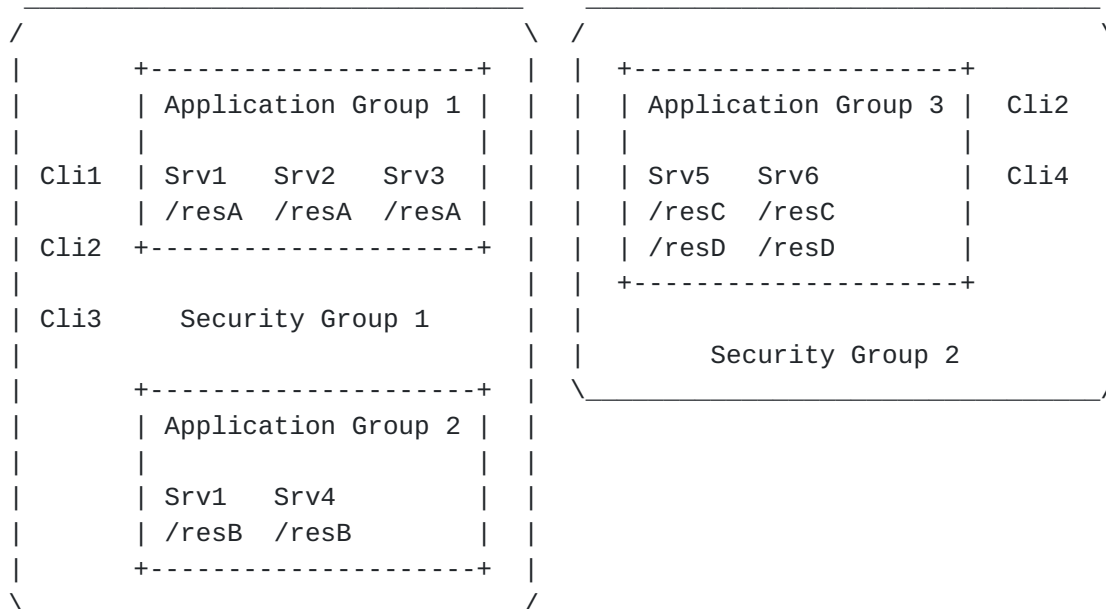


Figure 2: Deployment Example of Different Group Types

2.2. Group Configuration

The following defines how groups of different types are named, created, discovered and maintained.

2.2.1. Group Naming

A CoAP group is identified and named by the authority component in the Group URI, which includes host (possibly an IP multicast address literal) and an optional UDP port number. It is recommended to configure an endpoint with an IP multicast address literal, instead of a hostname, when configuring a CoAP group membership. This is because DNS infrastructure may not be deployed in many constrained networks. In case a group hostname is configured, it can be uniquely mapped to an IP multicast address via DNS resolution - if DNS client functionality is available in the endpoint being configured and the DNS service is supported in the network. Some examples of hierarchical CoAP group FQDN naming (and scoping) for a building control application are shown in [Section 2.2 of \[RFC7390\]](#).

An application group can be named in many ways through different types of identifiers, such as numbers, URIs or other strings. An application group name or identifier, if explicitly encoded in a CoAP request, is typically included in the path component or in the query component of a Group URI. It may also be encoded using the Uri-Host Option [\[RFC7252\]](#) in case application group members implement a virtual CoAP server specific to that application group. The

application group can then be identified by the value of the Uri-Host Option and each virtual server serves one specific application group. However, encoding the application group in the Uri-Host Option is not the preferred method because in this case the application group cannot be encoded in a Group URI, and also the Uri-Host Option is being used for another purpose than encoding the host part of a URI as intended by [\[RFC7252\]](#) - which is potentially confusing.

[Appendix A](#) of [\[I-D.ietf-core-resource-directory\]](#) shows an example registration of an application group into a Resource Directory (RD), along with the CoAP group it uses and the resources supported by the application group. In this example an application group identifier is not explicitly encoded in the RD nor in CoAP requests made to the group, but it implicitly follows from the CoAP group used for the request. So there is a one-to-one binding between the CoAP group and the application group. The "NoSec" security group is used.

A best practice for encoding application group into a Group URI is to use one URI path component to identify the application group and use the following URI paths component(s) to identify the resource within this application group. For example, `/<groupname>/res1` or `/base/<groupname>/res1/res2` conform to this practice. An application group identifier (like `<groupname>`) should be as short as possible when used in constrained networks.

A security group is identified by a stable and invariant string used as group name, which is generally not related with other kinds of group identifiers, specific to the chosen security solution. The "NoSec" security group name MUST be only used to represent the case of group communication without any security. It is typically characterized by the absence of any security group name, identifier, or security-related data structures in the CoAP message.

[2.2.2. Group Creation and Membership](#)

To create a CoAP group, a configuring entity defines an IP multicast address (or hostname) for the group and optionally a UDP port number in case it differs from the default CoAP port 5683. Then, it configures one or more devices as listeners to that IP multicast address, with a CoAP endpoint listening on the group's associated UDP port. These endpoints/devices are the group members. The configuring entity can be, for example, a local application with pre-configuration, a user, a software developer, a cloud service, or a local commissioning tool. Also, the devices sending CoAP requests to the group in the role of CoAP client need to be configured with the same information, even though they are not necessarily group members. One way to configure a client is to supply it with a CoAP Group URI. The IETF does not define a mandatory, standardized protocol to accomplish CoAP group creation. [\[RFC7390\]](#) defines an experimental

protocol for configuration of group membership for unsecured group communication, based on JSON-formatted configuration resources. For IPv6 CoAP groups, common multicast address ranges that are used to configure group addresses from are ff1x::/16 and ff3x::/16.

To create an application group, a configuring entity may configure a resource (name) or set of resources on CoAP endpoints, such that a CoAP request with Group URI sent by a configured CoAP client will be processed by one or more CoAP servers that have the matching URI path configured. These servers are the application group members.

To create a security group, a configuring entity defines an initial subset of the related security material. This comprises a set of group properties including the cryptographic algorithms and parameters used in the group, as well as additional information relevant throughout the group life-cycle, such as the security group name and description. This task MAY be entrusted to a dedicated administrator, that interacts with a Group Manager as defined in [Section 5](#). After that, further security materials to protect group communications have to be generated, compatible with the specified group configuration.

To participate in a security group, CoAP endpoints have to be configured with the group security material used to protect communications in the associated application/CoAP groups. The part of the process that involves secure distribution of group security material MAY use standardized communication with a Group Manager as defined in [Section 5](#). For unsecure group communication using the "NoSec" security group, any CoAP endpoint may become a group member at any time: there is no configuring entity that needs to provide security material for this group, as there is no security material for it. This means that group creation and membership cannot be tightly controlled for the "NoSec" group.

The configuration of groups and membership may be performed at different moments in the life-cycle of a device; for example during product (software) creation, in the factory, at a reseller, on-site during first deployment, or on-site during a system reconfiguration operation.

2.2.3. Group Discovery

It is possible for CoAP endpoints to discover application groups as well as CoAP groups, by using the RD-Groups usage pattern of the CoRE Resource Directory (RD), as defined in [Appendix A](#) of [\[I-D.ietf-core-resource-directory\]](#). In particular, an application group can be registered to the RD, specifying the reference IP multicast address, hence its associated CoAP group. The registration

is typically performed by a Commissioning Tool. Later on, CoAP endpoints can discover the registered application groups and related CoAP group, by using the lookup interface of the RD.

CoAP endpoints can also discover application groups by performing a multicast discovery query using the /.well-known/core resource. Such a request may be sent to a known CoAP group multicast address associated to application group(s), or to the All CoAP Nodes multicast address.

When secure communication is provided with Group OSCORE (see [Section 5](#)), the approach described in [\[I-D.tiloca-core-oscore-discovery\]](#) and also based on the RD can be used, in order to discover the security group to join.

In particular, the responsible OSCORE Group Manager registers its own security groups to the RD, as links to its own corresponding resources for joining the security groups [\[I-D.ietf-ace-key-groupcomm-oscore\]](#). Later on, CoAP endpoints can discover the registered security groups and related application groups, by using the lookup interface of the RD, and then join the security group through the respective Group Manager.

[2.2.4. Group Maintenance](#)

Maintenance of a group includes any necessary operations to cope with changes in a system, such as: adding group members, removing group members, changing group security material, reconfiguration of UDP port and/or IP multicast address, reconfiguration of the Group URI, renaming of application groups, splitting of groups, or merging of groups.

For unsecured group communication (see [Section 4](#)) i.e. the "NoSec" security group, addition/removal of CoAP group members is simply done by configuring these devices to start/stop listening to the group IP multicast address on the group's UDP port.

For secured group communication (see [Section 5](#)), the maintenance operations of the protocol Group OSCORE [\[I-D.ietf-core-oscore-groupcomm\]](#) MUST be implemented. When using Group OSCORE, CoAP endpoints participating in group communication are also members of a corresponding OSCORE security group, and thus share common security material. Additional related maintenance operations are discussed in [Section 5.1](#).

3. CoAP Usage in Group Communication

This section specifies the usage of CoAP in group communication, both unsecured and secured. This includes additional support for protocol extensions, such as Observe (see [Section 3.6](#)) and block-wise transfer (see [Section 3.7](#)).

How CoAP group messages are carried over various transport layers is the subject of [Section 3.8](#). Finally, [Section 3.9](#) covers the interworking of CoAP group communication with other protocols that may operate in the same network.

3.1. Request/Response Model

A CoAP client is an endpoint able to transmit CoAP requests and receive CoAP responses. Since the underlying UDP transport supports multiplexing by means of UDP port number, there can be multiple independent CoAP clients operational on a single host. On each UDP port, an independent CoAP client can be hosted. Each independent CoAP client sends requests that use the associated endpoint's UDP port number as the UDP source port of the request.

All CoAP requests that are sent via IP multicast MUST be Non-confirmable; see [Section 8.1 of \[RFC7252\]](#). The Message ID in an IP multicast CoAP message is used for optional message deduplication by both clients and servers, as detailed in [Section 4.5 of \[RFC7252\]](#).

A server sends back a unicast response to the CoAP group request - but the server MAY suppress the response for various reasons given in [Section 8.2 of \[RFC7252\]](#). This document adds the requirement that a server SHOULD suppress the response in case of error or in case there is nothing useful to respond, unless the application related to a particular resource requires such a response to be made for that resource. The unicast responses received by the CoAP client may be a mixture of success (e.g., 2.05 Content) and failure (e.g., 4.04 Not Found) codes, depending on the individual server processing results.

The CoAP No-Response Option [[RFC7967](#)] can be used by a client to influence the default response suppression on the server side. It is RECOMMENDED for a server to support this option only on selected resources where it is useful in the application context. If the option is supported on a resource, it MUST override the default response suppression of that resource.

Any default response suppression by a server SHOULD be performed consistently, as follows: if a request on a resource produces a particular Response Code and this response is not suppressed, then another request on the same resource that produces a response of the

same Response Code class is also not suppressed. For example, if a 4.05 Method Not Allowed error response code is suppressed by default on a resource, then a 4.15 Unsupported Content-Format error response code is also suppressed by default for that resource.

A CoAP client MAY repeat a multicast request using the same Token value and same Message ID value, in order to ensure that enough (or all) group members have been reached with the request. This is useful in case a number of group members did not respond to the initial request and the client suspects that the request did not reach these group members. However, in case one or more servers did receive the initial request but the response to that request was lost, this repeat does not help to retrieve the lost response(s) if the server(s) implement the optional Message ID based deduplication ([Section 4.5 of \[RFC7252\]](#)).

A CoAP client MAY repeat a multicast request using the same Token value and a different Message ID, in which case all servers that received the initial request will again process the repeated request since it appears within a new CoAP message. This is useful in case a client suspects that one or more response(s) to its original request were lost and the client needs to collect more, or even all, responses from group members, even if this comes at the cost of the overhead of certain group members responding twice (once to the original request, and once to the repeated request with different Message ID).

The CoAP client can distinguish the origin of multiple server responses by the source IP address of the UDP message containing the CoAP response and/or any other available application-specific source identifiers contained in the CoAP response payload or CoAP response options, such as an application-level unique ID associated to the server. If secure communication is provided with Group OSCORE (see [Section 5](#)), additional security-related identifiers in the CoAP response enable the client to retrieve the right security material for decrypting each response and authenticating its source.

While processing a response on the client, the source endpoint of the response is not matched to the destination endpoint of the request, since for a multicast request these will never match. This is specified in [Section 8.2 of \[RFC7252\]](#). It implies also that a server MAY respond from a UDP port number that differs from the destination UDP port number of the request, although a CoAP server normally SHOULD respond from the UDP port number that equals the destination port of the request - following the convention for UDP-based protocols. In case a single client has sent multiple group requests and concurrent CoAP transactions are ongoing, the responses received by that client are matched to an active request using only the Token

value. Due to UDP level multiplexing, the UDP destination port of the response MUST match to the client endpoint's UDP port value, i.e. to the UDP source port of the client's request.

For multicast CoAP requests, there are additional constraints on the reuse of Token values at the client, compared to the unicast case defined in [\[RFC7252\]](#) and updated by [\[I-D.ietf-core-echo-request-tag\]](#). Since for multicast CoAP the number of responses is not bound a priori, the client cannot use the reception of a response as a trigger to "free up" a Token value for reuse. Reusing a Token value too early could lead to incorrect response/request matching on the client, and would be a protocol error. Therefore, the time between reuse of Token values for different multicast requests MUST be greater than:

$$\text{MIN_TOKEN_REUSE_TIME} = (\text{NON_LIFETIME} + \text{MAX_LATENCY} + \text{MAX_SERVER_RESPONSE_DELAY})$$

where NON_LIFETIME and MAX_LATENCY are defined in [Section 4.8 of \[RFC7252\]](#). This specification defines MAX_SERVER_RESPONSE_DELAY as in [\[RFC7390\]](#), that is: the expected maximum response delay over all servers that the client can send a multicast request to. This delay includes the maximum Leisure time period as defined in [Section 8.2 of \[RFC7252\]](#). However, CoAP does not define a time limit for the server response delay. Using the default CoAP parameters, the Token reuse time MUST be greater than 250 seconds plus MAX_SERVER_RESPONSE_DELAY. A preferred solution to meet this requirement is to generate a new unique Token for every new multicast request, such that a Token value is never reused. If a client has to reuse Token values for some reason, and also MAX_SERVER_RESPONSE_DELAY is unknown, then using MAX_SERVER_RESPONSE_DELAY = 250 seconds is a reasonable guideline. The time between Token reuses is in that case set to a value greater than MIN_TOKEN_REUSE_TIME = 500 seconds.

When securing CoAP group communication with Group OSCORE [\[I-D.ietf-core-oscore-groupcomm\]](#), secure binding between requests and responses is ensured (see [Section 5](#)). Thus, a client may reuse a Token value after it has been freed up, as discussed above for the multicast case and considering a reuse time greater than MIN_TOKEN_REUSE_TIME. If an alternative security protocol for CoAP group communication is defined in the future which does not ensure secure binding between requests and responses, a client MUST follow the Token processing requirements as defined in [\[I-D.ietf-core-echo-request-tag\]](#).

Another method to more easily meet the above constraint is to instantiate multiple CoAP clients at multiple UDP ports on the same host. The Token values only have to be unique within the context of

a single CoAP client, so using multiple clients can make it easier to meet the constraint.

Since a client sending a multicast request with a Token T will accept multiple responses with the same Token T, it is possible in particular that the same server sends multiple responses with the same Token T back to the client. For example, this server might not implement the optional CoAP message deduplication based on Message ID; or it might be acting out of specification as a malicious, compromised or faulty server.

When this happens, the client normally processes at the CoAP layer each of those responses to the same request coming from the same server. If the processing of a response is successful, the client delivers this response to the application as usual.

Then, the application is in a better position to decide what to do, depending on the available context information. For instance, it might accept and process all the responses from the same server, even if they are not Observe notifications (i.e., they do not include an Observe option). Alternatively, the application might accept and process only one of those responses, such as the most recent one from that server, e.g. when this can trigger a change of state within the application.

3.2. Caching

CoAP endpoints that are members of a CoAP group MAY cache responses to a group request as defined in [Section 5.6 of \[RFC7252\]](#). In particular, these same rules apply to determine the set of request options used as "Cache-Key".

Furthermore, building on what is defined in [Section 8.2.1 of \[RFC7252\]](#):

- o A client sending a GET or FETCH group request over multicast MAY update a cache with the responses from the servers in the CoAP group. Then, the client uses both cached-still-fresh and new responses as the result of the group request.
- o A client sending a GET or FETCH group request over multicast MAY use a response received from a server, to satisfy a subsequent sent request intended to that server on the related unicast request URI. In particular, the unicast request URI is obtained by replacing the authority part of the request URI with the transport-layer source address of the cached response message.

- o A client MAY revalidate a cached response by making a GET or FETCH request on the related unicast request URI.

Note that, in the presence of proxies, doing any of the above (optional) unicast requests requires the client to distinguish the different responses to a group request, as well as to distinguish the different origin servers that responded. This in turn requires additional means to provide the client with information about the origin server of each response, as discussed in [Section 3.4.3](#).

The following subsections define the freshness model and validation model to use for cached responses, which update the models defined in [Section 5.6.1](#) and [Section 5.6.2 of \[RFC7252\]](#), respectively.

[3.2.1](#). Freshness Model

For caching at endpoints, the same freshness model relying on the Max-Age Option as defined in [Section 5.6.1 of \[RFC7252\]](#) applies.

For caching at proxies, the freshness model defined in [Section 3.4.3](#) of this specification applies.

[3.2.2](#). Validation Model

[Section 5.6.2 of \[RFC7252\]](#) defines a model to "validate" or "revalidate" responses stored in cache, hence enabling the suppression of responses that the client already has.

This relies on the ETag Option defined in [Section 5.10.6 of \[RFC7252\]](#), with its usage limited to exchanges between a CoAP client and one CoAP server. That is, [Section 8.2.1 of \[RFC7252\]](#) explicitly forbids using an ETag Option in requests sent over multicast, and leaves a mechanism to suppress responses for that case for further study.

This section provides such a model to "validate" or "revalidate" responses that the client already has cached. In particular, the group request can indicate entity-tag values separately for each CoAP server from which the client wishes to get a response revalidation, together with addressing information identifying that server.

To this end, this specification defines the new Multi-ETag Option. Operations related to this validation model and using the new option are defined in [Section 3.2.2.2](#) for the client side, and in [Section 3.2.2.3](#) for the server side.

The Multi-ETag Option has the properties summarized in Figure 3, which extends Table 4 of [\[RFC7252\]](#). The Multi-ETag Option is elective, safe to forward, part of the cache key, and repeatable.

The option is intended only for group requests, as directly sent to a CoAP group or to a CoAP proxy that forwards it to the CoAP group (see [Section 3.4](#)).

No.	C	U	N	R	Name	Format	Length	Default
TBD1				x	Multi-ETag	(*)	any	(none)

C=Critical, U=Unsafe, N=NoCacheKey, R=Repeatable

(*) See below.

Figure 3: The Multi-ETag Option.

The Multi-ETag Option has the same properties of the ETag Option defined in [Section 5.10.6 of \[RFC7252\]](#), but it differs in the format and length, as well as having a different reason for its repeatability.

Each occurrence of the Multi-ETag Option targets exactly one of the servers in the CoAP group, from which the client wishes to get a response revalidation. The option value is set to a CBOR sequence [\[RFC8742\]](#) composed of (1+M) elements, where:

- o The first element specifies the addressing information of the corresponding server, encoded as defined in [Section 3.2.2.1](#).

This mirrors the format of the Multicast-Signaling option defined in Section 3 of [\[I-D.tiloca-core-groupcomm-proxy\]](#). Thus, in the presence of a forward proxy supporting the mechanism defined in [\[I-D.tiloca-core-groupcomm-proxy\]](#), the client can seamlessly use the server addressing information obtained from the proxy, when this forwards back a response to a group request from that server.

- o The following M elements are CBOR byte strings, each of which has as value an entity-tag value that the client wants to try against the corresponding server.

The entity-tag values included in the Multi-ETag Option are subject to the same considerations for the entity-tag values used in an ETag Option (see [Section 5.10.6 of \[RFC7252\]](#)).

The Multi-ETag Option is of class E in terms of OSCORE processing (see [Section 4.1 of \[RFC8613\]](#)).

3.2.2.1. Encoding of Server Addressing Information

The first element of the CBOR sequence in the Multi-ETag Option value is set to the byte serialization of the CBOR array 'tp_info' defined in Section 2.2.1 of [\[I-D.tiloca-core-observe-multicast-notifications\]](#), including only the set of elements 'srv_addr'.

In turn, the set includes the integer 'tp_id' identifying the used transport protocol, and further elements whose number, format and encoding depend on the value of 'tp_id'.

When the Multi-ETag Option is used in group requests transported over UDP as in this specification, the 'tp_info' array includes the following elements, encoded as defined in Section 2.2.1.1 of [\[I-D.tiloca-core-observe-multicast-notifications\]](#).

- o 'tp_id': the CBOR integer with value 1 ("UDP"), from the "Value" column of the "Transport Protocol Identifiers" Registry defined in Section 14.4 of [\[I-D.tiloca-core-observe-multicast-notifications\]](#)
- o 'srv_host': a CBOR byte string, with value the unicast IP address of the server. This element is tagged and identified by the CBOR tag 260 "Network Address (IPv4 or IPv6 or MAC Address)".
- o 'srv_port': as a CBOR unsigned integer or the CBOR simple value Null. If it is a CBOR integer, it has as value the destination port number where to send individual requests intended to the server. This element MAY be present. If not included, the default port number 5683 is assumed.

The CDDL notation [\[RFC8610\]](#) provided below describes the 'tp_info' CBOR array using the format above.

```
tp_info = [  
    tp_id : 1,                ; UDP as transport protocol  
    srv_host : #6.260(bstr),   ; IP address where to reach the server  
    srv_port : uint / null     ; Port number where to reach the server  
]
```

3.2.2.2. Processing on the Client Side

Similar to what is defined in [Section 5.6.2 of \[RFC7252\]](#), the client may have one or more stored responses for a GET or FETCH group

request sent to the CoAP group, but cannot use any of them (e.g. because they are not fresh).

In that case, the client can send a GET or FETCH group request, in order to give the origin servers an opportunity both to select a stored response to be used, and to update its freshness. As in [\[RFC7252\]](#), this process is known as "validating" or "revalidating" the stored response.

When sending such a group request, the endpoint SHOULD include one Multi-ETag Option for each server it wishes to revalidate the corresponding response with. As defined in [Section 3.2.2](#), the Multi-ETag Option can include multiple entity-tag values, each applicable to a stored response from the corresponding server for that group request.

Specifically, in the same GET or FETCH group request:

- o The client MUST NOT include one or more ETag Option(s) together with one or more Multi-ETag Option(s).
- o The client MUST include only one Multi-ETag Option for each server it wishes to get a response revalidation from.
- o The client SHOULD limit the number of Multi-ETag Options, hence limiting the number of servers as intended target of the revalidation process, and SHOULD rather spread revalidation with different sets of servers over different group requests. Also, the client SHOULD limit the number of entity-tag values specified in each Multi-ETag Option, preferably indicating only one entity-tag value.

This allows for limiting the overall size of the group request. As a guideline, the server addressing information can be 9-24 bytes in size, while each entity-tag value can be 1-8 bytes in size. Thus, a single Multi-ETag Option can be up to $(24 + 8 * M)$ bytes in size, where M is the number of entity-tag values it includes.

A 2.03 (Valid) response indicates that the stored response identified by the entity-tag given in the response's ETag Option can be reused, after updating the stored response as described in [Section 5.9.1.3 of \[RFC7252\]](#). So the client can determine if any one of the stored representations from that server is current, without need to transfer the current resource representation again.

Any other Response Code indicates that none of the stored responses from that server, identified in the Multi-ETag Option of the group

request, are suitable. Instead, such response SHOULD be used to satisfy the request and MAY replace the stored response.

3.2.2.3. Processing on the Server Side

If a GET or FETCH request includes both one or more ETag Options together with one or more Multi-ETag Options, then the server MUST ignore all the included ETag and Multi-ETag Options.

The server MUST ignore any Multi-ETag Option which is malformed, or included in a request that is neither GET nor FETCH, or which specifies addressing information not matching with its own endpoint address.

The server considers only its pertaining Multi-ETag Option, i.e. specifying addressing information associated to its own endpoint. The server MUST ignore any pertaining Multi-ETag Option that occurs more than once.

If the pertaining Multi-ETag Option specifies the CBOR simple value Null for the 'srv_port' element of 'tp_info' (see [Section 3.2.2.1](#)), the server MUST assume the default port number 5683.

Then, the server can issue a 2.03 (Valid) response in place of a 2.05 (Content) response, if one of the entity-tag values from the pertaining Multi-ETag Option is the entity-tag for the current resource representation, i.e. it is valid. The 2.03 (Valid) response echoes this specific entity-tag within an ETag Option included in the response.

The inclusion of an ETag Option in a response works as defined in [Section 5.6.10.1 of \[RFC7252\]](#).

3.3. Port and URI Path Selection

A server that is a member of a CoAP group listens for CoAP messages on the group's IP multicast address, usually on the CoAP default UDP port 5683, or another non-default UDP port if configured. Regardless of the method for selecting the port number, the same port number MUST be used across all CoAP servers that are members of a CoAP group and across all CoAP clients performing the requests to that group.

The URI Path used in the request is preferably a path that is known to be supported across all group members. However there are valid use cases where a request is known to be successful only for a subset of the CoAP group, for example only members of a specific application group, while those group members for which the request is unsuccessful (for example because they are outside the application

group) either ignore the multicast request or respond with an error status code.

One way to create multiple CoAP groups is using different UDP ports with the same IP multicast address, in case the devices' network stack only supports a limited number of multicast address subscriptions. However, it must be taken into account that this incurs additional processing overhead on each CoAP server participating in at least one of these groups: messages to groups that are not of interest to the node are only discarded at the higher transport (UDP) layer instead of directly at the network (IP) layer. Also, a constrained network may be additionally burdened in this case with multicast traffic that is eventually discarded at the UDP layer by most nodes.

Port 5684 is reserved for DTLS-secured unicast CoAP and MUST NOT be used for any CoAP group communication.

For a CoAP server node that supports resource discovery as defined in [Section 2.4 of \[RFC7252\]](#), the default port 5683 MUST be supported (see [Section 7.1 of \[RFC7252\]](#)) for the "All CoAP Nodes" multicast group as detailed in [Section 3.8](#).

3.4. Proxy Operation

This section defines how proxies operate in a group communication scenario. In particular, [Section 3.4.1](#) defines operations of forward-proxies, [Section 3.4.2](#) defines operations of reverse-proxies, and [Section 3.4.3](#) defines operations of proxies that employ a cache for responses to group requests.

3.4.1. Forward-Proxies

CoAP enables a client to request a forward-proxy to process a CoAP request on its behalf, as described in [Section 5.7.2](#) and 8.2.2 of [\[RFC7252\]](#). For this purpose, the client specifies either the request group URI as a string in the Proxy-URI Option or it uses the Proxy-Scheme Option with the group URI constructed from the usual Uri-* Options. The forward-proxy then resolves the group URI to a destination CoAP group, multicasts the CoAP request, receives the responses and forwards all the individual (unicast) responses back to the client.

However, there are certain issues and limitations with this approach:

- o The CoAP client component that sent a unicast CoAP request to the proxy may be expecting only one (unicast) response, as usual for a CoAP unicast request. Instead, it receives multiple (unicast)

responses, potentially leading to fault conditions in the component or to discarding any received responses following the first one. This issue may occur even if the application calling the CoAP client component is aware that the forward-proxy is going to execute a CoAP group URI request.

- o Each individual CoAP response received by the client will appear to originate (based on its IP source address) from the CoAP Proxy, and not from the server that produced the response. This makes it impossible for the client to identify the server that produced each response, unless the server identity is contained as a part of the response payload or inside a CoAP option in the response.
- o The proxy does not know how many members there are in the CoAP group or how many group members will actually respond. Also, the proxy does not know for how long to collect responses before it stops forwarding them to the client. A CoAP client that is not using a Proxy might face the same problems in collecting responses to a multicast request. However, the client itself would typically have application-specific rules or knowledge on how to handle this situation, while an application-agnostic CoAP Proxy would typically not have this knowledge. For example, a CoAP client could monitor incoming responses and use this information to decide how long to continue collecting responses - which is something a proxy cannot do.

A forward-proxying method using this approach and addressing the issues raised above is defined in [[I-D.tiloca-core-groupcomm-proxy](#)].

An alternative solution is for the proxy to collect all the individual (unicast) responses to a CoAP group request and then send back only a single (aggregated) response to the client. However, this solution brings up new issues:

- o Like for the approach discussed above, the proxy does not know for how long to collect responses before sending back the aggregated response to the client. Analogous considerations apply to this approach too, both on the client and proxy side.
- o There is no default format defined in CoAP for aggregation of multiple responses into a single response. Such a format could be standardized based on, for example, the multipart content-format [[RFC8710](#)].

Due to the above issues, it is RECOMMENDED that a CoAP Proxy only processes a group URI request if it is explicitly enabled to do so. The default response (if the function is not explicitly enabled) to a group URI request is 5.01 Not Implemented. Furthermore, a proxy

SHOULD be explicitly configured (e.g. by allow-listing and/or client authentication) to allow proxied CoAP multicast requests only from specific client(s).

The operation of HTTP-to-CoAP proxies for multicast CoAP requests is specified in [Section 8.4](#) and 10.1 of [\[RFC8075\]](#). In this case, the "application/http" media type is used to let the proxy return multiple CoAP responses - each translated to a HTTP response - back to the HTTP client. Of course, in this case the HTTP client sending a group URI to the proxy needs to be aware that it is going to receive this format, and needs to be able to decode it into the responses of multiple CoAP servers. Also, the IP source address of each CoAP response cannot be determined anymore from the "application/http" response. The HTTP client still identify the CoAP servers by other means such as application-specific information in the response payload.

[3.4.2.](#) Reverse-Proxies

CoAP enables the use of a reverse-proxy, as an endpoint that stands in for one or more other server(s), and satisfies requests on behalf of these, doing any necessary translations (see [Section 5.7.3 of \[RFC7252\]](#)).

In a group communication scenario, a reverse-proxy can rely on its configuration and/or on information in a request from a client, in order to determine that the request has to be forwarded to a group of servers over IP multicast. For example, specific resources on the reverse-proxy could be allocated, each to a specific application group and/or CoAP group. Or alternatively, the application group and/or CoAP group in question could be encoded as URI path segments. The URI path encodings for a reverse-proxy may also use a URI mapping template as described in [Section 5.4 of \[RFC8075\]](#).

Furthermore, the reverse-proxy can actually stand in for (and thus prevent to directly reach) only the whole set of servers in the group, or also for each of those individual servers (e.g. if acting as firewall).

For a reverse-proxy that forwards a request to a group of servers over IP multicast, the same considerations as defined in [Section 5.7.3 of \[RFC7252\]](#) hold, with the following additions:

- o The three issues and limitations defined in [Section 3.4.1](#) for a forward proxy apply to a reverse-proxy as well, and have to be addressed, e.g. using the signaling method defined in [\[I-D.tiloca-core-groupcomm-proxy\]](#) or other means.

- o A reverse-proxy MAY have preconfigured time duration(s) that are used for the collecting of server responses and forwarding these back to the client. These duration(s) may be set as global configuration or resource-specific configurations. If there is such preconfiguration, then an explicit signaling of the time period in the client's request as defined in [\[I-D.tiloca-core-groupcomm-proxy\]](#) is not necessarily needed.
- o A client that is configured to access a reverse-proxy resource (i.e. one that triggers a CoAP group communication request) SHOULD be configured also to handle potentially multiple responses with the same Token value caused by a single request.

That is, the client needs to preserve the Token value used for the request also after the reception of the first response forwarded back by the proxy (see [Section 3.1](#)) and keep the request open to potential further responses with this Token. This requirement can be met by a combination of client implementation and proper proxied group communication configuration on the client.

- o A client might re-use a Token value in a valid new request to the reverse-proxy, while the reverse-proxy still has an ongoing group communication request for this client with the same Token value (i.e. its time period for response collection has not ended yet).

If this happens, the reverse-proxy MUST stop the ongoing request and associated response forwarding, it MUST NOT forward the new request to the group of servers, and it MUST send a 4.00 Bad Request error response to the client. The diagnostic payload of the error response SHOULD indicate to the client that the resource is a reverse-proxy resource, and that for this reason immediate Token re-use is not possible.

If the reverse-proxy supports the signalling protocol of [\[I-D.tiloca-core-groupcomm-proxy\]](#) it can include a Multicast-Signaling Option in the error response to convey the reason for the error in a machine-readable way.

For the operation of HTTP-to-CoAP reverse proxies, see the last paragraph of [Section 3.4.1](#) which applies also to this case.

[3.4.3. Caching](#)

A proxy that supports forwarding of group requests and that employs a cache maintains the following two types of cache entry.

- o The first type, "individual" cache entry, is associated to one server and stores one response from that server, regardless

whether it is a response to a unicast request or to a group request.

A hit to this entry would be produced by a matching request intended to that server, i.e. to the corresponding unicast URI.

When the response is a response to a unicast request to the server, the unicast URI is the same target URI used for the request.

When the response is a response to a group request to the CoAP group, the unicast URI is obtained by replacing the authority part of the group URI in the group request with the transport-layer source address and port number of the response message.

- o The second type, "aggregated" cache entry, is associated to the CoAP group, and stores all the responses that: the proxy has received as a response to a group request to that group; and that have been also forwarded back to the client that sent the group request.

A hit to this entry would be produced by a matching group request intended to the CoAP group, i.e. to the corresponding group URI.

When forwarding a group request to a CoAP group using the request's group URI and processing the responses, the proxy handles its cache entries as follows. The same applies if the proxy spontaneously re-sends a group request to the CoAP group, in order to refresh an aggregated cache entry after its expiration or invalidation.

1. For each response to the group request which is received and also forwarded back to the client:

- * The proxy creates or refreshes the individual cache entry associated to the origin server and for that response. That is, the response is stored in the individual cache entry, and the lifetime of the cache entry is set to the lifetime of the response, as indicated by the Max-Age Option if present, or as the default value of 60 seconds otherwise (see [Section 5.6.1 of \[RFC7252\]](#)). This cache entry becomes immediately usable to serve requests from clients.

- * The proxy adds the response to a temporary list L.

2. After stopping to forward the received responses back to the client:

- * The proxy creates an aggregated cache entry associated to the group for that group request, if not existing yet. In case of an existing entry to be refreshed, the proxy deletes all the responses stored in the entry.
- * The proxy stores all the responses from the list L in the aggregated cache entry.
- * The proxy sets the lifetime of the cache entry to the smallest lifetime among all the responses stored in the entry, determined in the same way as defined in step 1 above.
- * The proxy sets the aggregated cache entry as usable to serve group requests from clients.

When forwarding a request to an individual server using the associated unicast URI and processing its response, the proxy handles its cache entries as follows. The same applies if the proxy spontaneously re-sends a unicast request to a single server, in order to refresh an individual cache entry after its expiration or invalidation.

1. The proxy creates or refreshes the individual cache entry associated to the origin server and for that response. That is, the response is stored in the cache entry, and the lifetime of the cache entry is set to the lifetime of the response, as indicated by the Max-Age Option if present, or as the default value of 60 seconds otherwise (see [Section 5.6.1 of \[RFC7252\]](#)). This cache entry becomes immediately usable to serve requests from clients.
2. The proxy checks whether it has a non-expired and valid aggregated cache entry, such that a hit would be produced by a group request analogous to the forwarded unicast request.

That is, such group request would be intended to the group URI of the CoAP group associated to the aggregated cache entry, rather than intended to the unicast URI of the forwarded request.

3. If an aggregated cache entry is found at the previous step:
 - * The proxy stores the received response in the aggregated cache entry, possibly replacing an already stored instance of that response from that origin server.
 - * The proxy sets as new lifetime of the aggregated cache entry the minimum value between the current lifetime of the cache entry and the lifetime of the just-stored response, as

indicated by the Max-Age Option if present, or as the default value of 60 seconds otherwise (see [Section 5.6.1 of \[RFC7252\]](#)).

Note that a proxy embedded in a router can monitor network control messages, hence learning when a new server has joined a CoAP group and is listening to the multicast IP address of that CoAP group. This information could be used to guide the proxy in refreshing an aggregated cache entry, by sending a request to the CoAP group over the group URI before the entry expires, and thus storing also a response from the newly joined server.

Following the expiration or invalidation of a cache entry, as well as if wishing to refresh a cache entry, the proxy can directly interact with the servers in the CoAP group. To this end, it takes the role of a CoAP client as defined in [Section 3.2](#). In particular, the proxy can perform revalidation of responses to group requests by using the Multi-ETag Option, as defined in [Section 3.2.2](#).

As further discussed in [Section 5.2](#), additional means are required to enable cachability of responses at the proxy when communications in the group are secured with Group OSCORE [\[I-D.ietf-core-oscore-groupcomm\]](#).

3.4.3.1. Validation of Responses to a Group Request

A client can revalidate the full set of responses to a group request from the corresponding aggregated cache entry at the proxy. To this end, this specification defines the new Group-ETag Option.

The Group-ETag Option has the properties summarized in Figure 4, which extends Table 4 of [\[RFC7252\]](#). The Group-ETag Option is elective, safe to forward, part of the cache key, and repeatable.

The option is intended for group requests sent to a Forward-Proxy, as well as for the associated responses retrieved from the corresponding aggregated cache entry at the proxy.

No.	C	U	N	R	Name	Format	Length	Default
TBD2				x	Group-ETag	opaque	1-8	(none)

C=Critical, U=Unsafe, N=NoCacheKey, R=Repeatable

Figure 4: The Group-ETag Option.

The Group-ETag Option has the same properties of the ETag Option defined in [Section 5.10.6 of \[RFC7252\]](#).

The Group-ETag Option is of class U in terms of OSCORE processing (see [Section 4.1 of \[RFC8613\]](#)).

When providing 2.05 (Content) responses to a GET or FETCH group request from an aggregated cache entry, the proxy can include one Group-ETag Option, specifying the current entity-tag value associated to that cache entry. Each of such responses MUST NOT include more than one Group-ETag Option.

If the proxy supports this form of response revalidation, it MUST update the current entity-tag value associated to an aggregated cache entry, every time a response is added to that cache entry or replaces an already included response.

When sending a GET or FETCH group request to the proxy, to be forwarded to a CoAP group, the client can include one or more Group-ETag Option(s). Each option specifies one entity-tag value, as applicable to the aggregated cache entry for that group request.

In case the group request hits an aggregated cache entry and its current entity-tag value matches with one of the entity-tag value(s) specified in the Group-ETag option(s), then the proxy replies with a single 2.03 (Valid) response. This response has no payload and MUST include one Group-ETag Option, specifying the current entity-tag value of the aggregated cache entry.

That is, the 2.03 (Valid) response from the proxy indicates that the stored responses identified by the entity-tag given in the response's Group-ETag Option can be reused, after updating each of them as described in [Section 5.9.1.3 of \[RFC7252\]](#). In effect, the client can determine if any of the stored set of representations from the aggregated cache entry at the proxy is current, without needing to transfer any of them again.

Note that, if a client triggers the proxy to perform forwarding of a group request (i.e., there is no hit of an aggregated cache entry), this will result in a new aggregated cache entry created at the proxy. Then, the client cannot obtain an entity-tag value through a Group-ETag Option in any of the responses forwarded back by the proxy.

In fact, the proxy will only have an assigned entity-tag value to provide after all responses have been forwarded back to that client, which is the moment that the new aggregated cache entry is eventually created. However, when follow-up group requests from the same client

or different clients are served from this aggregated cache entry, the proxy can include a Group-ETag Option in each returned response, specifying the current entity-tag for the aggregated cache entry.

3.5. Congestion Control

CoAP group requests may result in a multitude of responses from different nodes, potentially causing congestion. Therefore, both the sending of IP multicast requests and the sending of the unicast CoAP responses to these multicast requests should be conservatively controlled.

CoAP [[RFC7252](#)] reduces IP multicast-specific congestion risks through the following measures:

- o A server may choose not to respond to an IP multicast request if there is nothing useful to respond to, e.g., error or empty response (see [Section 8.2 of \[RFC7252\]](#)).
- o A server should limit the support for IP multicast requests to specific resources where multicast operation is required ([Section 11.3 of \[RFC7252\]](#)).
- o An IP multicast request MUST be Non-confirmable ([Section 8.1 of \[RFC7252\]](#)).
- o A response to an IP multicast request SHOULD be Non-confirmable ([Section 5.2.3 of \[RFC7252\]](#)).
- o A server does not respond immediately to an IP multicast request and should first wait for a time that is randomly picked within a predetermined time interval called the Leisure ([Section 8.2 of \[RFC7252\]](#)).

Additional guidelines to reduce congestion risks defined in this document are as follows:

- o A server in a constrained network SHOULD only support group communication for resources that have a small representation (where the representation may be retrieved via a GET, FETCH or POST method in the request). For example, "small" can be defined as a response payload limited to approximately 5% of the IP Maximum Transmit Unit (MTU) size, so that it fits into a single link-layer frame in case IPv6 over Low-Power Wireless Personal Area Networks (6LoWPAN, see [Section 3.8.3](#)) is used on the constrained network.

- o A server SHOULD minimize the payload size of a response to a multicast GET or FETCH on `"/.well-known/core"` by using hierarchy in arranging link descriptions for the response. An example of this is given in [Section 5 of \[RFC6690\]](#).
- o A server MAY minimize the payload size of a response to a multicast GET or FETCH (e.g., on `"/.well-known/core"`) by using CoAP block-wise transfers [\[RFC7959\]](#) in case the payload is long, returning only a first block of the CoRE Link Format description. For this reason, a CoAP client sending an IP multicast CoAP request to `"/.well-known/core"` SHOULD support block-wise transfers. See also [Section 3.7](#).
- o A client SHOULD be configured to use CoAP groups with the smallest possible IP multicast scope that fulfills the application needs. As an example, site-local scope is always preferred over global scope IP multicast if this fulfills the application needs. Similarly, realm-local scope is always preferred over site-local scope if this fulfills the application needs.

[3.6](#). Observing Resources

The CoAP Observe Option [\[RFC7641\]](#) is a protocol extension of CoAP, that allows a CoAP client to retrieve a representation of a resource and automatically keep this representation up-to-date over a longer period of time. The client gets notified when the representation has changed. [\[RFC7641\]](#) does not mention whether the Observe Option can be combined with CoAP multicast.

This section updates [\[RFC7641\]](#) with the use of the Observe Option in a CoAP multicast GET request, and defines normative behavior for both client and server. Consistent with [Section 2.4 of \[RFC8132\]](#), it is also possible to use the Observe Option in a CoAP multicast FETCH request.

Multicast Observe is a useful way to start observing a particular resource on all members of a CoAP group at the same time. Group members that do not have this particular resource or do not allow the GET or FETCH method on it will either respond with an error status - 4.04 Not Found or 4.05 Method Not Allowed, respectively - or will silently suppress the response following the rules of [Section 3.1](#), depending on server-specific configuration.

A client that sends a multicast GET or FETCH request with the Observe Option MAY repeat this request using the same Token value and the same Observe Option value, in order to ensure that enough (or all) members of the CoAP group have been reached with the request. This is useful in case a number of group members did not respond to the

initial request. The client MAY additionally use the same Message ID in the repeated request to avoid that group members that had already received the initial request would respond again. Note that using the same Message ID in a repeated request will not be helpful in case of loss of a response message, since the server that responded already will consider the repeated request as a duplicate message. On the other hand, if the client uses a different, fresh Message ID in the repeated request, then all the group members that receive this new message will typically respond again, which increases the network load.

A client that has sent a multicast GET or FETCH request with the Observe Option MAY follow up by sending a new unicast CON request with the same Token value and same Observe Option value to a particular server, in order to ensure that the particular server receives the request. This is useful in case a specific group member, that was expected to respond to the initial group request, did not respond to the initial request. In this case, the client MUST use a Message ID that differs from the initial multicast message.

Furthermore, consistent with [Section 3.3.1 of \[RFC7641\]](#) and following its guidelines, a client MAY at any time send a new multicast GET or FETCH request with the same Token value and same Observe Option value as the original request. This allows the client to verify that it has an up-to-date representation of an observed resource and/or to re-register its interest to observe a resource.

In the above client behaviors, the Token value is kept identical to the initial request to avoid that a client is included in more than one entry in the list of observers ([Section 4.1 of \[RFC7641\]](#)).

Before repeating a request as specified above, the client SHOULD wait for at least the expected round-trip time plus the Leisure time period defined in [Section 8.2 of \[RFC7252\]](#), to give the server time to respond.

A server that receives a GET or FETCH request with the Observe Option, for which request processing is successful, SHOULD respond to this request and not suppress the response. A server that adds a client to the list (as a new entry) of observers for a resource due to an Observe request MUST respond to this request and not suppress it.

A server SHOULD have a mechanism to verify liveness of its observing clients and the continued interest of these clients in receiving the observe notifications. This can be implemented by sending notifications occasionally using a Confirmable message. See

[Section 4.5 of \[RFC7641\]](#) for details. This requirement overrides the regular behavior of sending Non-Confirmable notifications in response to a Non-Confirmable request.

A client can use the unicast cancellation methods of [Section 3.6 of \[RFC7641\]](#) and stop the ongoing observation of a particular resource on members of a CoAP group. This can be used to remove specific observed servers, or even all servers in the group (using serial unicast to each known group member). In addition, a client MAY explicitly deregister from all those servers at once, by sending a multicast GET or FETCH request that includes the Token value of the observation to be cancelled and includes an Observe Option with the value set to 1 (deregister). In case not all the servers in the CoAP group received this deregistration request, either the unicast cancellation methods can be used at a later point in time or the multicast deregistration request MAY be repeated upon receiving another observe response from a server.

For observing a group of servers through a CoAP-to-CoAP proxy, the limitations stated in [Section 3.4](#) apply. The method defined in [\[I-D.tiloca-core-groupcomm-proxy\]](#) enables group communication including resource observation through proxies and addresses those limitations.

3.7. Block-Wise Transfer

[Section 2.8 of \[RFC7959\]](#) specifies how a client can use block-wise transfer (Block2 Option) in a multicast GET request to limit the size of the initial response of each server. Consistent with [Section 2.5 of \[RFC8132\]](#), the same can be done with a multicast FETCH request.

The client has to use unicast for any further request, separately addressing each different server, in order to retrieve more blocks of the resource from that server, if any. Also, a server (member of a targeted CoAP group) that needs to respond to a multicast request with a particularly large resource can use block-wise transfer (Block2 Option) at its own initiative, to limit the size of the initial response. Again, a client would have to use unicast for any further requests to retrieve more blocks of the resource.

A solution for multicast block-wise transfer using the Block1 Option is not specified in [\[RFC7959\]](#) nor in the present document. Such a solution would be useful for multicast FETCH/PUT/POST/PATCH/iPATCH requests, to efficiently distribute a large request payload as multiple blocks to all members of a CoAP group. Multicast usage of Block1 is non-trivial due to potential message loss (leading to missing blocks or missing confirmations), and potential diverging block size preferences of different members of the CoAP group.

3.8. Transport

In this document only UDP is considered as a transport protocol, both over IPv4 and IPv6. Therefore, [\[RFC8323\]](#) (CoAP over TCP, TLS, and WebSockets) is not in scope as a transport for group communication.

3.8.1. UDP/IPv6 Multicast Transport

CoAP group communication can use UDP over IPv6 as a transport protocol, provided that IPv6 multicast is enabled. IPv6 multicast MAY be supported in a network only for a limited scope. For example, [Section 3.9.2](#) describes the potential limited support of RPL for multicast, depending on how the protocol is configured.

For a CoAP server node that supports resource discovery as defined in [Section 2.4 of \[RFC7252\]](#), the default port 5683 MUST be supported as per [Section 7.1](#) and 12.8 of [\[RFC7252\]](#) for the "All CoAP Nodes" multicast group. An IPv6 CoAP server SHOULD support the "All CoAP Nodes" groups with at least link-local (2), admin-local (4) and site-local (5) scopes. An IPv6 CoAP server on a 6LoWPAN node (see [Section 3.8.3](#)) SHOULD also support the realm-local (3) scope.

Note that a client sending an IPv6 multicast CoAP message to a port that is not supported by the server will not receive an ICMPv6 Port Unreachable error message from that server, because the server does not send it in this case, per [Section 2.4 of \[RFC4443\]](#).

3.8.2. UDP/IPv4 Multicast Transport

CoAP group communication can use UDP over IPv4 as a transport protocol, provided that IPv4 multicast is enabled. For a CoAP server node that supports resource discovery as defined in [Section 2.4 of \[RFC7252\]](#), the default port 5683 MUST be supported as per [Section 7.1](#) and 12.8 of [\[RFC7252\]](#), for the "All CoAP Nodes" IPv4 multicast group.

Note that a client sending an IPv4 multicast CoAP message to a port that is not supported by the server will not receive an ICMP Port Unreachable error message from that server, because the server does not send it in this case, per [Section 3.2.2 of \[RFC1122\]](#).

3.8.3. 6LoWPAN

In 6LoWPAN [\[RFC4944\]](#) [\[RFC6282\]](#) networks, IPv6 packets (up to 1280 bytes) may be fragmented into smaller IEEE 802.15.4 MAC frames (up to 127 bytes), if the packet size requires this. Every 6LoWPAN IPv6 router that receives a multi-fragment packet reassembles the packet and refragments it upon transmission. Since the loss of a single fragment implies the loss of the entire IPv6 packet, the performance

in terms of packet loss and throughput of multi-fragment multicast IPv6 packets is typically far worse than the performance of single-fragment IPv6 multicast packets. For this reason, a CoAP request sent over multicast in 6LoWPAN networks SHOULD be sized in such a way that it fits in a single IEEE 802.15.4 MAC frame, if possible.

On 6LoWPAN networks, multicast groups can be defined with realm-local scope [[RFC7346](#)]. Such a realm-local group is restricted to the local 6LoWPAN network/subnet. In other words, a multicast request to that group does not propagate beyond the 6LoWPAN network segment where the request originated. For example, a multicast discovery request can be sent to the realm-local "All CoAP Nodes" IPv6 multicast group (see [Section 3.8.1](#)) in order to discover only CoAP servers on the local 6LoWPAN network.

[3.9](#). Interworking with Other Protocols

[3.9.1](#). MLD/MLDv2/IGMP/IGMPv3

CoAP nodes that are IP hosts (i.e., not IP routers) are generally unaware of the specific IP multicast routing/forwarding protocol being used in their network. When such a host needs to join a specific (CoAP) multicast group, it requires a way to signal to IP multicast routers which IP multicast address(es) it needs to listen to.

The MLDv2 protocol [[RFC3810](#)] is the standard IPv6 method to achieve this; therefore, this method SHOULD be used by members of a CoAP group to subscribe to its multicast IPv6 address, on IPv6 networks that support it. CoAP server nodes then act in the role of MLD Multicast Address Listener. MLDv2 uses link-local communication between Listeners and IP multicast routers. Constrained IPv6 networks that implement either RPL (see [Section 3.9.2](#)) or MPL (see [Section 3.9.3](#)) typically do not support MLDv2 as they have their own mechanisms defined for subscribing to multicast groups.

The IGMPv3 protocol [[RFC3376](#)] is the standard IPv4 method to signal multicast group subscriptions. This SHOULD be used by members of a CoAP group to subscribe to its multicast IPv4 address on IPv4 networks.

The guidelines from [[RFC6636](#)] on the tuning of MLD for mobile and wireless networks may be useful when implementing MLD in constrained networks.

3.9.2. RPL

RPL [[RFC6550](#)] is an IPv6 based routing protocol suitable for low-power, lossy networks (LLNs). In such a context, CoAP is often used as an application protocol.

If only RPL is used in a network for routing and its optional multicast support is disabled, there will be no IP multicast routing available. Any IPv6 multicast packets in this case will not propagate beyond a single hop (to direct neighbors in the LLN). This implies that any CoAP group request will be delivered to link-local nodes only, for any scope value ≥ 2 used in the IPv6 destination address.

RPL supports (see [Section 12 of \[RFC6550\]](#)) advertisement of IP multicast destinations using Destination Advertisement Object (DAO) messages and subsequent routing of multicast IPv6 packets based on this. It requires the RPL mode of operation to be 3 (Storing mode with multicast support).

In this mode, RPL DAO can be used by a CoAP node that is either an RPL router or RPL Leaf Node, to advertise its CoAP group membership to parent RPL routers. Then, RPL will route any IP multicast CoAP requests over multiple hops to those CoAP servers that are group members.

The same DAO mechanism can be used to convey CoAP group membership information to an edge router (e.g., 6LBR), in case the edge router is also the root of the RPL Destination-Oriented Directed Acyclic Graph (DODAG). This is useful because the edge router then learns which IP multicast traffic it needs to pass through from the backbone network into the LLN subnet, and which traffic not. In LLNs, such ingress filtering helps to avoid congestion of the resource-constrained network segment, due to IP multicast traffic from the high-speed backbone IP network.

3.9.3. MPL

The Multicast Protocol for Low-Power and Lossy Networks (MPL) [[RFC7731](#)] can be used for propagation of IPv6 multicast packets throughout a defined network domain, over multiple hops. MPL is designed to work in LLNs and can operate alone or in combination with RPL. The protocol involves a predefined group of MPL Forwarders to collectively distribute IPv6 multicast packets throughout their MPL Domain. An MPL Forwarder may be associated to multiple MPL Domains at the same time. Non-Forwarders will receive IPv6 multicast packets from one or more of their neighboring Forwarders. Therefore, MPL can be used to propagate a CoAP multicast request to all group members.

However, a CoAP multicast request to a group that originated outside of the MPL Domain will not be propagated by MPL - unless an MPL Forwarder is explicitly configured as an ingress point that introduces external multicast packets into the MPL Domain. Such an ingress point could be located on an edge router (e.g., 6LBR). Methods to configure which multicast groups are to be propagated into the MPL Domain could be:

- o Manual configuration on each ingress MPL Forwarder.
- o MLDv2 protocol, which works only in case all CoAP servers joining a group are in link-local communication range of an ingress MPL Forwarder. This is typically not the case on mesh networks.
- o A new/custom protocol to register multicast groups at an ingress MPL Forwarder. This could be for example a CoAP-based protocol offering multicast group subscription features similar to MLDv2.

4. Unsecured Group Communication

CoAP group communication can operate in CoAP NoSec (No Security) mode, without using application-layer and transport-layer security mechanisms. The NoSec mode uses the "coap" scheme, and is defined in [Section 9 of \[RFC7252\]](#). The conceptual "NoSec" security group as defined in [Section 2.1](#) is used for unsecured group communication. Before using this mode of operation, the security implications ([Section 6.1](#)) must be well understood.

5. Secured Group Communication using Group OSCORE

The application-layer protocol Object Security for Constrained RESTful Environments (OSCORE) [[RFC8613](#)] provides end-to-end encryption, integrity and replay protection of CoAP messages exchanged between two CoAP endpoints. These can act both as CoAP Client as well as CoAP Server, and share an OSCORE Security Context used to protect and verify exchanged messages. The use of OSCORE does not affect the URI scheme and OSCORE can therefore be used with any URI scheme defined for CoAP.

OSCORE uses COSE

[[I-D.ietf-cose-rfc8152bis-struct](#)][[I-D.ietf-cose-rfc8152bis-algs](#)] to perform encryption operations and protect a CoAP message carried in a COSE object, by using an Authenticated Encryption with Associated Data (AEAD) algorithm. In particular, OSCORE takes as input an unprotected CoAP message and transforms it into a protected CoAP message transporting the COSE object.

OSCORE makes it possible to selectively protect different parts of a CoAP message in different ways, while still allowing intermediaries (e.g., CoAP proxies) to perform their intended functionalities. That is, some message parts are encrypted and integrity protected; other parts are only integrity protected to be accessible to, but not modifiable by, proxies; and some parts are kept as plain content to be both accessible to and modifiable by proxies. Such differences especially concern the CoAP options included in the unprotected message.

Group OSCORE [[I-D.ietf-core-oscore-groupcomm](#)] builds on OSCORE, and provides end-to-end security of CoAP messages exchanged between members of an OSCORE group, while fulfilling the same security requirements.

In particular, Group OSCORE protects CoAP requests sent over IP multicast by a CoAP client, as well as multiple corresponding CoAP responses sent over IP unicast by different CoAP servers. However, the same security material can also be used to protect CoAP requests sent over IP unicast to a single CoAP server in the OSCORE group, as well as the corresponding responses.

Group OSCORE ensures source authentication of all messages exchanged within the OSCORE group, by means of two possible methods.

The first method, called group mode, relies on digital signatures. That is, sender devices sign their outgoing messages using their own private key, and embed the signature in the protected CoAP message.

The second method, called pairwise mode, relies on a symmetric key, which is derived from a pairwise shared secret computed from the asymmetric keys of the message sender and recipient. This method is intended for one-to-one messages sent in the group, such as all responses individually sent by servers, as well as requests addressed to an individual server.

A Group Manager is responsible for managing one or multiple OSCORE groups. In particular, the Group Manager acts as repository of public keys of group members; manages, renews and provides security material in the group; and handles the join process of new group members.

As defined in [[I-D.ietf-ace-oscore-gm-admin](#)], an administrator entity can interact with the Group Manager to create OSCORE groups and specify their configuration (see [Section 2.2.2](#)). During the lifetime of the OSCORE group, the administrator can further interact with the Group Manager, in order to possibly update the group configuration and eventually delete the group.

As recommended in [[I-D.ietf-core-oscore-groupcomm](#)], a CoAP endpoint can join an OSCORE group by using the method described in [[I-D.ietf-ace-key-groupcomm-oscore](#)] and based on the ACE framework for Authentication and Authorization in constrained environments [[I-D.ietf-ace-oauth-authz](#)].

A CoAP endpoint can discover OSCORE groups and retrieve information to join them through their respective Group Managers by using the method described in [[I-D.tiloca-core-oscore-discovery](#)] and based on the CoRE Resource Directory [[I-D.ietf-core-resource-directory](#)].

If security is required, CoAP group communication as described in this specification MUST use Group OSCORE. In particular, a CoAP group as defined in [Section 2.1](#) and using secure group communication is associated to an OSCORE security group, which includes:

- o All members of the CoAP group, i.e. the CoAP endpoints configured (also) as CoAP servers and listening to the group's multicast IP address on the group's UDP port.
- o All further CoAP endpoints configured only as CoAP clients, that send (multicast) CoAP requests to the CoAP group.

[5.1. Secure Group Maintenance](#)

As part of group maintenance operations (see [Section 2.2.4](#)), additional key management operations are required for an OSCORE group, depending on the security requirements of the application (see [Section 6.2](#)). Specifically:

- o Adding new members to a CoAP group or enabling new client-only endpoints to interact with that group require also that each of such members/endpoints join the corresponding OSCORE group. By doing so, they are securely provided with the necessary cryptographic material. In case backward security is needed, this also requires to first renew such material and distribute it to the current members/endpoints, before new ones are added and join the OSCORE group.
- o In case forward security is needed, removing members from a CoAP group or stopping client-only endpoints from interacting with that group requires removing such members/endpoints from the corresponding OSCORE group. To this end, new cryptographic material is generated and securely distributed only to the remaining members/endpoints. This ensures that only the members/endpoints intended to remain are able to continue participating in secure group communication, while the evicted ones are not able to.

The key management operations mentioned above are entrusted to the Group Manager responsible for the OSCORE group [[I-D.ietf-core-oscore-groupcomm](#)], and it is RECOMMENDED to perform them according to the approach described in [[I-D.ietf-ace-key-groupcomm-oscore](#)].

5.2. Caching of Responses at Proxies

When using Group OSCORE to protect communications end-to-end between a client and multiple servers in the group, it is normally not possible for an intermediary proxy to cache protected responses.

In fact, when starting from the same plain CoAP message, different clients generate different protected requests to send on the wire. This prevents different clients to generate potential cache hits, and thus makes response caching at the proxy pointless.

5.2.1. Using Deterministic Requests to Achieve Cachability

For application scenarios that require secure group communication, it is still possible to achieve cachability of responses at proxies, by using the approach defined in [[I-D.amsuess-core-cachable-oscore](#)] which is based on Deterministic Requests protected with the pairwise mode of Group OSCORE. This approach is limited to group requests that are safe (in the RESTful sense) to process and do not yield side effects at the server. As for any protected group request, it requires the clients and all the servers in the CoAP group to have already joined the correct OSCORE group.

Starting from the same plain CoAP request, this allows different clients in the OSCORE group to deterministically generate a same request protected with Group OSCORE, which is sent to the proxy for being forwarded to the CoAP group. The proxy can now effectively cache the resulting responses from the servers in the CoAP group, since the same plain CoAP request will result again in the same Deterministic Request and thus will produce a cache hit.

When caching of Group OSCORE secured responses is enabled at the proxy, the same as defined in [Section 3.4.3](#) applies, with respect to cache entries and their lifetimes.

Note that different Deterministic Requests result in different cache entries at the proxy. This includes the case where different plain group requests differ only in their set of Multi-ETag Options.

That is, even though the servers would produce the same plain CoAP responses in reply to the different Deterministic Requests, those will result in different protected responses to each respective

Deterministic Request, and hence in different cache entries at the proxy.

Thus, given a plain group request, a client needs to reuse the same set of Multi-ETag Options, in order to send that group request as a Deterministic Request that can actually produce a cache hit at the proxy. However, while this would prevent the caching at the proxy to be inefficient and unnecessarily redundant, it would also limit the flexibility of end-to-end response revalidation for a client.

5.2.2. Validation of Responses

When directly interacting with the servers in the CoAP group to refresh its cache entries, the proxy cannot rely on response revalidation anymore. In fact, responses protected with Group OSCORE cannot have 2.03 (Valid) as Outer Code. Response revalidation remains possible end-to-end between the client and the servers in the group, by means of including inner ETag Option(s) or inner Multi-ETag Option(s).

Finally, it is not possible for a client to revalidate responses to a group request from an aggregated cache entry at the proxy, by using the outer Group-ETag Option as defined in [Section 3.4.3.1](#). In fact, that would require the proxy to respond with an unprotected 2.03 (Valid) response potentially. However, success responses have to be protected with Group OSCORE, so cannot have 2.03 (Valid) as Outer Code.

6. Security Considerations

This section provides security considerations for CoAP group communication using IP multicast.

6.1. CoAP NoSec Mode

CoAP group communication, if not protected, is vulnerable to all the attacks mentioned in [Section 11 of \[RFC7252\]](#) for IP multicast.

Thus, for sensitive and mission-critical applications (e.g., health monitoring systems and alarm monitoring systems), it is NOT RECOMMENDED to deploy CoAP group communication in NoSec mode.

Without application-layer security, CoAP group communication SHOULD only be deployed in applications that are non-critical, and that do not involve or may have an impact on sensitive data and personal sphere. These include, e.g., read-only temperature sensors deployed in non-sensitive environments, where the client reads out the values

but does not use the data to control actuators or to base an important decision on.

Discovery of devices and resources is a typical use case where NoSec mode is applied, since the devices involved do not have yet configured any mutual security relations at the time the discovery takes place.

6.2. Group OSCORE

Group OSCORE provides end-to-end application-level security. This has many desirable properties, including maintaining security assurances while forwarding traffic through intermediaries (proxies). Application-level security also tends to more cleanly separate security from the dynamics of group membership (e.g., the problem of distributing security keys across large groups with many members that come and go).

For sensitive and mission-critical applications, CoAP group communication MUST be protected by using Group OSCORE as specified in [[I-D.ietf-core-oscore-groupcomm](#)]. The same security considerations from Section 10 of [[I-D.ietf-core-oscore-groupcomm](#)] hold for this specification.

6.2.1. Group Key Management

A key management scheme for secure revocation and renewal of group security material, namely group rekeying, should be adopted in OSCORE groups. In particular, the key management scheme should preserve backward and forward security in the OSCORE group, if the application requires so (see Section 3.1 of [[I-D.ietf-core-oscore-groupcomm](#)]).

Group policies should also take into account the time that the key management scheme requires to rekey the group, on one hand, and the expected frequency of group membership changes, i.e. nodes' joining and leaving, on the other hand.

In fact, it may be desirable to not rekey the group upon every single membership change, in case members' joining and leaving are frequent, and at the same time a single group rekeying instance takes a non-negligible time to complete.

In such a case, the Group Manager may consider to rekey the group, e.g., after a minimum number of nodes has joined or left the group within a pre-defined time interval, or according to communication patterns with predictable time intervals of network inactivity. This would prevent paralyzing communications in the group, when a slow rekeying scheme is used and frequently invoked.

This comes at the cost of not continuously preserving backward and forward security, since group rekeying might not occur upon every single group membership change. That is, most recently joined nodes would have access to the security material used prior to their join, and thus be able to access past group communications protected with that security material. Similarly, until the group is rekeyed, most recently left nodes would preserve access to group communications protected with the retained security material.

6.2.2. Source Authentication

Both the group mode and the pairwise mode of Group OSCORE ensure source authentication of messages exchanged by CoAP endpoints through CoAP group communication.

To this end, outgoing messages are either countersigned by the message sender endpoint with its own private key (group mode), or protected with a symmetric key, which is in turn derived using the asymmetric keys of the message sender and recipient (pairwise mode).

Thus, both modes allow a recipient CoAP endpoint to verify that a message has actually been originated by a specific and identified member of the OSCORE group.

[Appendix F](#) of [\[I-D.ietf-core-oscore-groupcomm\]](#) discusses a number of cases where a recipient CoAP endpoint may skip the verification of countersignatures in messages protected with the group mode, possibly on a per-message basis. However, this is NOT RECOMMENDED. That is, a CoAP endpoint receiving a message secured with the group mode of Group OSCORE SHOULD always verify the countersignature.

6.2.3. Countering Attacks

As discussed below, Group OSCORE addresses a number of security attacks mentioned in [Section 11 of \[RFC7252\]](#), with particular reference to their execution over IP multicast.

- o Since Group OSCORE provides end-to-end confidentiality and integrity of request/response messages, proxies in multicast settings cannot break message protection, and thus cannot act as man-in-the-middle beyond their legitimate duties (see [Section 11.2 of \[RFC7252\]](#)). In fact, intermediaries such as proxies are not assumed to have access to the OSCORE Security Context used by group members. Also, with the notable addition of countersignatures for the group mode, Group OSCORE protects messages using the same procedure as OSCORE (see [Sections 8.1 and 8.3 of \[I-D.ietf-core-oscore-groupcomm\]](#)), and especially processes

CoAP options according to the same classification in U/I/E classes.

- o Group OSCORE protects against amplification attacks (see [Section 11.3 of \[RFC7252\]](#)), which are made e.g. by injecting (small) requests over IP multicast from the (spoofed) IP address of a victim client, and thus triggering the transmission of several (much bigger) responses back to that client. In fact, upon receiving a request over IP multicast as protected with Group OSCORE in group mode, a server is able to verify whether the request is fresh and originates from the alleged sender in the OSCORE group, by verifying the countersignature included in the request using the public key of that sender (see Section 8.2 of [\[I-D.ietf-core-oscore-groupcomm\]](#)). Furthermore, as also discussed in Section 8 of [\[I-D.ietf-core-oscore-groupcomm\]](#), it is recommended that servers failing to decrypt and verify an incoming message do not send back any error message.
- o Group OSCORE limits the impact of attacks based on IP spoofing also over IP multicast (see [Section 11.4 of \[RFC7252\]](#)). In fact, requests and corresponding responses sent in the OSCORE group can be correctly generated only by legitimate group members.

Within an OSCORE group, the shared symmetric-key security material strictly provides only group-level authentication (see Section 10.1 of [\[I-D.ietf-core-oscore-groupcomm\]](#)). However, source authentication of messages is also ensured, both in the group mode by means of countersignatures (see Sections [8.1](#) and [8.3](#) of [\[I-D.ietf-core-oscore-groupcomm\]](#)), and in the pairwise mode by using additionally derived pairwise keys (see Sections [9.1](#) and [9.3](#) of [\[I-D.ietf-core-oscore-groupcomm\]](#)). Thus, recipient endpoints can verify a message to be originated by the alleged, identifiable sender in the OSCORE group.

Note that the server may additionally rely on the Echo Option for CoAP described in [\[I-D.ietf-core-echo-request-tag\]](#), in order to verify the aliveness and reachability of the client sending a request from a particular IP address.

- o Group OSCORE does not require group members to be equipped with a good source of entropy for generating security material (see [Section 11.6 of \[RFC7252\]](#)), and thus does not contribute to create an entropy-related attack vector against such (constrained) CoAP endpoints. In particular, the symmetric keys used for message encryption and decryption are derived through the same HMAC-based HKDF scheme used for OSCORE (see [Section 3.2 of \[RFC8613\]](#)). Besides, the OSCORE Master Secret used in such derivation is securely generated by the Group Manager responsible for the OSCORE

group, and securely provided to the CoAP endpoints when they join the group.

- o Group OSCORE prevents to make any single group member a target for subverting security in the whole OSCORE group (see [Section 11.6 of \[RFC7252\]](#)), even though all group members share (and can derive) the same symmetric-key security material used in the OSCORE group (see Section 10.1 of [[I-D.ietf-core-oscore-groupcomm](#)]). In fact, source authentication is always ensured for exchanged CoAP messages, as verifiable to be originated by the alleged, identifiable sender in the OSCORE group. This relies on including a countersignature computed with a node's individual private key (in the group mode), or on protecting messages with a pairwise symmetric key, which is in turn derived from the asymmetric keys of the sender and recipient CoAP endpoints (in the pairwise mode).

6.3. Replay of Non-Confirmable Messages

Since all requests sent over IP multicast are Non-confirmable, a client might not be able to know if an adversary has actually captured one of its transmitted requests and later re-injected it in the group as a replay to the server nodes. In fact, even if the servers sent back responses to the replayed request, the client would typically not have a valid matching request active anymore so this attack would not accomplish anything in the client.

If Group OSCORE is used, such a replay attack on the servers is prevented, since a client protects every different request with a different Sequence Number value, which is in turn included as Partial IV in the protected message and takes part in the construction of the AEAD cipher nonce. Thus, a server would be able to detect the replayed request, by checking the conveyed Partial IV against its own replay window in the OSCORE Recipient Context associated to the client.

This requires a server to have a synchronized, up to date view of the sequence number used by the client. If such synchronization is lost, e.g. due to a reboot, or suspected so, the server should use one of the methods described in [Appendix E](#) of [[I-D.ietf-core-oscore-groupcomm](#)], such as the one based on the Echo Option for CoAP described in [[I-D.ietf-core-echo-request-tag](#)], in order to (re-)synchronize with the client's sequence number.

6.4. Use of CoAP No-Response Option

When CoAP group communication is used in CoAP NoSec (No Security) mode (see [Section 4](#)), the CoAP No-Response Option [[RFC7967](#)] could be misused by a malicious client to evoke as much responses from servers

to a multicast request as possible, by using the value '0' - Interested in all responses. This even overrides the default behaviour of a CoAP server to suppress the response in case there is nothing of interest to respond with. Therefore, this option can be used to perform an amplification attack.

A proposed mitigation is to only allow this option to relax the standard suppression rules for a resource in case the option is sent by an authenticated client. If sent by an unauthenticated client, the option can be used to expand the classes of responses suppressed compared to the default rules but not to reduce the classes of responses suppressed.

6.5. 6LoWPAN

In a 6LoWPAN network, a multicast IPv6 packet may be fragmented prior to transmission. A 6LoWPAN Router that forwards a fragmented packet may have a relatively high impact on the occupation of the wireless channel and may locally experience high memory load due to packet buffering. For example, the MPL [[RFC7731](#)] protocol requires an MPL Forwarder to store the packet for a longer duration, to allow multiple forwarding transmissions to neighboring Forwarders. If one or more of the fragments are not received correctly by an MPL Forwarder during its packet reassembly time window, the Forwarder discards all received fragments and at a future point in time it needs to receive again all the packet fragments (this time, possibly from another neighboring MPL Forwarder).

For these reasons, a fragmented IPv6 multicast packet is a possible attack vector in a Denial of Service (DoS) amplification attack. See [Section 11.3 of \[RFC7252\]](#) for more details on amplification. To mitigate the risk, applications sending multicast IPv6 requests to 6LoWPAN hosted CoAP servers SHOULD limit the size of the request to avoid 6LoWPAN fragmentation of the request packet. A 6LoWPAN Router or multicast forwarder SHOULD deprioritize forwarding for multi-fragment 6LoWPAN multicast packets. Also, a 6LoWPAN Border Router SHOULD implement multicast packet filtering to prevent unwanted multicast traffic from entering a 6LoWPAN network from the outside. For example, it could filter out all multicast packet for which there is no known multicast listener on the 6LoWPAN network. See [Section 3.9](#) for protocols that allow multicast listeners to signal which groups they would like to listen to.

6.6. Wi-Fi

In a home automation scenario using Wi-Fi, Wi-Fi security should be enabled to prevent rogue nodes from joining. The Customer Premises Equipment (CPE) that enables access to the Internet should also have

its IP multicast filters set so that it enforces multicast scope boundaries to isolate local multicast groups from the rest of the Internet (e.g., as per [[RFC6092](#)]). In addition, the scope of IP multicast transmissions and listeners should be site-local (5) or smaller. For site-local scope, the CPE will be an appropriate multicast scope boundary point.

[6.7.](#) Monitoring

[6.7.1.](#) General Monitoring

CoAP group communication can be used to control a set of related devices: for example, simultaneously turn on all the lights in a room. This intrinsically exposes the group to some unique monitoring risks that devices not in a group are not as vulnerable to. For example, assume an attacker is able to physically see a set of lights turn on in a room. Then the attacker can correlate an observed CoAP group communication message to the observed coordinated group action - even if the CoAP message is (partly) encrypted.

This will give the attacker side-channel information to plan further attacks (e.g., by determining the members of the group some network topology information may be deduced).

[6.7.2.](#) Pervasive Monitoring

A key additional threat consideration for group communication is pervasive monitoring [[RFC7258](#)]. CoAP group communication solutions that are built on top of IP multicast need to pay particular heed to these dangers. This is because IP multicast is easier to intercept compared to IP unicast. Also, CoAP traffic is typically used for the Internet of Things. This means that CoAP multicast may be used for the control and monitoring of critical infrastructure (e.g., lights, alarms, HVAC, electrical grid, etc.) that may be prime targets for attack.

For example, an attacker may attempt to record all the CoAP traffic going over a smart grid (i.e., networked electrical utility) and try to determine critical nodes for further attacks. For example, the source node (controller) sends out CoAP group communication messages which easily identifies it as a controller. CoAP multicast traffic is inherently more vulnerable compared to unicast, as the same packet may be replicated over many more links, leading to a higher probability of packet capture by a pervasive monitoring system.

One mitigation is to restrict the scope of IP multicast to the minimal scope that fulfills the application need. See the congestion control recommendations in the last bullet of

[Section 3.5](#) to minimize the scope. Thus, for example, realm-local IP multicast scope is always preferred over site-local scope IP multicast if this fulfills the application needs.

Even if all CoAP multicast traffic is encrypted/protected, an attacker may still attempt to capture this traffic and perform an off-line attack in the future.

7. IANA Considerations

This document has the following actions for IANA.

7.1. CoAP Option Numbers Registry

IANA is asked to enter the following option numbers to the "CoAP Option Numbers" registry defined in [\[RFC7252\]](#) within the "CoRE Parameters" registry.

Number	Name	Reference
TBD1	Multi-ETag	[This document]
TBD2	Group-ETag	[This document]

8. References

8.1. Normative References

[I-D.ietf-core-echo-request-tag]

Amsuess, C., Mattsson, J., and G. Selander, "CoAP: Echo, Request-Tag, and Token Processing", [draft-ietf-core-echo-request-tag-12](#) (work in progress), February 2021.

[I-D.ietf-core-oscore-groupcomm]

Tiloca, M., Selander, G., Palombini, F., Mattsson, J., and J. Park, "Group OSCORE - Secure Group Communication for CoAP", [draft-ietf-core-oscore-groupcomm-11](#) (work in progress), February 2021.

[I-D.ietf-cose-rfc8152bis-algs]

Schaad, J., "CBOR Object Signing and Encryption (COSE): Initial Algorithms", [draft-ietf-cose-rfc8152bis-algs-12](#) (work in progress), September 2020.

[I-D.ietf-cose-rfc8152bis-struct]

Schaad, J., "CBOR Object Signing and Encryption (COSE): Structures and Process", [draft-ietf-cose-rfc8152bis-struct-15](#) (work in progress), February 2021.

[I-D.tiloca-core-observe-multicast-notifications]

Tiloca, M., Hoeglund, R., Amsuess, C., and F. Palombini, "Observe Notifications as CoAP Multicast Responses", [draft-tiloca-core-observe-multicast-notifications-05](#) (work in progress), February 2021.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", [RFC 3376](#), DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.

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

- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", [RFC 6690](#), DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", [RFC 7641](#), DOI 10.17487/RFC7641, September 2015, <<https://www.rfc-editor.org/info/rfc7641>>.
- [RFC7959] Bormann, C. and Z. Shelby, Ed., "Block-Wise Transfers in the Constrained Application Protocol (CoAP)", [RFC 7959](#), DOI 10.17487/RFC7959, August 2016, <<https://www.rfc-editor.org/info/rfc7959>>.
- [RFC8075] Castellani, A., Loreto, S., Rahman, A., Fossati, T., and E. Dijk, "Guidelines for Mapping Implementations: HTTP to the Constrained Application Protocol (CoAP)", [RFC 8075](#), DOI 10.17487/RFC8075, February 2017, <<https://www.rfc-editor.org/info/rfc8075>>.
- [RFC8132] van der Stok, P., Bormann, C., and A. Sehgal, "PATCH and FETCH Methods for the Constrained Application Protocol (CoAP)", [RFC 8132](#), DOI 10.17487/RFC8132, April 2017, <<https://www.rfc-editor.org/info/rfc8132>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", [RFC 8610](#), DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.
- [RFC8613] Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security for Constrained RESTful Environments (OSCORE)", [RFC 8613](#), DOI 10.17487/RFC8613, July 2019, <<https://www.rfc-editor.org/info/rfc8613>>.
- [RFC8742] Bormann, C., "Concise Binary Object Representation (CBOR) Sequences", [RFC 8742](#), DOI 10.17487/RFC8742, February 2020, <<https://www.rfc-editor.org/info/rfc8742>>.

8.2. Informative References

[Californium]

Eclipse Foundation, "Eclipse Californium", March 2019, <<https://github.com/eclipse/californium/tree/2.0.x/californium-core/src/main/java/org/eclipse/californium/core>>.

[Go-OCF] Open Connectivity Foundation (OCF), "Implementation of CoAP Server & Client in Go", March 2019, <<https://github.com/go-ocf/go-coap>>.

[I-D.amsuess-core-cachable-oscore]

Amsuess, C. and M. Tiloca, "Cachable OSCORE", [draft-amsuess-core-cachable-oscore-01](#) (work in progress), February 2021.

[I-D.ietf-ace-key-groupcomm-oscore]

Tiloca, M., Park, J., and F. Palombini, "Key Management for OSCORE Groups in ACE", [draft-ietf-ace-key-groupcomm-oscore-10](#) (work in progress), February 2021.

[I-D.ietf-ace-oauth-authz]

Seitz, L., Selander, G., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "Authentication and Authorization for Constrained Environments (ACE) using the OAuth 2.0 Framework (ACE-OAuth)", [draft-ietf-ace-oauth-authz-37](#) (work in progress), February 2021.

[I-D.ietf-ace-oscore-gm-admin]

Tiloca, M., Hoeglund, R., Stok, P., Palombini, F., and K. Hartke, "Admin Interface for the OSCORE Group Manager", [draft-ietf-ace-oscore-gm-admin-02](#) (work in progress), February 2021.

[I-D.ietf-core-coap-pubsub]

Koster, M., Keranen, A., and J. Jimenez, "Publish-Subscribe Broker for the Constrained Application Protocol (CoAP)", [draft-ietf-core-coap-pubsub-09](#) (work in progress), September 2019.

[I-D.ietf-core-resource-directory]

Amsuess, C., Shelby, Z., Koster, M., Bormann, C., and P. Stok, "CoRE Resource Directory", [draft-ietf-core-resource-directory-26](#) (work in progress), November 2020.

[I-D.tiloca-core-groupcomm-proxy]

Tiloca, M. and E. Dijk, "Proxy Operations for CoAP Group Communication", [draft-tiloca-core-groupcomm-proxy-03](#) (work in progress), February 2021.

[I-D.tiloca-core-oscore-discovery]

Tiloca, M., Amsuess, C., and P. Stok, "Discovery of OSCORE Groups with the CoRE Resource Directory", [draft-tiloca-core-oscore-discovery-08](#) (work in progress), February 2020.

[RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", [RFC 6092](#), DOI 10.17487/RFC6092, January 2011, <<https://www.rfc-editor.org/info/rfc6092>>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., 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](#), DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.

[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](#), DOI 10.17487/RFC6636, May 2012, <<https://www.rfc-editor.org/info/rfc6636>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

[RFC7346] Droms, R., "IPv6 Multicast Address Scopes", [RFC 7346](#), DOI 10.17487/RFC7346, August 2014, <<https://www.rfc-editor.org/info/rfc7346>>.

[RFC7390] Rahman, A., Ed. and E. Dijk, Ed., "Group Communication for the Constrained Application Protocol (CoAP)", [RFC 7390](#), DOI 10.17487/RFC7390, October 2014, <<https://www.rfc-editor.org/info/rfc7390>>.

[RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", [RFC 7731](#), DOI 10.17487/RFC7731, February 2016, <<https://www.rfc-editor.org/info/rfc7731>>.

- [RFC7967] Bhattacharyya, A., Bandyopadhyay, S., Pal, A., and T. Bose, "Constrained Application Protocol (CoAP) Option for No Server Response", [RFC 7967](#), DOI 10.17487/RFC7967, August 2016, <<https://www.rfc-editor.org/info/rfc7967>>.
- [RFC8323] Bormann, C., Lemay, S., Tschafenig, H., Hartke, K., Silverajan, B., and B. Raymor, Ed., "CoAP (Constrained Application Protocol) over TCP, TLS, and WebSockets", [RFC 8323](#), DOI 10.17487/RFC8323, February 2018, <<https://www.rfc-editor.org/info/rfc8323>>.
- [RFC8710] Fossati, T., Hartke, K., and C. Bormann, "Multipart Content-Format for the Constrained Application Protocol (CoAP)", [RFC 8710](#), DOI 10.17487/RFC8710, February 2020, <<https://www.rfc-editor.org/info/rfc8710>>.

Appendix A. Use Cases

To illustrate where and how CoAP-based group communication can be used, this section summarizes the most common use cases. These use cases include both secured and non-secured CoAP usage. Each subsection below covers one particular category of use cases for CoRE. Within each category, a use case may cover multiple application areas such as home IoT, commercial building IoT (sensing and control), industrial IoT/control, or environmental sensing.

A.1. Discovery

Discovery of physical devices in a network, or discovery of information entities hosted on network devices, are operations that are usually required in a system during the phases of setup or (re)configuration. When a discovery use case involves devices that need to interact without having been configured previously with a common security context, unsecured CoAP communication is typically used. Discovery may involve a request to a directory server, which provides services to aid clients in the discovery process. One particular type of directory server is the CoRE Resource Directory [[I-D.ietf-core-resource-directory](#)]; and there may be other types of directories that can be used with CoAP.

A.1.1. Distributed Device Discovery

Device discovery is the discovery and identification of networked devices - optionally only devices of a particular class, type, model, or brand. Group communication is used for distributed device discovery, if a central directory server is not used. Typically in distributed device discovery, a multicast request is sent to a particular address (or address range) and multicast scope of

interest, and any devices configured to be discoverable will respond back. For the alternative solution of centralized device discovery a central directory server is accessed through unicast, in which case group communication is not needed. This requires that the address of the central directory is either preconfigured in each device or configured during operation using a protocol.

In CoAP, device discovery can be implemented by CoAP resource discovery requesting (GET) a particular resource that the sought device class, type, model or brand is known to respond to. It can also be implemented using CoAP resource discovery ([Section 7 of \[RFC7252\]](#)) and the CoAP query interface defined in [Section 4 of \[RFC6690\]](#) to find these particular resources. Also, a multicast GET request to /.well-known/core can be used to discover all CoAP devices.

[A.1.2.](#) Distributed Service Discovery

Service discovery is the discovery and identification of particular services hosted on network devices. Services can be identified by one or more parameters such as ID, name, protocol, version and/or type. Distributed service discovery involves group communication to reach individual devices hosting a particular service; with a central directory server not being used.

In CoAP, services are represented as resources and service discovery is implemented using resource discovery ([Section 7 of \[RFC7252\]](#)) and the CoAP query interface defined in [Section 4 of \[RFC6690\]](#).

[A.1.3.](#) Directory Discovery

This use case is a specific sub-case of Distributed Service Discovery (Appendix A.1.2), in which a device needs to identify the location of a Directory on the network to which it can e.g. register its own offered services, or to which it can perform queries to identify and locate other devices/services it needs to access on the network. [Section 3.3 of \[RFC7390\]](#) shows an example of discovering a CoRE Resource Directory using CoAP group communication. As defined in [\[I-D.ietf-core-resource-directory\]](#), a resource directory is a web entity that stores information about web resources and implements REST interfaces for registration and lookup of those resources. For example, a device can register itself to a resource directory to let it be found by other devices and/or applications.

A.2. Operational Phase

Operational phase use cases describe those operations that occur most frequently in a networked system, during its operational lifetime and regular operation. Regular usage is when the applications on networked devices perform the tasks they were designed for and exchange of application-related data using group communication occurs. Processes like system reconfiguration, group changes, system/device setup, extra group security changes, etc. are not part of regular operation.

A.2.1. Actuator Group Control

Group communication can be beneficial to control actuators that need to act in synchrony, as a group, with strict timing (latency) requirements. Examples are office lighting, stage lighting, street lighting, or audio alert/Public Address systems. Sections [3.4](#) and [3.5](#) of [\[RFC7390\]](#) show examples of lighting control of a group of 6LoWPAN-connected lights.

A.2.2. Device Group Status Request

To properly monitor the status of systems, there may be a need for ad-hoc, unplanned status updates. Group communication can be used to quickly send out a request to a (potentially large) number of devices for specific information. Each device then responds back with the requested data. Those devices that did not respond to the request can optionally be polled again via reliable unicast communication to complete the dataset. The device group may be defined e.g. as "all temperature sensors on floor 3", or "all lights in wing B". For example, it could be a status request for device temperature, most recent sensor event detected, firmware version, network load, and/or battery level.

A.2.3. Network-wide Query

In some cases a whole network or subnet of multiple IP devices needs to be queried for status or other information. This is similar to the previous use case except that the device group is not defined in terms of its function/type but in terms of its network location. Technically this is also similar to distributed service discovery (Appendix A.1.2) where a query is processed by all devices on a network - except that the query is not about services offered by the device, but rather specific operational data is requested.

A.2.4. Network-wide / Group Notification

In some cases a whole network, or subnet of multiple IP devices, or a specific target group needs to be notified of a status change or other information. This is similar to the previous two use cases except that the recipients are not expected to respond with some information. Unreliable notification can be acceptable in some use cases, in which a recipient does not respond with a confirmation of having received the notification. In such a case, the receiving CoAP server does not have to create a CoAP response. If the sender needs confirmation of reception, the CoAP servers can be configured for that resource to respond with a 2.xx success status after processing a notification request successfully.

A.3. Software Update

Multicast can be useful to efficiently distribute new software (firmware, image, application, etc.) to a group of multiple devices. In this case, the group is defined in terms of device type: all devices in the target group are known to be capable of installing and running the new software. The software is distributed as a series of smaller blocks that are collected by all devices and stored in memory. All devices in the target group are usually responsible for integrity verification of the received software; which can be done per-block or for the entire software image once all blocks have been received. Due to the inherent unreliability of CoAP multicast, there needs to be a backup mechanism (e.g. implemented using CoAP unicast) by which a device can individually request missing blocks of a whole software image/entity. Prior to multicast software update, the group of recipients can be separately notified that there is new software available and coming, using the above network-wide or group notification.

Appendix B. Document Updates

RFC EDITOR: PLEASE REMOVE THIS SECTION.

B.1. Version -02 to -03

- o Multiple responses from same server handled at the application.
- o Clarifications about issues with forward-proxies.
- o Operations for reverse-proxies.
- o Caching of responses at proxies.
- o Client-Server response revalidation, with Multi-ETag Option.

- o Client-Proxy response revalidation, with the Group-ETag Option.

B.2. Version -01 to -02

- o Clarified relation between security groups and application groups.
- o Considered also FETCH for requests over IP multicast.
- o More details on Observe re-registration.
- o More details on Proxy intermediaries.
- o More details on servers changing port number in the response.
- o Usage of the Uri-Host Option to indicate an application group.
- o Response suppression based on classes of error codes.

B.3. Version -00 to -01

- o Clarifications on group memberships for the different group types.
- o Simplified description of Token reuse, compared to the unicast case.
- o More details on the rationale for response suppression.
- o Clarifications of creation and management of security groups.
- o Clients more knowledgeable than proxies about stopping receiving responses.
- o Cancellation of group observations.
- o Clarification on multicast scope to use.
- o Both the group mode and pairwise mode of Group OSCORE are considered.
- o Updated security considerations.
- o Editorial improvements.

Acknowledgments

The authors sincerely thank Christian Amsuess, Carsten Bormann, Thomas Fossati, Rikard Hoeglund and Jim Schaad for their comments and feedback.

The work on this document has been partly supported by VINNOVA and the Celtic-Next project CRITISEC; and by the H2020 project SIFIS-Home (Grant agreement 952652).

Authors' Addresses

Esko Dijk
IoTconsultancy.nl
_____\n
Utrecht
Netherlands

Email: esko.dijk@iotconsultancy.nl

Chonggang Wang
InterDigital
1001 E Hector St, Suite 300
Conshohocken PA 19428
United States

Email: Chonggang.Wang@InterDigital.com

Marco Tiloca
RISE AB
Isafjordsgatan 22
Kista SE-16440 Stockholm
Sweden

Email: marco.tiloca@ri.se

