Constrained RESTful Environments WG (core)

Chairs:

Jaime Jiménez <jaime@iki.fi>

Carsten Bormann <cabo@tzi.org>

Mailing List:

core@ietf.org

Jabber:

core@jabber.ietf.org

- **We assume people have read the drafts**
- **Meetings serve to advance difficult issues by making good use of face-to-face communications**
- **We work as individuals and <u>try to be nice to each other</u>**
- **Note Well: Be aware of the IPR principles, according to RFC 8179 and its updates**

★Blue sheets
★Scribe(s)

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- BCP 9 (Internet Standards Process)
- BCP 25 (Working Group processes)
- BCP 25 (Anti-Harassment Procedures)
- BCP 54 (Code of Conduct)
- BCP 78 (Copyright)
- BCP 79 (Patents, Participation)
- https://www.ietf.org/privacy-policy/ (Privacy Policy)

Agenda Bashing

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- **10:00–10:10 Intro, Agenda, Status**
- **10:10–10:25 CoRECONF (IP)**
- **10:25–10:35 OSCORE groupcomm (MT)**
- **10:35–10:45 OSCORE discovery (MT)**
- **10:45–11:00 Observe multicast notifications (MT)**
- **11:00–11:15 Groupcomm bis (ED)**
- **11:15–11:35 SenML in IESG (etch, units) (AK)**
- **11:35–11:45 SenML data ct (AK)**
- **11:45–12:00 SenML base prefix (AK)**

All times are in time-warped SGT (UTC+08:00)

Friday (90 min)

- **12:20–12:25 Intro, Agenda**
- **12:25–13:10 CoRE applications (KH)**
- **13:10–13:50 Flextime**

Hallway discussions and side meetings

- CoRE Applications:
  Tuesday 17:00..18:40, Butterworth

Other document status

RFC-Editor's Queue: draft-ietf-core-multipart-ct-04

In IESG processing:

* draft-ietf-core-hop-limit-07: Approved-announcement to be sent::Point Raised - writeup needed
* draft-ietf-core-resource-directory-23: AD Evaluation::Revised I-D Needed
* (discuss later under SenML cluster)
  * draft-ietf-core-senml-etch-05: IESG Evaluation::Revised I-D Needed
  * draft-ietf-core-senml-more-units-03: Waiting for Writeup

* In Post-WGLC processing:
  * draft-ietf-core-stateless-03 (author to address recent input)
  * draft-ietf-core-echo-request-tag-08 (shepherd writeup needed)

Other document status (2)

Expired, otherwise ready for WGLC:

- draft-ietf-core-dev-urn-03

In WG adoption call:

- draft-bormann-core-corr-clar (in limbo)

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

CORECONF

Andy Bierman
Michel Veillette
Peter van der Stok
Alexander Pelov
Ivaylo Petrov

draft-ietf-core-sid status update

- No new changes
- The SID generation tool have been merged with pyang as requested during the last IETF

draft-ietf-core-sid next step

- Should be ready for Working Group Last Call

draft-ietf-core-yang-cbor status update

- Changes between v10 and v11
  - Updated the CBOR tag from 42 to 47
  - YANG data templates encoding now MUST follow rules in sec 4.2 instead of SHOULD (in sec 5)
  - Clarified delta usages
  - Referencing now RFC8610 for CBOR diagnostic notation.
  - A number of editorial changes to improve readability
    - Updated the Terminology section with the latest template
    - In the examples ‘+’ will not be needed due to explicit reference SID with value 0
  - Improved examples
    - Added previously missing examples for leaf, anyxml and other encodings
    - Example corrections and no longer pointing to obsoleted RFCs

draft-ietf-core-yang-cbor next step

- Should be ready for Working Group Last Call

draft-veillette-core-yang-library status update

- Changed name after WG adoption (draft-veillette-core-yang-library-05 => draft-ietf-core-yang-library-00)
- No new changes

draft-ietf-core-yang-library next step

- Should be ready for Working Group Last Call

draft-ietf-core-comi status update

- v07 and v08
  - References RFC8613 as possible security
  - Changed a SHOULD to a MUST when using user-ordered lists (sec 4.3.1)
  - Clarified sec 4.3.2 that the key should be present when list element is created
  - Updated reference to core-yang-library
  - Added IETF COPYRIGHT to yang description
  - A number of other editorial changes to improve readability
    - Fixed forgotten updates of CoMI to CORECONF
    - Format use abstract path references, not the recommended values.
    - Removed '+' sign from examples

draft-ietf-core-comi next step

- Should be ready for Working Group Last Call

Questions and answers

Thank you!

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

Group OSCORE - Secure Group Communication for CoAP

draft-ietf-core-oscore-groupcomm-06

Marco Tiloca, RISE
Göran Selander, Ericsson
Francesca Palombini, Ericsson
Jiye Park, Universität Duisburg-Essen

IETF 106, CoRE WG, Singapore, November 20th, 2019

Selected updates (1/2)

- Processed review from Ludwig (thanks!)
- Group ID MUST be unique under the same Group Manager
- Responsibilities of the Group Manager
  - Validate consistency of public keys (format, parameters, …)
- Handling of replied/repeated responses on clients
  - Now moved to *draft-dijk-core-groupcomm-bis*
  - At most 1 fresh response from each server, except for Notifications

Selected updates (2/2)

- Improved encoding of *external_aad*

- Application-specific and decoupled from message reception
  - Derivation of a Recipient Context
  - Retrieval of a public key

- Clarifications on group rekeying
  - Downsides for short-term retaining an old Security Context

Ongoing

- What countersignature algorithm(s)? Need feedback
  - Signature size vs. computing speed
  - ECDSA, EDDSA w/ Ed25519 (now MTI)

- More detailed considerations on rekeying
  - Request protected with old Ctx && Response protected with new Ctx
  - In this case, MUST include the Group ID in the response
  - For notifications, MUST include in the first after rekeying, MAY in the next ones

- Remove IANA registries on signature params and key params
  - Point at the lately extended registries in COSE-bis

Next steps

- Adopt latest comments from Jim
  - From a recent review and follow-up discussions (thanks!)
  - More details on key rollover, also for Observe notifications

- Remove IANA registries on alg/key parameters
  - Refer to the new registries in COSE-bis instead

26

- Move to WGLC ?

Thank you!

Comments/questions?

https://github.com/core-wg/oscore-groupcomm

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

Discovery of OSCORE Groups with the CoRE Resource Directory

draft-tiloca-core-oscore-discovery-04

Marco Tiloca, RISE
Christian Amsüss
Peter van der Stok

IETF 106, CoRE WG, Singapore, November 20th, 2019

Recap

- A newly deployed device:
  - May not know the OSCORE groups and their Group Manager (GM)
  - May have to wait GMs to be deployed or OSCORE groups to be created

- Use the CoRE Resource Directory (RD):
  - Discover an OSCORE group and retrieve information to join it
  - CoAP Observe supports early discovery and changes in group information

30

- Use resource lookup, to retrieve especially:
  - The identifier of the OSCORE group
  - A pointer to the resource at the GM for joining the group

Workflow overview

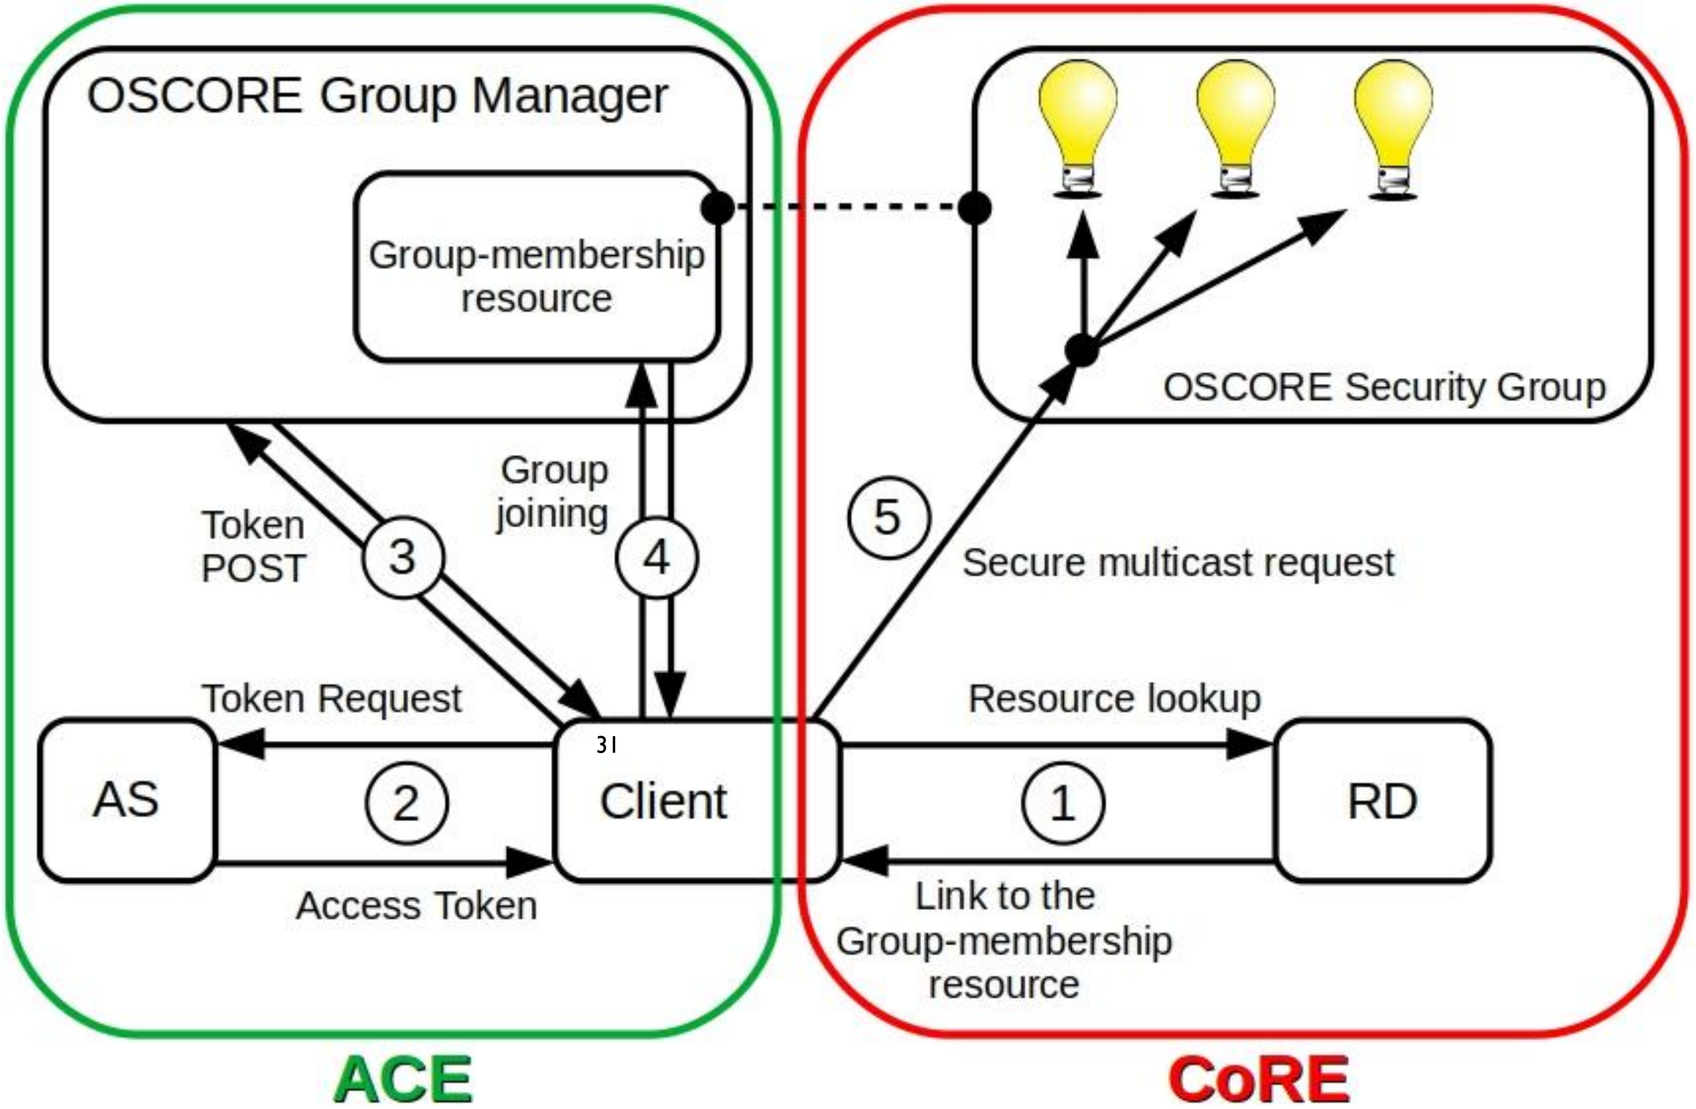

(1)

This document

(2)(3)(4)

ace-key-groupcomm-oscore

(5)

core-oscore-groupcomm

Updates from -03 (1/2)

- Approach reviewed at a design workshop (Stockholm) and CoRE interim

- Clarified rationale
  1. Use description and links of resources to discover OSCORE groups
  2. The problem becomes finding those links and descriptions
  3. Use the Resource Directory to discover links, hence OSCORE groups

- "Group-membership resource" at the Group Manager
  - Used to be "Join[32] resource", now it offers more services

Updates from -03 (2/2)

- ‘sec-gp’ parameter
  - Invariant plain name of the OSCORE group
  - Not related anymore to a (zeroed-epoch) OSCORE Group ID

- Target attributes for COSE parameters, e.g. ‘cs_alg’
  - Optional early hints on how the OSCORE group works
  - Values now taken from the ‘Value’ column in IANA registries
  - Those values MUST be unique, unlike in the ‘Name’ column (SHOULD)

33

- Updated examples
  - Including step-by-step lightweight installation scenario (BACnet)

Open points

- Registration of link target attributes
  - Mandatory: *sec-gp* and *app-gp*
  - Optional: *cs_alg* , *cs_crv* , *cs_kty* , *cs_enc* , *alg* , *hkdf*
  - A new registry will come with a *core-attributes* document

- One more optional target attribute?
  - URI of the Authorization Server associated to the GM
  - The client can avoid an unauthorized access at the GM

Summary and next steps

- Main updates
  - Clarified rationale and encoding of target attributes
  - Simpler invariant group name for OSCORE groups

- Outcome from IETF 104 [1]
  - “Time to start reading it in order to decide for WGA”
  - People volunteered to review (Jim, Carsten, Bill, Klaus)

35

- Way forward
  - Process reviews as they come

[1] https://etherpad.ietf.org/p/notes-ietf-104-core?useMonospaceFont=true

Thank you!

Comments/questions?

https://gitlab.com/crimson84/draft-tiloca-core-oscore-discovery

Backup

Application & Security Groups

- Application group
  - Defined in {RD} and reused as is
  - Set of CoAP endpoints sharing a pool of resources
  - Registered and looked up just as per Appendix A of {RD}

- OSCORE Security Group
  - Set of CoAP endpoints sharing a common Group OSCORE Security Context
  - A GM registers the group-membership resources for accessing its groups

38

Application vs. Security Groups

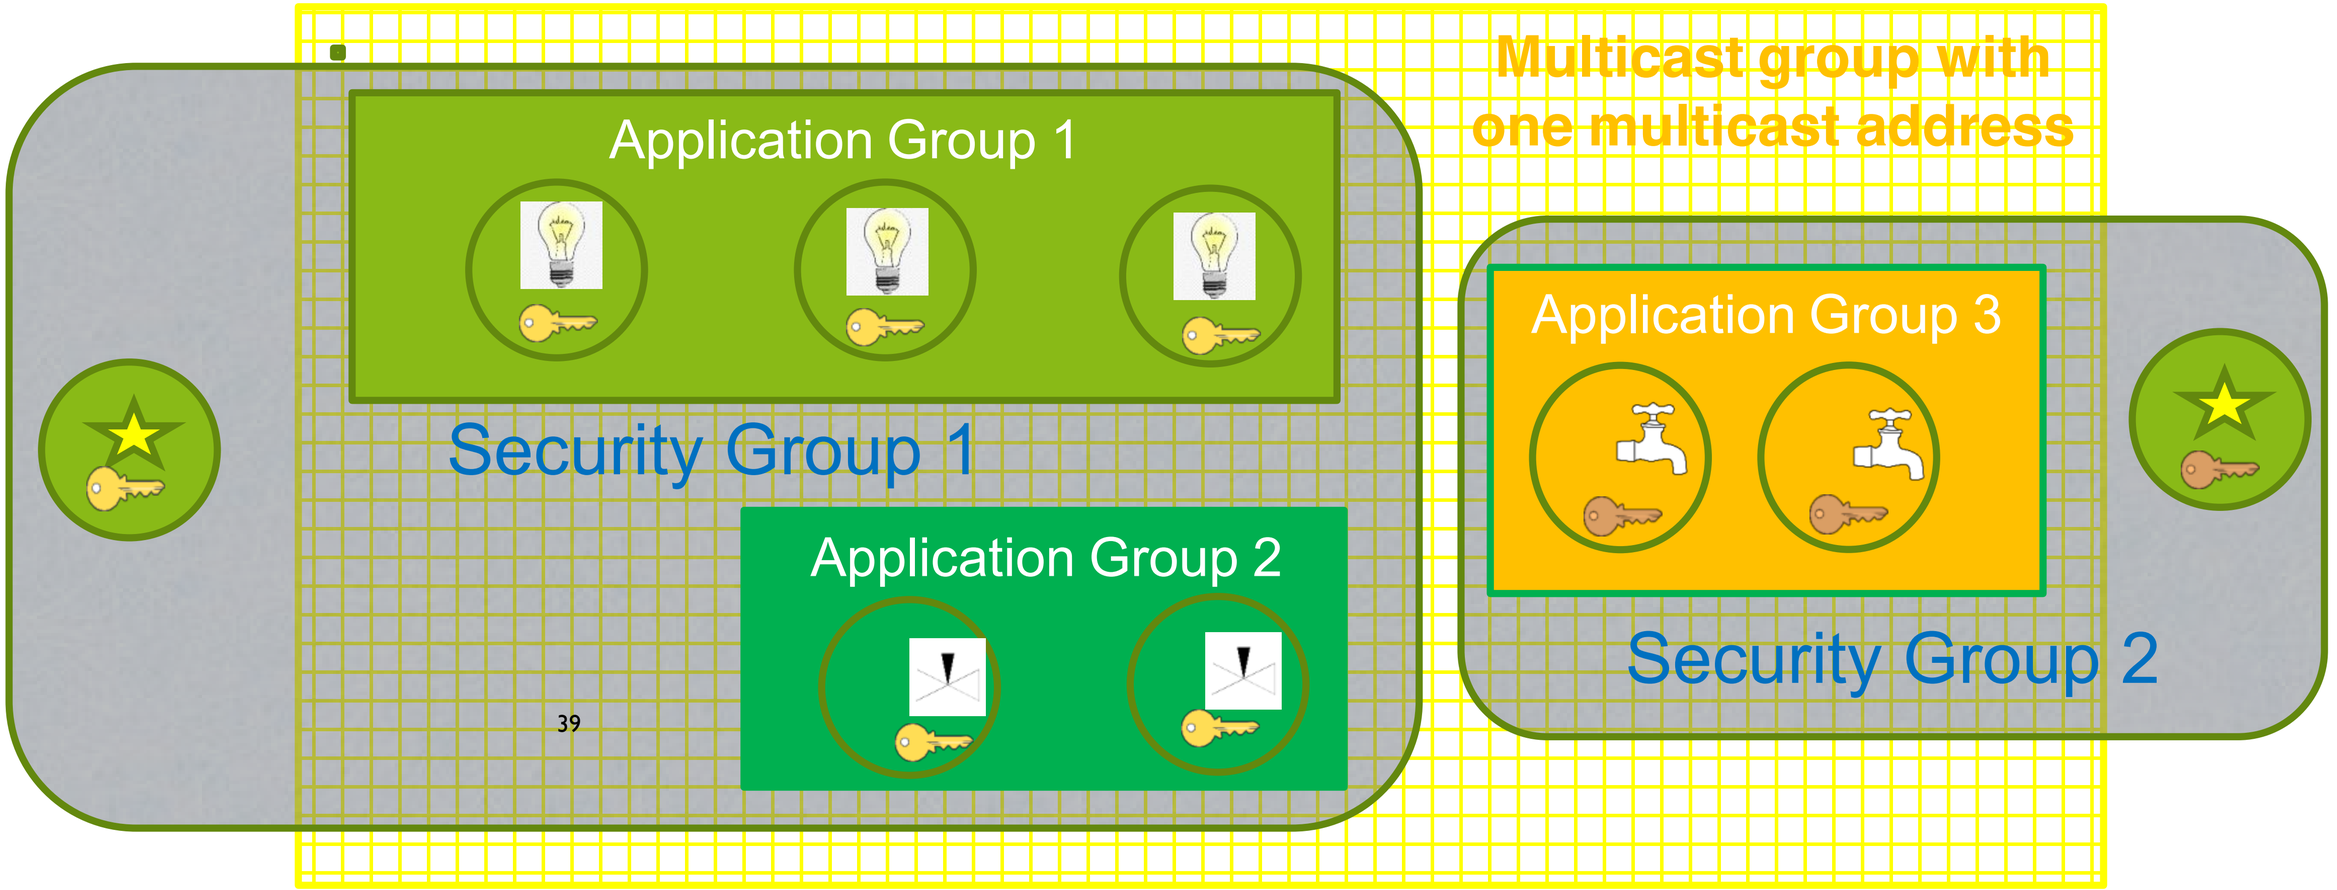

Client of application group

Different key sets

Resources for given function

Alg/key related parameters

- New optional parameters for a registered join resource
  - (*)(**) *cs_alg* : countersignature algorithm, e.g. “EdDSA”
  - (*) *cs_crv* : countersignature curve (if applicable), e.g. “Ed25519”
  - (*) *cs_kty* : countersignature key type, e.g. “OKP”
  - (*) *cs_kenc* : encoding of public keys, e.g. “COSE_Key”
  - (**) *alg* : AEAD algorithm
  - (**) *hkdf* : HKDF algorithm

- Benefits for a joining node, when discovering the OSCORE group
  - (*) No need to ask the GM or to have a trial-and-error when joining the group
  - (**) Decide whether to join the group or not, based on supported the algorithms

Registration

- The GM registers itself with the RD
  - MUST include all its join resources, with their link attributes
  - New ‘rt’ value “osc.j” in the CoRE Parameters registry

```
Request: GM -> RD

Req: POST coap://rd.example.com/rd?ep=gm1
Content-Format: 40
Payload:
</group-oscore/feedca570000>;ct=41;rt="core.osc.mbr";
                             sec-gp="feedca570000";app-gp="group1";
                             cs_alg="-8";cs_crv="6";cs_kty="1";
                             cs_kenc="1"

Response: RD -> GM

Res: 2.01 Created
Location-Path: /rd/4521
```

41

Discovery (1/2)

- The device performs a resource lookup at the RD
  - Known information: name of the Application Group, i.e. “group1”
  - Need to know: OSCORE Group Identifier; Join resource @ GM; Multicast IP address
  - ‘*app-gp*’ → Name of the Application Group, acting as tie parameter in the RD

```
Req: GET coap://rd.example.com/rd-lookup/res
  ?rt=core.osc.mbr&app-gp=group1

Response: RD -> Joining node

Res: 2.05 Content
Payload:
<coap://[2001:db8::ab]/group-oscore/feedca570000>;rt="core.osc.mbr";
     sec-gp="feedca570000";app-gp="group1";
     cs_alg="-8";cs_crv="6";cs_kty="1";
     cs_kenc="1";anchor="coap://[2001:db8::ab]"
```

Discovery (2/2)

- The device performs an endpoint lookup at the RD
  - Still need to know the Multicast IP address
  - 'ep' // Name of the Application Group, value from '*app-gp*'
  - 'base' // Multicast IP address used in the Application Group

```
Request: Joining node -> RD

Req: GET coap://rd.example.com/rd-lookup/ep
  ?et=core.rd-group&ep=group1

Response: RD ->[43] Joining node

Res: 2.05 Content
Payload:
</rd/501>;ep="group1";et="core.rd-group";
          base="coap://[ff35:30:2001:db8::23]"
```

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

Observe Notifications as CoAP Multicast Responses

draft-tiloca-core-observe-multicast-notifications-01

Marco Tiloca, RISE
Rikard Höglund, RISE
Christian Amsüss
Francesca Palombini, Ericsson

IETF 106, CoRE WG, Singapore, November 20th, 2019

Recap

- Observe notifications as multicast responses
  - Many clients observe the same resource on a server S
  - Improved performance due to multicast delivery
  - Multicast responses are not defined. Token binding? Security?

- Practical use case
  - Pub-Sub scenario
  - Many clients subscribe to a same topic on the Broker
  - Better performance
  - Subscribers are clients only

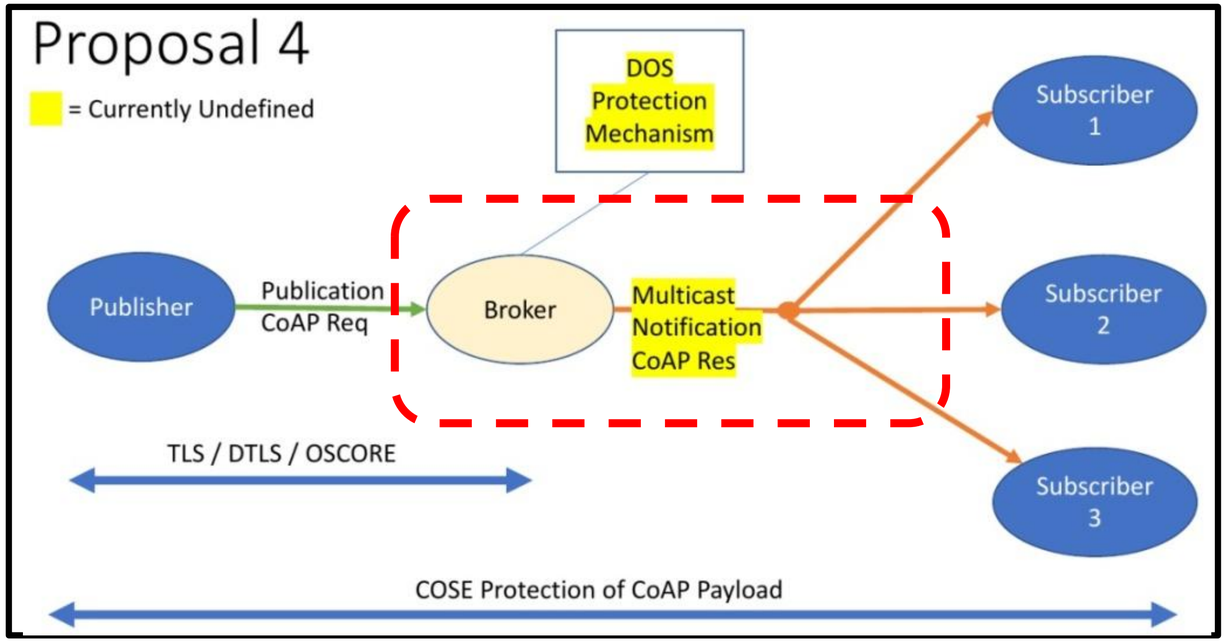

From the Hallway Discussion @ IETF 104

Contribution

- Define Observe notifications as multicast responses

- Management and enforcement of a common Token space
  - The Token space belongs to the group
  - The group entrusts the management to the server
  - All clients in a group observation use the same Token value

47

- Use of Group OSCORE to protect multicast notifications
  - The server aligns all clients of an observation on a same *external_aad*
  - All notifications for a resource are protected with that *external_aad*

Assumptions

- Clients have previously discovered the resource to access
- The server knows the IP multicast address where to send notifications
- If Group OSCORE is used to secure multicast notifications
  - The server has previously joined the right OSCORE group
- The server provides the clients with other required information

New design

- Compared to v -00
  - Revised simpler approach (no new CoAP options; no reserved Token range)
  - Re-shaped through a design workshop (Stockholm) and a CoRE interim

- The server can start a group observation for a resource, e.g.
  1. With no observers yet, a traditional registration request comes from a first client
  2. With many traditional observations, all clients are shifted to a group observation

- Phantom observation request
  - Generated inside the server, it does not hit the wire
  - Like if sent by the group, from the multicast IP address of the group
  - Multicast notifications are responses to this phantom request

Server side

1. Build a GET phantom request; Observe option set to 0

2. Choose a value T, from the Token space for messages …
   - … coming from the multicast IP address and addressed to target resource

3. Process the phantom request
   - As coming from the group and its IP multicast address
   - As addressed to the target resource

50

4. Hereafter, use T as token value for the group observation

5. Store the phantom request, with no reply right away

Interaction with clients

- The server sends to new/shifted clients an ***error response*** with
  - '*address*': IP multicast address where notifications are sent to
  - '*registr*': byte serialization of the phantom request
  - ' *res*': current representation of the target resource

- When the value of the target resource changes
  - The server sends an Observe notification to the IP multicast address
  - The notification has the Token value T of the phantom request

51

- When getting the error response, a client
  - Configures an observation from an endpoint associated to the multicast IP address
  - Accepts observe notifications with Token value T, sent to that multicast IP address

C1 registration

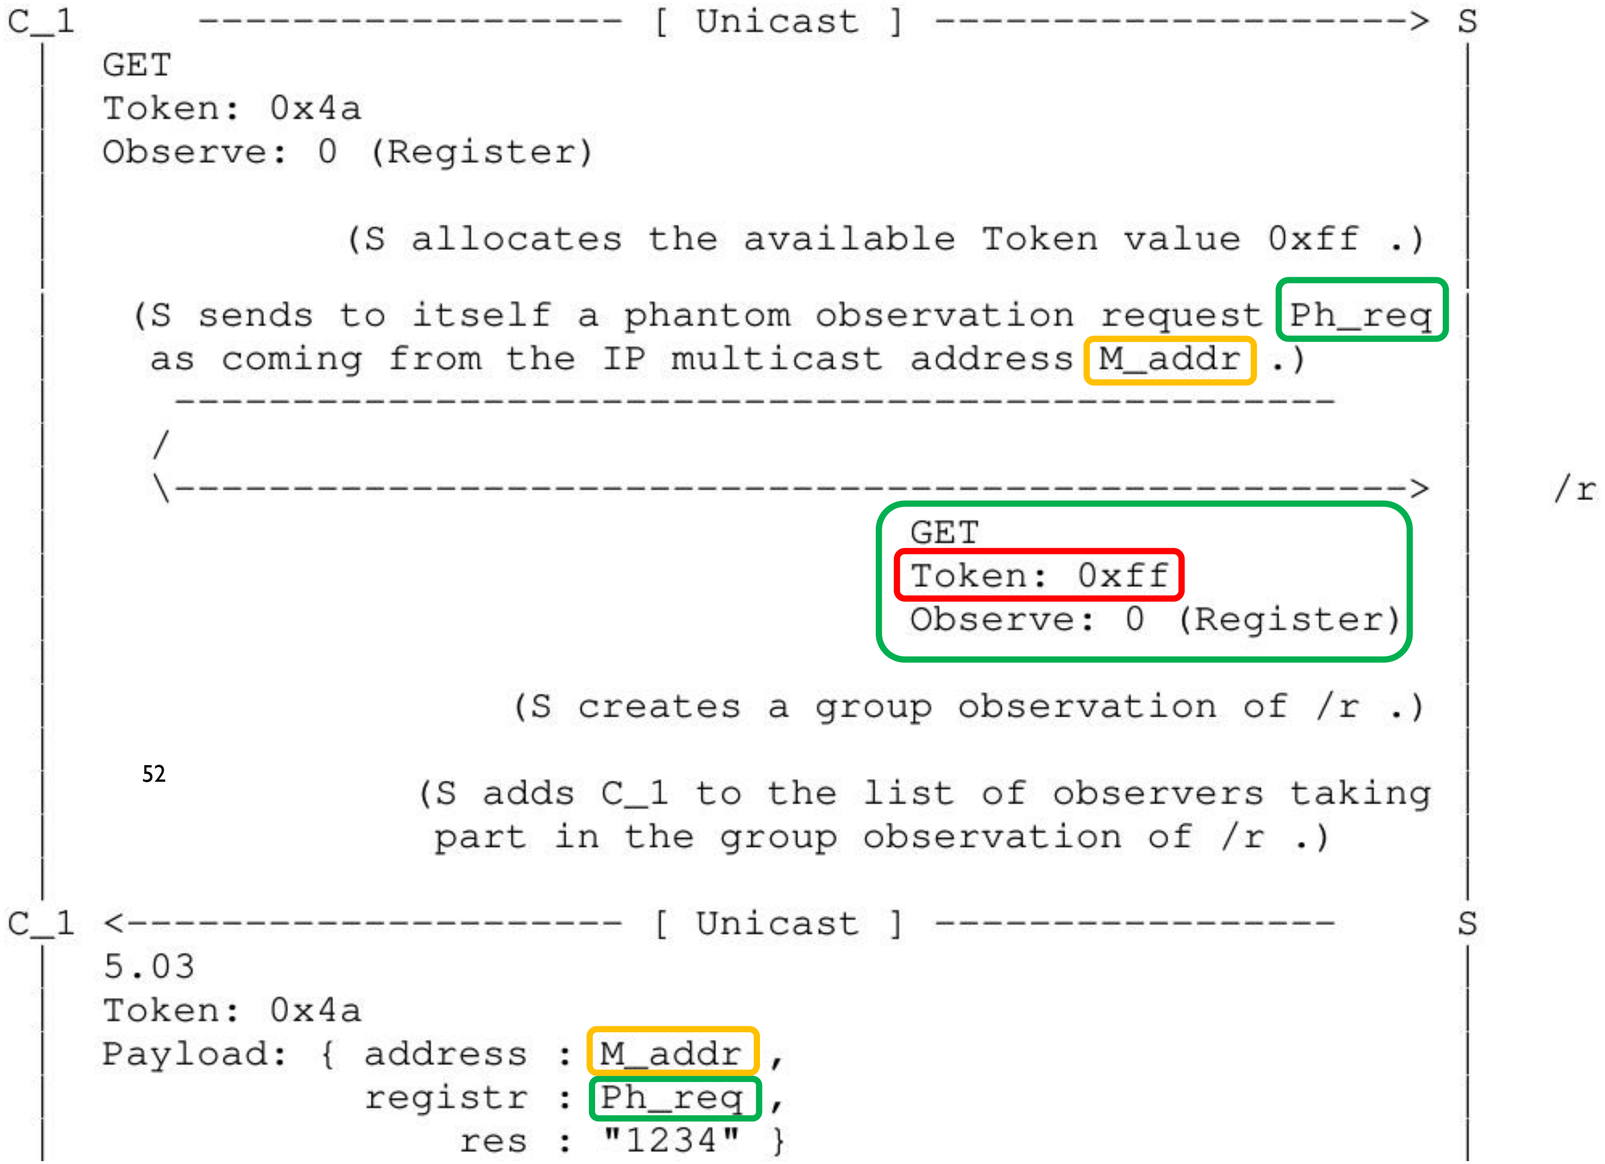

52

C2 registration

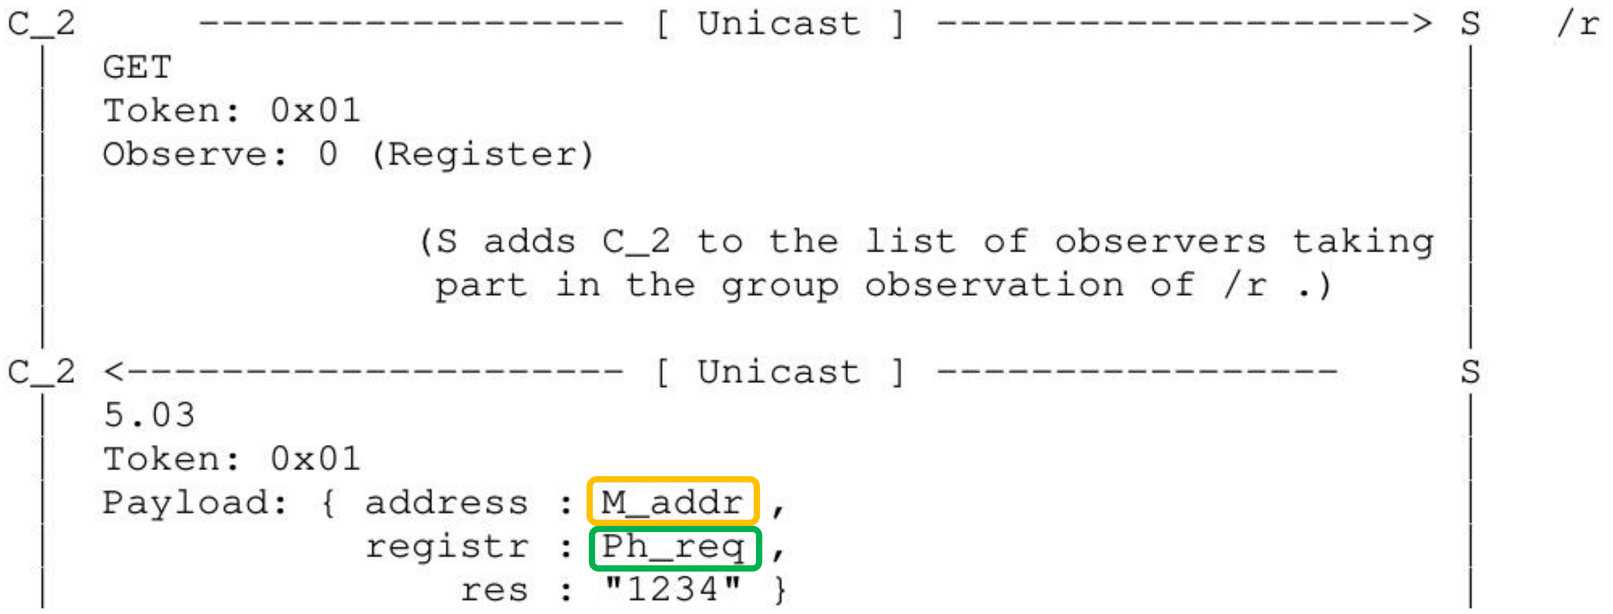

53

Multicast notification

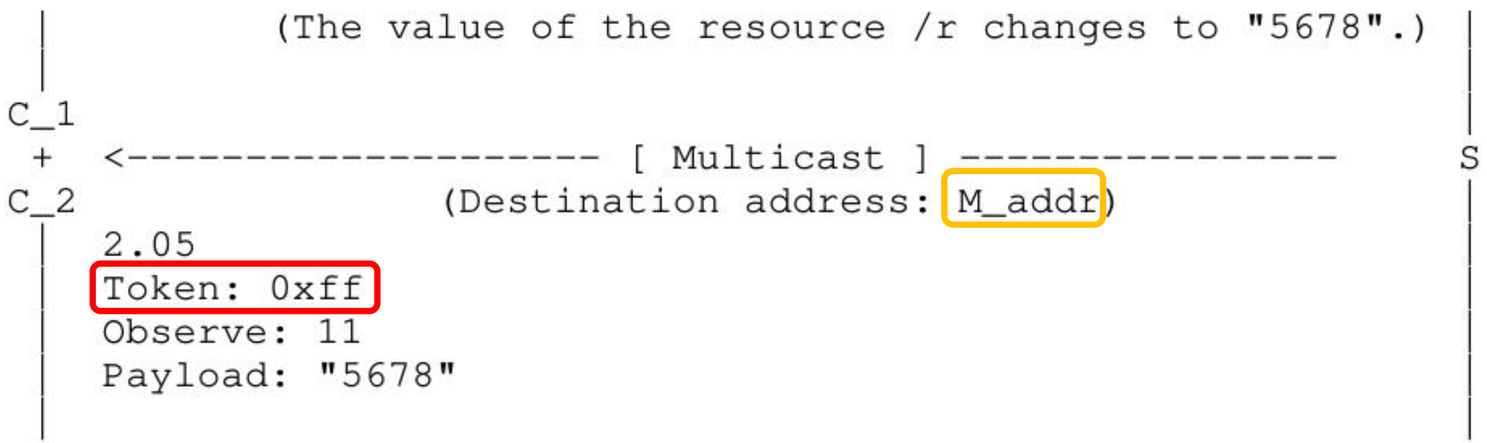

- Same Token value of the Phantom Request

- Enforce binding between
  - Every multicast notification for the target resource
  - The (group) observation that each client takes part in

Security with Group OSCORE

- The phantom request is protected with Group OSCORE
  - ***x*** : the Sender ID ('kid') of the Server in the OSCORE group
  - ***y*** : the current SN value ('piv') used by the Server in the OSCORE group
  - Note: the Server consumes the value ***y*** and does not reuse it as SN in the group

- To secure/verify all multicast notifications, the OSCORE *external_aad* is built with:
  - 'req_kid' = ***x***
  - 'req_piv' = ***y***

55

- The phantom request is still included in the informative response
  - Each client retrieves ***x*** and ***y*** from the OSCORE option

Security with Group OSCORE

- In the error response, the server can ***optionally*** specify also:
  - '*join-uri*' : link to the Group Manager to join the OSCORE group
  - '*sec-gp*' : name of the OSCORE group
  - '*as-uri*' : link to the ACE Authorization Server associated to the Group Manager
  - '*cs-alg*' : countersignature algorithm
  - '*cs-crv*' : countersignature curve
  - '*cs-kty*' : countersignature key type
  - '*cs-kenc*' : countersignature key encoding
  - '*alg*' : AEAD algorithm
  - '*hkdf*' : HKDF algorithm

MUST

MAY

56

- Clients can still discover the OSCORE group through other means
  - E.g., using the CoRE Resource Directory, as in *draft-tiloca-core-oscore-discovery*

C1 registration w/ security

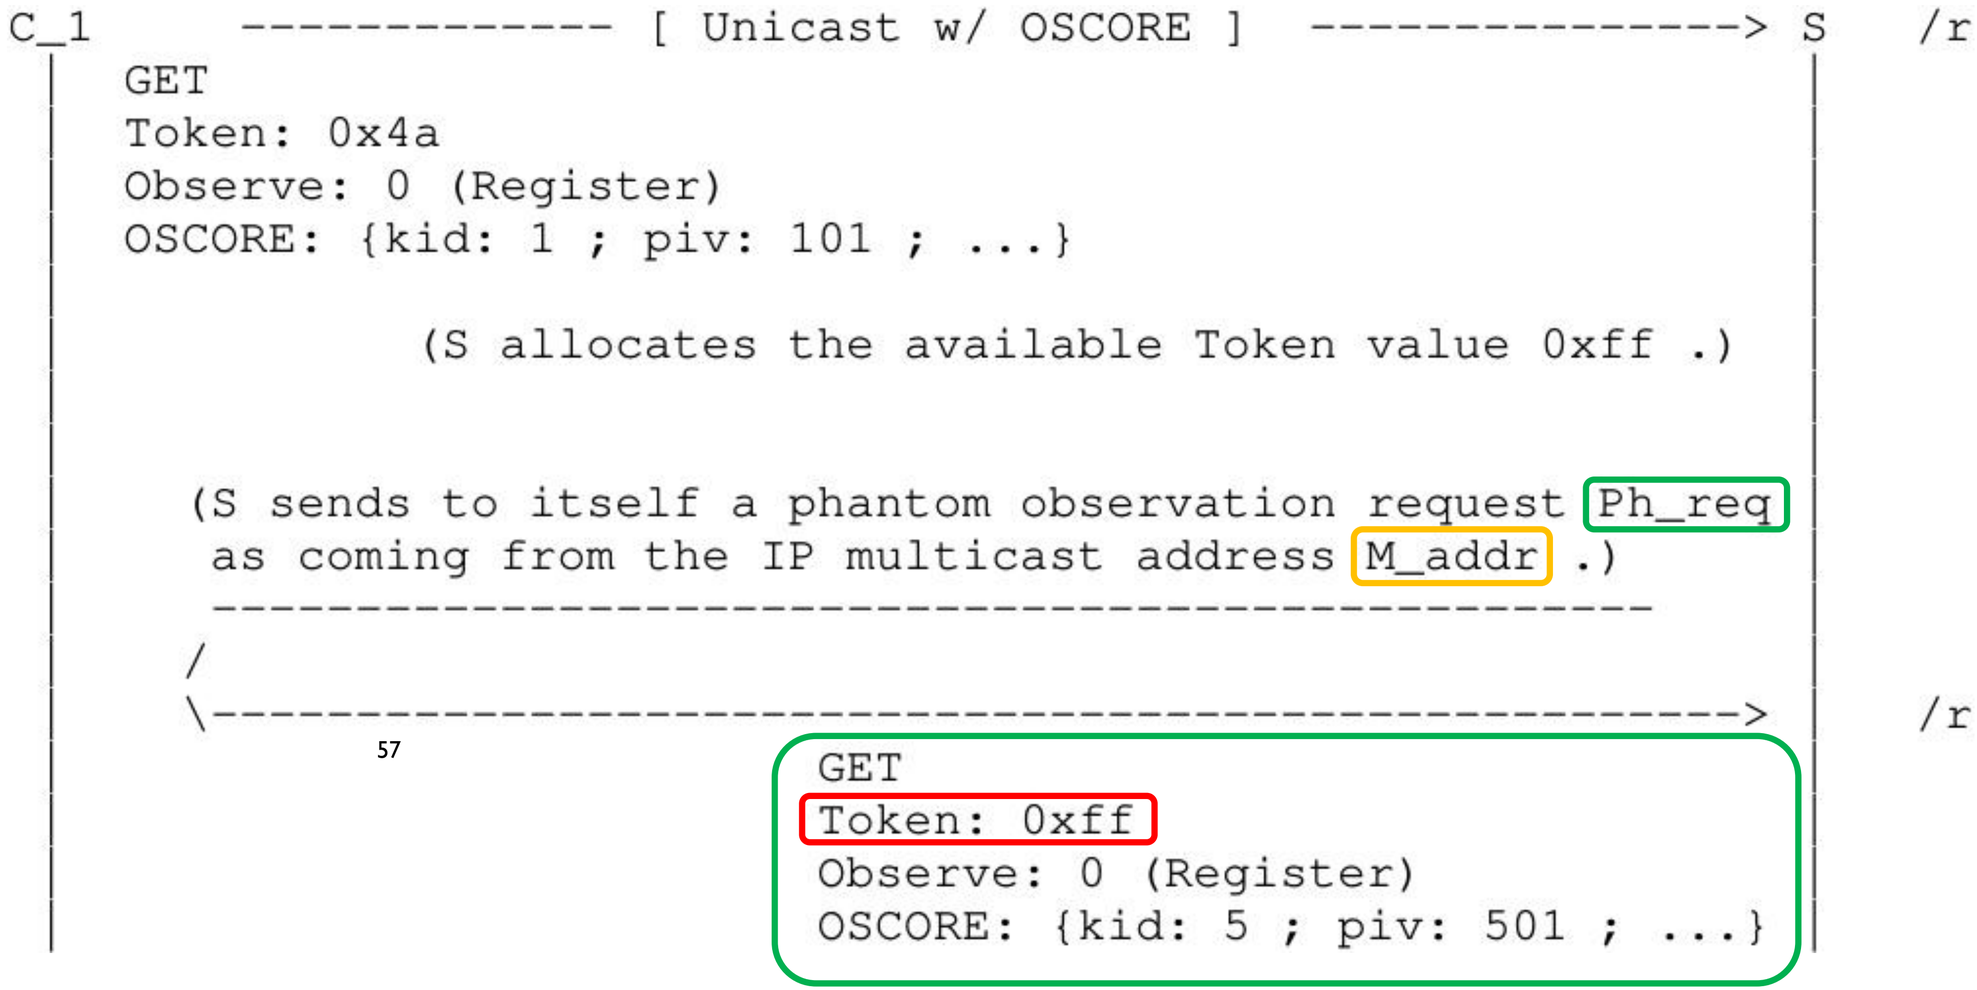

C1 registration w/ security

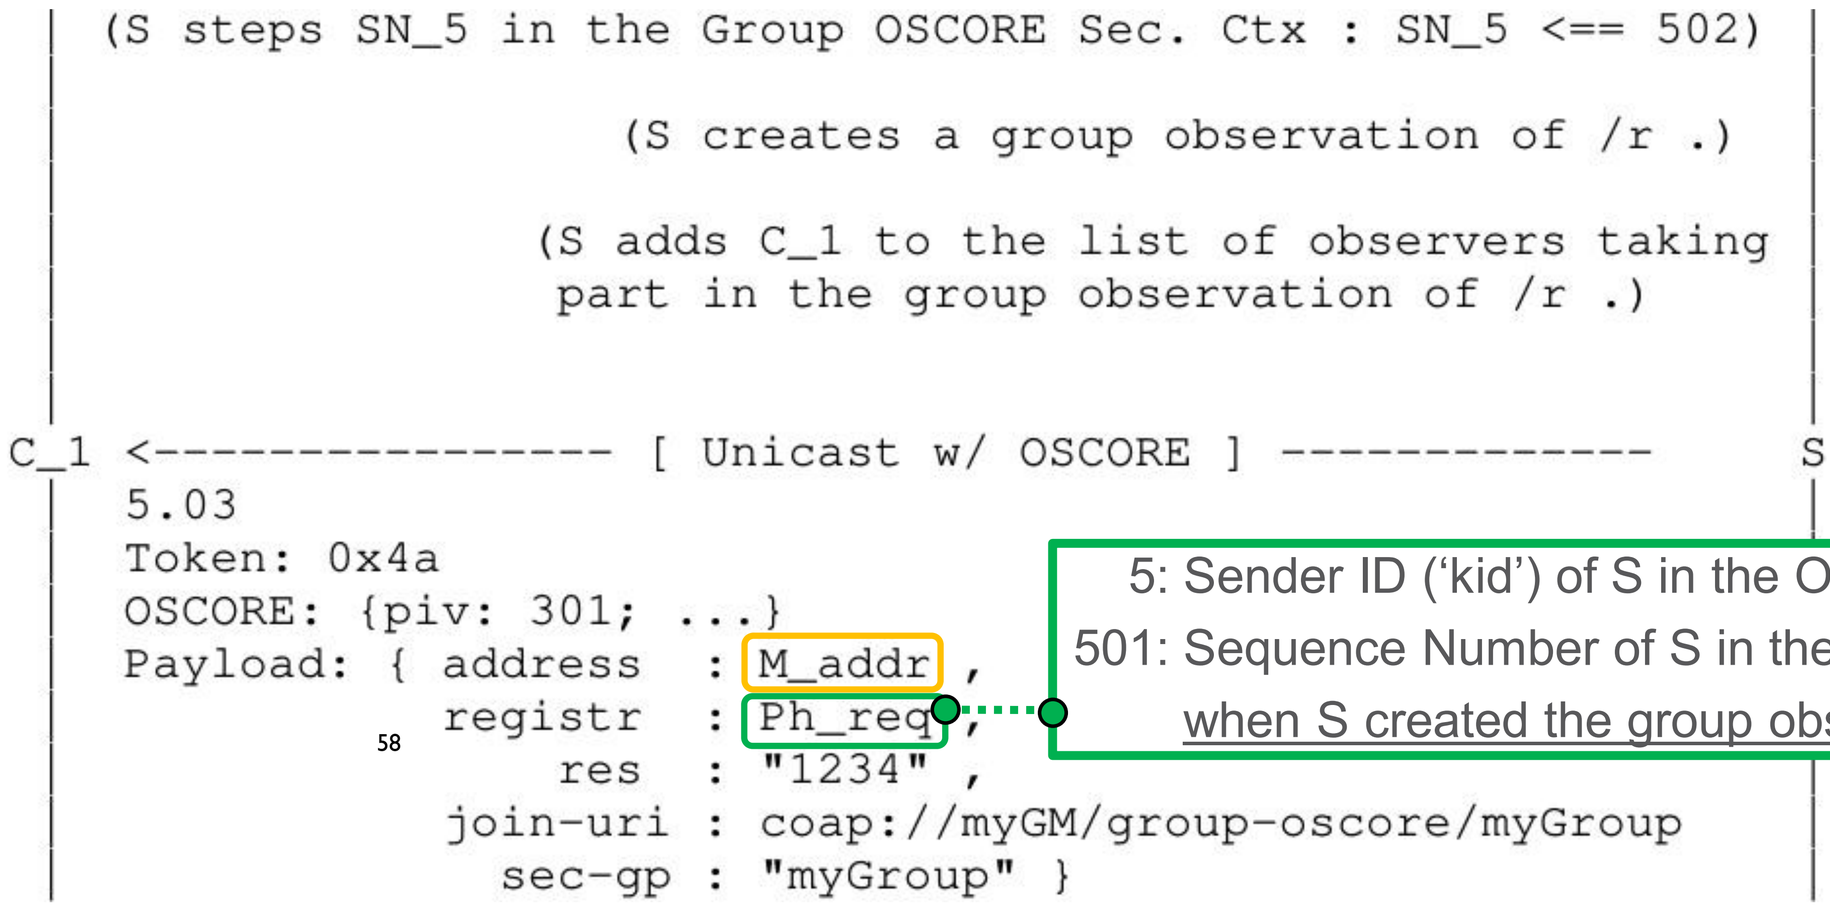

C2 registration w/ security

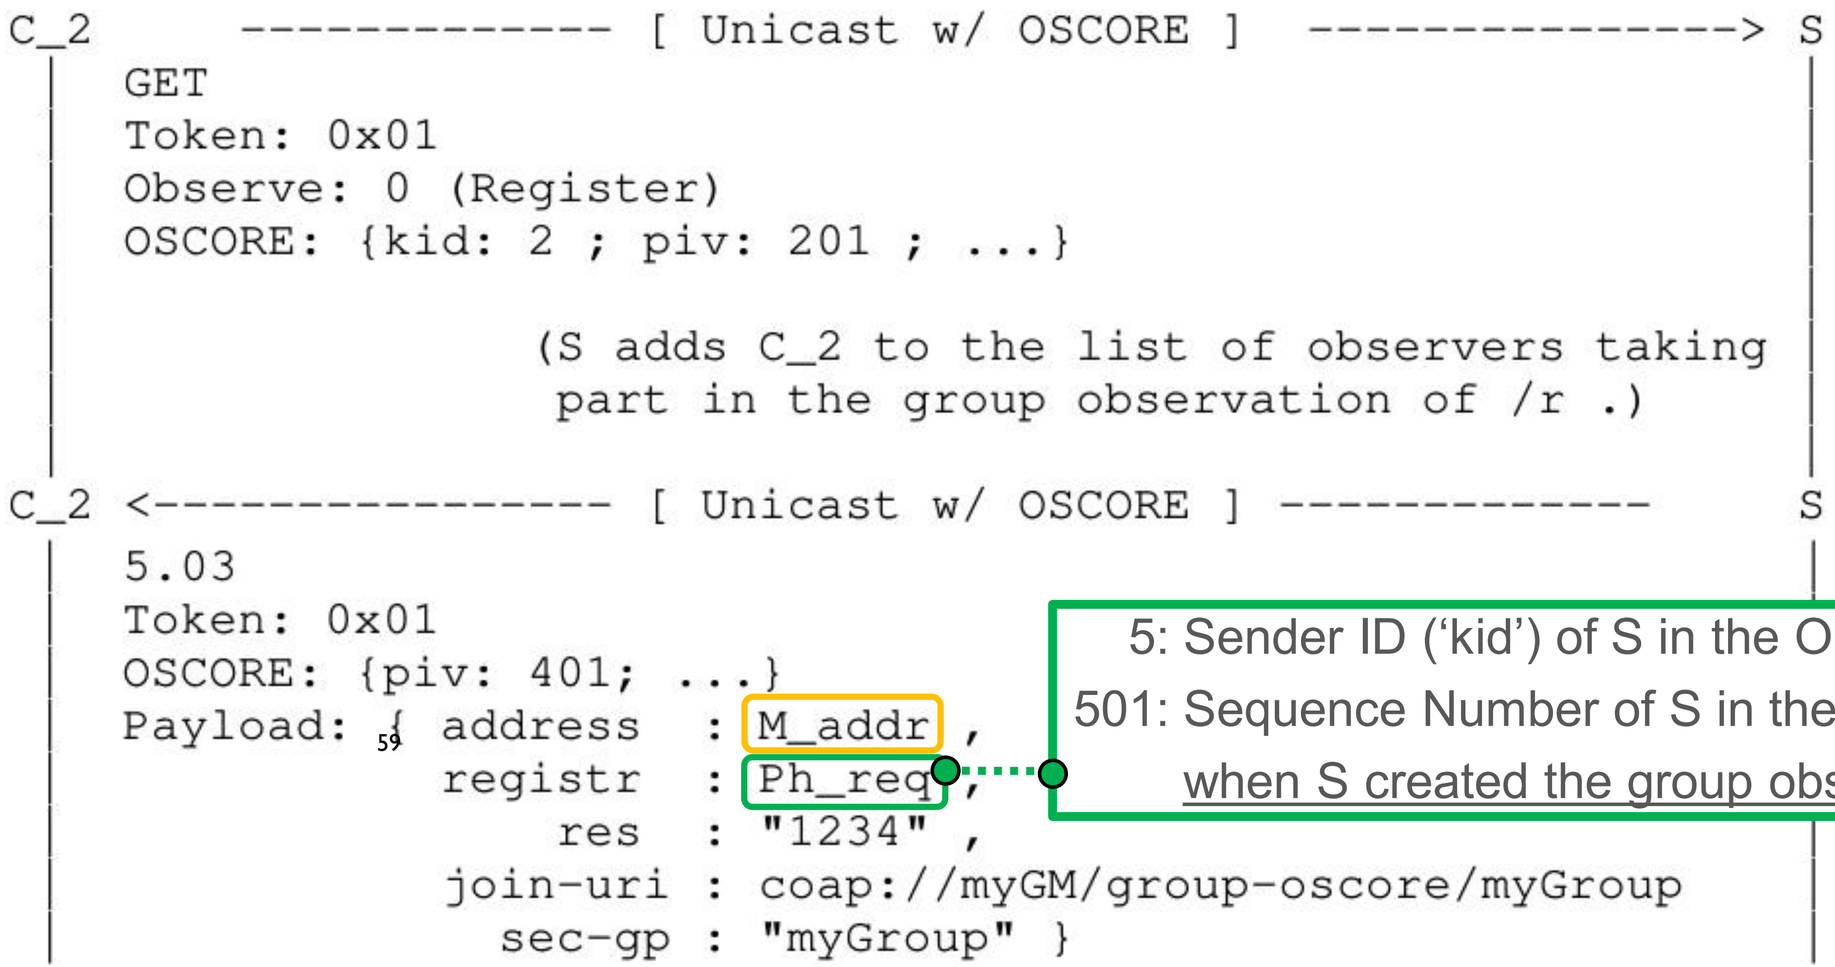

59

5: Sender ID ('kid') of S in the OSCORE group

501: Sequence Number of S in the OSCORE group when S created the group observation

Multicast notification w/ security

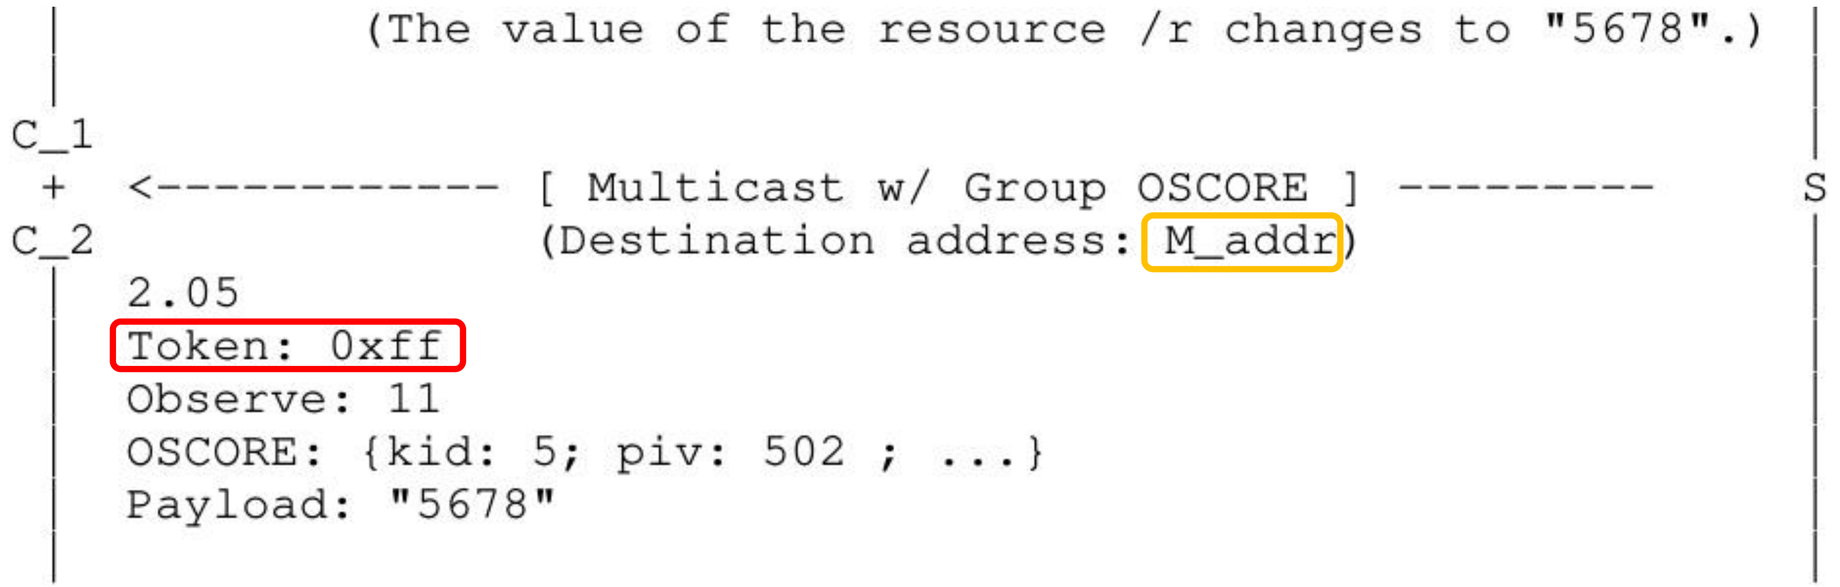

- When encrypting and signing the multicast notification:
  - The OSCORE *external_aad* has 'req_kid' = 5 and 'req_iv' = 501
  - Same for all following notifications for the same resource

- Enforce secure binding between
  - Every multicast notification for the target resource
  - The (group) observation that each client takes part in

Summary

- Multicast notifications to all clients observing a resource
  - The Server is entrusted to manage the Token space for the group
  - All notifications are (securely) bound to the group observation

- Benefits
  - Better performance when many clients observe a same resource
  - In pub-sub scenarios, subscribers can be only clients

61

- Revised simpler approach, compared to v -00

- Need for document reviews

Thank you!

Comments/questions?

https://gitlab.com/crimson84/draft-tiloca-core-observe-responses-multicast

Backup

Example with security

- {C_1, S} – OSCORE Initial status
  - C_1 : Sender ID ‘kid’ = 1; Sequence Number SN_1 = 101
  - S : Sender ID ‘kid’ = 3; Sequence Number SN_3 = 301
- {C_2, S} – OSCORE Initial status
  - C_2 : Sender ID ‘kid’ = 2; Sequence Number SN_2 = 201
  - S : Sender ID ‘kid’ = 4; Sequence Number SN_4 = 401
- {S} – Initial status[64] in the OSCORE group
  - Group ID ‘kid_context’ = “feedca57ab2e”
  - S: Sender ID ‘kid’ = 5; Sequence Number SN_5 = 501

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

Group Communication for the Constrained Application Protocol (CoAP)

draft-dijk-core-groupcomm-bis-02

Esko Dijk, IoTconsultancy.nl
Chonggang Wang, InterDigital
Marco Tiloca, RISE

IETF 106, CoRE WG, Singapore, November 20th, 2019

Goal (updated from IETF 105)

- Intended normative successor of experimental RFC 7390 (if approved)
  - As a Standards Track document
  - Obsoletes RFC 7390, except for the experimental RESTful protocol
- Be standard reference for implementations now based on RFC 7390, e.g.:
  - “Eclipse Californium 2.0.x” (Eclipse Foundation)
  - “Implementation of CoAP Server & Client in Go” (OCF)
- What’s in scope?
  - CoAP group communication over UDP/IP, including latest developments (Observe/Blockwise/Security …)
  - Unsecured CoAP or group-OSCORE-secured communication
  - Principles for secure group configuration
  - Use cases (Appendix A)

Groupcomm-bis-02: process view

- Updated with all pending reviewers' comments
  - thanks to the reviewers!
- "Copied over" and updated more of RFC 7390 content
- Closed open "TBD" items
  - Multicast transport, internetworking with other protocols
- Closed open GitLab issues

Groupcomm-bis-02: content view

- Clarified CoAP Request/Response model (2.2.1)
  - Clarified server response suppression
  - No-Response option (RFC 7967) RECOMMENDED
  - Repeated request, same or different Message ID
  - At most one response per request per server (except notifications)

- Security considerations added

- Fixes & clarifications

Observe (RFC 7641)

- -02 updates RFC 7641
  - With Multicast GET + Observe Option usage
- Unreliable request transport
  - Client may repeat the request
  - Token / Message ID usage defined for repeats
- Clarified server response suppression rules
  - Specifically for “multicast GET + Observe”
  - Server needs to verify liveness of client using occasional CON observe notifications

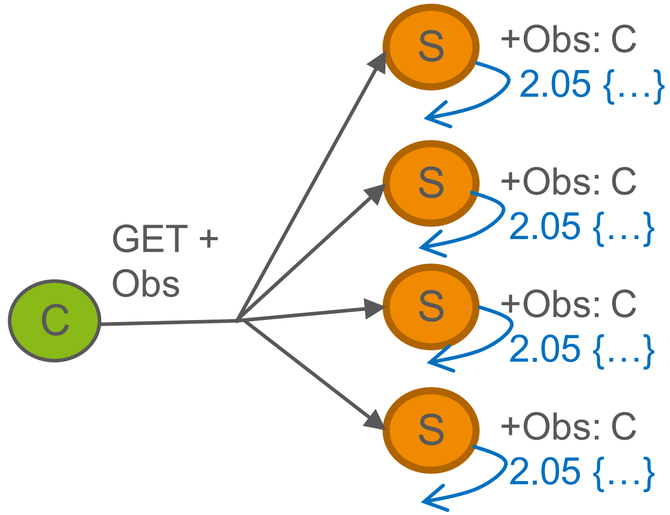

Block-wise Transfer (RFC 7959)

- -02 does not update RFC 7959 anymore
- Only considers multicast GET + Block2 as per RFC 7959
- Removed solution Block1 + Multicast
  - Turns out to be too complex to specify in scope of Section 2.2.6 …
  - Separate I-D would be possible, if there is interest

Next steps

- Fix issues found in -02
- Include -02 review comments (thanks Jim Schaad)
- Await -02 review (thanks Thomas Fossati)
- Test selected functions in CoAP implementations
  - E.g. “Observe + multicast” extension of RFC 7641
- Propose adoption by the CoRE WG!

Thank you!

Comments/questions?

https://gitlab.com/crimson84/draft-groupcomm-bis

Motivation (backup slide)

- RFC 7390 was published in 2014
  - CoAP functionalities available by then were covered
  - No group security solution was available to indicate
  - It is an Experimental document (started as Informational)
- What has changed?
  - More CoAP functionalities have been developed (Block-Wise, Observe)
  - RESTful interface for membership configuration is not really used
  - Group OSCORE provides group end-to-end security for CoAP
- Practical considerations
  - Group OSCORE clearly builds on RFC 7390 normatively
  - However, it can refer RFC 7390 only informationally

“Group” concept (backup slide)

- **Distinguish *types* of groups** (identifiers for group type:)
  - CoAP group: network level → multicast-address + port
  - OSCORE group (‘security group’) → Group ID (Gid) + Master Secret
  - Application group: application level → <any application-specific ID>
- *To do in -03: relations between group types to be detailed*
- Example of group relations:

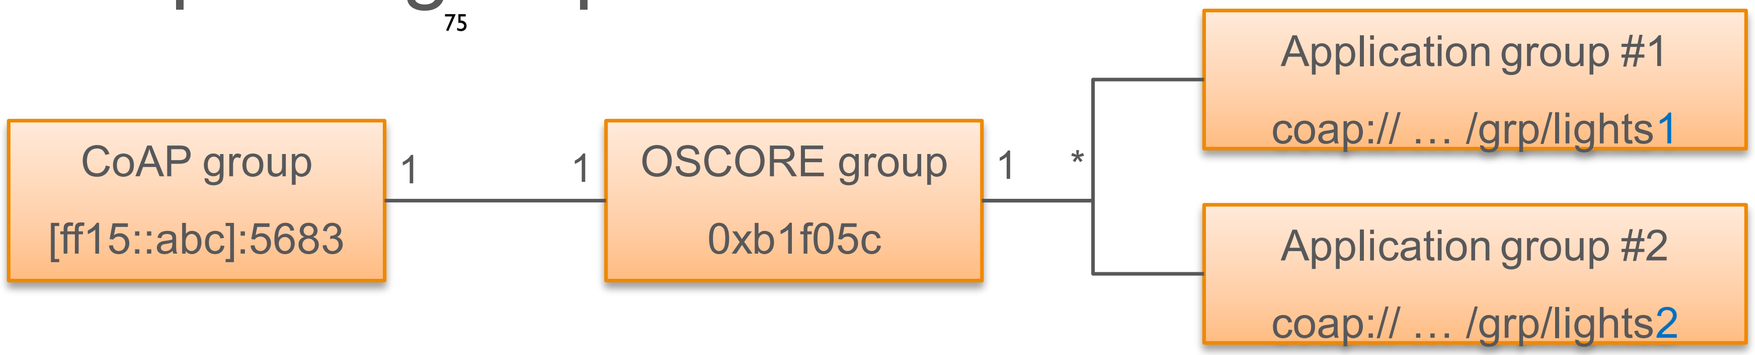

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

FETCH & PATCH with SenML

draft-ietf-core-senml-etch

IETF 106

Status

- In IESG review; revised draft needed
- Latest changes: see Github
- Addressed comments from IoT dir review (Matthias), IESG comments from Warren, Barry & Alissa, and IESG discuss from Roman
  - Re-wrote Fragment ID section; added examples
  - Clarified: Fetch needs at least one Record and Name
  - Clarified: Patch Records MUST contain Value or Sum
  - Clarified: CoAP provides the security
  - Clarified: iPATCH and PATCH are equivalent here
  - Return empty Pack when no matches to FETCH
  - IANA registration clarifications
  - Editorial fixes
- TBD: Discuss from Adam & Comment from Ben

FETCH/PATCH Unit Selector

- Currently name (always) and time (optionally) used to select Records
  - Enables to select Record with name & time and Patch the other fields
- Some use cases have same name & time, but different Unit
  - e.g., lat/lon or V and A
- Proposal: use name **and** {time, unit} to select Records
  - Exact match: Fetch Record without Unit would not match Record with Unit

Patch Record order matters

- A Patch Record may change the Pack so that Pack becomes valid/invalid for later Patch Records
  - For example, remove same record twice
- Solution suggested by Adam:
  - If Patch Record matches more than one SenML Record: error
    - TBD: multi-record Patch?
  - If one Patch Record fails, the state of the Pack is not changed
  - MUST apply Patch Records sequentially

Multi-record Patch?

- Currently: Patch selector matching multiple Records is error (but text is ambiguous)
- Alternative: same Patch operation is applied to all matching Records
  - For example, "delete all Records with name X"

Additional Units for SenML

(Slides donated by Ari, thank you)

SenML Units Registry today

- Registry of short strings to represent units of measurement
  - e.g., "m" for meters and "s" for seconds
- Restrictive registration policy to facilitate interoperability
  - Unscaled SI units and "a few more"
  - Only one unit for each kind of measurement (no “km” or “miles” as we have “m”)
- However, many derived and other units in practical use today
  - OMA SpecWorks IPSO/LwM2M models use a richer set:
    "ms" for time values, also "kWh", "dBm", etc.
  - Would need to **change a large amount of existing models** to use unscaled SI
  - Many use cases have a "natural" scaled/offset unit (e.g., "ms" for time or "um" for particle size); having to use exponent every time brings extra cost

Proposal: **secondary** registry

- Another SenML IANA (sub) registry for units with different rules
  - No completely “new” units: must be based on a primary SenML unit
  - Secondary registry describes translation rules for conversions to the primary set, e.g.:

"Relaxed unit"	SenML unit	scale	offset
km	m	1000	0
dBm	dBW	1	-30

- Discussed with OMA SpecWorks, consensus that this is a good way forward
- Good thing: Conversion can be entirely automatic
  (would need retrieval API in IANA, though)

Issue: How mandatory is implementing units from secondary registry?

- Currently **use** is NOT RECOMMENDED:
  - “SenML packs MAY, but SHOULD NOT, use secondary units in place of SenML units, where the exception of the "SHOULD NOT" lies in the context of specific existing data models that are based on these secondary units.”
- Some commenters do not like “MAY, but SHOULD NOT”
  - But that is exactly intended: use only if you have to
- Some commenters worry about creating two SenMLs, one where the secondary registry is **implemented** and one where it isn’t
- But RFC 8428 is clearly updated to now include secondary registry
- ➔ Be more clear, then

SenML More Units (more considerations)

IETF 106

Cullen, Ari, Carsten

Proposed expert review guidance for "table 2"

- Unit IDs subset of (ALPHA+DIGIT+...)
  - no quotes; character set same as names?
- Unit has to be found in existing scientific literature or specification
- If same unit as something that exists (same conversion)
  we ask if really necessary to duplicate; if yes, we do
- We try to give the string to most common use in scientific literature (check potential conflicts)
- OK to have two things for same SI unit, e.g., reactive/apparent power
- Naming syntax for things like "events per hour per square meter"
  - 1/h/m2? Double slashes OK?

Using "u" or (new) "u2" field

- "u" proposal: both table 1 and 2 units go into the existing "u" field
  - Software written only expecting table 1 will start receiving many things not expected from table 2
  - Could be useful to have clear indication on wire which table you are using
- "u2" proposal: units from either table can go into a new field called "u2" (or similar) but only table 1 units can go to "u"
  - Complexities with different fields for similar things (especially when combining SenML Packs from different sources)
- Need to do analysis and think of corner cases. Feedback welcome!
  - More discussion Friday?

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- 10:00–10:10 Intro, Agenda, Status
- 10:10–10:25 CoRECONF (IP)
- 10:25–10:35 OSCORE groupcomm (MT)
- 10:35–10:45 OSCORE discovery (MT)
- 10:45–11:00 Observe multicast notifications (MT)
- 11:00–11:15 Groupcomm bis (ED)
- 11:15–11:35 SenML in IESG (etch, units) (AK)
- 11:35–11:45 SenML data ct (AK)
- 11:45–12:00 SenML base prefix (AK)

SenML Data Value Content-Format Indication

draft-ietf-core-senml-data-ct-01

Ari Keränen

IETF 106

Examples

```
{"n":"nfc-reader", "vd":"gmNmb28YKg", "ct":"60"}
```

```
{"n":"nfc-reader-42",
 "vd":"H4sIAA+dmFwAAzMx0jEZMAQALnH8Yn0AAAA",
 "ct":"text/csv@gzip"}
```

Changes (as agreed at IETF105)

- Using same field and string values for both content-format (numbers) and content-type and -coding (strings)
- Mandatory to understand "ct" field ("ct_", and "bct_")
- Example ct values
  - "60" (CoAP Content-Format for "application/cbor")
  - "0" (CoAP Content-Format for "text/plain" with parameter "charset=utf-8")
  - "application/json@deflate" (JSON with "deflate" as Content-Coding - equivalent to "11050" CoAP Content-Format identifier)
  - "text/csv" (CSV Content-Type)
  - "text/csv@gzip" (CSV with "gzip" as Content-Coding)

Mixing b and _ fields: what are the resolution rules?

1)

```
[
 {"bfoo_":42, "n":"t1", "v":1},
 {            "n":"t2", "v":2}
 {"foo": 1,   "n":"t3", "v":3}
]
```

2)

```
[
 {"bfoo_":42, "n":"t1", "v":1},
 {            "n":"t2", "v":2}
 {"foo_": 1,  "n":"t3", "v":3}
]
```

3)

```
[
 {"bfoo":42, "n":"t1", "v":1},
 {           "n":"t2", "v":2}
 {"foo_": 1, "n":"t3", "v":3}
]
```

4)

```
[
 {"bfoo":42, "n":"t1", "v":1},
 {           "n":"t2", "v":2}
 {"foo": 1,  "n":"t3", "v":3}
]
```

Resolution rules with mandatory fields

- Current text: *"ct_" field overrides the "ct" field. Using both "ct" and "ct_" in the same Record is NOT RECOMMENDED*
- Unfortunately: it is more complicated than you think

- TBD; feedback welcome. But not specific to this draft.

94

All times are in time-warped SGT (UTC+08:00)

Wednesday (120 min)

- **10:00–10:10 Intro, Agenda, Status**
- **10:10–10:25 CoRECONF (IP)**
- **10:25–10:35 OSCORE groupcomm (MT)**
- **10:35–10:45 OSCORE discovery (MT)**
- **10:45–11:00 Observe multicast notifications (MT)**
- **11:00–11:15 Groupcomm bis (ED)**
- **11:15–11:35 SenML in IESG (etch, units) (AK)**
- **11:35–11:45 SenML data ct (AK)**
- **11:45–12:00 SenML base prefix (AK)**

SenML Base Name Prefix Indication

draft-keranen-core-senml-base-prefix-00

Ari Keränen

IETF 106

SenML Base Name Prefix Indication

- SenML uses globally unique names (e.g., IPv6 address as prefix)
  - facilitates information exchange across systems
- Result: "long" names and hence potentially "large" (for extremely constrained networks) amount of data to be transmitted
- Usually out-of-band information available for the prefix: IP address, request URI, TLS identity, etc.
- Proposal: indicate which out-of-band info to use in the SenML Pack
  - Convert out-of-band info to regular base name after constrained hop
- IPR #3662

Example

```
[ {"bn":"2001:db8:1234:5678::1",
   "n":"/temp", "u":"Cel", "v":25.2},
  {"n":"/humi", "u":"%RH", "v":30} ]
```

```
[ {"bpi":1,
   "n":"/temp", "u":"Cel", "v":25.2},
  {"n":"/humi", "u":"%RH", "v":30} ]
```

98

Proposed bpi values

- IP address
- IP address & port
- Request base URI
- Public key fingerprint (RFC 6920 URL Segment Format)
- TLS PSK Identity
- CoRE RD endpoint

Status

- draft-tschofenig-core-senml-lbn addressing same problem
  - Replaces "base name" with "local base name"
  - Relaxes requirements for global uniqueness

- Next steps: closer look at use cases
  - LwM2M uses SenML but apparently no interop issues between clients and servers today (basically: informally do this without indicating bpi)

100

Constrained RESTful Environments WG (core)

Chairs:

Jaime Jiménez <jaime@iki.fi>

Carsten Bormann <cabo@tzi.org>

Mailing List:

core@ietf.org

Jabber:

core@jabber.ietf.org

- **We assume people have read the drafts**
- **Meetings serve to advance difficult issues by making good use of face-to-face communications**
- **Note Well: Be aware of the IPR principles, according to RFC 8179 and its updates**

üBlue sheets
üScribe(s)

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- BCP 9 (Internet Standards Process)
- BCP 25 (Working Group processes)
- BCP 25 (Anti-Harassment Procedures)
- BCP 54 (Code of Conduct)
- BCP 78 (Copyright)
- BCP 79 (Patents, Participation)
- https://www.ietf.org/privacy-policy/ (Privacy Policy)

All times are in time-warped SGT (UTC+08:00)

Friday (90 min)

- **12:20–12:25 Intro, Agenda**
- **12:25–13:10 CoRE applications (KH)**
- **13:10–13:50 Flextime**

SenML More Units (but not fields?)

Ari, Cullen, Carsten

IETF 106

Options

- Indicate if using units from 2nd registry ("* option")
- Use units from 2nd registry as such

Why need to indicate?

- Some SenML systems use the Unit field for routing SenML records
  - "Is this information relevant to me"
  - Names discovered dynamically based on incoming records
- Assumption based on SenML RFC: application that does, e.g., speed needs to look for only "m/s"
  - Not "km/h" or "furlongs per fortnight"

- Useful to indicate that with some units this may no longer be true
  - Raise error instead of silently discard

Proposal

- Prefix units from Table 2, when used in SenML Record Unit field, with a special character not used in the units
  - Asterisk ("*") seems like a safe bet
  - Example: `"u": "*km/h"`
- Considerations
  - Yes, it's one more byte (but if you really care about the last byte, probably use some form of extra compression?)
  - Breaking the principle of least astonishment
  - Not very elegant

Resource Directory

draft-ietf-core-resource-directory

Zach Shelby, Michael Koster, Carsten Bormann, Peter van der Stok,
Christian Amsüss

2019-11-22

Status

- -21 etc. addressed WGLC comments
- -23 just received AD review

Open points for -24

- Is RD-DNS-SD a normative reference?

  ```
  The following RD discovery mechanisms are recommended:

  o In managed networks, [RDAO or anycast].

  o The use of DNS facilities is described in [RD-DNS-SD].
  ```

- Editorial and markup changes

All times are in time-warped SGT (UTC+08:00)

Friday (90 min)

- **12:20–12:25 Intro, Agenda**
- **12:25–13:10 CoRE applications (KH)**
- **13:10–13:50 Flextime**

Sub-agenda

Report on Side-Meeting

* Problem Details for CoAP APIs
* Extensibiliy and Code Points
* Error Cases in CoAP API Specifications

Update on CoRAL

* Update on draft-ietf-core-coral-01 (quick)
* Update on draft-ietf-core-href-01 (quick)

Update on CoRAL-based Applications

* Update on draft-hartke-t2trg-coral-pubsub-00
* Update on draft-hartke-t2trg-coral-reef-03 (quick)
* Update on draft-hartke-t2trg-data-hub-05 (quick)
* Update on draft-tiloca-ace-oscore-gm-admin-01 to be

Planning; discuss timeline

Completing CoRAL: A timeline

- Step 1: Attain stable version
  - Address the https://github.com/core-wg/coral/issues
    - Sequentially, pull out each non-trivial issue to mailing list as “issue of the day”
  - Proposed target: IETF107 (March 2020, Vancouver)
- Step 2: Validate, implement, review
  - Obtain input from a wider group of implementers and experts; address issues that emerge
  - Proposed target: IETF108 (July 2020, Madrid)
- Step 3: WGLC mid-2020

Flextime