Constrained RESTful Environments WG (core)

Chairs:

Jaime Jiménez <jaime.jimenez@ericsson.com>

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)

Agenda Bashing

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP — moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

All times are in time-warped EDT (UTC–04:00)

Thursday (60 min)

- **18:10–18:15 Intro, Agenda**
- **18:15–18:20 DOTS heads-up (DOTS chairs)**
- **18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)**
- **18:34–18:46 Housekeeping cluster (AK, CB)**
- **18:46–18:58 Other WG drafts (MK) /candidates (BS)**
- **18:58–19:10 FASOR: Alternative Congestion Control**

Advertisements

- DNSSD: Thu 09:30..12:00 Duluth
- (see also cluster agenda on mailing list)

- OCF/T2TRG coordination call Wed 11..12 (please ask chairs)

draft-ietf-core-links-json: Status

- JSON version of 6690-to-be — avoid need for another parser
  - Started Feb 2012, added CBOR variants mid-2015
- Focus was: roundtrippable with RFC 6690
  - Inherit limitations of RFC 6690 (e.g., percent-encoding)
- Submitted to IESG on 2017-04-02: Lots of feedback
- Re-focus:
  - Still cover all of RFC 6690
  - Be more general, don’t inherit the limitations
- Lots more input from CorE-RD, W3C WoT TDir work, related concepts in OCF spec
- Discussions will go on in hallways this week

draft-ietf-core-cocoa: Status

- Submitted to IESG
  - Responsible AD here: Mirja Kühlewind (TSV AD)
  - Great AD feedback
- Authors need to generate new version (this week?)
- Should go though normal process then

- CoCoA is not the end-all of congestion control work for CoAP
- Proposed new work: draft-jarvinen-core-fasor
  (Thu, if we have time)

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP -- moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

Object Security for Constrained RESTFUL Environments

OSCORE

draft-ietf-core-object-security-13

Göran Selander, Ericsson
John Mattsson, Ericsson
Francesca Palombini, Ericsson
Ludwig Seitz, RISE SICS

IETF 102, CoRE WG, Montreal, Jul 16, 2018

Status (v-13)

- Main changes: Clarifications and further details based on the comments received by IESG and other Post LC reviews
- In particular in the new Appendix D – Overview of security properties
- Increased protection of certain CoAP options and motivation for lack of protection of certain options
- Additional clarifications and simplifications of processing

12

- Up-to-date comments on the wiki: https://github.com/core-wg/oscoap/wiki

V-13 Changes In Detail

- Observe is now additionally Inner, which enables the endpoints to verify each others intent and simplifies the specification, at the cost of making some of the proxy processing out of scope. Observe processing is separated.
- No-Response is now essentially Inner, following a review byJim Schaad
- Uri-Host/Port processing is clarified in a separate subsection
- A corresponding change of the analysis of unprotected header fields was made in appendix D

V-13 Changes In Detail

- HTTP processing updated based on comments from Martin Thomson
- CoAP-to-CoAP Forwarding Proxy description is expanded
- ID Context added to the security context and key derivation. Such a parameter was already in use by Group OSCORE and 6TiSCH Minimal Security and they can now apply this in a common way
- Updated deployment examples, test vectors (appendices B and C), and references

Next Steps

- Update based on recent review comments
- Continue IESG evaluation
- Interop-testing the next version

15

Secure group communication for CoAP

draft-ietf-core-oscore-groupcomm-02

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

IETF 102, CoRE WG, Montreal, July 16th, 2018

Updates from -01 (1/3)

- Major revision:
  - Based on discussions at IETF 101
  - Aligned with latest *draft-ietf-core-object-security*

- Section 1.1 – “Terminology”
  - Removed “Multicaster” and “Listener”
  - Now simply “Client” and “Server”, or “Sender” and “Recipient”
  - The old “Pure listener” is now called “Silent server”

17

- Section 2 – “OSCORE Security Context”
  - Group Identifier (Gid) stored as the “ID Context”
  - “ID context” defined in *draft-ietf-core-object-security*

Updates from -01 (2/3)

› Section 3 – “The COSE Object”
  – Format of ‘external_aad’ consistent with *draft-ietf-core-object-security*

› Section 4 – “Message Processing”
  – Major rewriting for plain alignment with *draft-ietf-core-object-security*
  – Now pointing at exact steps of the OSCORE message processing
  – Only the Gid is used for context retrieval, regardless the IP address

18

› Section 7 – “Security Considerations”
  – Section 7.2 – “Uniqueness of (key, nonce)” // The same holds from OSCORE
  – Section 7.3 – “Collision of Group Identifiers” // Not impairing security

Updates from -01 (3/3)

- Appendix C – “Example of Group Identifier Format”
  - Clarified practical implications in case of collisions
  - A recipient may go for trial & error, until the right context is found
  - Favorable to discourage collisions with appropriate Gid sizes
  - Thanks to Esko Dijk for the good discussion!

- Appendix D.2 – “Provisioning and retrieval of public keys”
  - Updates for alignment with *draft-palombini-ace-key-groupcomm* [19]

- See full list of updates in Appendix G.1

Implementation

- Plans for a Java version in Californium
  - Build on the current OSCORE implementation

- OSRAM Innovation
  - Developed in C
  - MediaTek LinkIt Smart 7688
  - Aligned with individual submission at IETF99

- Proof-of-concept for Contiki OS
  - Wismote (MSP430; TI CC2520)
  - SmartRF (MSP430; TI CC2538)
  - Aligned with individual submission at IETF99
  - https://github.com/tdrlab/mcast

Related activity

› *draft-tiloca-ace-oscoap-joining-04*
  – Referred by Appendix D.3

› Join an OSCORE group using the ACE framework
  – Joining node → Client
  – Group Manager → Resource Server
  – Message formats aligned with *draft-palombini-ace-key-groupcomm*

21

› Renaming for consistency
  – “Multicaster” → ”Requester” , as in *oscore-groupcomm*
  – “Pure listener” is the “silent server” of *oscore-groupcomm*
  – Kept “Listener” and “Pure listener” to avoid confusion with ACE roles

Thank you!

Comments/questions?

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

Support for group comm.

- draft-ietf-core-oscore-groupcomm-02
  - The Sender Context stores the endpoint's public-private key pair
  - The Recipient Context stores the public key associated to the endpoint from which messages are received
  - Recipient Contexts are derived at runtime

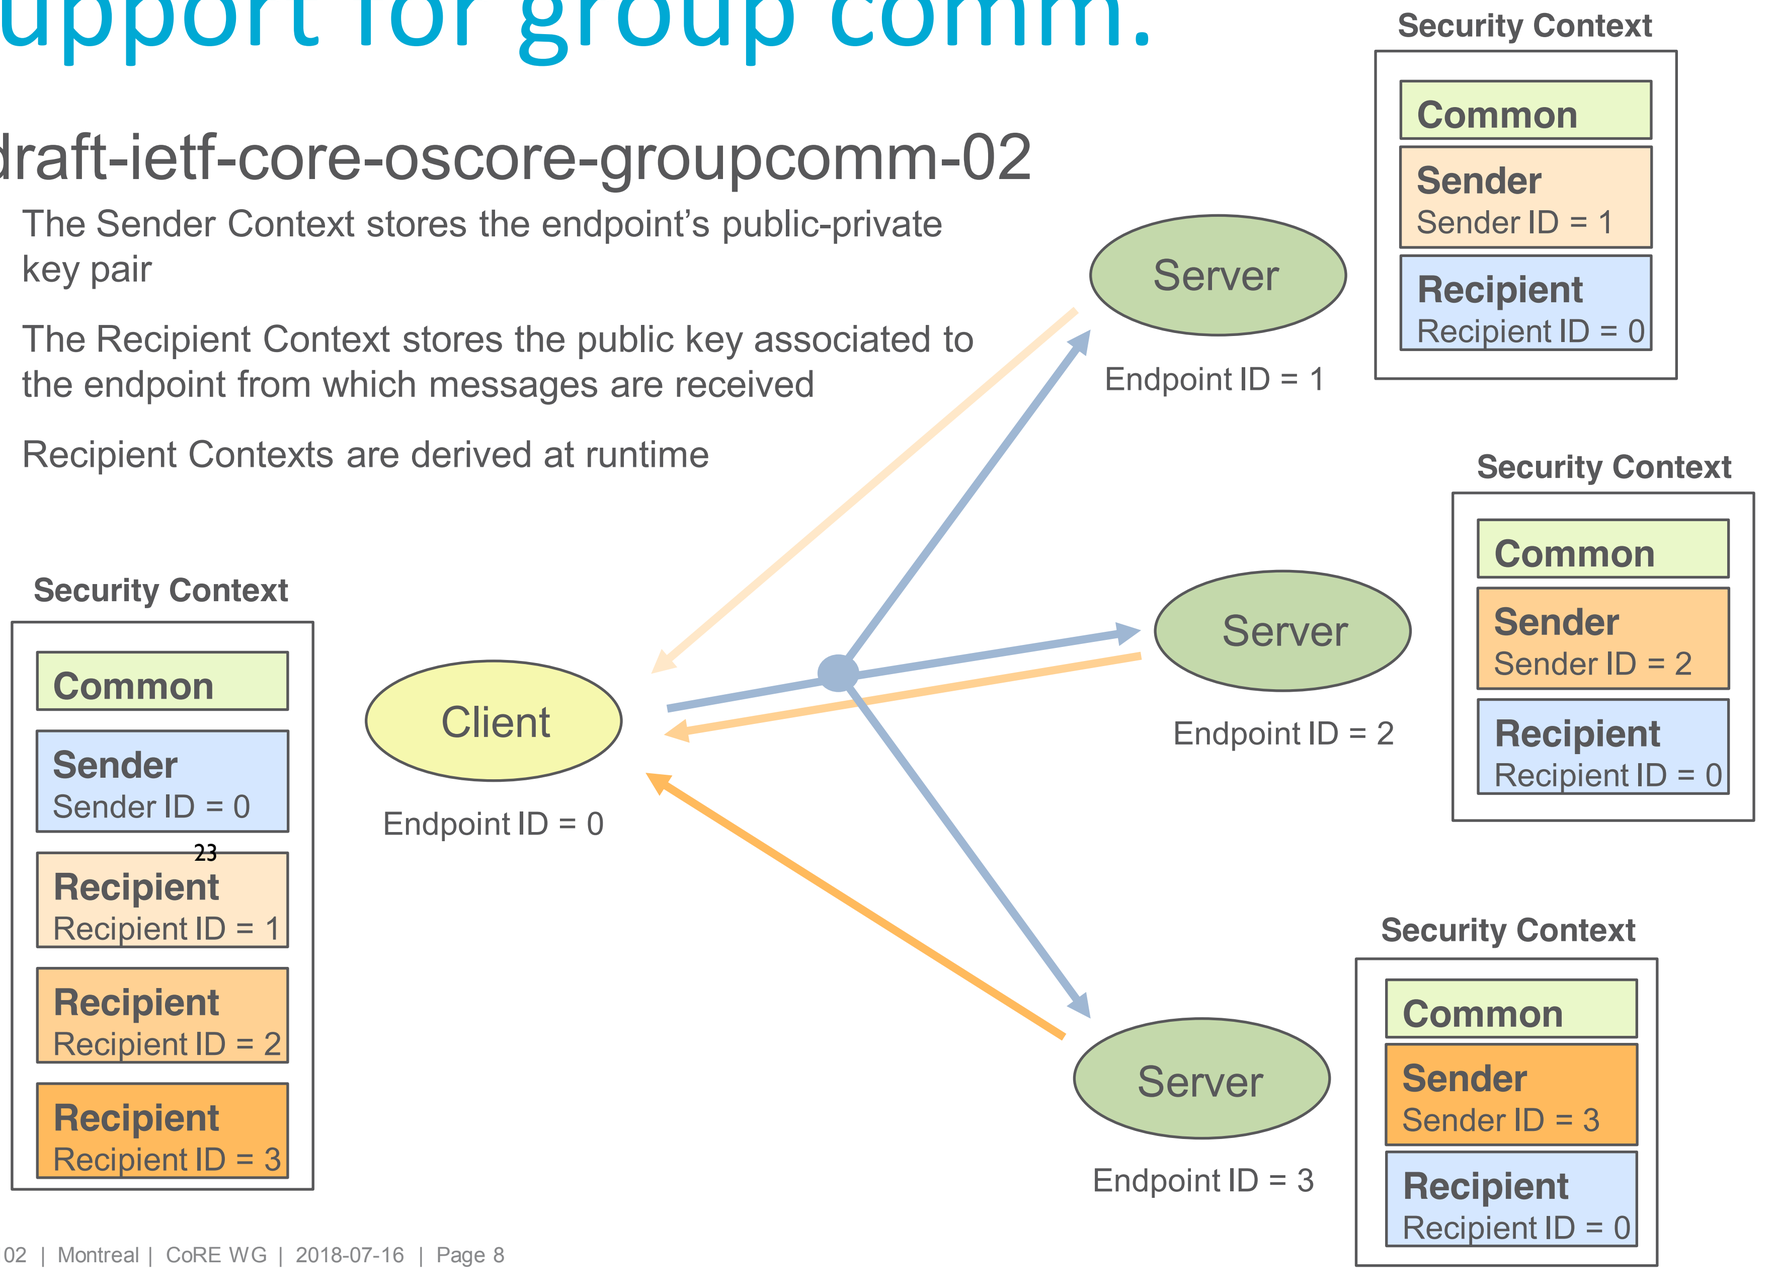

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP -- moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

Resource Directory

Peter van der Stok, Carsten Bormann, Michael Koster
Christian Amsuess

25

IETF 102 - CoRE Working Group

URI

URI syntax: scheme://authority/path/?query#fragment

URI reference is a URI or relative reference (no scheme component)

scheme://authority part is needed as prefix to relative reference

Resolving a URI reference against Base URI results in target URI
RFC8288

26

Relative references available in /.well-known/core
"hosts" relation from RFC6690 links scheme://authority part
to relative references

Maintain link semantics from host to RD

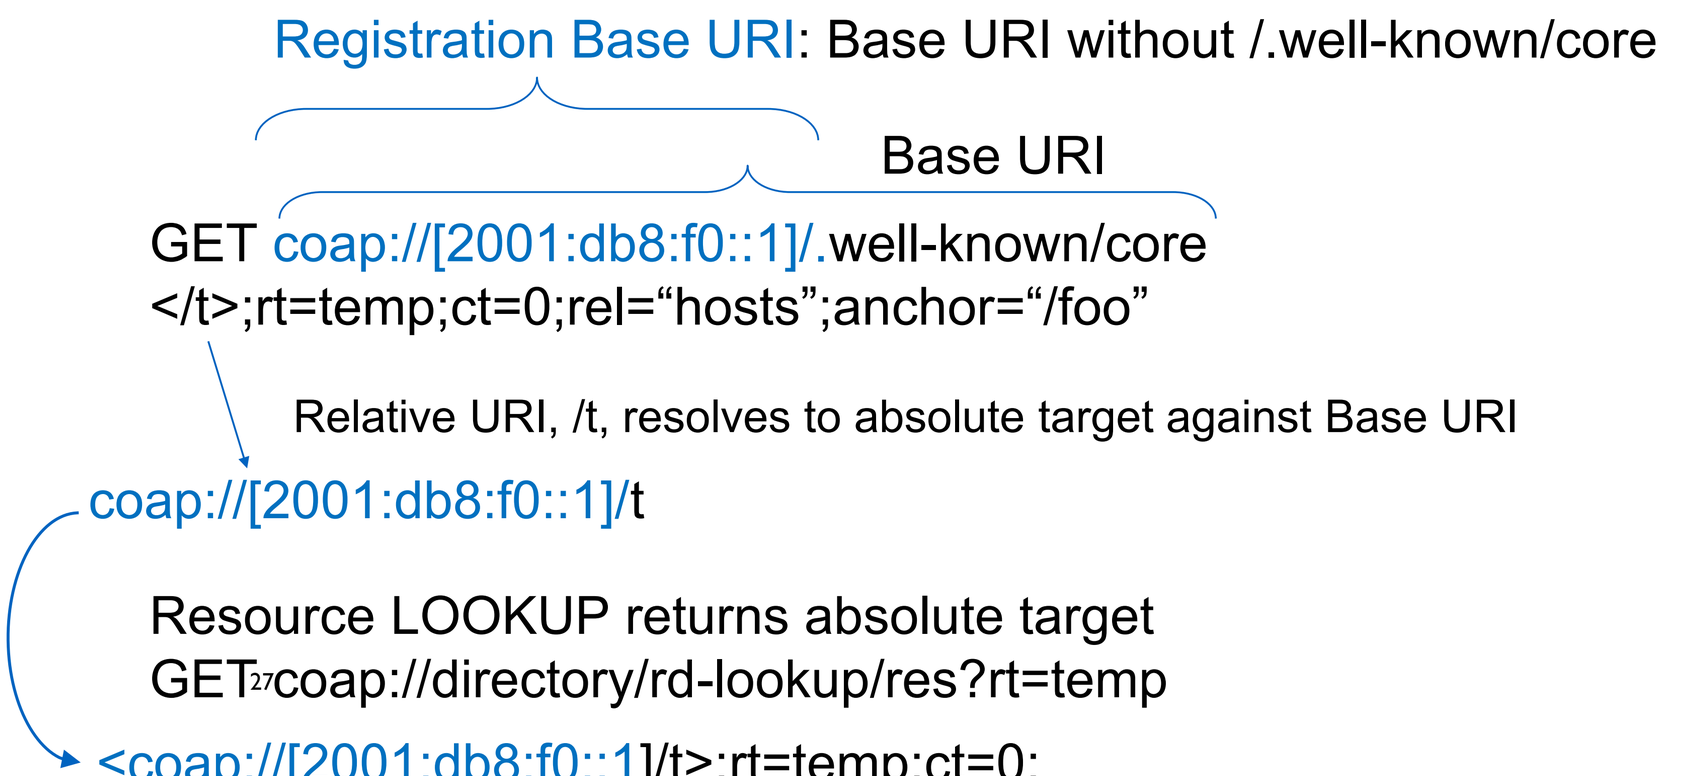

The link context is:

- Value of the anchor=context parameter in link specification
- With no anchor=, context is the base URI

Registration Base URI

Registration Base URI:

- Base URI with /.well-known/core stripped
- Value of base=Registration Base URI in link specification

Stored in Resource directory Registration

IN LOOKUP:

- Registration Base URI prefixed to relative reference
  28 to return absolute reference
- Otherwise absolute reference is returned

RFC 6690 and RFC 8288

RFC6690: anchor is used as Base URI against which relative target is resolved
RFC8288: anchor is immaterial to resolution

RFC6690: without anchor, context is target URI with paths stripped off.
RFC8288: context is given by Base URI

Modernized Link format to avoid ambiguities:

- Relative target URI always resolved against Base URI[29]
- Anchor= context
- When no anchor, Base URI is context

Other improvements to RD text

- domain -> sector (maintained d=)
- con= -> base= (registration context -> registration base URI)
- rt-types: core.rd-ep and core.rd-gp introduced
- Simple registration more concrete and reworded
- Lookup: return of resolved references.
- It not exposed in lookup (ambiguous result)
- Registration update clarified

30

TODO

- React to reviews (thanks for the many we received Jim)
- Remove ambiguous unclear text

31

WGLC

Yes, please,

We think that no structural changes are needed any more

Discovery Mapping

CoRE Link Format <-> DNS-SD RRs

draft-ietf-core-rd-dns-sd

Kerry Lynn, Peter van der Stok, Michael Koster, Christian Amsüss
2018-07-16, IETF 102 CoRE WG, Montréal

Why? (Use Cases)

- Support alternate methods of discovery in heterogeneous environments (e.g. HTTPS clients and CoAPS servers)
- Support hierarchical discovery in large environments (e.g. many K's of points)
  - DNS-SD for coarse-grained discovery
  - CoRE Link Format for fine-grained discovery
- Discovery bootstrapping (i.e. locating Resource Directories)

DNS-Based Service Discovery [RFC6763]

- A conventional use of existing DNS RRs and message formats to support service discovery:

DNS Resource Record	Binding
PTR	<ServiceType> to service instance name
SRV	Service instance name to host, port (end-point)
TXT	Arbitrary key=value pairs (e.g. "path=/lamp/1”)
A, AAAA	Host name to IP address

- Expand the definition of *service* to include REST API entry point (e.g. in multi-function devices)
- Service instance names are of the form:

  <Instance>.<ServiceType>.<Domain>

New/Required Link Target Attributes

- exp, hint that information about this resource should be exported
- ins=, instance name in UTF-8 format
- rt=, resource type (federated namespace?)
- if=, semantic tag or link to interface description

Link-format to DNS-SD mapping

Link Format	DNS-SD
Resource Instance (ins=)	<Instance>
Resource Type (rt=)	<ServiceType>
<uri>	TXT path=/{relativeURI}
Interface Description (if=)	TXT if={anyURI}
Other attribute (key=value)	TXT key=value

TBD:

- Domain name (the DNS zone where the records are created)
- Host name (if it doesn't already exist) for naming AAAA RRs

Link Format -> DNS-SD Example

CoRE query

REQ: GET coap://[ff02::1]/.well-known/core?exp

RES: 2.05 "Content" (from [fdfd::1234]:5678)
</sensors/temp/1>;exp;ct=50;rt="oic.r.temperature";
ins="indoorTemp"; if="oic.if.s",

Resulting RRs

_oic._udp.example.com. IN PTR indoorTemp._oic._udp…

r-temperature._sub._oic._udp… IN PTR indoorTemp._oic._udp…
indoorTemp._oic._udp… IN TXT txtver=1
indoorTemp._oic._udp… IN TXT path=/sensors/temp/1
indoorTemp._oic._udp… IN TXT if=oic.if.s
indoorTemp._oic._udp… IN SRV 0 0 5678 node1234...
node1234.example.com. IN AAAA fdfd::1234

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP -- moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

Marketing message: "CoRECONF"

Note:
You can mix and match
(to a certain extent)

CoMI update

draft-ietf-core-comi-03

Andy Bierman

Michel Veillette

Peter van der Stok

Alexander Pelov <a@ackl.io>

Draft status

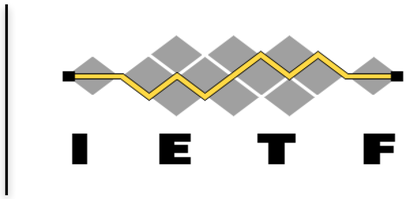

Actions from last time:
- Official Hackathon @ IETF 102

Draft	Version
ietf-core-yang-cbor	6
ietf-core-sid	4
ietf-core-comi	3
veillette-core-yang-library	3

Draft status

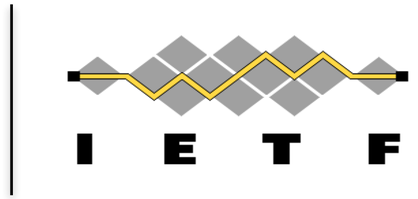

Actions from last time:
- Official Hackathon @ IETF 102

Draft	Version	
ietf-core-yang-cbor	6	**CORECONF** Like NETCONF & RESTCONF
ietf-core-sid	4	
ietf-core-comi	3	
veillette-core-yang-library	3	

The YANG protocol family

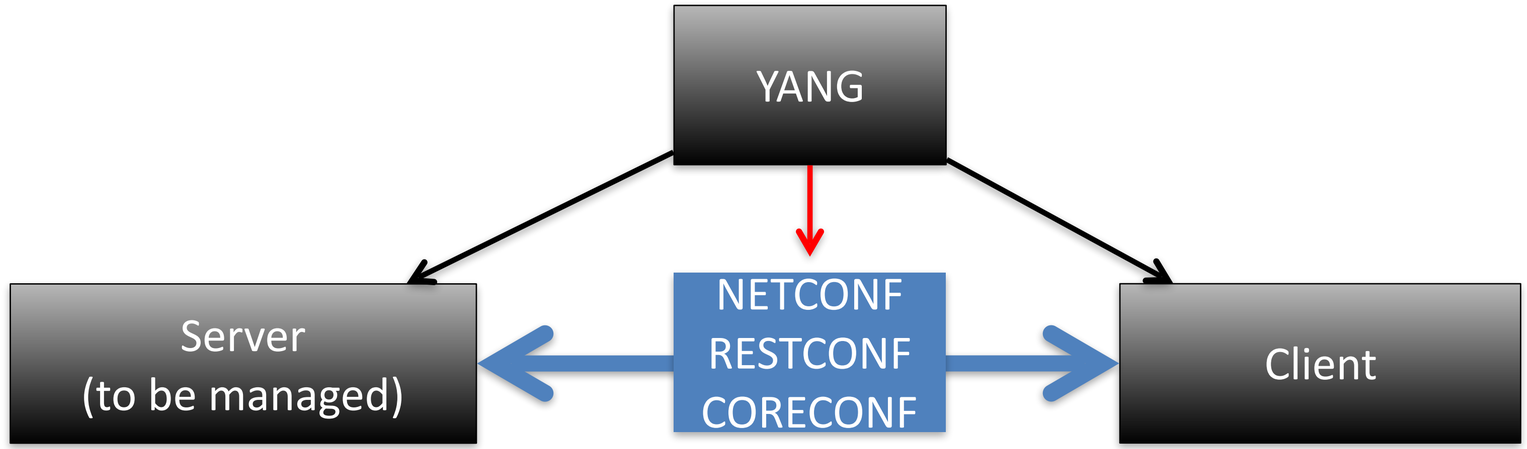

XML	JSON	CBOR
RPC	REST / RPC	REST / RPC
HTTP	HTTP	CoAP (CoMI)
TCP	TCP	UDP
IP	IP	IP
NETCONF	RESTCONF	CORECONF

What we have today

- Example SID Registry
  - http://comi.space

- Existing implementations
  - GoLang: server + client
  - C: server + client
  - 2 more partial proprietary implementations

- Interoperability
  - Virtual interop @ Hackathon IETF100 (FETCH with ietf-system) – existing implementations
  - Hackathon IETF101 – Semantic interoperability
  - Example implementation (client+server) accessible for everyone
    - F-Interop

Hackathon IETF 102

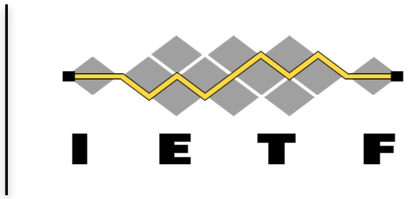

What we wanted to achieve

- **Open-source Python-based examples**
  - **Help people boot-strap implementations**
- Full open-source Python implementation
  - Client
- **Document our work**

What got done

- Developed base examples working on various OS (Lin/Mac)
- Clearly identified development process for CoMI
  - **Independent development of YANG-CBOR & CoAP**
  - Compatible with commercial / open-source NETCONF/RESTCONF servers
  - Identified next steps for a C implementation
- Started YDK-based CoMI client implementation

https://etherpad.tools.ietf.org/p/comi
https://github.com/Acklio/pycomi

YANG-CBOR + SID

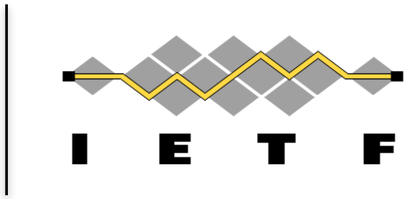

- Reviews
  - Juergen Schoenwaelder
  - Robert Wilton
- Minor changes / improvements suggestion
- One more significant
  - Always return top node, so that delta SIDs can be resolved unambiguously by only looking at the payload

Top node

1721

```
REQ: GET example.com/c/a5

RES: 2.05 Content (Content-Format: application/yang-value+cbor)
{
    +2 : "2014-10-26T12:16:51Z",   / current-datetime SID 1723 /
    +1 : "2014-10-21T03:00:00Z"    / boot-datetime SID 1722 /
}
```

Existing:

Pros:

- more compact

Cons:

- requires additional processing step
- may render debugging more difficult

```
REQ: GET example.com/c/a5

RES: 2.05 Content (Content-Format: application/yang-value+cbor)
{
  1721 : {
    +2 : "2014-10-26T12:16:51Z",   / current-datetime SID 1723 /
    +1 : "2014-10-21T03:00:00Z"    / boot-datetime SID 1722 /
  }
}
```

a5 in URI-safe Base64

Proposed:

Pros:

- Easier debugging
- Straightforward processing

Cons:

- 4-5 bytes more / response

Conclusion

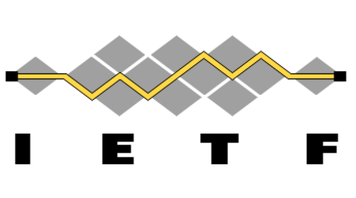

- YANG-CBOR + SID ready to ship after this IETF
  - Application in RESTCONF, CORECONF
  - Two reviews from NETMOD
  - WGLC

- Same for CoMI
  - One or two reviews from CORE are welcome
    - During WGLC?

- Action points IETF 103
  - Hackathon for open-source implementation
  - YANG of Things BOF

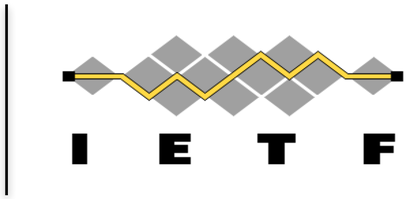

Thanks!

Concise YANG Telemetry

(on adding YANG Datastore Subscription & YANG Subscribed Notifications Capabilities to CoRECONF/CoMI)

@IETF 102 July'18

Henk Birkholz {henk.Birkholz@sit.fraunhofer.de}

&

Eric Voit {evoit@cisco.com}

Datastore Subscriptions & YANG (the thing formally called Push)

- Once Notifications were just about “Control Plane”…
- Now, they can have a variety of characteristics, have a “hard-coded” format… composing Events, Alarms or maybe even Incidents (currently exploring that scope) OR they can be about changes of Data Node Value of your favorite YANG Datastore
- Also, they now provide the capabilities to convey security-related information, diffusing in the Security Area domain (featuring levels of visibility and resilient subscriptions)
- I.e. there is an early draft to look at: https://datatracker.ietf.org/doc/draft-birkholz-yang-core-telemetry/

SID+keys really make things easier

- CoAP operations on a CoMI store that enable have the potential of actually being lightweight, resilient and intuitive
- E.g. a subscription on a datastore using a subtree expression could be realized simply using a Get+Observe on a SID in /c that is representing an intermediary node of a module
- YANG RPC can be used via POST/iPATCH. The response including a new key (subscription-id) that will also be populating stream resource /s as a sub-resource
- There is chance (currently exploring this option) to create a concise filter expression that is not a… naive transformation of an XPATH expression

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP -- moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

Draft-ietf-core-dev-urn-02

Arkko, Jennings & Shelby

A Uniform Resource Name (URN) namespace for hardware device identifiers.

Potentially useful in applications such as in sensor data streams and storage, or equipment inventories.

Complements other similar identifiers NIs (RFC 6920), UUIDs (RFC 4122), IMEIs (RFC 7254) etc. Supports, e.g., MAC and EUI-64, identifiers.

urn:dev:mac:0024befffe804ff1

Version -02

- For aligning the usage across the world:
    - **Folded in the “urn:dev:os:” and “urn:dev:ops:” sub-branches from OMA** LwM2M specifications
    - Three levels of “private” device identifiers
- Other changes made as a consequence of the above:
    - **Changed the “org:” sub-branch** to use “-“, not “:” to separate the PEN and the rest of the identifier (to align with the above)
    - A few other **syntax changes**, including allowing %-encoding

The Private Device Identifier Spaces

- Three levels of “private” device identifiers
- My organisation (org:), my serial number (os:), my product and serial number (ops:)

  urn:dev:org:32473-blaablaa

  urn:dev:ops:32473-Refrigerator-12345

  urn:dev:ops:32473-Refrigerator-12345

Questions

- The **unification** with suggested OMA types seems necessary — **do we agree**?
- However, OMA used OUIs, not PEN numbers
    - Easy if you already have an OUI, but otherwise acquiring one is costly, **change to PEN**?
- The OMA and IETF draft syntax style for os/ops/org was different, which leads to another desired change
- Do we have **usage of org/os/ops that would be affected**?

Constrained RESTful Environments WG (core)

Chairs:

Jaime Jiménez <jaime.jimenez@ericsson.com>

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 EDT (UTC–04:00)

Thursday (60 min)

- **18:10–18:15 Intro, Agenda**
- **18:15–18:20 DOTS heads-up (DOTS chairs)**
- **18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)**
- **18:34–18:46 Housekeeping cluster (AK, CB)**
- **18:46–18:58 Other WG drafts (MK) /candidates (BS)**
- **18:58–19:10 FASOR: Alternative Congestion Control**

All times are in time-warped EDT (UTC–04:00)

Monday (120 min)

- **15:50–16:00 Intro, Agenda, Status**
- **16:00–16:15 Up for WGLC soon: CoRECONF (AP -- moved)**
- **16:15–16:55 Post-WGLC: OSCORE (GS)**
- **16:55–17:35 Near-WGLC: RD/DNS-SD (PV, KL)**
- **17:35–17:50 Approved: SenML + related (JA, CB, AK)**

FETCH & PATCH with SenML

IETF 102, Montréal, CA

draft-keranen-senml-fetch-01

Ari Keränen & Mojan Mohajer

Updates since -00

- Re-using the base SenML media types (no need to register new ones)
- Wild-card feature left for future documents
- Focus on iPATCH instead of PATCH
- Security considerations: single FETCH/(i)PATCH can impact multiple resources; should be careful with access control
- Appending and deleting with iPATCH (next slide)

Add/Append/Replace/Delete with (i)PATCH

- Add: when no existing record with matching name the Patch record is added
  - Need to clarify that time is not mandatory
- Append: name matches but different time
- Replace: name (and time if in the target and patch records) matches
- Delete: match like above but with value set to null
  - Base SenML does not have null values so this should work
- Considerations
  - No need for op-codes. If later need, we can define new media type.
  - Can't add a time to a Record without time with a single Patch operation

TBD

- Clarify PATCH operations
- Rename "FETCH/PATCH Record/Pack" to "Fetch/Patch Record/Pack" to differentiate from the PATCH/iPATCH methods
- Ready for WG adoption?

IANA registry maintenance for SenML

- The usual fare.
- Except:
  - Every new field name needs a change to the XML schema
  - This then needs a new name for reference from EXI (“a” now)
- Who does this work?
- Most registrants are not interested in EXI
  - Example: LWM2M registration of “vlo”
- What the draft says: accumulate changes
  - The next new registrant that cares about EXI does all the changes so far
  - Weirdness: the schema in effect at any time could be in an obscure document…

All times are in time-warped EDT (UTC–04:00)

Thursday (60 min)

- 18:10–18:15 Intro, Agenda
- 18:15–18:20 DOTS heads-up (DOTS chairs)
- 18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)
- 18:34–18:46 Housekeeping cluster (AK, CB)
- 18:46–18:58 Other WG drafts (MK) /candidates (BS)
- 18:58–19:10 FASOR: Alternative Congestion Control

Introducing DDoS Open Threat Signaling WG (DOTS)

Thursday July 19, 2018

IETF 102, Montreal

Roman Danyliw (DOTS co-chair)

DOTS Architecture

(simplified)

[DOTS-REQUIREMENTS]
[DOTS-ARCHITECTURE]

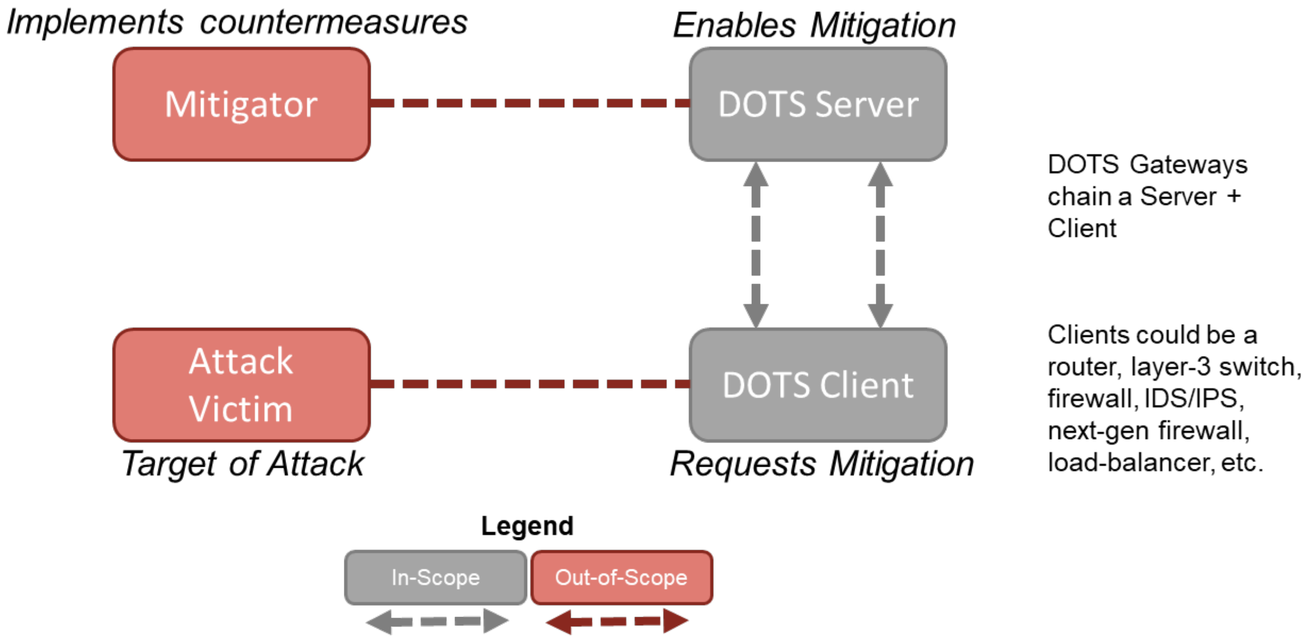

DOTS Protocols

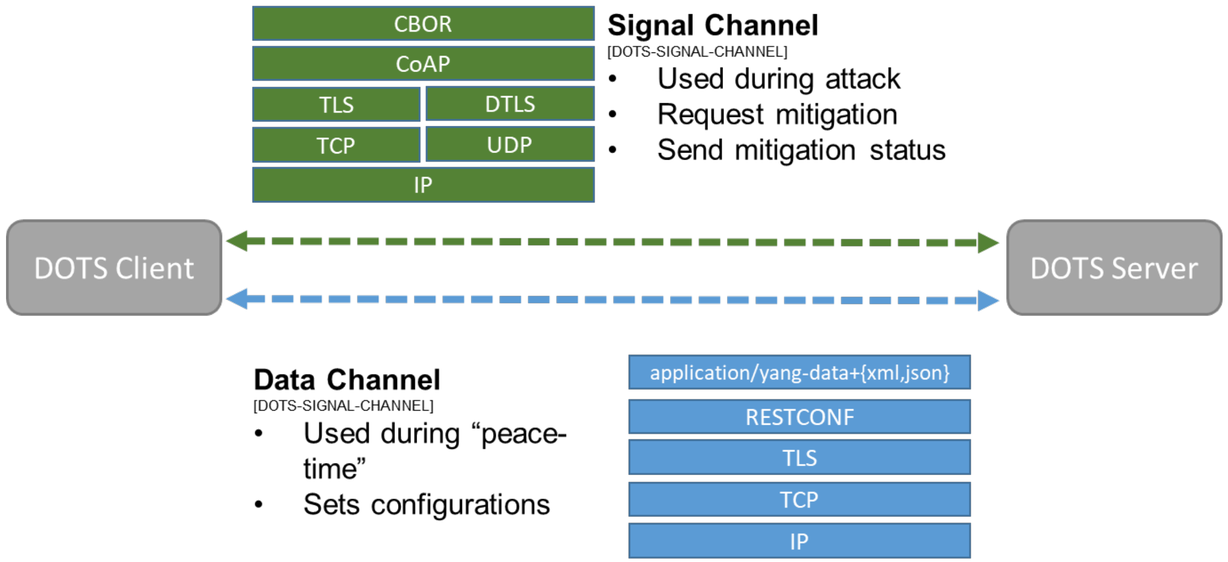

Properties of the Signal Channel [dots-signal-channel]

- (Section 3) CoAP chosen because of (a) expectation of packet loss, (b) support for non-confirmable messaging and (c) Small message overhead
- CoAP session established in peace-time
- (Section 3) DOES NOT use default 5684 port to allow for differentiated behavior in environments where both DOTS gateway and an IoT gated are present (per RFC7452)
- (Section 3) Uses “coaps” or “coaps+tcp” URI scheme
- (Section 3) To avoid fragmentation, follows Section 4.6 of RFC7252
- (Section 4.2) DOTS servers uses “/.well-known/dots”
- (Section 4.3) Uses Happy Eyeballs per RFC8305
- (Section 4.4) For mitigation requests during attack uses PUT, GET and DELETE methods; non-confirmable
- (Section 4.5) DOTS client can negotiate, configure and retrieve session configurations (e.g., heartbeat-interval; # of mission heartbeats, maximum retransmission, transmission timeout value, etc.)
- (Section 4.7) Heartbeat mechanism to distinguish between idle, disconnected and defunct

References

[DOTS-ARCHITECTURE] Distributed-Denial-of-Service Open Threat Signaling (DOTS) Architecture. https://datatracker.ietf.org/doc/draft-ietf-dots-architecture/

[DOTS-DATA-CHANNEL] Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification. https://datatracker.ietf.org/doc/draft-ietf-dots-data-channel/

[DOTS-REQUIREMENTS] Distributed Denial of Service (DDoS) Open Threat Signaling Requirements. https://datatracker.ietf.org/doc/draft-ietf-dots-requirements/

[DOTS-SIGNAL-CHANNEL] Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification. https://datatracker.ietf.org/doc/draft-ietf-dots-signal-channel/

[DOTS-USE-CASES] Use cases for DDoS Open Threat Signaling https://datatracker.ietf.org/doc/draft-ietf-dots-use-cases/

All times are in time-warped EDT (UTC–04:00)

Thursday (60 min)

- 18:10–18:15 Intro, Agenda
- 18:15–18:20 DOTS heads-up (DOTS chairs)
- 18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)
- 18:34–18:46 Housekeeping cluster (AK, CB)
- 18:46–18:58 Other WG drafts (MK) /candidates (BS)
- 18:58–19:10 FASOR: Alternative Congestion Control

Stateless Forward Proxies

Mališa Vučinić

Thomas Watteyne

Carsten Bormann

Göran Selander

Klaus Hartke

CoRE WG

IETF 102/Montreal

Stateful Forward Proxy

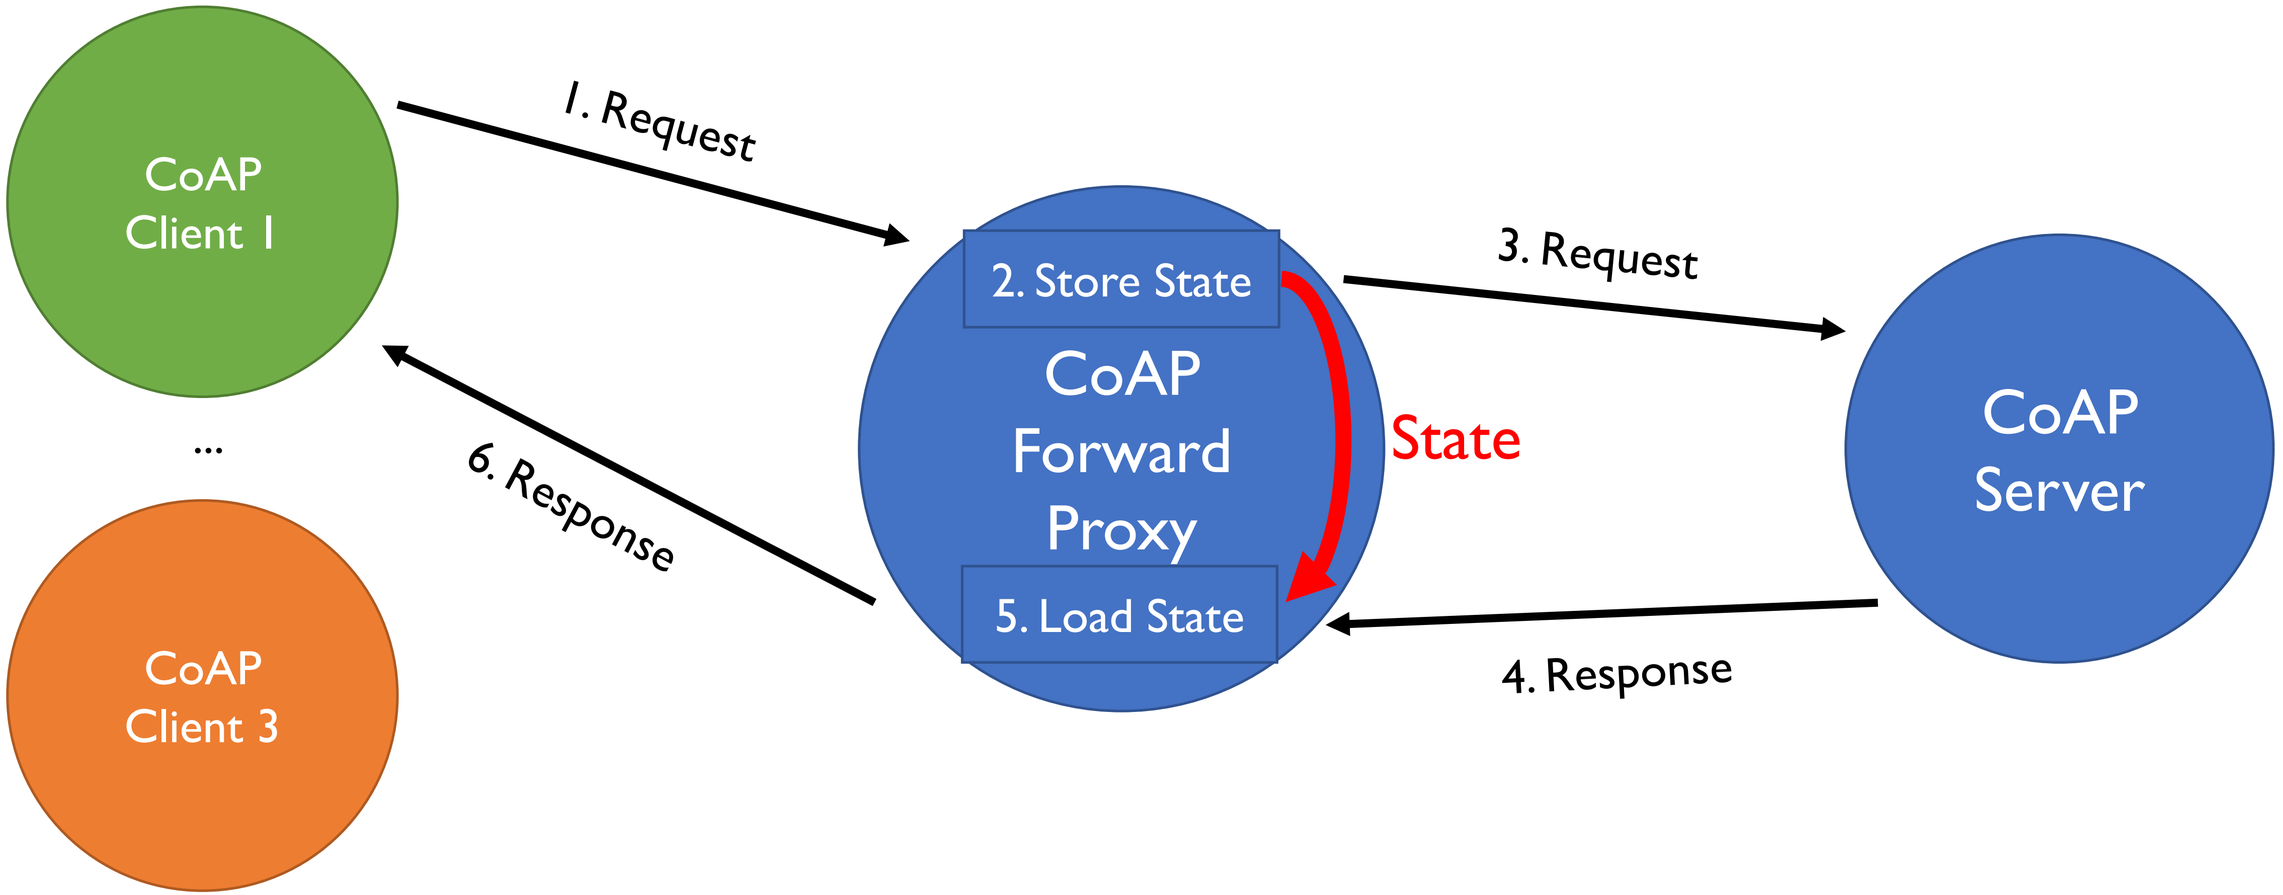

Stateless Forward Proxy

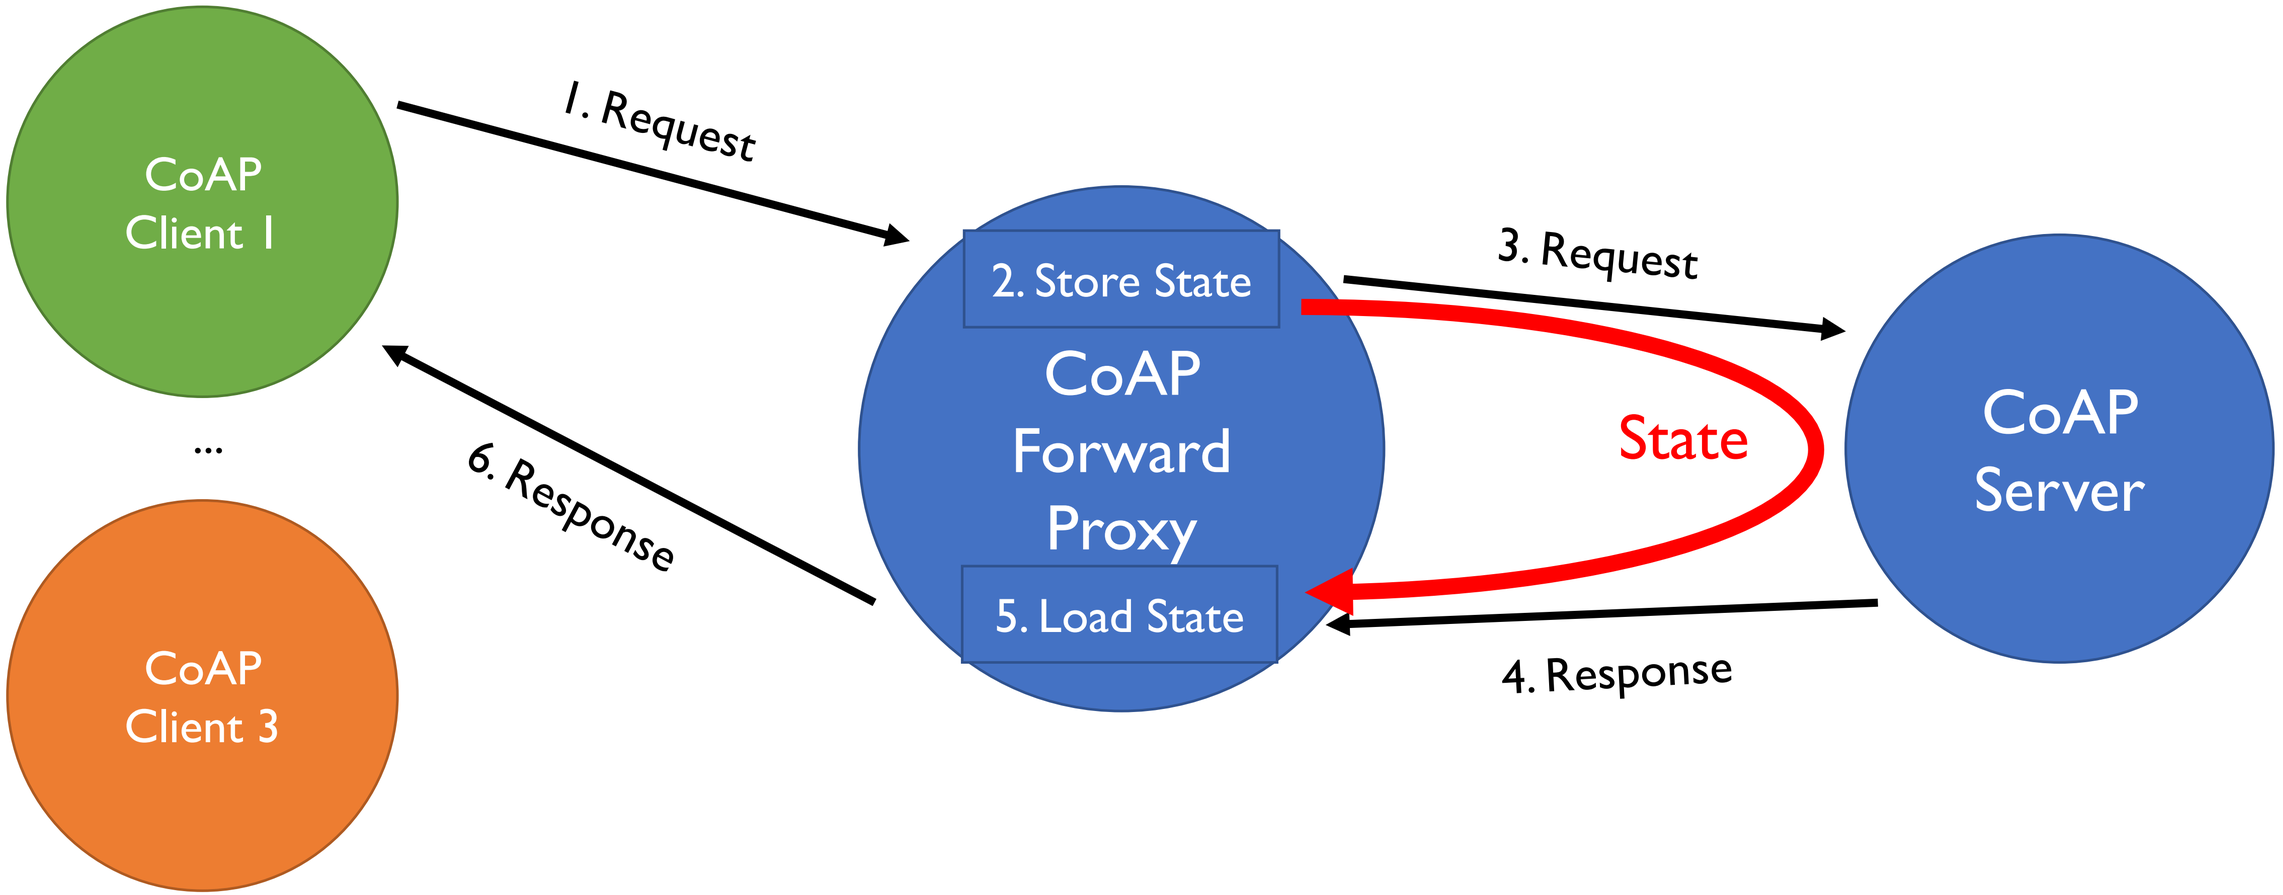

First attempt: Stateless-Proxy Option

Critical & Safe to Forward

Problems:

- There are two tokens in a message
- The option is not always echoed back (Intermediaries and servers that do not implement the option will not include it in responses they generate)
- The option cannot be not part of the cache key (Clients would receive the option value from another client) and cannot be part of the cache key (This would break caching since cache keys would never be the same)

These problems cannot be solved without requiring the next hop to have support for the option when it's used

Second attempt: Second-Token Option

Critical & **Not** Safe to Forward

Problems:

- There are still two tokens in a message
- It's ugly

Second attempt: Second-Token Option

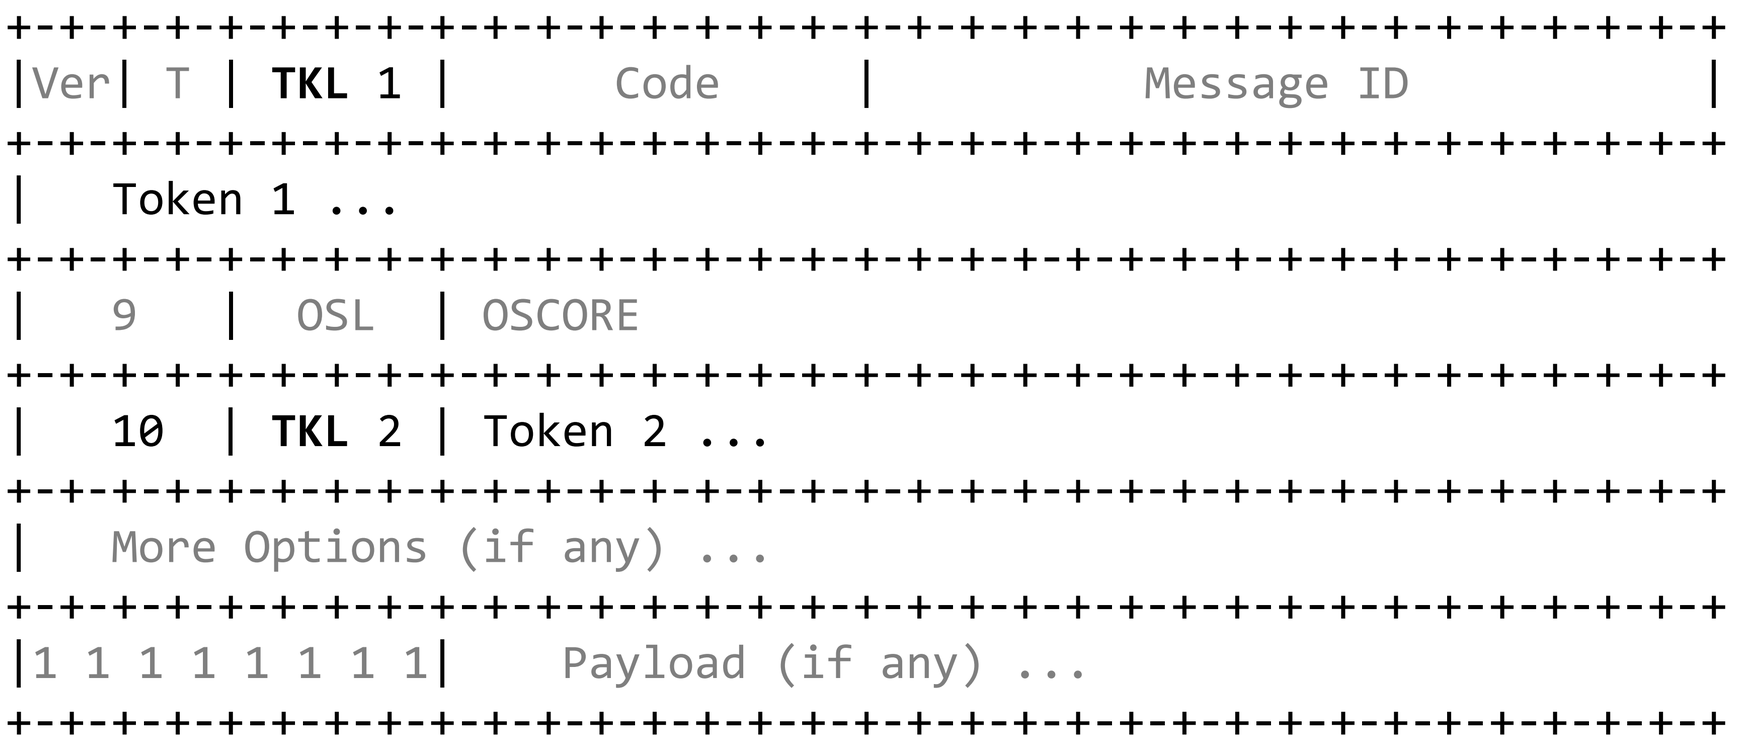

Proposal: Extending the Token Length

- Should have been done in 2013 already

Proposal: Extending the Token Length

0 to ~~8~~ 12 bytes

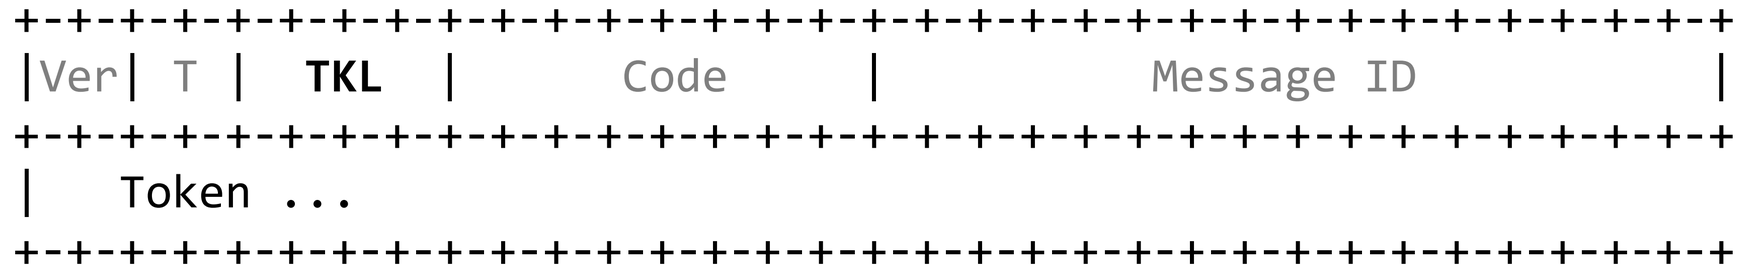

13 to 268 bytes

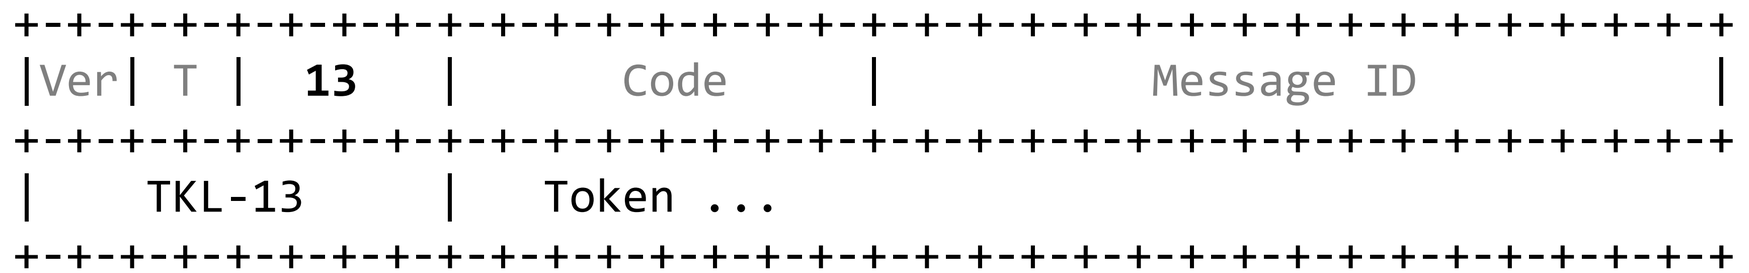

269 to 65804 bytes

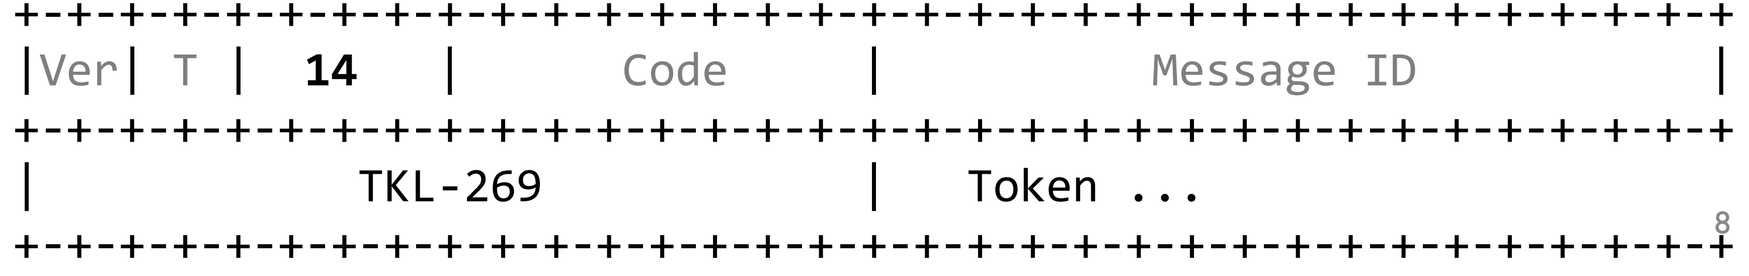

Proposal: Extending the Token Length

- Caveat:
  A client (or intermediary in the role of a client) needs to perform an stateful request with extended token length to the next hop first to discover support before it can be stateless

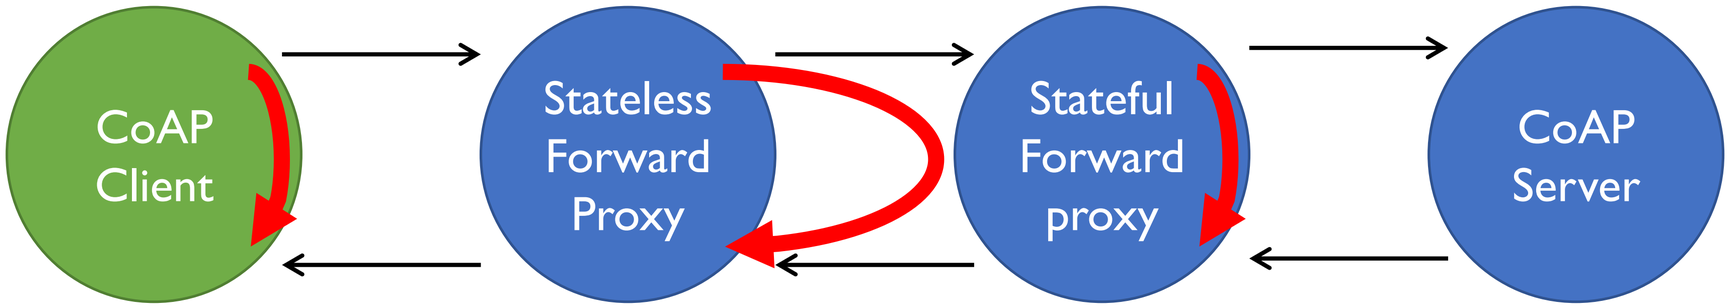

Next Steps

- Write a draft
- Adopt it as a WG document
- Have a WGLC around IETF 103/Bangkok

Thursday (60 min)

- 18:10–18:15 Intro, Agenda
- 18:15–18:20 DOTS heads-up (DOTS chairs)
- 18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)
- 18:34–18:46 Housekeeping cluster (AK, CB)
- 18:46–18:58 Other WG drafts (MK) /candidates (BS)
- 18:58–19:10 FASOR: Alternative Congestion Control

Too Many Requests Response Code for CoAP

IETF 102, Montréal, CA

draft-ietf-core-too-many-reqs-02

Ari Keränen <ari.keranen@ericsson.com>

Background

- CoAP client can cause overload in server with too frequent requests
- How can server tell client to back off
- HTTP error code 429 “Too many requests”
- Solution: register 4.29 for CoAP
  - With MaxAge to indicate when it’s OK to request again
- Originally part of CoAP Pub/sub Broker draft; also OCF interest

Changes since IETF 101

- Added a hint that action payloads can be used by the server to guide clients about next actions
- Instead of only "same request" also "similar requests" can be suppressed with too-many-requests response code
  - "Client SHOULD NOT repeat similar request until Max-Age times out"

Same vs. Similar request

- Input from Abhijan B: extends use to e.g., stream transfer pattern use cases (see T2TRG STP draft)
- "same request": same method and target resource
- "similar request": same method and **related** target resource
  - E.g., resources are part of same collection
- Up to application what is "similar enough"
  - Could be part of application specification
  - Future documents may define action payloads to guide client on this

draft-ietf-core-multipart-ct

- Continuation of draft-fossati-multipart-ct of 2012 vintage:
  - Join request/response bodies into a single combined one
  - keep information about the constituent content-formats
- 2018: Ported to the CBOR age
- ```
  multipart-core = [* multipart-part]
  ```
- ```
  multipart-part = (type: uint .size 2, part: bytes / null)
  ```
- Use case: Needed by EST-over-coaps

- Are we done?

draft-bormann-core-proactive-ct

- There is a threshold for using CoAP in place of HTTP:
  - Get the content-format numbers for the media types needed
- There are < 2000 media types, > 65000 content format numbers
- Why don’t we just register them **proactively**?
  - Deliberately wasting some hundreds of code points, just in case.
- Draft contains proposed procedure, and discussion of limitations
- Where it doesn’t work, no change from today.
- Where it works, can use CoAP out of the box with existing media types

- Do we want to do this? (If yes, is the draft ready for adoption?)

Thursday (60 min)

- 18:10–18:15 Intro, Agenda
- 18:15–18:20 DOTS heads-up (DOTS chairs)
- 18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)
- 18:34–18:46 Housekeeping cluster (AK, CB)
- 18:46–18:58 Other WG drafts (MK) /candidates (BS)
- 18:58–19:10 FASOR: Alternative Congestion Control

draft-ietf-core-interfaces-12
draft-ietf-core-dynlink-06

Bill Silverajan, Julian Zhu, Michael Koster

Status of core-interfaces

- Used in the OCF Specification
- Editorial changes made for clarity
  - First section discusses resource collections
  - Second section discusses interface descriptions
- Content types for interfaces rectified
- Draft updated with the new SenML format
- Draft is ready for WGLC

Status of core-dynlink

- OCF using Dynlink in many use cases, e.g. Rules, Events, Push model, Direct Device-to-Cloud
- OMA LWM2M uses Observation Attributes
- All remaining issues in github are being closed
- Major cleanup performed, but 1 more revision necessary to organise the document better
- More examples needed, particularly for band and observation attribute interactions
- Responding to current reviews on core-ml
- Draft will be ready for WGLC once these are done (in a few weeks)

CoAP Pub/Sub

IETF 102

Status and Recent Changes

- Addressed all outstanding comments from Jim S. and the mailing list
  - Sorry for the delay…
- More cleanup and clarification
- Went through the issues list and closed or deferred all but 2 issues

Remaining Issues

- 2 issues from Github
- What happens when a client tries to publish to a topic that exists
  - Currently specify 4.03 Unacceptable
  - 4.09 Conflict is proposed, based on HTTP 409 semantics
- How should the broker respond when the data are stale
  - Currently specify HTTP 204 semantics
  - Propose a new code 2.07 with HTTP 202 semantics

Roadmap

- Final edit pass and resolution of last 2 issues
- One more WG review?
- WGLC candidate in a couple of weeks or after WG review

CoAP Protocol Negotiation

draft-silverajan-core-coap-protocol-negotiation-09

Bill Silverajan TUT
Mert Ocak Ericsson

Changes between -08 and -09

- Based on Jim Schaad's review
  - Clarified usage of the 'tt' lookup parameter
  - 'tt' parameter value given as URI scheme (eg coaps+tcp) instead of CoAP transport (eg tcp)
  - Alternative-Transport option usage updated to correspond to changes for 'tt'
- Updated Resource Directory examples
  - Because "con" has been supplanted by "base"

Thursday (60 min)

- 18:10–18:15 Intro, Agenda
- 18:15–18:20 DOTS heads-up (DOTS chairs)
- 18:20–18:34 Stateless-Proxy option (6TiSCH -- moved)
- 18:34–18:46 Housekeeping cluster (AK, CB)
- 18:46–18:58 Other WG drafts (MK) /candidates (BS)
- 18:58–19:10 FASOR: Alternative Congestion Control

FASOR Retransmission Timeout and Congestion Control Mechanism

draft-jarvinen-core-fasor

Ilpo Järvinen[*], Iivo Raitahila[*], Zhen Cao[†] and Markku Kojo[*]

[*]University of Helsinki [†]Huawei

core @ IETF-102
July 19, 2018

Introduction and Objectives

- FASOR (Fast-Slow RTO) balances between the contradictory goals in handling random loss and congestion
  - Triggers RTO fast in case of random losses
  - Triggers RTO slow enough to handle congestion
- In IoT deployments, congestion expected to occur mainly due to large number of parallel devices
  - Test such extreme congestion scenarios now rather than later
- Unlike default CoAP and CoCoA, FASOR is not vulnerable to Congestion collapse
  - But still outperforms them in cases with random losses

Problem with Current CoAP RTO Management

- Karn's algorithm: exponential backoff and keep the backed off RTO until unambiguous RTT sample acquired
- CoAP CC algorithms: exponential backoff but DO NOT retain the backed off RTO
- Default CoAP and CoCoA prone to Congestion collapse*
  - Unnecessary retransmissions occur persistently if RTT $>$ RTO with the default congestion control algorithm
  - CoCoA not safe either but more complicated
    - Weak estimator hacks around the lack of retaining the backed off RTO (but RTO only updated if $<$3 rexmits were made)
    - Inflated RTT that triggers 3+ rexmits still causes the collapse
- Lack of retaining RTO good for random losses though

*

I. Järvinen, I. Raitahila, L. Pesola, Z. Cao, and M. Kojo, "Experimental Results with Default CoAP, CoCoA and CoAP over TCP RTO Management & Congestion Control," in *Proceedings of IETF101 / core WG*, Mar. 2018

I. Järvinen, I. Raitahila, Z. Cao, and M. Kojo, "Is CoAP Congestion Safe?," in *Proceedings of the Applied Networking Research Workshop 2018 (ANRW'18)*, July 2018

FASOR (Fast-Slow RTO) in Nutshell

- FASOR (Fast-Slow RTO)* tries to find a good middle ground
  - Try to improve random loss
  - . . . but still handles congestion safely, including unnecessary rexmits
- Two ways to calculate RTO
  - FastRTO (normal RTO)
  - New SlowRTO
- New back off logic

---

* I. Järvinen, M. Kojo, I. Raitahila, and Z. Cao, "Fast-Slow Retransmission and Congestion Control Algorithm for CoAP," Internet Draft, June 2018. Work in progress

I. Järvinen, I. Raitahila, Z. Cao, and M. Kojo, "FASOR Retransmission Timeout and Congestion Control Mechanism for CoAP," in *Proceedings of IEEE Globecom 2018*, Dec. 2018. To appear

FastRTO and SlowRTO

- FastRTO $\approx$ RFC 6298 RTT/RTO computation
  - Initialization of RTTVAR changed to R/2K
    - Lowers RTO for short exchanges
- SlowRTO analogous to Karn's algorithm keeping RTO until unambiguous RTT sample
  - Measured when retransmissions were made as the time elapsed from the original copy
  - Multiplied by a factor to allow load growth (1.5 by default)
  - More conservative than Karn's algorithm

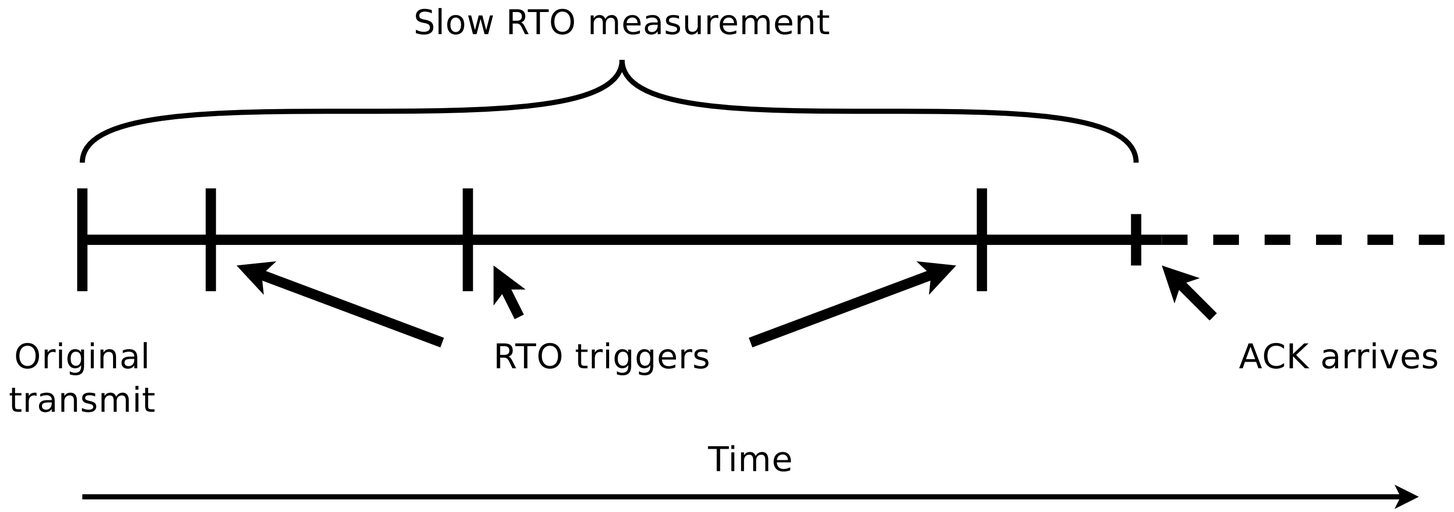

FASOR Back Off Logic

- Modify 2-state RTO logic of Karn's algorithm by adding a new state and modify back off series:

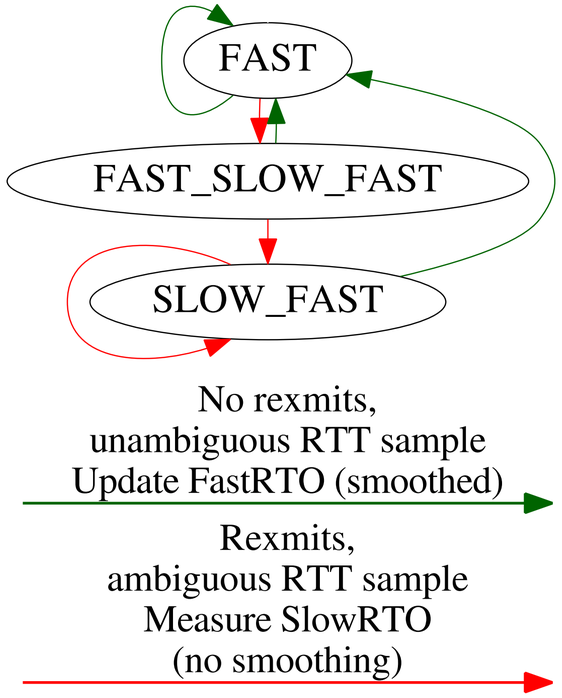

Back Off Series
FastRTO, FastRTO*2^1, FastRTO*2^2, ...
FastRTO, max(SlowRTO, FastRTO*2), FastRTO*2^1, FastRTO*2^2, ...
SlowRTO, FastRTO, FastRTO*2^1, FastRTO*2^2, ...

FASOR States

- FAST
  - "Normal" RTO series with exponential back off
  - When network state is not dubious
- FAST_SLOW_FAST
  - Probe first with FastRTO
    - Helps random loss cases to retransmit quickly
  - If no response and RTO expires, use SlowRTO as conservative back off
    - Allow draining unnecessary retransmissions from network
    - Due to lack of response so far, the sender cannot know if unnecessary retransmissions occurred or not
    - Safe and conservative option taken
  - If still more RTOs trigger, continue with the Fast RTO based exponential back off
- SLOW_FAST
  - Start with SlowRTO to acquire an unambiguous RTT sample with high probability

Optional Features

- Token/option variant
  - Encodes ordinal number of the transmissions for the request message to either token or option
  - Receiver echos the ordinal number back unchanged
  - Removes retransmission ambiguity problem
  - Allows accurate RTT estimation also with retransmitted messages

Test Setup

- Bottleneck BW: 30 kbps
- Base RTT $\approx$ 660 msecs
- Workload
  - A flow: a series of short-lived clients perform 50 request-responses exchanges in total
  - CC state reset after 1 to 10 message exchanges (new short-lived client starts)
  - Response payload: 60 bytes
- Test scenarios
  - Heavy congestion and bufferbloat
    - Up to 400 parallel flows
    - Varying buffer size, including infinite buffer (1410000 bytes)
    - RTT $\approx$ 10 secs (for 400 clients + infinite buffer)
    - Error-free link
  - Random losses
    - 10 parallel flows
    - No congestion
    - 2-state error model: 0%/50% (medium) or 2%/80% (high) packet error rate

Results with Heavy Congestion and Bufferbloat

FCT

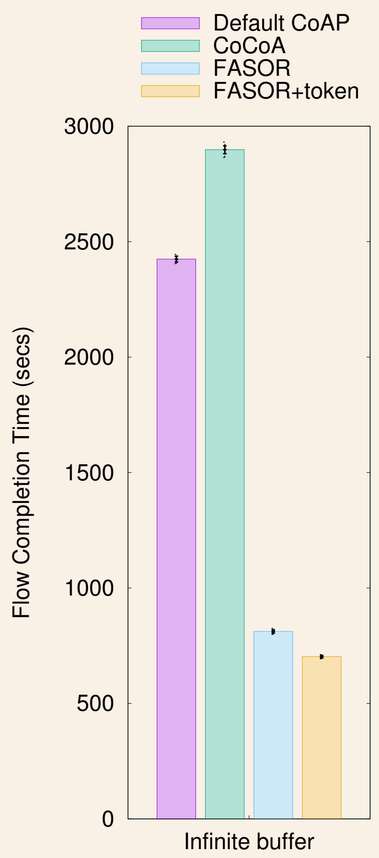

Unnec. Rexmits

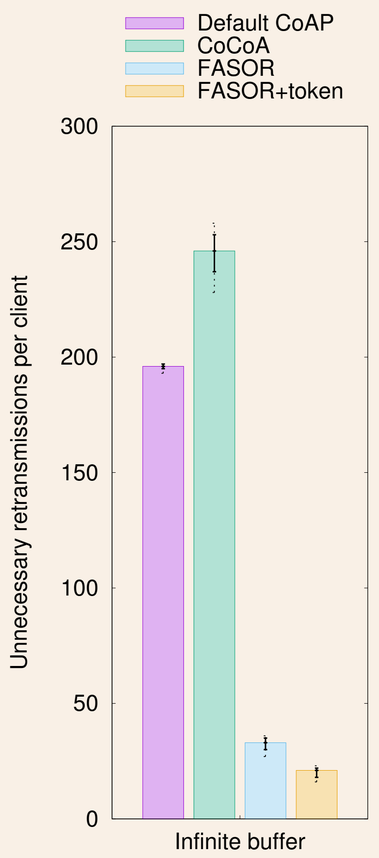

Observations

- FCT for Default CoAP and CoCoA long due to unnecessary rexmits
- Reduction in median with FASOR
  - FCT: 67%-76%
  - Unnecessary rexmits: 83%-91%
- Some unnecessary rexmits unavoidable when new client starts
- Similar pattern visible also in RTT

Results with Random Loss

FCT

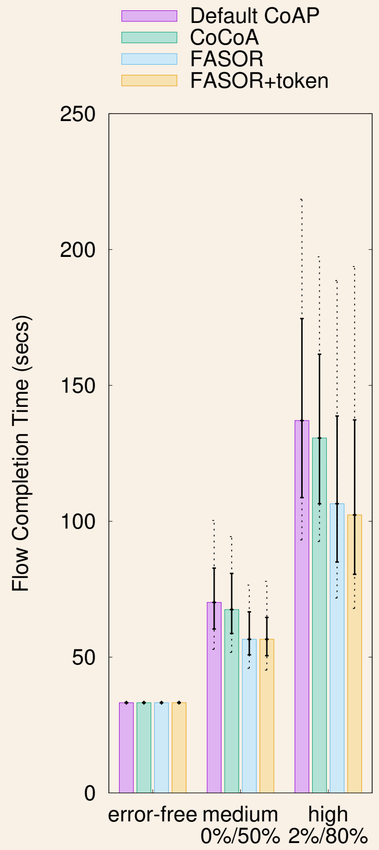

Expired RTOs

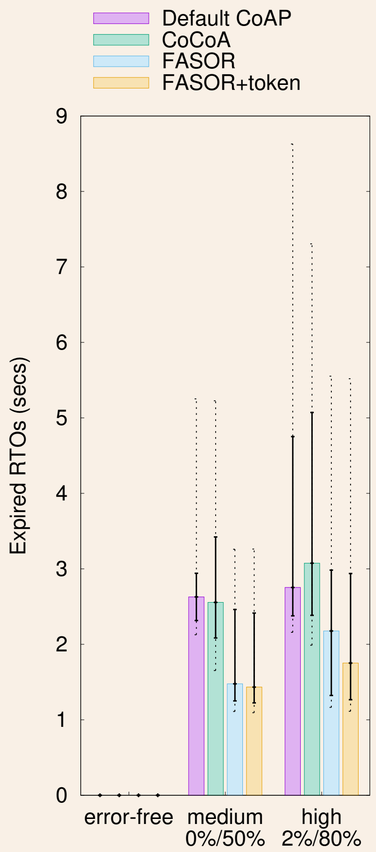

Observations

- Median of the FCT shorter with FASOR:
  - medium: 16%-19%
  - high: 19%-25%
- FASOR is able to lower RTO value despite the challenging short-lived clients
- CoCoA's weak estimator measures random loss noise on ambiguous RTT samples
  - Its RTO values increase instead of converging towards the real RTT ($\approx$ 660 msecs)

Work Items under Consideration

- FAST_SLOW_FAST back off series may currently be more aggressive than that of FAST state
  - A more conservative version has small but measurable performance impact
- Test with a dithering algorithm that is more similar to the standard dithering algorithm
  - Currently the specification matches with our current implementation
  - Dithering mostly orthogonal to the other parts of FASOR algorithm

Concluding Remarks

- FASOR achieves good balance between handling random losses efficiently and responding to congestion adequately in contrast to the other CC proposals
- Despite handling congestion safely, FASOR outperforms both default CoAP and CoCoA in cases with random losses
  - Making default CoAP and CoCoA congestion safe will have significant negative impact on their performance
  - Therefore, the performance gap is likely to become even larger
- Complexity of FASOR algorithm is comparable to that of CoCoA
- We believe FASOR would be beneficial for the ecosystem
  - Is there interest in this WG to work on this?

Backup Slides

Backup Slides

- "Continuous" workload: 50 request-replies; does not reset CC state after 1 to 10 exchanges
- "Random" workload: 50 request-replies; CC state reset after 1 to 10 exchanges
- "Fullbackoff" variants* are congestion safe versions of default CoAP and CoCoA adding retaining RTO similar to Karn's algorithm

---

*

I. Järvinen, I. Raitahila, Z. Cao, and M. Kojo, "Is CoAP Congestion Safe?," in *Proceedings of the Applied Networking Research Workshop 2018 (ANRW'18)*, July 2018

Backup Slides: Fullbackoff Variants

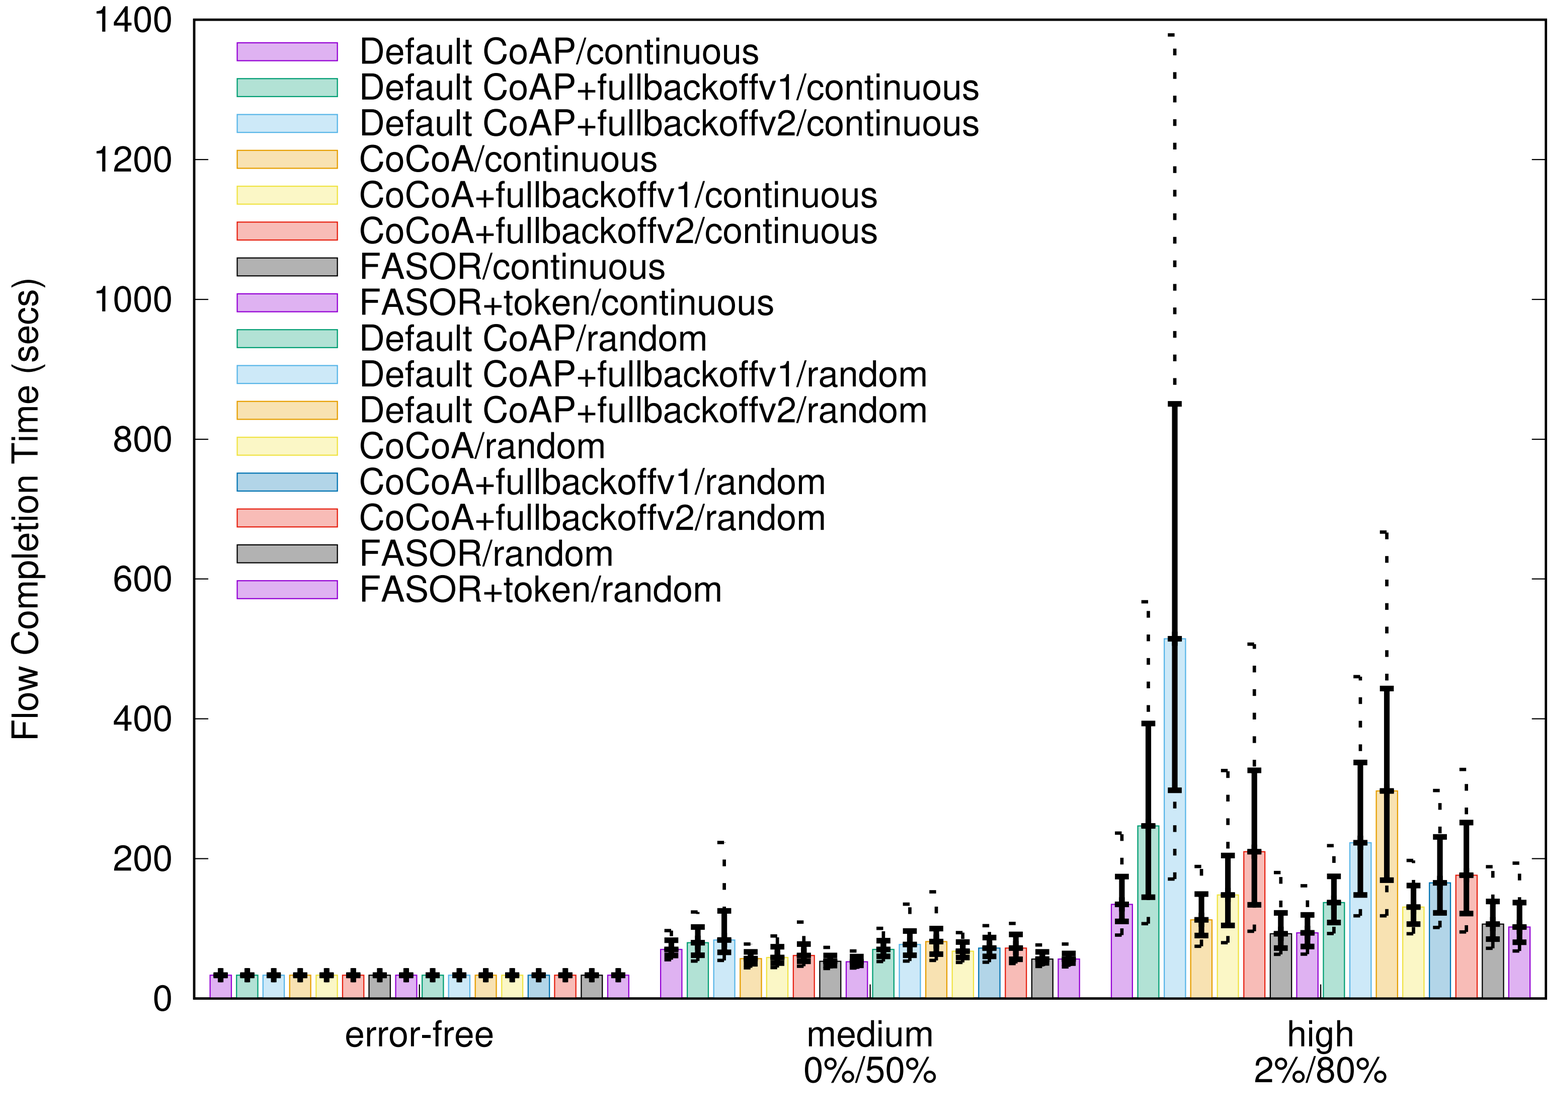

Backup Slides: 100 Parallel Flows

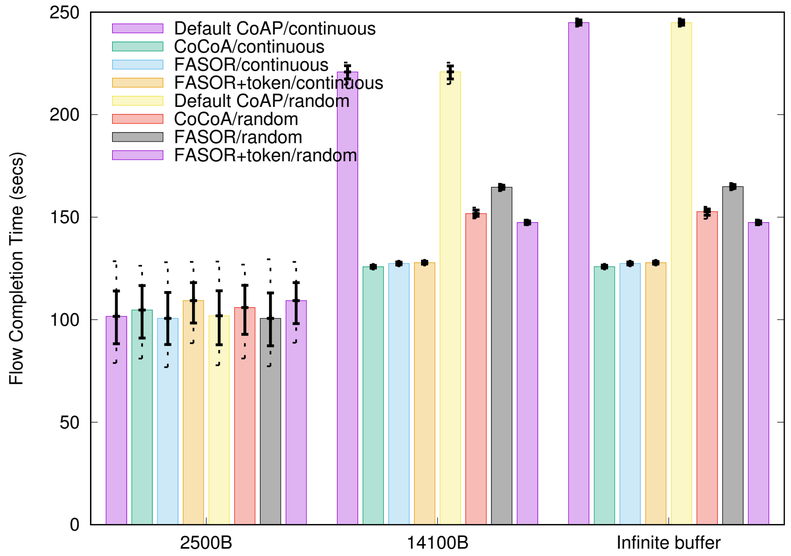

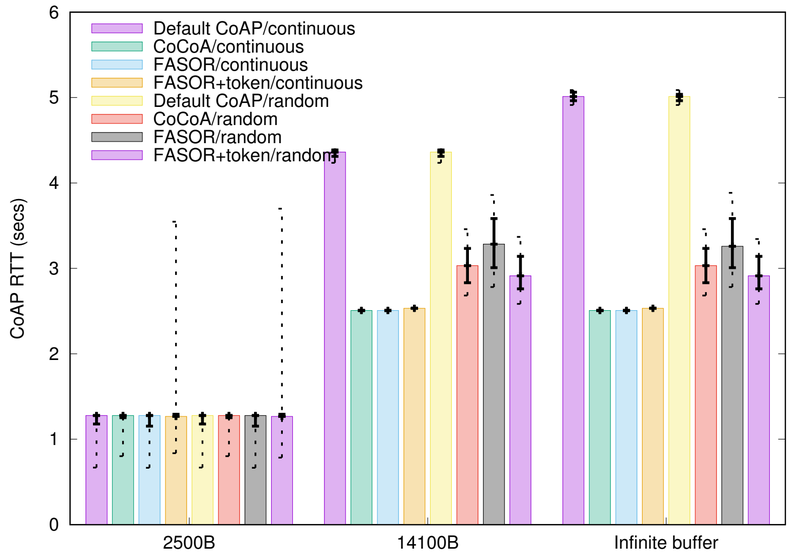

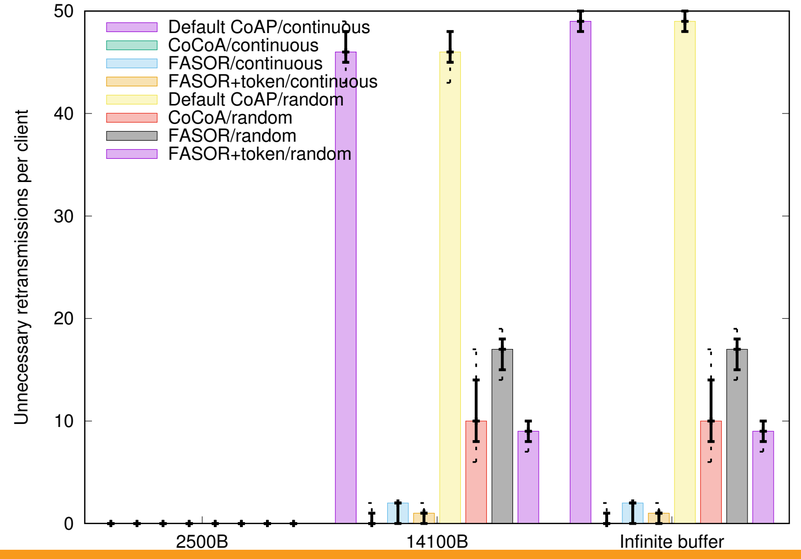

Backup Slides: 200 Parallel Flows

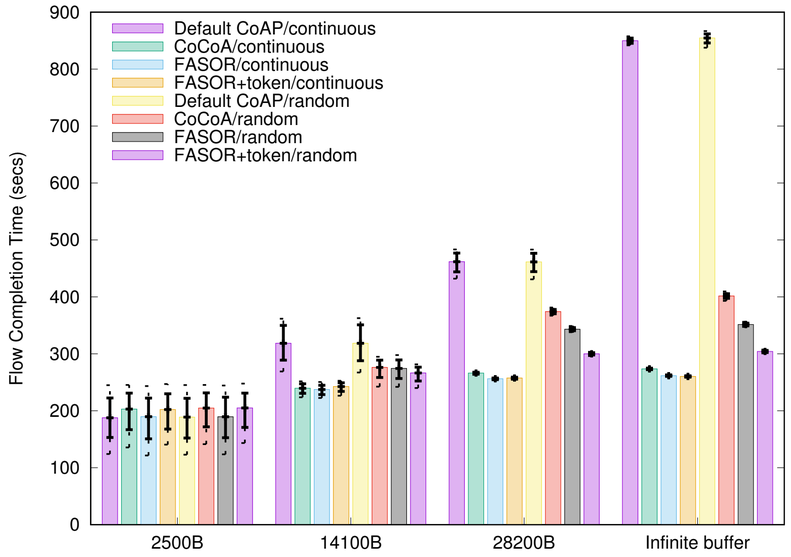

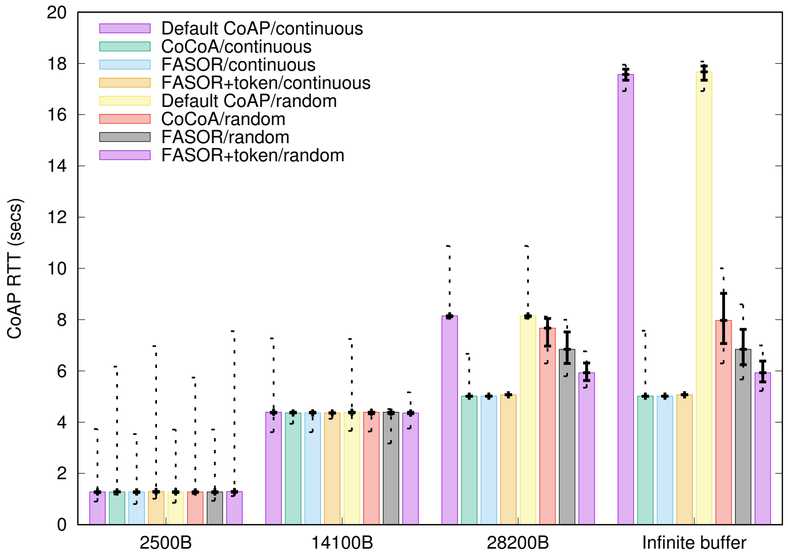

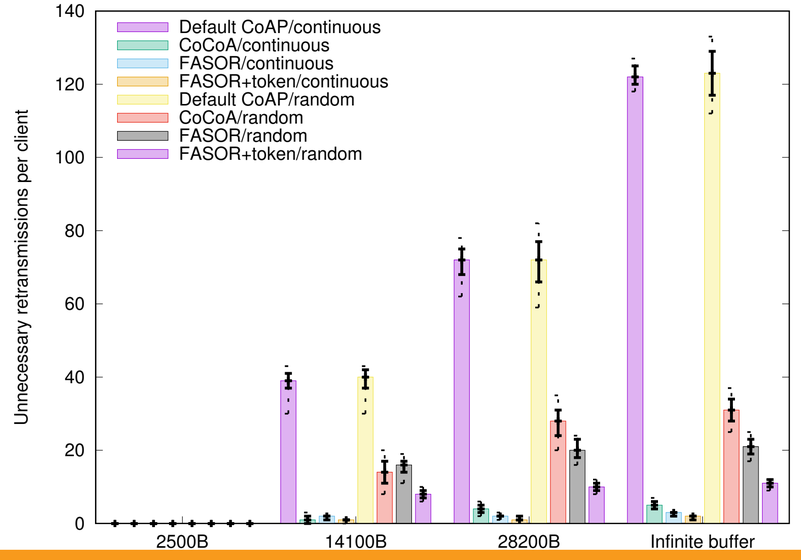

Backup Slides: 400 Parallel Flows