Constrained RESTful Environments WG (core)

Chairs:

Andrew McGregor <andrewmcgr@gmail.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 3979 and its updates**

- Blue sheets
- Scribe(s): http://tools.ietf.org/wg/core/minutes

http://www.ietf.org/about/note-well.html

Note Well

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

- The IETF plenary session
- The IESG, or any member thereof on behalf of the IESG
- Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices
- Any IETF working group or portion thereof
- Any Birds of a Feather (BOF) session
- The IAB or any member thereof on behalf of the IAB
- The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice. Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

Agenda Bashing

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

Milestones (from WG charter page)

http://datatracker.ietf.org/wg/core/charter/

Document submissions to IESG:

- **Done** CoAP protocol specification with mapping to HTTP Rest API to IESG
- **Oct 2013** Blockwise transfers in CoAP to IESG
- **Done** Observing Resources in CoAP to IESG
- **Done** Group Communication for CoAP to IESG
- **Jan 2014** BP for HTTP-CoAP Mapping Impl to IESG
- **Jan 2014** CoRE Link Collections in JSON to IESG
- **May 2014** CoRE Interfaces to IESG
- **Dec 2099** HOLD (date TBD) Constrained security bootstrapping specification to IESG

RFC 7390

~~draft-ietf-core-groupcomm-25~~

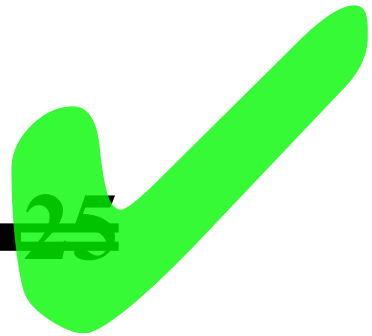

- **Was approved 2014-09-17**
  - **Reclassified as "EXPERIMENTAL" (was Informational) due to normative protocol content**
- **Published 2014-10-30**

- **More can be done on making multicast more useful**
  - **Security? (DICE, object security?)**
  - **Routing? (ROLL: MPL, others?)**
  - **Reliability? (ROLL: MPL, others?)**
  - **Integration into CoAP (e.g., multicast notifications?)**

Observe

- **draft-ietf-core-observe-14 (2014-07-20) submitted for Standards-Track (Proposed Standard)**

  Status:
  DISCUSSes are cleared (-16),
  working on text changes for COMMENTs

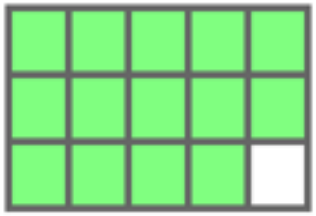

- **Some of the interesting COMMENTs may instead turn into text changes in draft-ietf-lwig-coap — we need to pay more attention to documenting implementation information!**

WG documents

- draft-ietf-core-block — 3rd WGLC any time now
  - should be fully cooked
- draft-ietf-core-http-mapping
  - WGLC very soon
- draft-ietf-core-links-json
  - still waiting for more implementation experience?
- draft-ietf-core-resource-directory
  - charter work needed, to resume activity!
- draft-ietf-core-interfaces
  - to resume activity!

Group 1: WG docs

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

Guidelines for HTTP-CoAP Mapping Implementations

Angelo Castellani, Salvatore Loreto, Akbar Rahman, Thomas Fossati, Esko Dijk

IETF-91, Nov 2014

http://tools.ietf.org/html/draft-ietf-core-http-mapping-05

Main Changes (from IETF-90 Toronto)

- Changes from ietf-04 to ietf-05:
  - Addressed Ticket #366 (Mapping of CoRE Link Format payloads to be valid in HTTP Domain?)
  - Addressed Ticket #375 (Add requirement on mapping of CoAP diagnostic payload)
  - Addressed comment from Yusuke (on warning of potential loss of meta data in certain transcoding schemes)
    - "… What I'd appreciate if the document can mention such specific case are out of the scope of the document and the implementor should be responsible and be aware of the side-effect (may not specific to EXI cases). In particular, schema-informed EXI will be mapped to application/exi without any schema information OR application/octet-stream. In-band metadata is lost in the case and the proxy should be aware of the metadata required to manage the metadata (schema, or any other application-specific data) …"
  - Various editorial improvements

Reverse Cross-Protocol Proxy Deployment Scenario

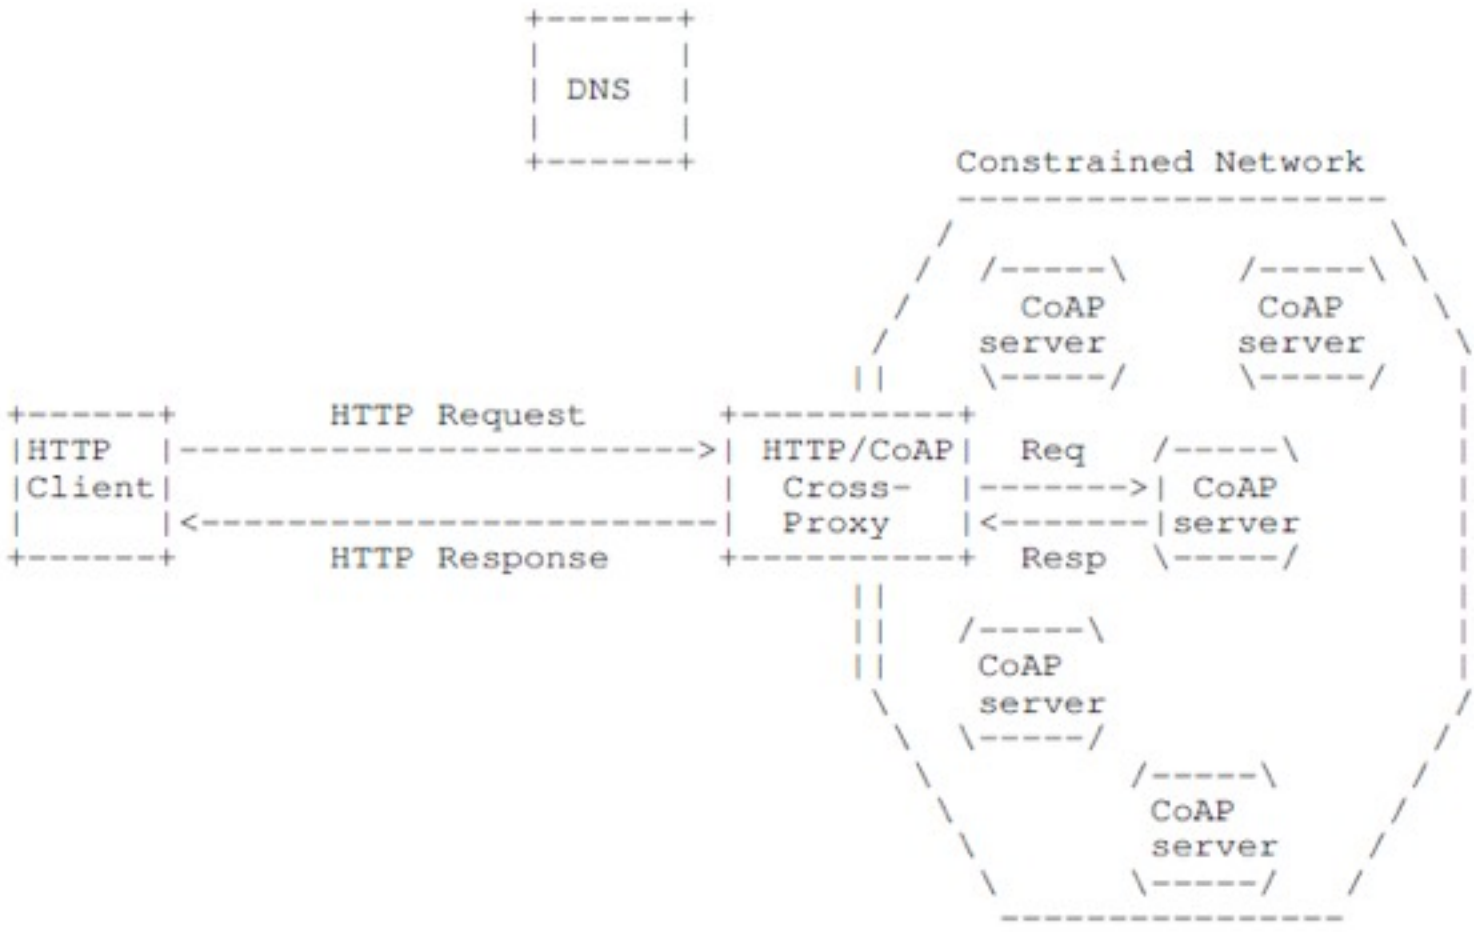

Reminder: Focus of I-D is reverse HTTP-CoAP (HC) Cross Proxy

Closed Ticket #366 (1/3)

- Ticket #366 (Mapping of CoRE Link Format payloads to be valid in HTTP Domain?)
  - Addressed in new Section 6.3.3.2 (Content Transcoding- CORE Link Format):

    - "The CoRE Link Format [RFC6690] is a set of links (i.e., URIs and their formal relationships) which is carried as content payload in a CoAP response. These links usually include CoAP URIs that might be translated by the HC proxy to the correspondent HTTP URIs using the implemented URI mapping function (see Section 5). Such a process would inspect the forwarded traffic and attempt to re-write the body of resources with an application/link-format media type, mapping the embedded CoAP URIs to their HTTP counterparts. Some potential issues with this approach are:

      1. Tampering with payloads is incompatible with resources that are integrity protected (although this is a problem with transcoding in general).
      2. The HC proxy needs to fully understand [RFC6690] syntax and semantics, otherwise there is a inherent risk to corrupt the payloads.

    - Therefore, CoRE Link Format payload should only be transcoded at the risk and discretion of the proxy implementer."

Ticket #366 resolution details (2/3)

- Choice made: we do not mandate that a CoRE Link Format payload is mapped to contain only HTTP-links. Default is no payload mapping.
- Reasons:
  1. Typically a HC Proxy does not know a-priori all the CoAP servers it needs to map. Hence including all servers' entries in /.well-known/core , which would be the cleanest solution, is most often not feasible.
     - *Also gives issues of keeping HC Proxy /.well-known/core 'fresh'*
  2. A HTTP Client that is capable of mapping itself a CoAP URI to a HTTP URI can interpret Link Format payloads and do the URI translation itself.
  3. A HTTP Client not capable of mapping CoAP URIs to HTTP URIs, is unlikely to fetch `http://<proxy>/<path>/.well-known/core` and use this in the RFC 6690 manner. Following RFC 6690, it would only fetch `http://<proxy>/.well-known/core` and none of the Link Format descriptions of the CoAP servers would be there, so no mapping needed.

Ticket #366 resolution details (3/3)

- Final Thought: In the draft we could advise Link Format mapping for the case that CoAP servers' Link Format content is served in `http://<proxy>/.well-known/core`, in order to comply with RFC 6690
  - I.E. Add an informative reminder (in section 6.3.3.2) to developers that to comply with RFC 6690 the Link Format actually should be translated correctly if the developer wants to include the CoAP server "/.well-known/core" into the HC Proxy "/.well-known/core"
    - (see also section 5.4, Discovery of "core.hc")
    - If this translation is omitted, the Link Format on the HC Proxy would be incorrect and violate RFC 6690 rules

Closed Ticket #375

- Closed Ticket #375 (Add requirement on mapping of CoAP diagnostic payload)
  - Addressed in new Section 6.3.3.3 (Diagnostic Messages):

    - "CoAP responses may, in certain error cases, contain a diagnostic message in the payload explaining the error situation, as described in Section 5.5.2 of [RFC7252]. In this scenario, the CoAP response diagnostic payload MUST NOT be returned as the regular HTTP payload (message body). Instead, the CoAP diagnostic payload should be used as the HTTP reason-phrase (of the HTTP status line as defined in Section 3.1.2 of [RFC7230]) without any alterations."

Addressed Yusuke's comment

- Addressed Yusuke's comment (on warning of potential loss of meta data in certain transcoding schemes)
  - Addressed in new last paragraph of Section 6.3.3.1 (Content Transcoding, General):

    - "However, it should be noted that in certain cases, transcoding can lose information in a non-obvious manner. For example, encoding an XML document using schema-informed EXI encoding leads to a loss of information when the destination does not know the exact schema version used by the encoder. So whenever the HC Proxy transcodes an application/ XML to application/EXI in-band meta data could be lost. Therefore, the implementer should always carefully verify such lossy payload transformations before triggering the transcoding."

Next Steps

- Any other issues that the WG thinks we need to solve for HTTP-CoAP reverse proxies?
  - Editors will do an editorial pass if there is no other technical issues

And/or

- Are we ready for WGLC?

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

CoRE Resource Directory

draft-ietf-core-resource-directory-02

Z. Shelby, C. Bormann

Changes in this version

- #369 – Web app catalogue use case addition
- #370 – Integrate the DNS-SD mapping
  - Together with Kerry Lynn
- #371 – DDoS security consideration
- #372 – Example section for use cases
  - Help from the WG still needed on examples
- #373 –Registration update interface now using POST
- Added text on endpoint identification and authorization
- Error code 4.04 added to Registration Update and Delete interfaces
- Made 63 byte size a SHOULD for endpoint name and endpoint type parameters

Registration update interface

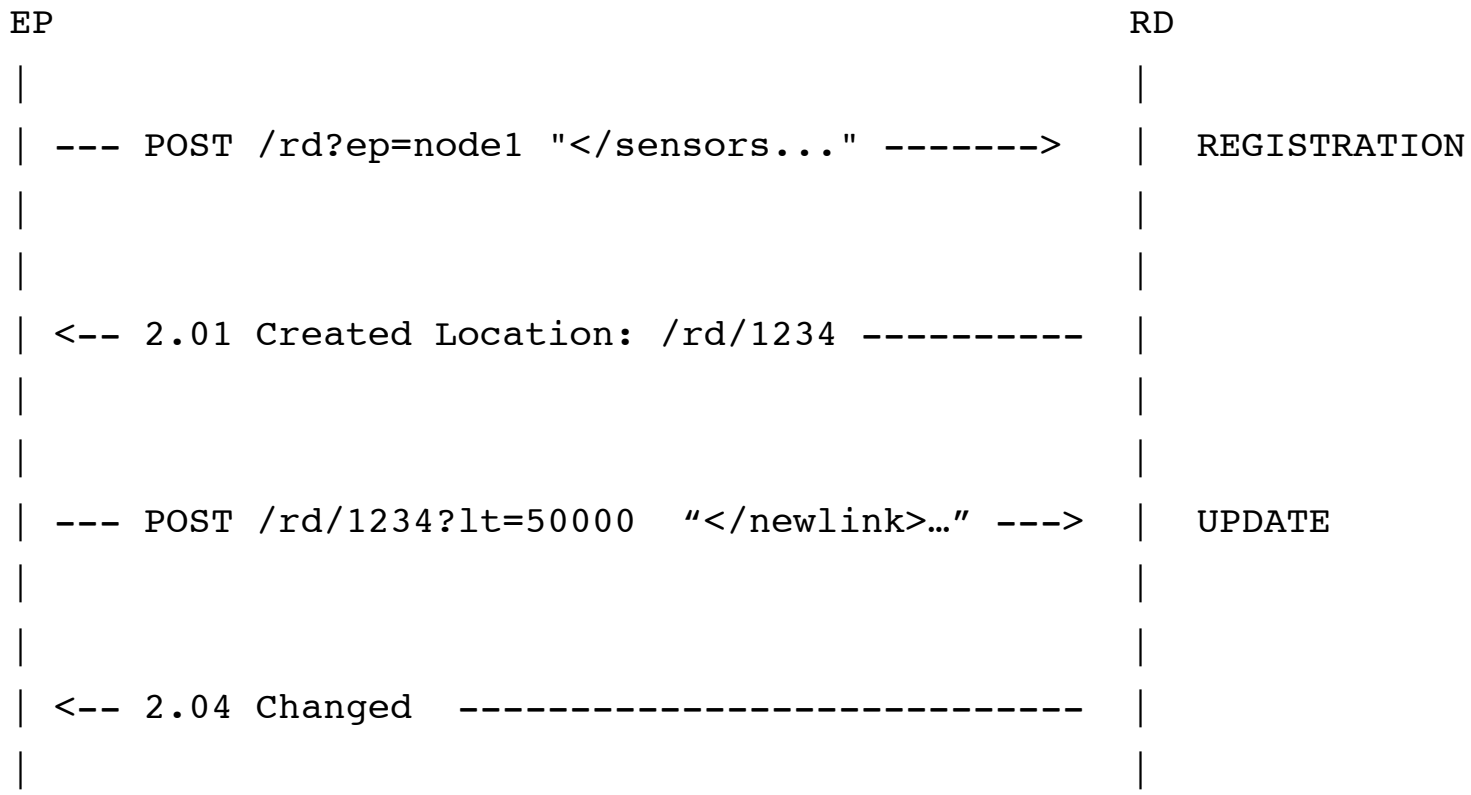

When are we done?

1. Do one more editing round for -03
2. Known issues for -03
   - Link maintenance collection interface needed at /{+location} after registration (see next slide)
   - Advanced examples needed (Volunteer?)
3. Get at least 3 expert reviews, at least one from OMA
4. And add this to our charter ☺

Link Collection Interface

- Goal
  - For an endpoint (or installation tool) to manage collections of links after registration
- Interface
  - REST interface at /{+location} (returned in registration)
  - Get link collection
  - Add a link
  - Remove a link
  - Replace a link
- Design proposals welcome, but let's not make this an excuse to invent a patch mechanism unless really useful

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

Authentication and authorisation for the Resource Directory

Bert Greevenbosch, bert.greevenbosch@huawei.com

Abstract

- draft-greevenbosch-ace-resource-directory aims to describe how ACE can be used to allow authorised access to the Resource Directory.
- The document is informative.
- The document currently considers issues about authorised access with the resource directory, and distillates requirements from these issues.
- https://datatracker.ietf.org/doc/draft-greevenbosch-ace-resource-directory/

Recap of resource directory

- The resource directory is used as a discovery mechanism for CoAP endpoints.
- The CoAP endpoints are often resource servers.
- When registering themselves with the resource directory, they provide information about the resources they host.
  - They use the CoAP link format for this.
- The resource directory retains the information, and makes it available to other CoAP endpoints.
- https://datatracker.ietf.org/doc/draft-ietf-core-resource-directory/

Architecture

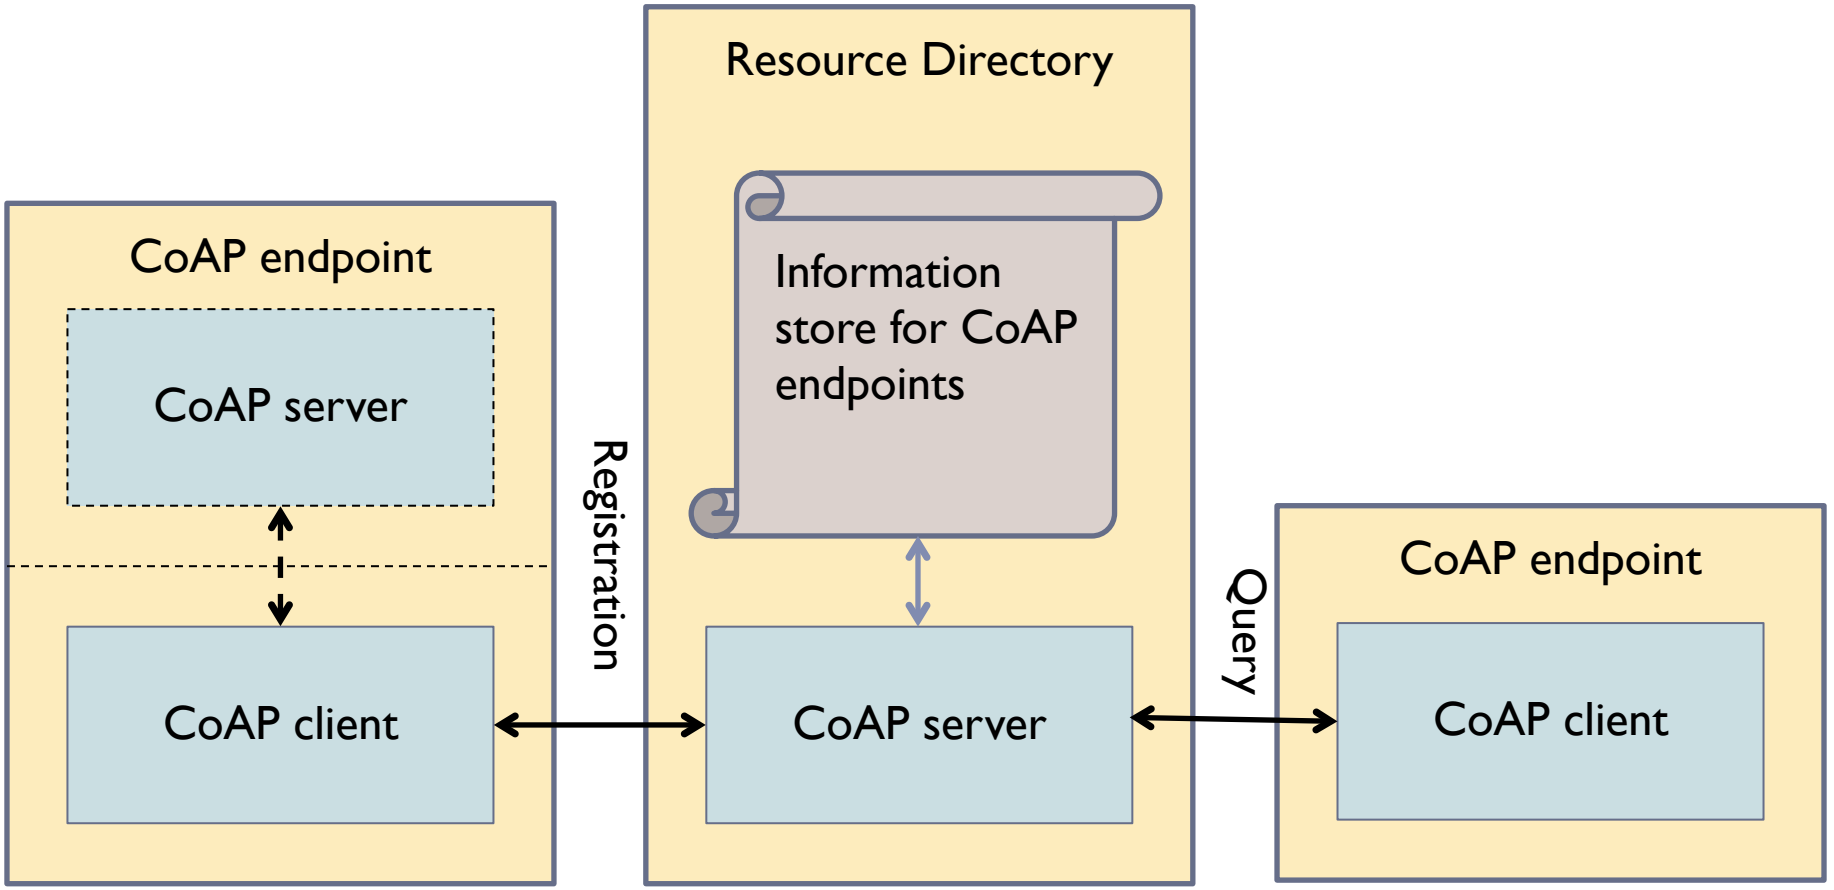

Observations

- The resource directory acts as a CoAP server.
- From the resource directory's perspective, the CoAP endpoints all act as CoAP clients.
  - In reality, the main function of certain endpoints may be CoAP servers.
  - CoAP endpoints that only contain CoAP server functionality, require a different CoAP endpoint to register with the RD on their behalf.
- Ideally, authentication and authorisation should be performed as usual for other CoAP server/client pairs.

Requirements for authentication and authorisation

- REQ-1 The RD should be able to perform authentication and authorisation from a CoAP server's point of view.
- REQ-2 The RD should be able to authenticate and verify authorisation of CoAP clients.
- REQ-3 The RD should be able to authenticate and authorise delegation clients.
- REQ-4 CoAP servers should be able to authenticate and authorise delegation clients.

Requirements for registration

- REQ-5 The endpoint should be authenticated, such that it cannot spoof other endpoints.
- REQ-6 The endpoint should only be able to provide, change or delete registration information of other endpoints if it is authorised to do so.
- REQ-7 If the endpoint is member of a domain, it should be possible to ensure the true membership of that domain.

Requirements for querying

- REQ-8 The RD should be able to grant different access rights to different clients.
- REQ-9 If the different areas are implemented through different URIs, it should be possible for the RD to approve or block access to related areas by providing access rights to specific URIs.
- *REQ-10 (TBD) Do we want to specify specific access rights to URI- queries?*

Conclusion

- The Resource Directory can be considered as just another CoAP server.
- Authentication is required for registering, querying and delegating.
- The Resource Directory can be used as a test drive for the authentication and authorisation work.

Thank you!
Discussion...

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

draft-ietf-core-links-json-00.txt

- RFC 6690 (link-format) documents are somewhat foreign to many web app developers
  - would prefer to have them in JSON format
- There is no standard way to represent link-format documents in applications
  - but everyone knows how to handle JSON

→ Define a standard JSON translation for link-format

```
</sensors>;ct=40;title="Sensor Index",
</sensors/temp>;rt="temperature-c";if="sensor",
</sensors/light>;rt="light-lux";if="sensor",
<http://www.example.com/sensors/t123>
  ;anchor="/sensors/temp";rel="describedby",
</t>;anchor="/sensors/temp";rel="alternate"
```

➔

```
[{"href":"/sensors","ct":"40","title":"Sensor Index"},
  {"href":"/sensors/temp","rt":"temperature-c","if":"sensor"},
  {"href":"/sensors/light","rt":"light-lux","if":"sensor"},
  {"href":"http://www.example.com/sensors/t123",
   "anchor":"/sensors/temp","rel":"describedby"},
  {"href":"/t","anchor":"/sensors/temp","rel":"alternate"}]
```

Potential Issue: How to update

- Structure: Array of links
- RD update might
  - add links: trivial
  - change links: replace on href as key?
  - remove links (how to indicate this?)
- draft-ietf-appsawg-json-merge-patch was defined to solve problems like this
  - but does not fit: only can update object (map), not array
- restructure links-json as an object (map)?
  - are hrefs really unique in real-life link sets???

Group 2: other docs

Tuesday

All times are in time-warped HST

- **15:20–15:30 Intro**
- **15:30–15:45 HTTP mapping (SL)**
- **15:45–16:25 Resource Directory (ZS)**
- **16:25–16:45 ACE for Resource Directory (BG)**
- **16:45–16:47 Links-JSON (chairs)**
- **16:47–16:49 Core-Interfaces (chairs)**
- **16:49–17:04 alt trans: DTLS on SMS (HT)**
- **17:04–17:20 CoAP-PubSub (MK)**

DTLS over SMS

IETF 91 – Honolulu, Hawaii

Thomas Fossati, Hannes Tschofenig

Overview

- **Use case**: Securing CoAP messaging transmitted over SMS in machine-to-machine deployments.
- **Characteristics of this transport**:
  - Substantial and highly variable latency
  - Limited bandwidth (140 bytes)
- **History**: Work started in context of the OMA LWM2M and swapped over to the IETF.

Challenges for DTLS

- **Latency**
  Setting appropriate timeout values is critical to avoid spurious retransmissions.
- **Fragmentation**
  Since TLS retransmission applies to flights, a single lost T-PDU (or one that has been delayed a bit too much) implies retransmission of all T-PDUs in the same flight.

Work in Progress

MultiTech SocketModem

- **Prototype implementation**
  - PSK can be done in 4-6 T-PDUs (i.e. one per flight)
  - X.509 rarely completes successfully, and when it does it can take minutes.
  - Evaluation of handshake with raw public keys.
- **Standardization activities**
  - Level of interest: others exploring similar deployments?
  - What would be the best group to discuss this topic? (CORE, DICE, …)

Flextime

- **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 3979 and its updates**

Blue sheets
Scribe(s)

http://www.ietf.org/about/note-well.html

Note Well

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

- The IETF plenary session
- The IESG, or any member thereof on behalf of the IESG
- Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices
- Any IETF working group or portion thereof
- Any Birds of a Feather (BOF) session
- The IAB or any member thereof on behalf of the IAB
- The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice. Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

How 6TiSCH is using CoAP for management

Thomas Watteyne

IEEE802.15.4e Time Synchronized Channel Hopping

- Nodes are synchronized
- Communication follows a schedule
- Schedule gives tunable trade-off between
  - packets/second
  - latency
  - robustness

} …and energy consumption

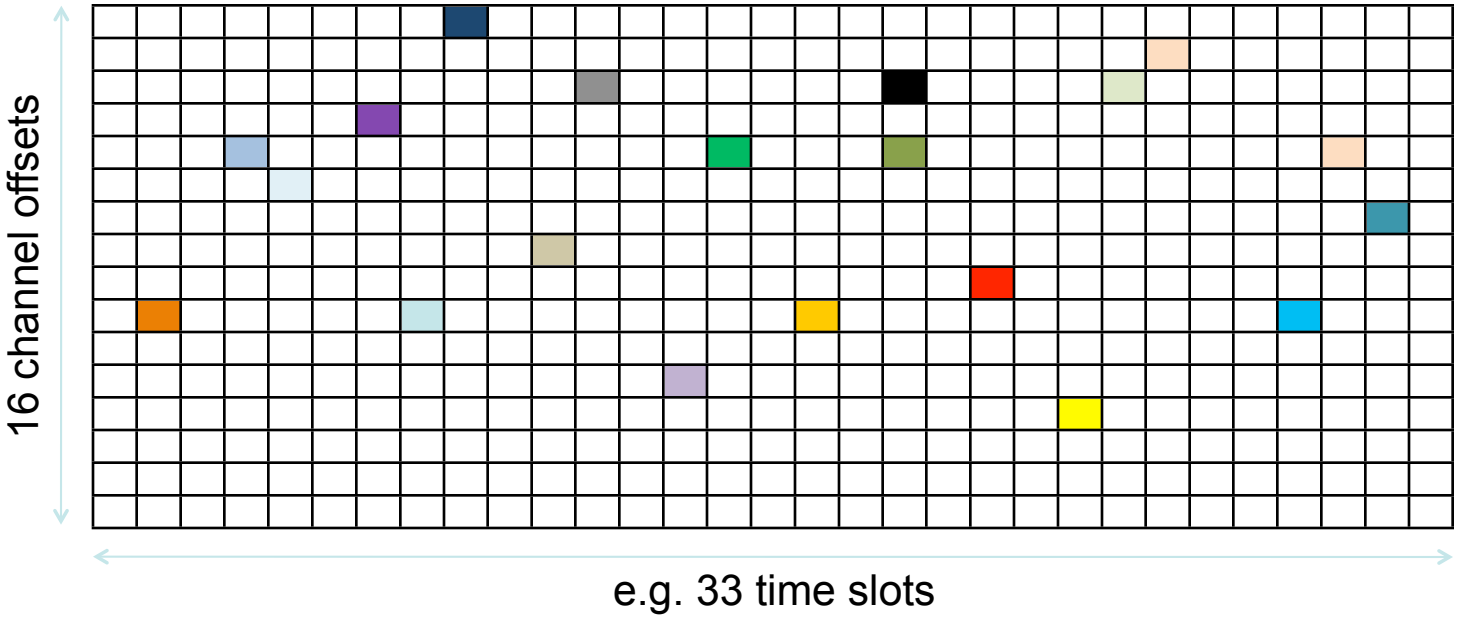

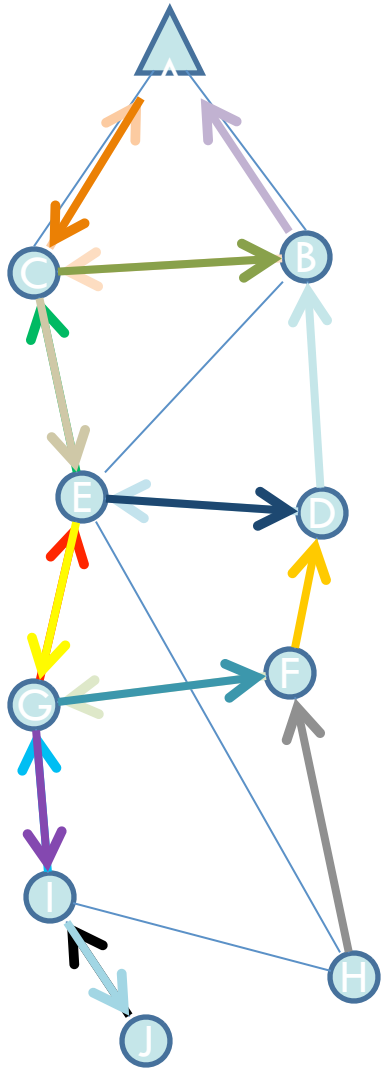

→ 6TiSCH defines mechanisms to build/maintain the TSCH schedule

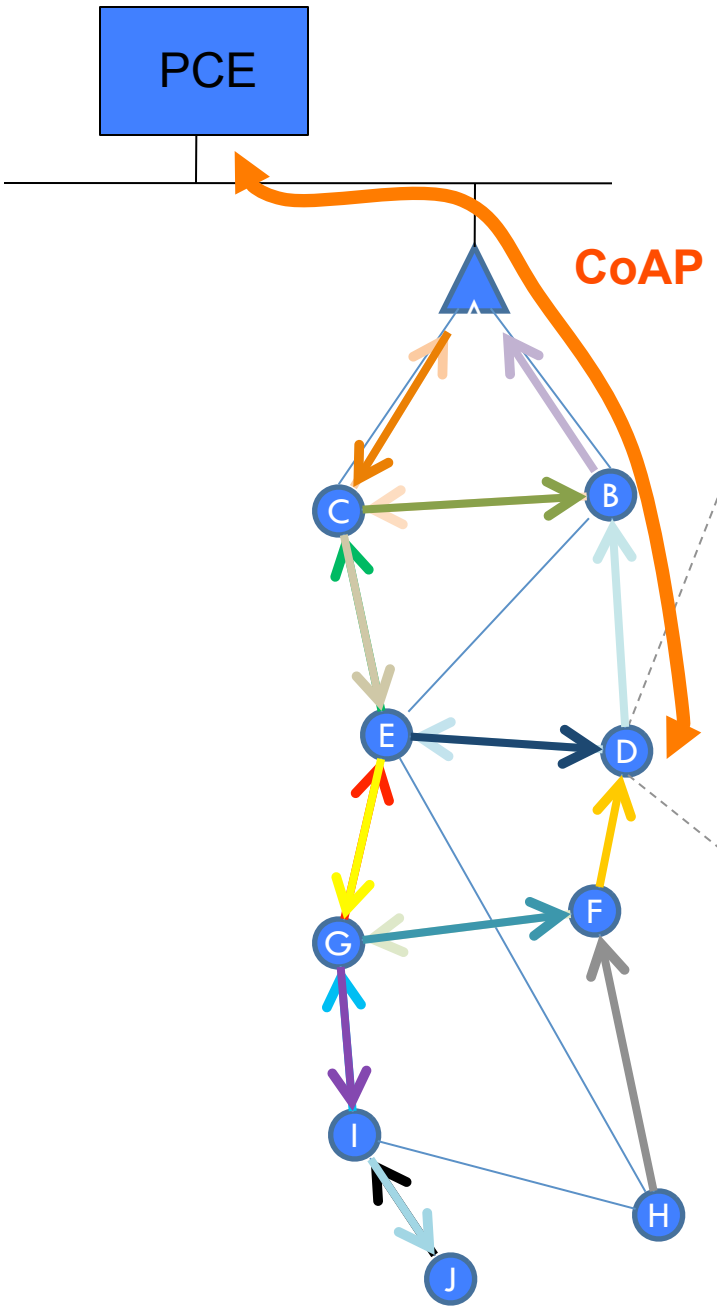

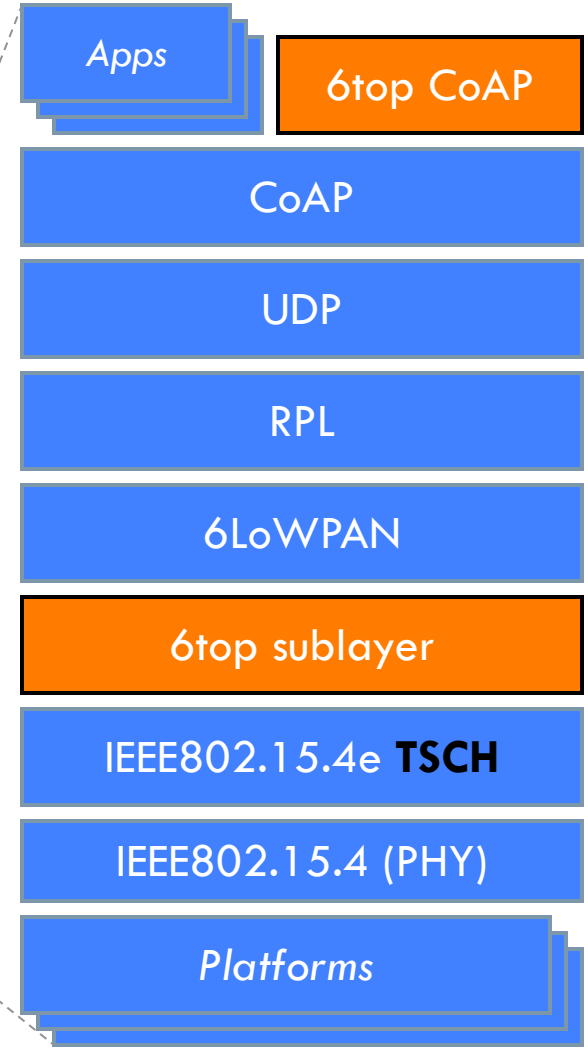

"centralized"

- RESTful access to resources → CoAP
- observe changes to resources → observe
- secure → DTLS/DICE
- fine-grained authorazation → ACE

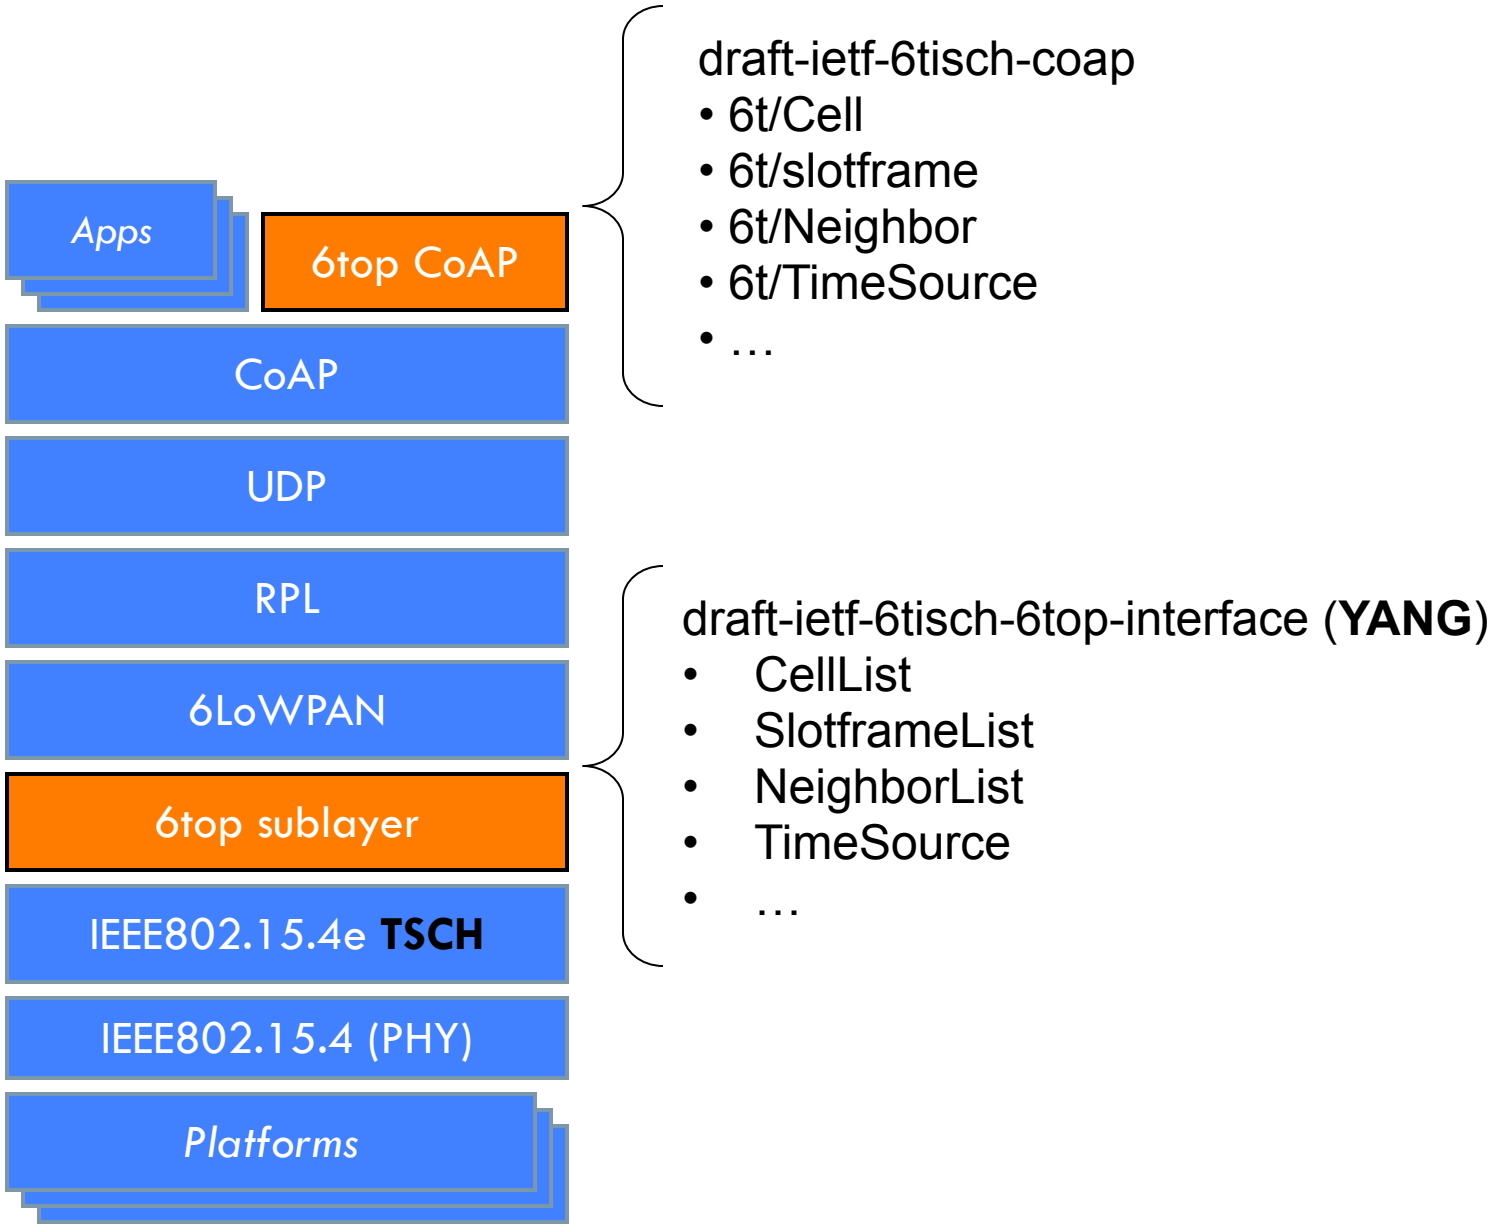

"centralized"

Open Question

How does draft-ietf-6tisch-coap compare to

- draft-vanderstok-core-comi-05
- “constrained RESTconf”

?

How to introduce some “version” in the CoAP interface?

- /6t/ve to discover the version
- /6t/1/, /6t/2/ to prefix different versions

“centralized”

draft-wang-6tisch-6top-coapie-00

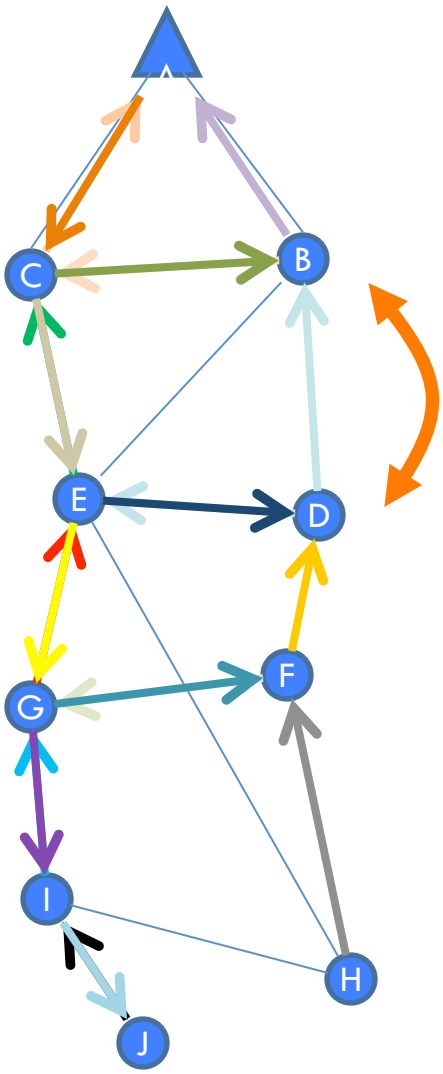

CoAP

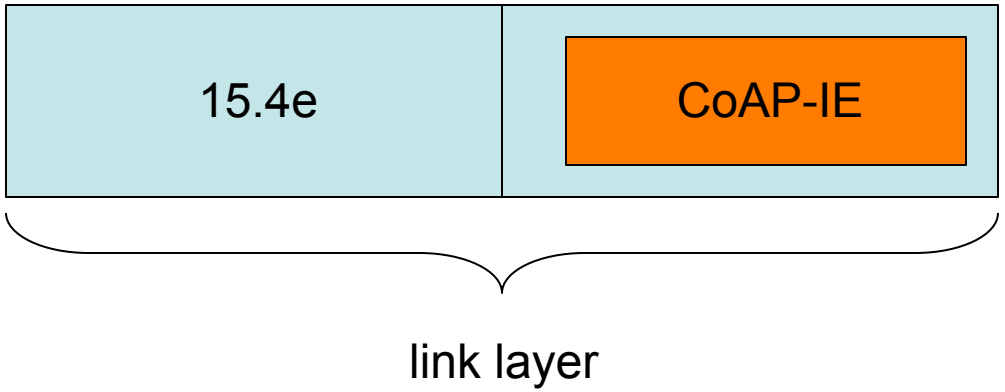

6TiSCH WG Meeting

- 0900-1030 HST Thursday Morning Session I
- Hibiscus room
- https://datatracker.ietf.org/meeting/91/agenda/6tisch/

CoRE working group

CoAP Management Interface
draft-vanderstok-core-comi-05

P. van der Stok, B. Greevenbosch, A. Bierman, J. Schoenwalder

November 12, 2014

Motivation

Provide transport over CoAP between clients and “reduced resource” servers

to access standardized object (written in SMI or YANG) to:

- Do statistics (e.g. fragmentation percentage in LoWPAN packets)
- Initialize parameters (e.g. DIOIntervalMin in RPL)

With the wish to:

- Provide small payloads and transport overhead
- Single transport interface and security protocol for all applications

Including management

History

Earlier versions discussed replacement of SNMP by CoAP to access MIBs

- Evolved from XML/EXI payload to JSON/CBOR payload
- Addition of name string to number conversion

Existence of RESTconf and SMI to YANG algorithm led to current version 5

- Aligned with RESTconf (which uses http)
- JSON/CBOR payload
- Hash for name strings
- Discovery
- Access granularity
- Dependence on Block and Observe

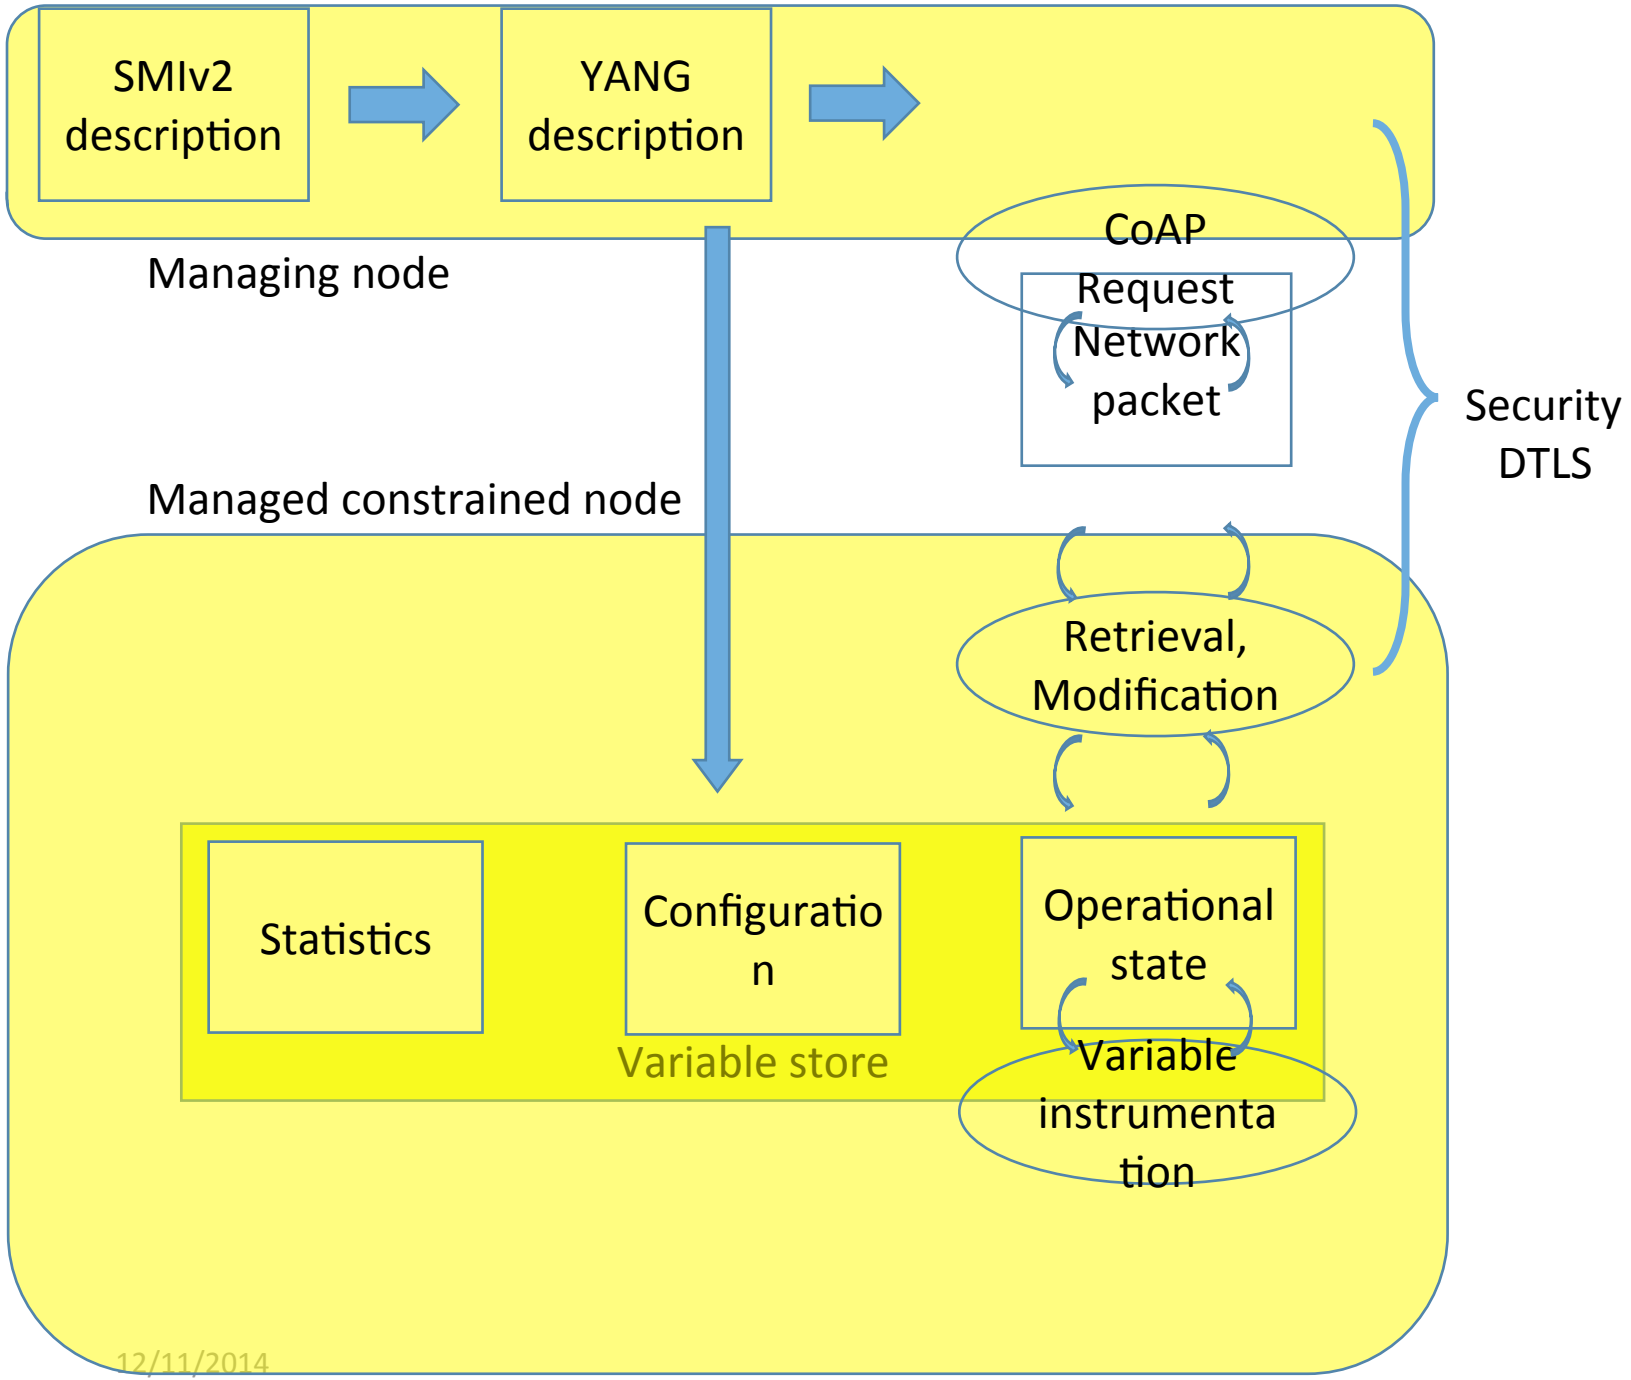
SMIv2 description
YANG description
Managing node
CoAP Request Network packet
Security DTLS
Managed constrained node
Retrieval, Modification
Statistics
Configuratio n
Operational state
Variable store
Variable instrumenta tion

Profile of CoMI Function set

name	path	rt	Data type
Management	/mg	Core.mg	n/a
data	/mg/data	Core.mg.data	Application/cbor
Module set URI	/mg/moduri	Core.mg.moduri	Application/cbor
YANG Hash Infor	/mg/yang-hash	Core/mg.yang-hash	Application/cbor

YANG module and MIB

URI in request specifies the data envelope

```
REQ: GET example.com/mg/data/system-state/clock/current-
datetime

   RES: 2.05 Content (Content-Format: application/cbor)
   {
       "current-datetime" : "2014-10-26T12:16:31Z"
   }

REQ: GET example.com/mg/data/system-state/clock

   RES: 2.05 Content (Content-Format: application/cbor)
   {
      "clock" : {
         "current-datetime" : "2014-10-26T12:16:51Z",
         "boot-datetime" : "2014-10-21T03:00:00Z"
         "timezone" : {
            "timezone-location" : "Europe/Stockholm",
            "timezone-utc-offset" : -60
         }
      }
   }
```

?Select query specifies items within data envelope

Not mentioned in CoMI-05, but functionality available in CoMI-04

```
 GET //example.com/mg/data/system-state/clock
      ?select=/current-datetime, /timezone/
timezone-location
RES: 2.05 Content (Content-Format: application/cbor)
   {
         "current-datetime" : "2014-10-26T12:16:51Z",
         "timezone" : {
            "timezone-location" : "Europe/Stockholm",
         }
   }
```

Within envelop /data, possible to select from multiple modules

Payload encoding

In I-D text, examples are all described with JSON payload
In the packet, CBOR is transported, name strings are hashed

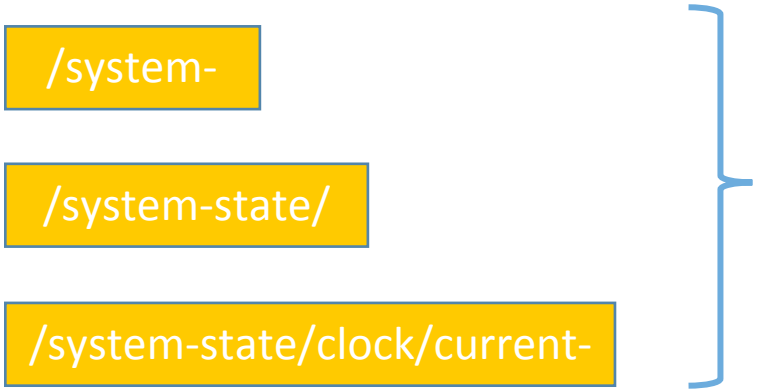

On 1700+ YANG objects, the average YANG identifier local-name length is between 9 and 10 bytes.
Proposed Hash function: murmur3

No Hash collision detected yet.

Encoding discussion

Required: Hash code for given object identical over all servers.
Note:

- Clashes between servers are irrelevant
    - Client can use same hash for different objects on different servers
- Clash on a server needs to be notified with rehash

Alternatives payload formats to be investigated (e.g.)

- ZIP (seems to be less performant than CBOR)
- JSON with hash (no CBOR)
- Etc…
- Other hash

Discovery

Discovery information exported to “YANG module” servers
Managed server only knows hash codes,
names are present on “YANG module” server

/Moduri provides link to external ietf-yang-library module

Managed servers can add data items, identified by hash code, to /.well-known/core with their rt value

WHY continue ?

Discontinuation reasons:

- SNMP exists
  - Small payload
  - Small footprint
- Industry will not easily switch to support of “another” protocol
  - Current routers support SNMP or NETCONF servers.

Continuation reasons:

- RESTful access to MIBs and YANG modules
  - One homogeneous CoAP interface application development (including mgmt)
- CoMI addresses a new “small device” market
- SNMP security footprint is large, and not needed with CoAP security present
- CoMI aims at small payloads without information loss

Where continue ?

Title says it all:

- CoRE Management Interface (CoMI)
  - CoAP developed in CoRE is central for CoMI
  - Motivated by wishes of CoRE wg members (and related market)
- Access to standardized management data definitions (YANG modules and MIBs)
  - Leaves the data definitions unchanged.
- CoMI provides subset of SNMP and NETCONF functionality
  - Does not affect SNMP or NETCONF

Proposal: **CoRE WG adoption**

HOW to continue ?

Payload formats:

- XML/EXI have been excluded earlier
- JSON or CBOR
- Name string size reduction

Conformance with RESTConf syntax:

- URI path names
- ?Query formats

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

CoAP over TCP and TLS

Results from ad-hoc group 2014-11-11

Carles Gomez, Carsten Bormann, Hannes Tschofenig, Matthias Kovatsch,
Michael Koster, Mikko Saarnivala, Robert Cragie, Simon Lemay

Many Options

- Many of them described in draft-bormann-core-coap-tcp-01
- One workable solution described in draft-tschofenig-core-coap-tcp-tls-01
- Ad-hoc meeting between draft authors and others after CoRE Meeting Tuesday
  - We skipped the social for this!

Connection model

- What happens in a connection, stays in the connection
  - CoAP endpoint bound to the connection
  - No transfer of e.g. observation relationships between connections
- Hard to define a continuation relationship for “resilient model”

Delimiting

- Fixed length Length field
  - (not: self-delimiting; Minion-style scanning)
- efficient, independent from packet parsing
- chunking not required (≤ 1152 bytes)

Length of length field

- Fixed 2-byte

- 64 KiB is enough for everyone (≤ 1152)
- Stay compatible with UDP max size
  (no reserved bit)

4-byte header?

- Clear message type, remove message-ID
- Fields have no purpose (TCP does reliability)
  - Sending them anyway is interop disaster
- “Message type” bits might be used as future extension point

a) length first

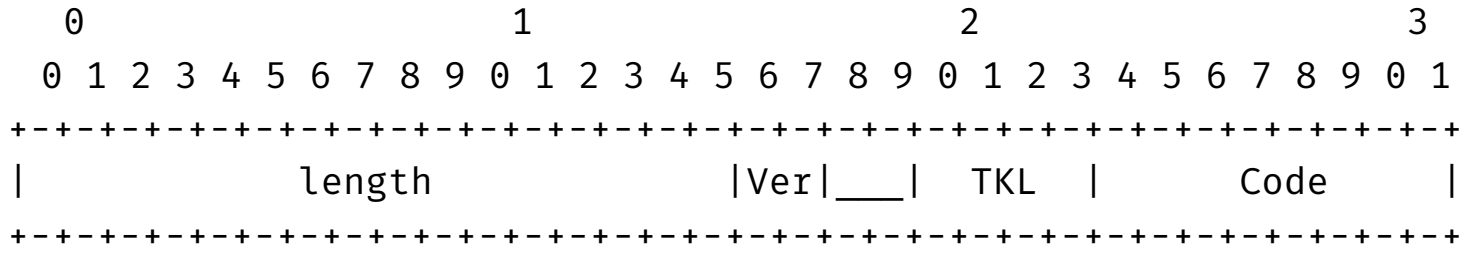

- length includes (2-byte) CoAP header
- can use lengths 0 and 1 for some special applications (e.g., keepalive)
- some existing libraries support this directly (e.g., Erlang)

b) replace message-ID

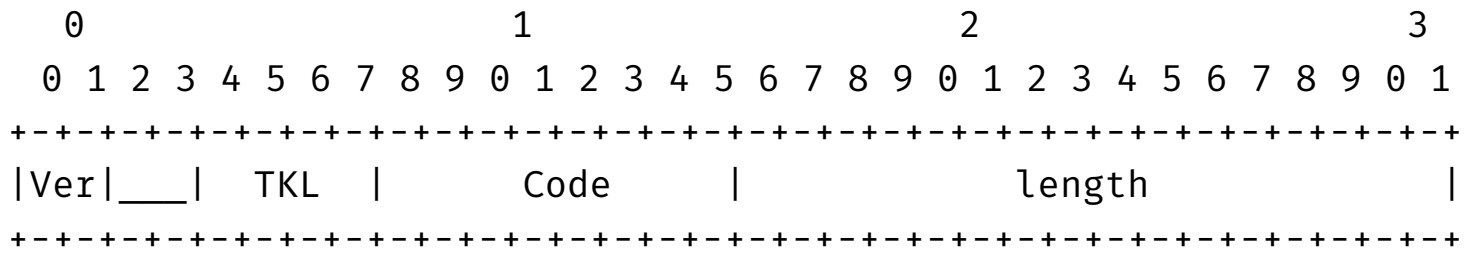

- (length measures Options + Payload)
- “feels like CoAP”
- Can use reserved bits for format extension points

Rendezvous, URI

- TCP: Port 5683, URI: `coap+tcp://`
  TLS: ALPN "`coap1`", URI: `coaps+tcp://`

- (no "signatures"/"magic strings")

Other observations

- "Pipelining" (no strict lock-step)
  - Token provides response matching
- Instead of RST, close connection

draft-silverajan-core-coap-alternative-transports

Main Points

Guidance draft outlining:

- Objective and Usage Scenarios
- Alt. Transports for CoAP message exchanges
  - Payload and Encapsulation considerations
  - Impact on CoAP Header Fields, Message Types, Mechanisms
  - Single Transports vs Multiple Alt. Transports
- Exposing CoAP resources over Alt Transports
  - URI design goals and format
  - Rationale for discarding other URI formats (Appendix)

Way Forward:

- Alternative transports URI format WG discussion seems to have concluded on including transport in scheme:

`coap+transport://..`

coap+tcp://...
coap+sms://...
coap+ws://...

- The draft does NOT allocate any specific schemes by itself

- **Adopt as WG document for Informational RFC track?**

CoAP SMS Transport

draft-becker-core-coap-sms-gprs-05

Markus Becker, mab@comnets.uni-bremen.de
Kepeng Li, likepeng@huawei.com
Koojana Kuladinithi, koo@comnets.uni-bremen.de
Thomas Poetsch, thp@comnets.uni-bremen.de

Recap - Motivation

- In M2M communication, IP connectivity is not **always** supported by the constrained end-points
  - Power saving
  - Coverage (GPRS, 3G, LTE)
- SMS based communication is almost **always supported**
- OMA uses SMS as an alternative transport in OMA-TS-LightweightM2M

Changes in -04, 05

- Changed from draft-04 to draft-05:
  - Updated Options.
  - Adapted URI scheme.
- Changed from draft-03 to draft-04:
  - Removed USSD and GPRS related parts.
  - Removed section 5: Examples
  - Removed section 14: Proxying Considerations
  - Added more block size considerations.
  - Added more concatenated SMS considerations.
  - Rewrote encoding scheme section;
  - 7 bit encoding only.

Way Forward

- Related work
  - draft-fossati-dtls-over-gsm-sms

- Ready for Adoption?

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

TATA CONSULTANCY SERVICES
Experience certainty.

TATA

The No-Response option for CoAP

draft-tcs-coap-no-response-option-07

CoRE WG meeting @ IETF 91

Abhijan Bhattacharyya, Soma Bandyopadhyay, Arpan Pal
TCS Innovation Labs

From meeting #90, Toronto

- Addressed all the comments received so far
  - No new comments received during Toronto meeting
  - Received some supporting nods regarding WG adoption in the mailing list
- One last important aspect was resolving the issue of token-reuse (section 5.2)
  - How the client to decide a suitable time for retiring a token and reuse it
    - NON request with No-Response has no reverse path – no token matching

A little recap …

- A TOKEN_REUSE_TIME is defined similar to Section 2.5 of RFC 7390

```
TOKEN_REUSE_TIME = NON_LIFETIME + MAX_SERVER_RESPONSE_DELAY + MAX_LATENCY
```

- Similar interpretation for multicast
- Unicast specific modifications:
  - MAX_SERVER_RESPONSE_DELAY : simply the expected maximum response delay from the (single) server to which client sent the request
- If not possible for the client to get a reasonable estimate of MAX_SERVER_RESPONSE_DELAY then client SHOULD use a unique token for the request with No-Response to be safe – clarified in the present draft

The option as it looks now

Number	C	U	N	R	Name	Format	Length	Default
92			X		No-Response	uint	1	0

Option Properties

Requested based on option properties

Value	Binary Representation	Description
0	00000000	Suppress all responses (same as empty value).
2	00000010	Allow 2.xx success responses.
8	00001000	Allow 4.xx client errors.
16	00010000	Allow 5.xx server errors.

Option values

Thank you

Are we there yet?

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

Stable Identifiers for CoAP Endpoints

draft-li-core-coap-node-id-option-01
Kepeng Li, likepeng@huawei.com
Gengyu Wei, weigengyu@bupt.edu.cn

draft-kleine-core-coap-endpoint-id-01
Oliver Kleine, kleine@itm.uni-luebeck.de

Problem Statement

- During an observation relationship, client or server IP address may be changed, it becomes difficult to correlate notifications with observe request.
- If a client sends a multicast observation request to a group URI, and receive different notifications from different servers. Client can't determine that several responses (update notifications) come from the same server if that server has changed its IP address.

Problem Illustration (Unicast)

- Client sends Request to start observation
- At some point Server sends Responses (Update Notifications) with new IP address
- Client can not relate Responses with initial Request
- Update Notifications after Servers IP address change are lost
- The same holds for raw CoAP (without observation) if the Server changes its IP between reception of a Request and sending of a Response

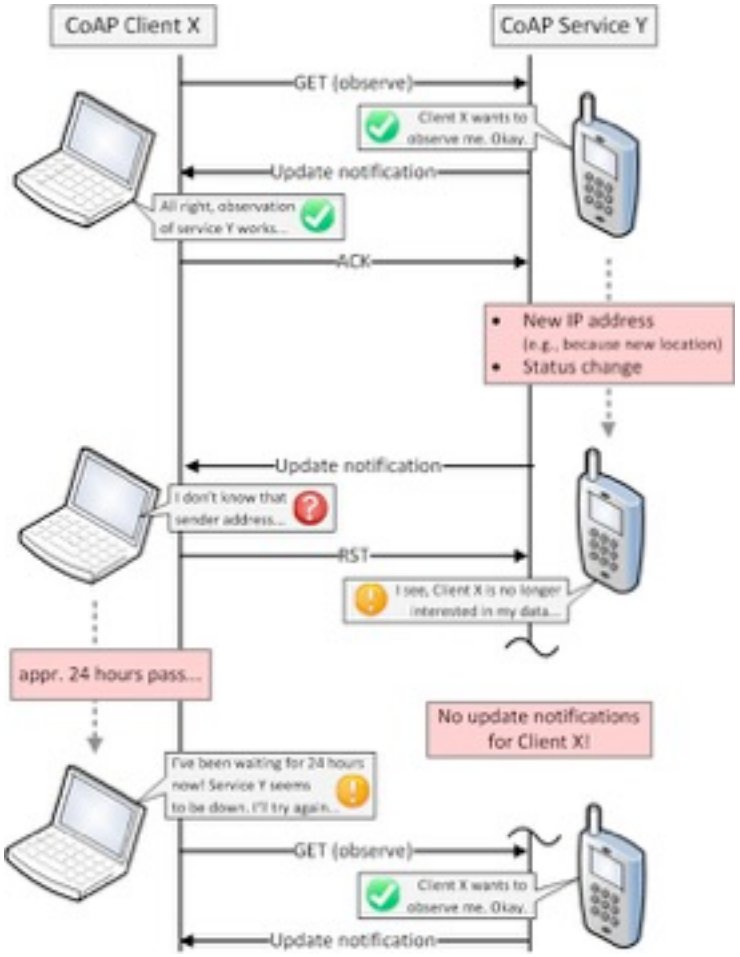

Stateless Solution: NodeId Option

Type	C	U	N	R	Name	Format	Length	Default
TBD	-	-	-	-	NodeId	string	1-255 B	(none)

- Client or server can use self-generated NodeId.
- Endpoint Identifier specified in Resource Directory draft can be reused.
- NodeId option can be used to identify endpoints, and correlate request with notification responses.

Stateful Solution: CoAP Endpoint

- Client assigns an ID to server (value of EID1 option)
- Server repeats assigned ID in every follow-up message (value of EID2 option)
- Assigned IDs need to be unique per client
  - Different servers are assigned different IDs
  - One server may be assigned different IDs from different clients
- Lifetime of IDs: Duration of conversation, e.g.
  - Request/Reponse pair
  - Request/Update Notifications

Endpoint ID Illustration

Open Issues

- Does the problem exist if DTLS channel supported?
- For the stateless solution, how to guarantee the uniqueness of the NodeId?
- For the stateful solution, how to reduce message overhead and reduce state memory for EndpointId?

CoAP Endpoint Unit Identification for Multiple Sensor and Actuator in a Node

draft-hong-core-coap-endpoint-unit-id-01

Yong-Geun Hong (ETRI)

core WG Meeting@IETF 91 – Honolulu, USA
2014.11.12

Scenario of unit Id

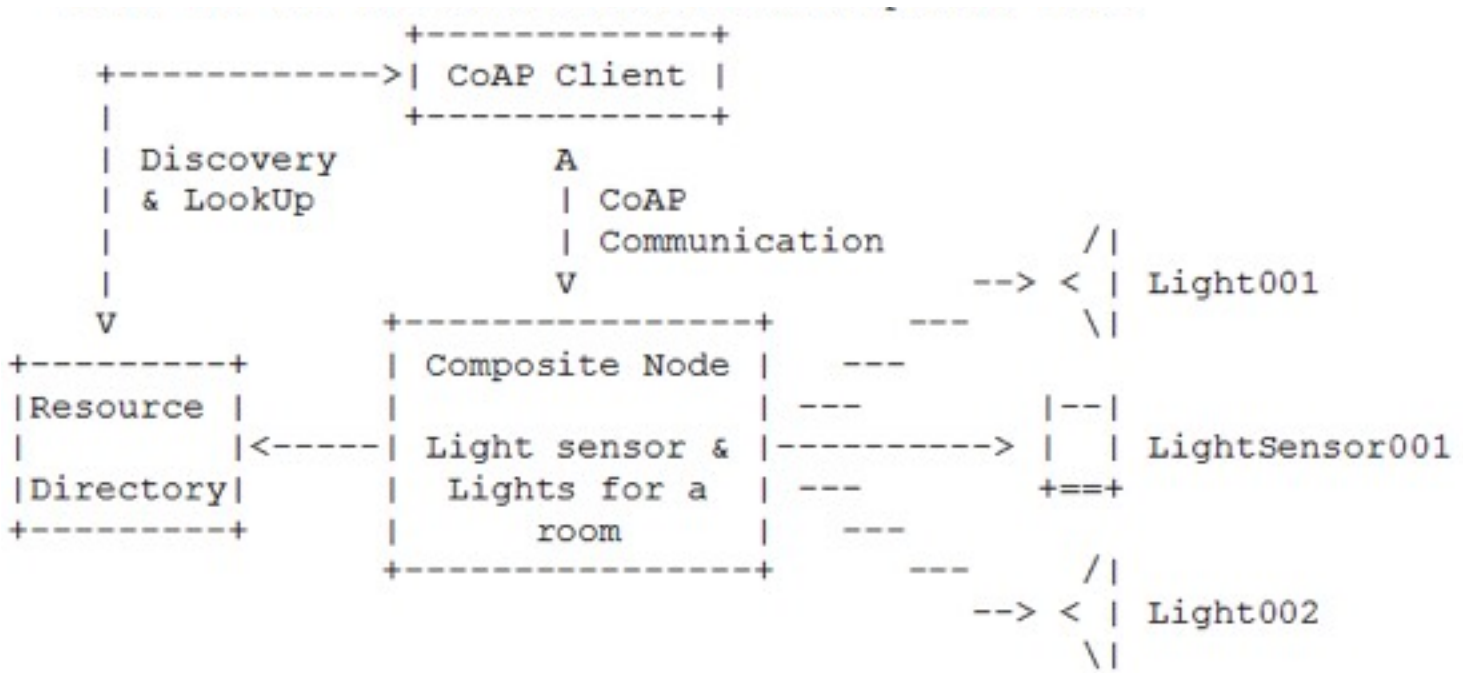

Figure 2: Multiple UnitID based composite CoAP node interaction use Case

Benefits of unit id

- Multiple resources can be uniquely identified by using just multiple unit IDs.
- Single CoAP message can be used to control multiple unit resources by using special characters in conjunction with multi-ID CoAP protocol.
- The reduction in message transmission results in reduced traffic and hence energy conservation in constrained resources.

CoAP Client Server Interaction (Multiple Units)

Figure 8: CoAP based client server interaction (Endpoint multiple Unit ID

Prototype Development Environment

- Prototype development environment for CoAP endpoint unit identification

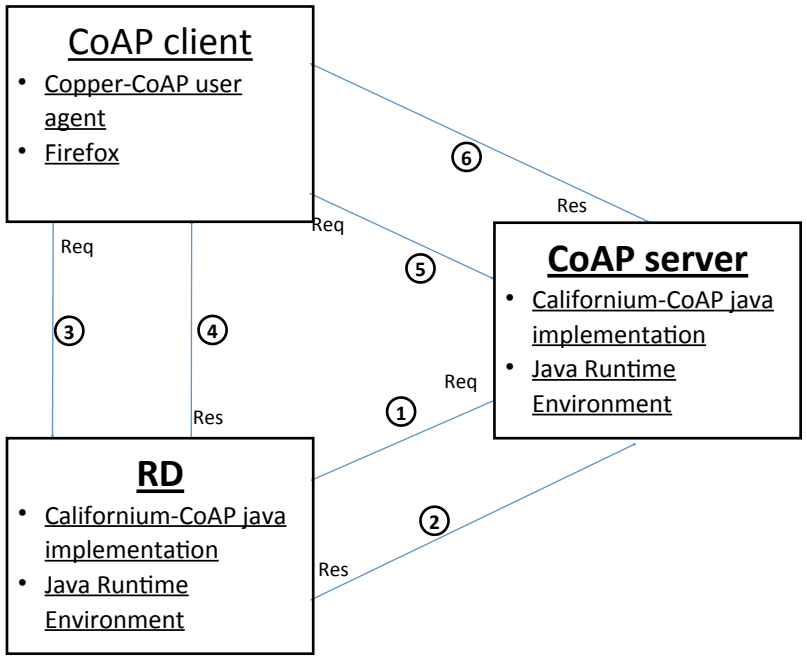

1. Req: POST coap://{rd-ip:port}/rd?ep=node1
   Payload:
   <unitID001>;ct=50;rt="temperature-c";if="sensor",
   <unitID002>;ct=50;rt="light-lux";if="sensor"
   ---
2. Res: 2.01 Created
   ==========
3. Req: GET coap://{rd-ip:port}/rd-lookup/res?rt=temperature
   ---
4. Res: 2.05 Content
   <coap://{node-ip:port}/node1/unitID001>
   ==========
5. Req: GET coap:// {node-ip:port}/node1/unitID001
   ---
6. Res: 2.05 Content
   {"unit_id":"unitID001","data":"22.5 C"}
   ==========
5. Req: GET coap://{node-ip:port}/node1?unit_size=2
   ---
6. Res: 2.05 Content
   {"node1":[
   {"unit_id":"unitID001","data":"22.5 C"},
   {"unit_id":"unitID002","data":"1000 LUX"}]}
   ==========

Prototype Development Results

- Prototype development results for CoAP endpoint unit identification

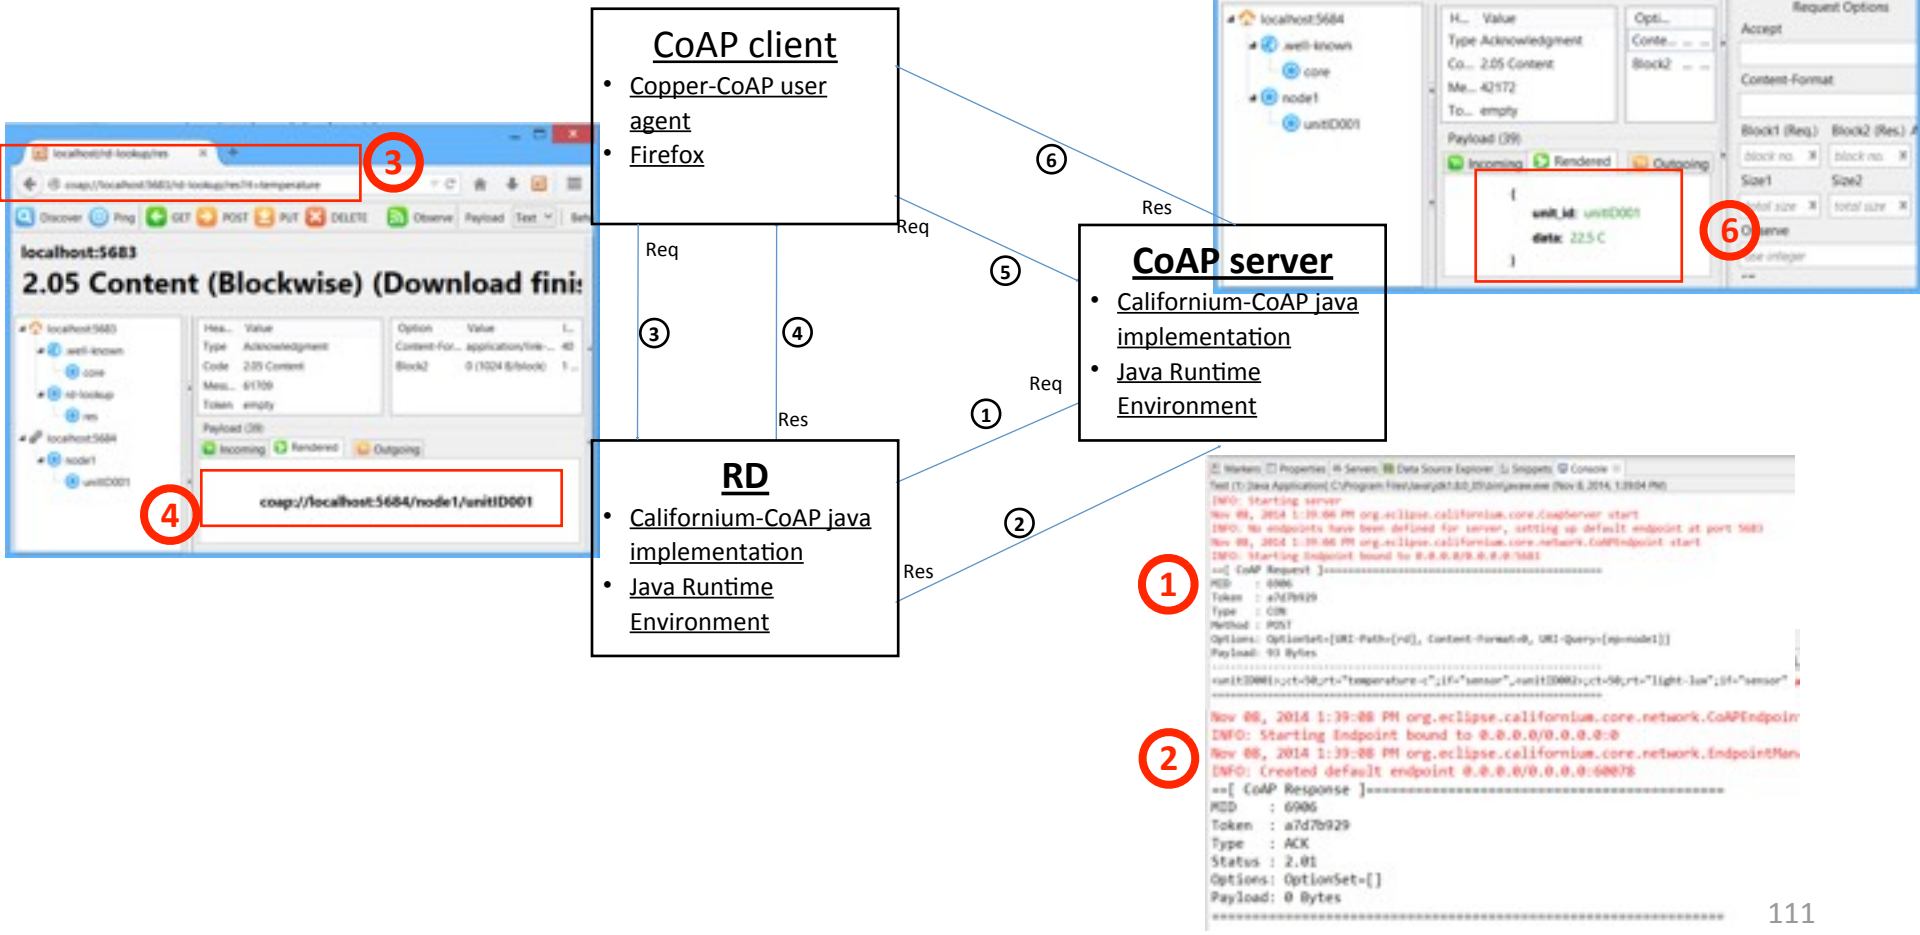

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

CoAP Patience Option

-draft-li-core-coap-patience-option-05

Kepeng Li

Bert Greevenbosch

Esko Dijk

Salvatore Loreto

Motivation

- Client sends a unicast request with a Patience option, to indicate the maximum time the client is prepared to wait for a response.
- It can avoid that the server wastes resources by sending a response which already exceeds the set patience timeout of the client.

Patience Option

No.	C	U	N	R	Name	Format	Length	Default
28			x		Patience	integer	1-2 B	none

- The value of the Patience option is measured in seconds.
- The range is from 1 second to 2^16 seconds, that is, 65535 seconds, around 18 hours.

Recommendation

- Ready for adoption?

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

Update on CoAP Simple Congestion Control/Advanced (CoCoA)

draft-bormann-core-cocoa-02

Carsten Bormann – Universität Bremen TZI
cabo@tzi.org
August Betzler, Carles Gomez, Ilker Demirkol
Universitat Politècnica de Catalunya (UPC)/Fundació i2cat
carlesgo@entel.upc.edu

UPC members have been partly supported by the Spanish Government through project TEC2012-32531, and FEDER

Current status

- Good feedback from IETF'90
  - "Questions" collected in draft-bormann-core-cc-qq-00
  - Suggestion: comparison with alternative RTT/RTO algorithms for CONs
    - Minimalistic RTT-based algorithms
    - RTT/RTO enhancements designed for TCP
      - Improve RFC 2988 (basis of 6298)
      - Linux RTO, Peak-Hopper
- Today presentation: overview of experiment results
  - GPRS scenario
- Version -03 in progress
  - Terminology alignment (e.g. with RFC 6298)
  - Clarifications

Considered RTO algorithms

- Default CoAP
  - RTO randomly chosen from the [2, 3] s interval
    - Insensitive to RTT
- CoCoA
  - Strong and weak estimator
    - CoCoA-S: strong only
  - Includes Variable Backoff Factor (VBF) and aging
- Basic RTO
  - RTO randomly chosen from [last_RTT, 1.5*last_RTT]
  - Also uses weak RTTs
- Linux RTO
  - Reduces contribution of variance to the RTO when RTT decreases
  - Avoids RFC 2988 RTO getting too close to the RTT
- Peak-Hopper RTO
  - Short history and long history estimator
  - Maximum of the two estimators

Running code

- cocoa-02 has been implemented for Californium
  - CoAP implementation for unconstrained platforms
  - Optional *CongestionControlLayer*

- Californium with CoCoA is now publicly available
  - https://github.com/eclipse/californium
    - cf-cocoa example
    - org.eclipse.californium.core.network.stack.congestioncontrol
      - Includes the RTO algorithms considered in these slides!

GPRS scenario

- Two PCs
  - One PC running multiple instances of Californium (Cf) client
    - Sending CON requests
    - cocoa-02

Constant traffic scenario

- Successful exchanges per time unit
  - New CON sent once the previous one is ACKed

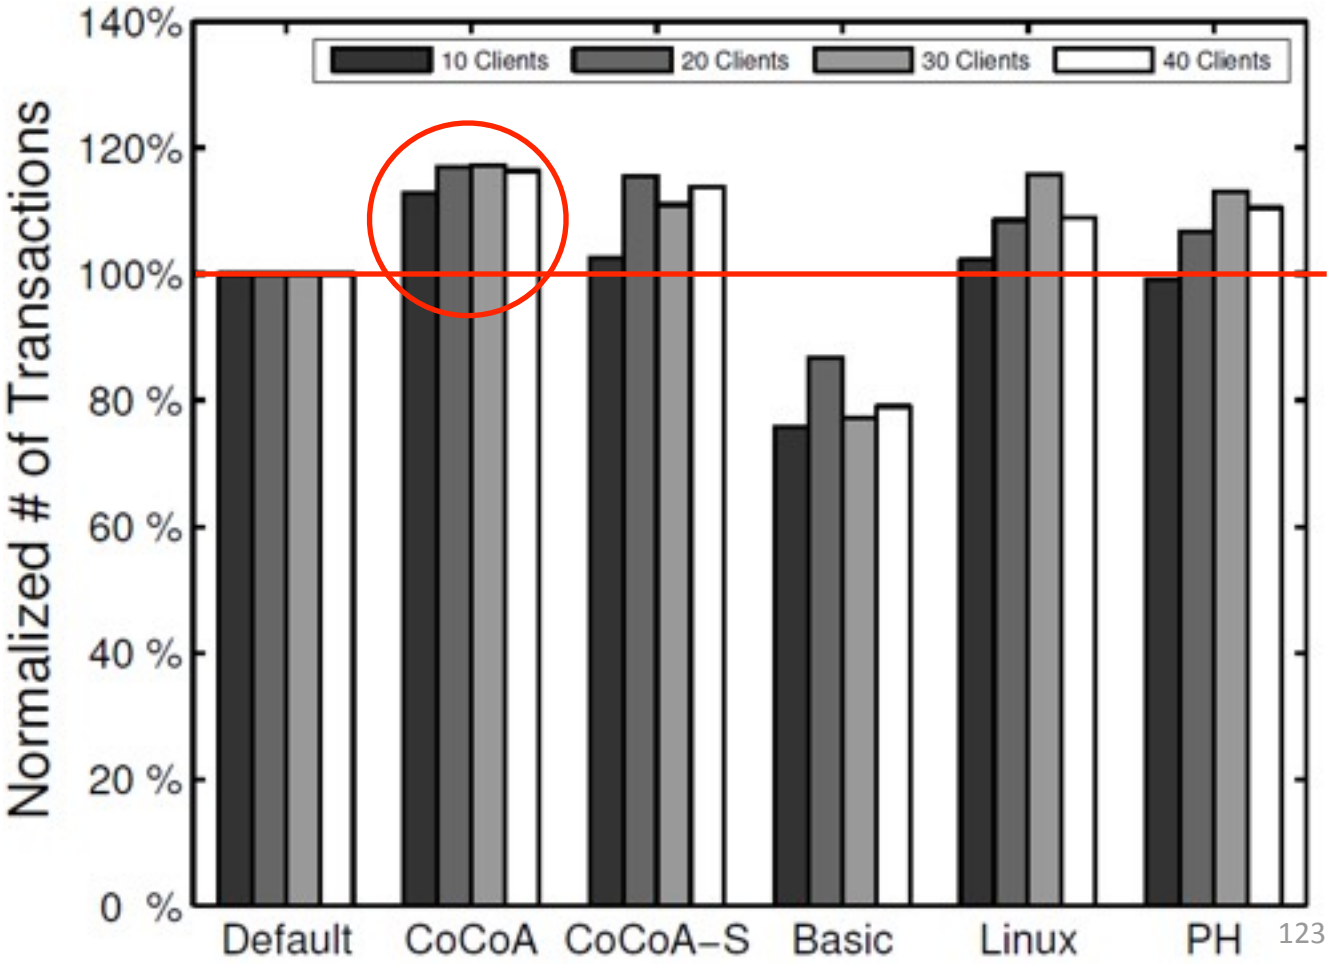

123

Constant traffic scenario

- Retry ratio

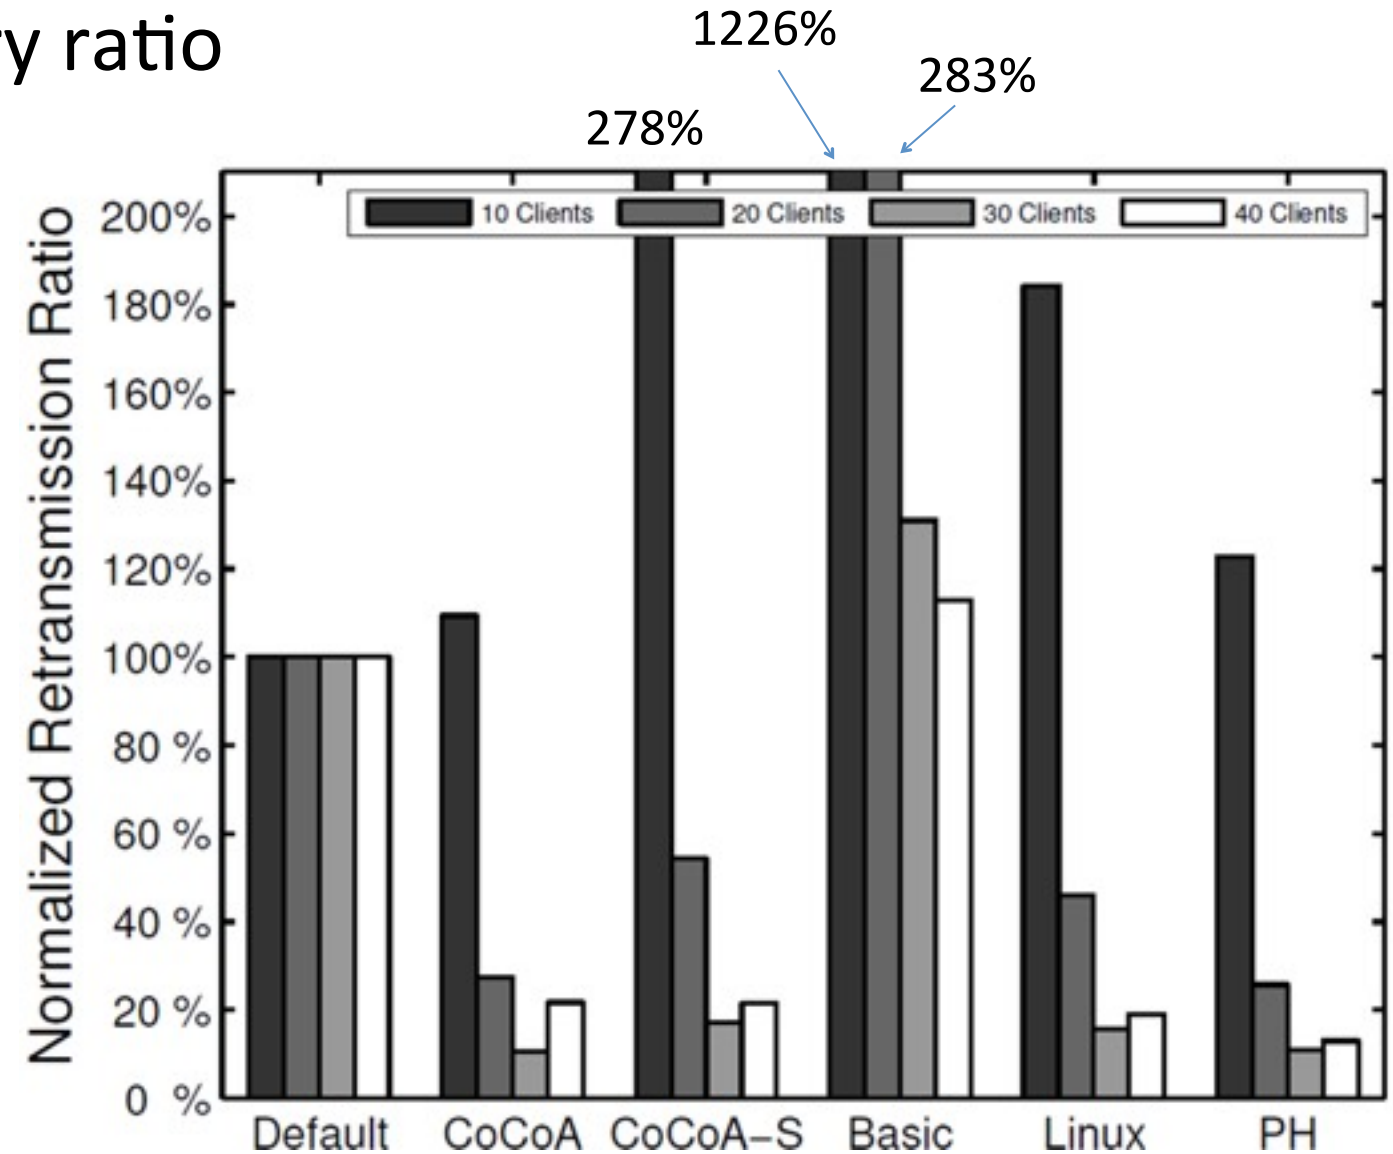

Constant traffic scenario

- Initial RTO

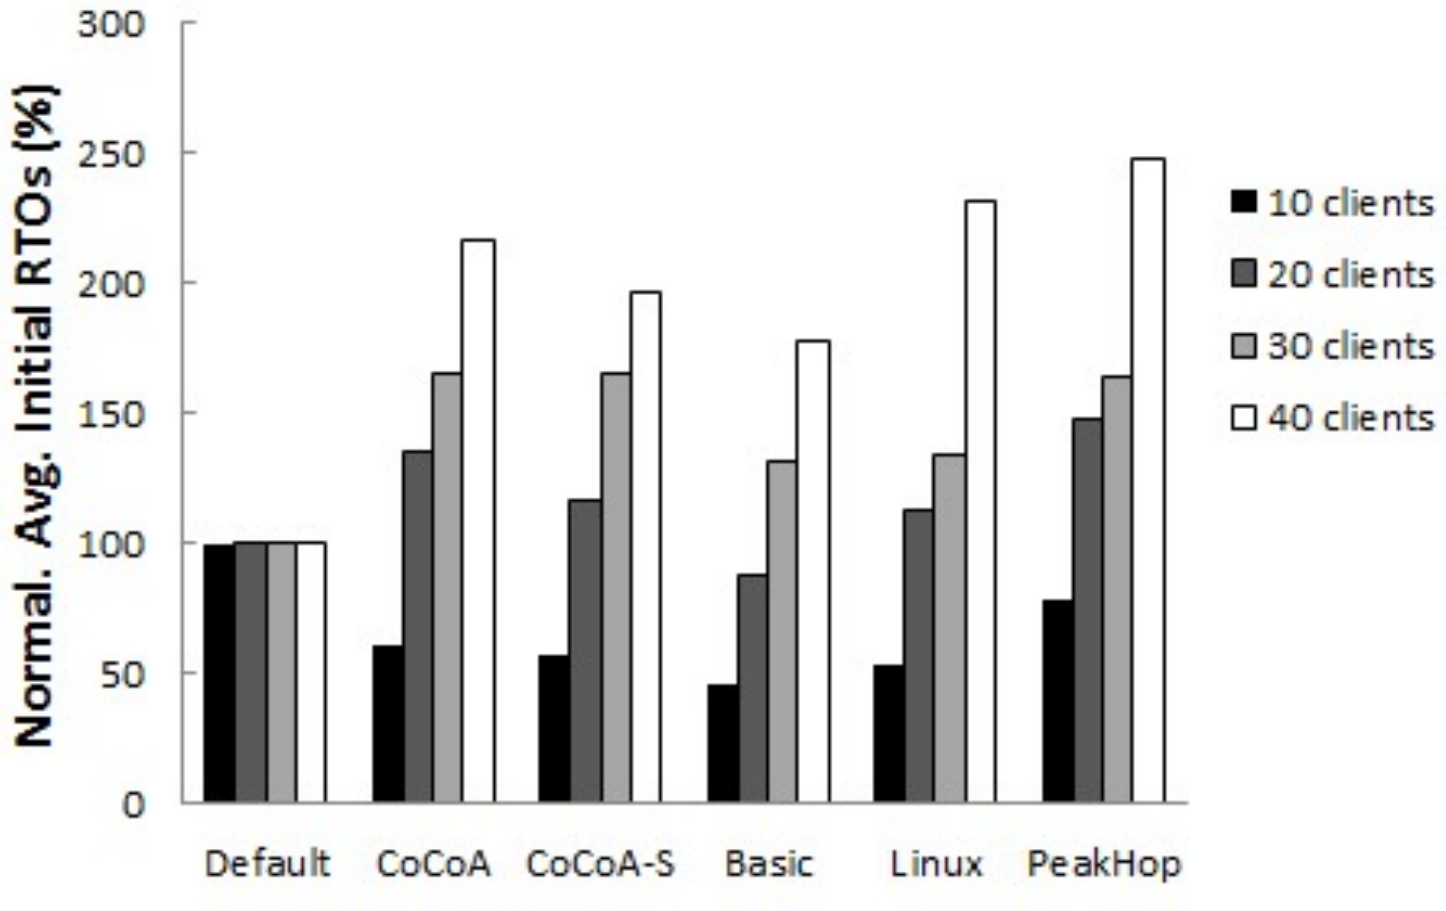

Burst traffic scenario

- Settling time (ST)
  - 80% of the requests generated in a burst are served

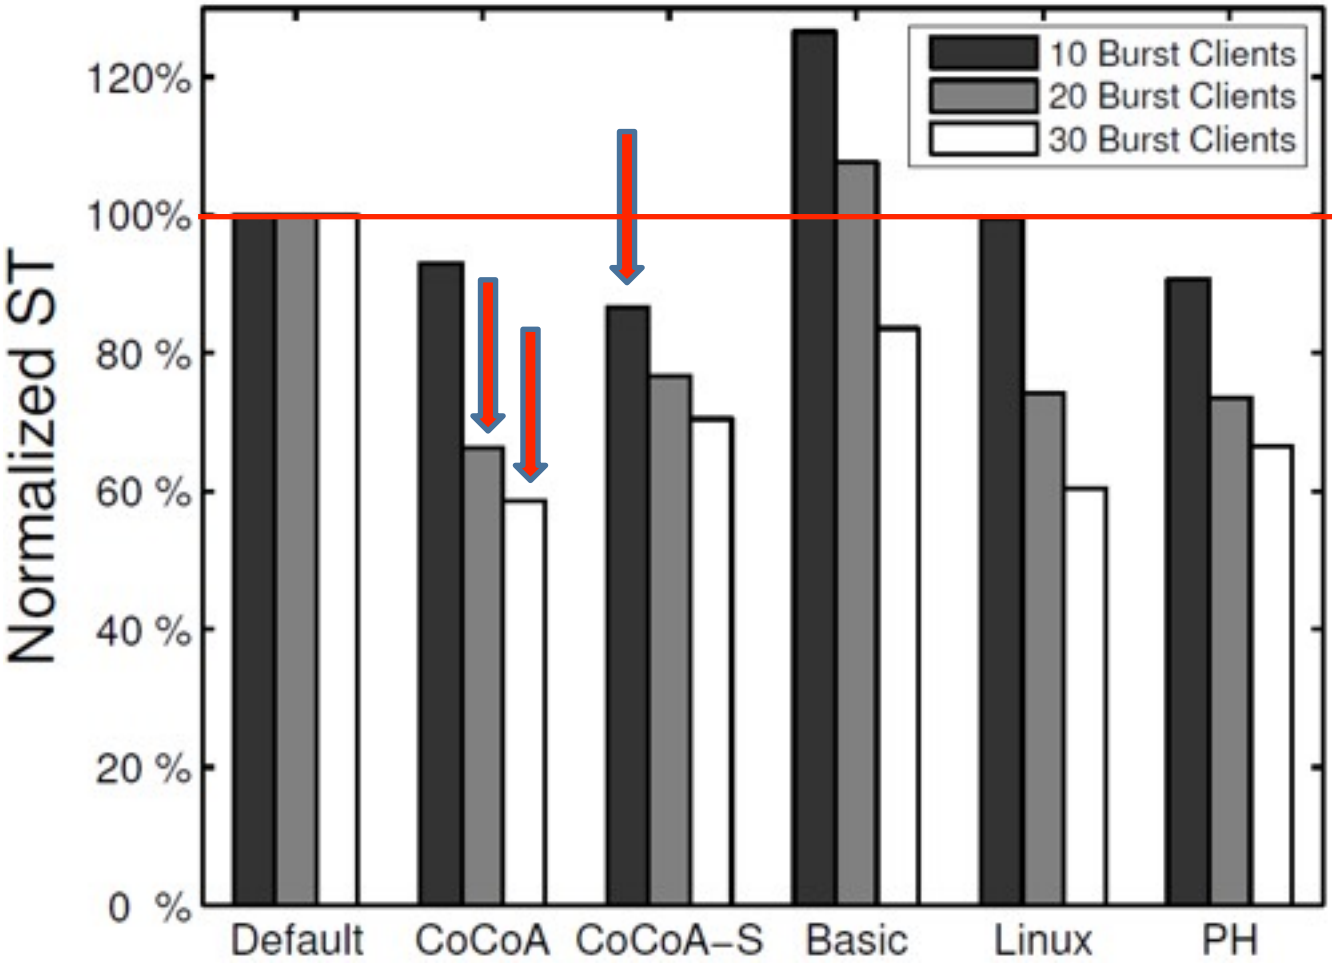

Observations (I)

- CoCoA outperforms default CoAP
  - Good use of RTT samples
  - Higher amount of work done (16-18%)
  - Retry ratio decrease by a factor up to 10
  - Settling time decrease by up to 40% after a request burst
- CoCoA-S
  - Often good performance in congested scenarios (vs default CoAP)
    - But high number of retries in low congestion scenario!
  - A bit less conservative than CoCoA (no weak RTTs/ RTO)

Observations (II)

- Too simplistic RTT-sensitive approaches underperform default CoAP
  - Basic RTO considers only the last RTT sample
  - Not enough safety margin (RTO vs actual RTT)
  - Huge amount of (too early) retries
- TCP-oriented RTO algorithms generally outperform default CoAP
  - Good use of RTT samples
- CoCoA performs similarly or slightly better than TCP-oriented designs
  - Trade-off: conservative vs aggressive
  - Differences include the VBF, contribution of RTT drop to the RTO, weak RTTs and dithering

Future work

- Extend prior study on the Flocklab testbed
  - CoCoA shown to outperform default CoAP
    - 30 Tmote Sky motes
    - Contiki OS, Erbium CoAP implementation
    - ContikiMAC, Null RDC
  - Include alternative RTO algorithms
    - Basic RTO, Linux RTO, Peak Hopper

Call to Action

- Please experiment with Californium+cocoa-02
  - Publicly available
- Please implement cocoa-02
- Please provide feedback before IETF'92

~~Tuesday~~ Wednesday

All times are in time-warped HST

- ~~15:20–15:30 Intro~~
- ~~15:30–15:45 HTTP mapping (SL)~~
- ~~15:45–16:25 Resource Directory (ZS)~~
- ~~16:25–16:45 ACE for Resource Directory (BG)~~
- ~~16:45–16:47 Links-JSON (chairs)~~
- ~~16:47–16:49 Core-Interfaces (chairs)~~
- ~~16:49–17:04 alt trans: DTLS on SMS (HT)~~
- **17:04–17:20 CoAP-PubSub (MK)**

CoAP-PubSub: Publish-Subscribe Extensions For CoAP

Extensions to enable a publish-subscribe interaction model between CoAP endpoints and CoAP services with asynchronous notifications and supporting sleeping and partially reachable endpoints

Earlier Draft was called CoAP-MQ – New Name, Same Goals and Objectives

- Provide support for very simple nodes, battery powered and energy harvesting nodes, sleeping and partially reachable endpoints
- Cover the use case requirements for Mirror Server and Sleepy Node drafts
- Simplify and Clarify material from the previous draft (work in progress)
- The new name is more descriptive of the main idea in this draft
- Publish/Subscribe pattern is a good mechanism to support simple nodes and sleeping endpoints

Terminology

- Topic – a string, similar to a path, that uniquely identifies an item being subscribed or published
- CoAP-PubSub Broker – A server node that stores information published to it, referenced by topic, and published said information to all subscribed entities
- PubSub Client Endpoint – an endpoint node that uses CoAP-PubSub to subscribe to and publish items to a broker

Design Overview

- Uses CoRE Resource Directory to register PubSub endpoints
- Uses new core.pubsub server attribute and core.pubsub registration parameters
- Topics are analogous to resource paths
- CoAP-PubSub Broker becomes the origin server
- PubSub endpoint nodes use only client-initiated transactions

Architecture

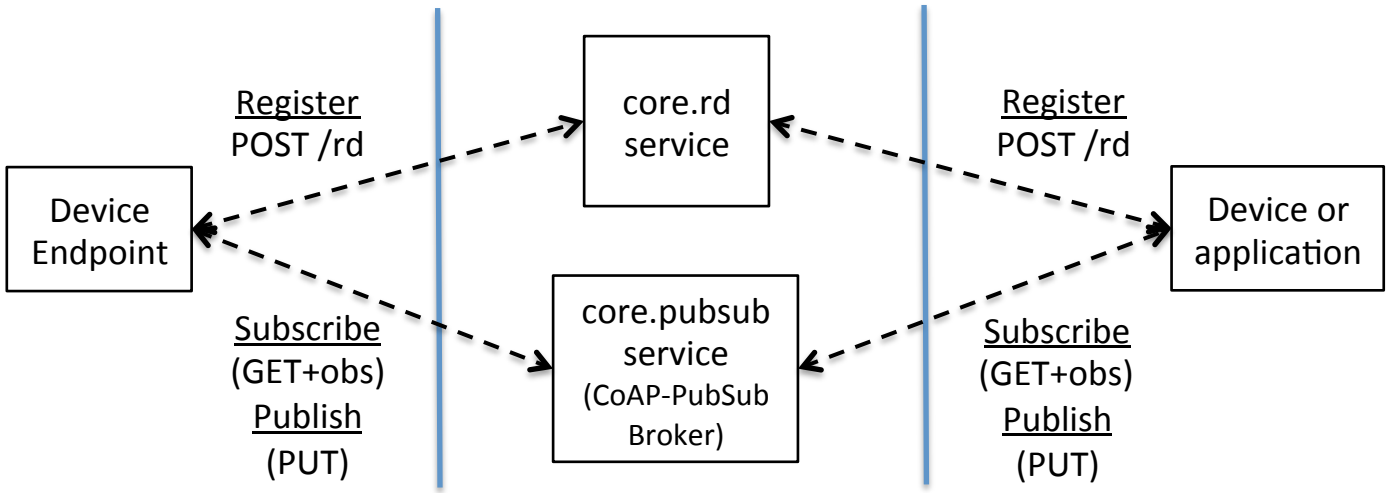

CoAP-PubSub Broker Discovery and Client Registration Using Resource Directory Service

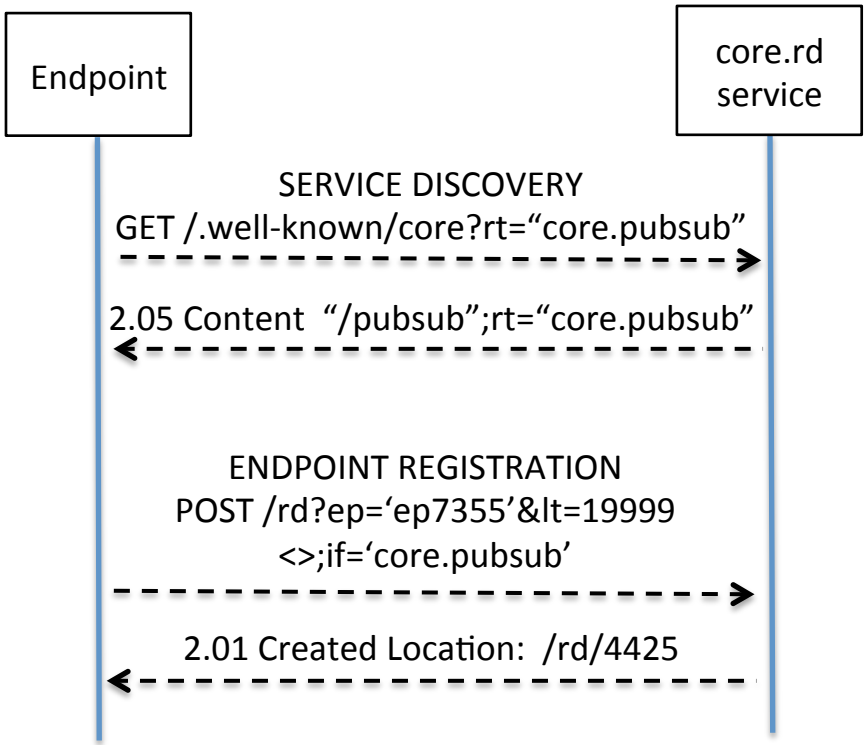

Client endpoint Publishes To CoAP-PubSub Broker

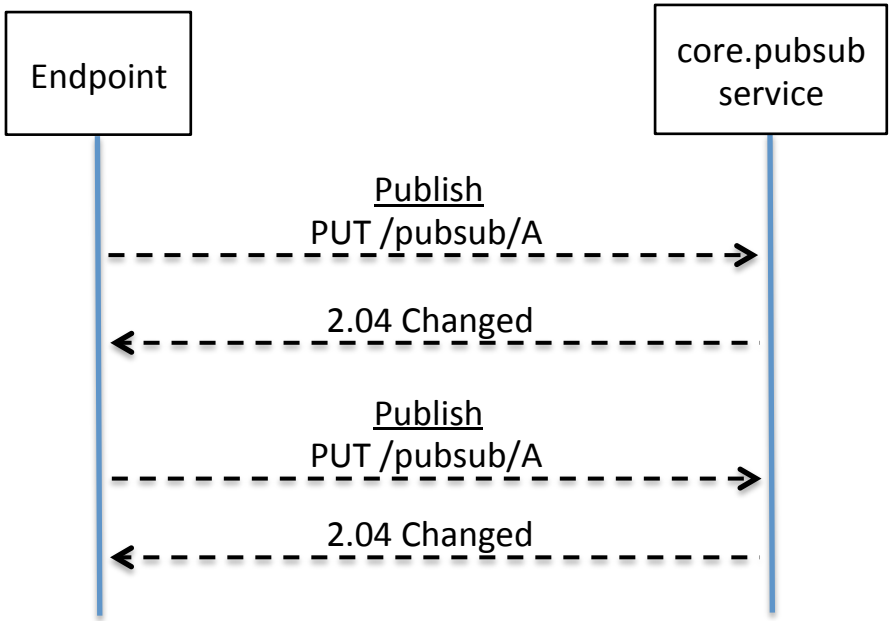

Sleeping Client endpoint Publishes To CoAP-PubSub Broker

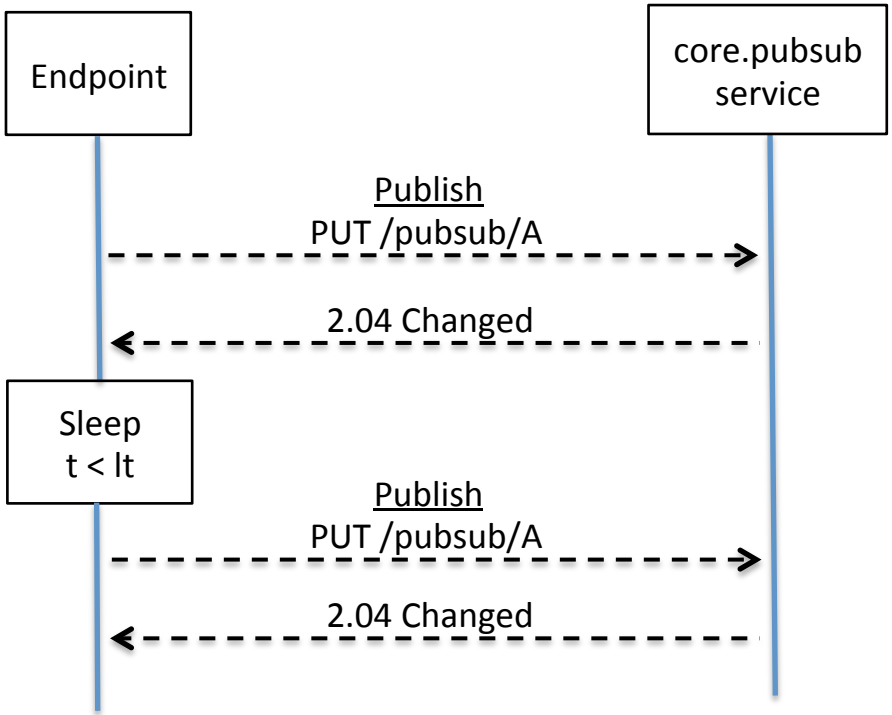

Client Endpoint Subscribes To CoAP-PubSub, Broker Publishes To EP

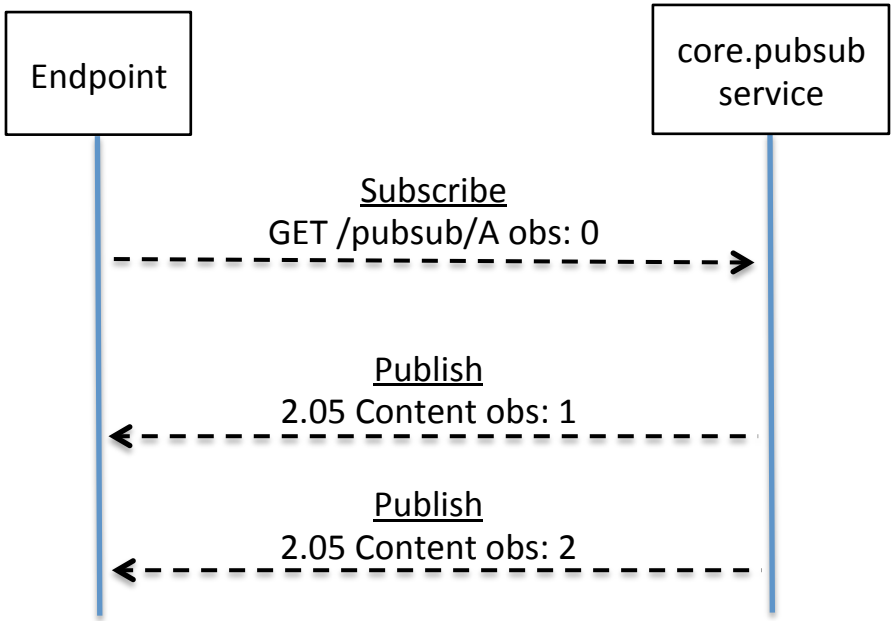

Client Endpoint Subscribes To CoAP-PubSub, Broker Publishes To Sleeping EP

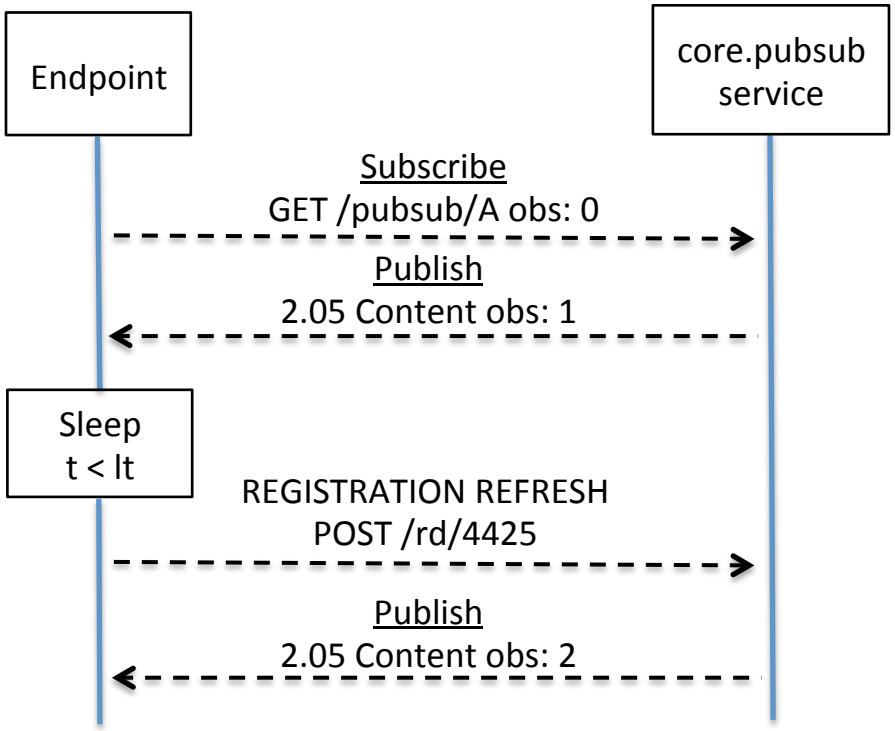

Multiple Publishers and Subscribers(2)

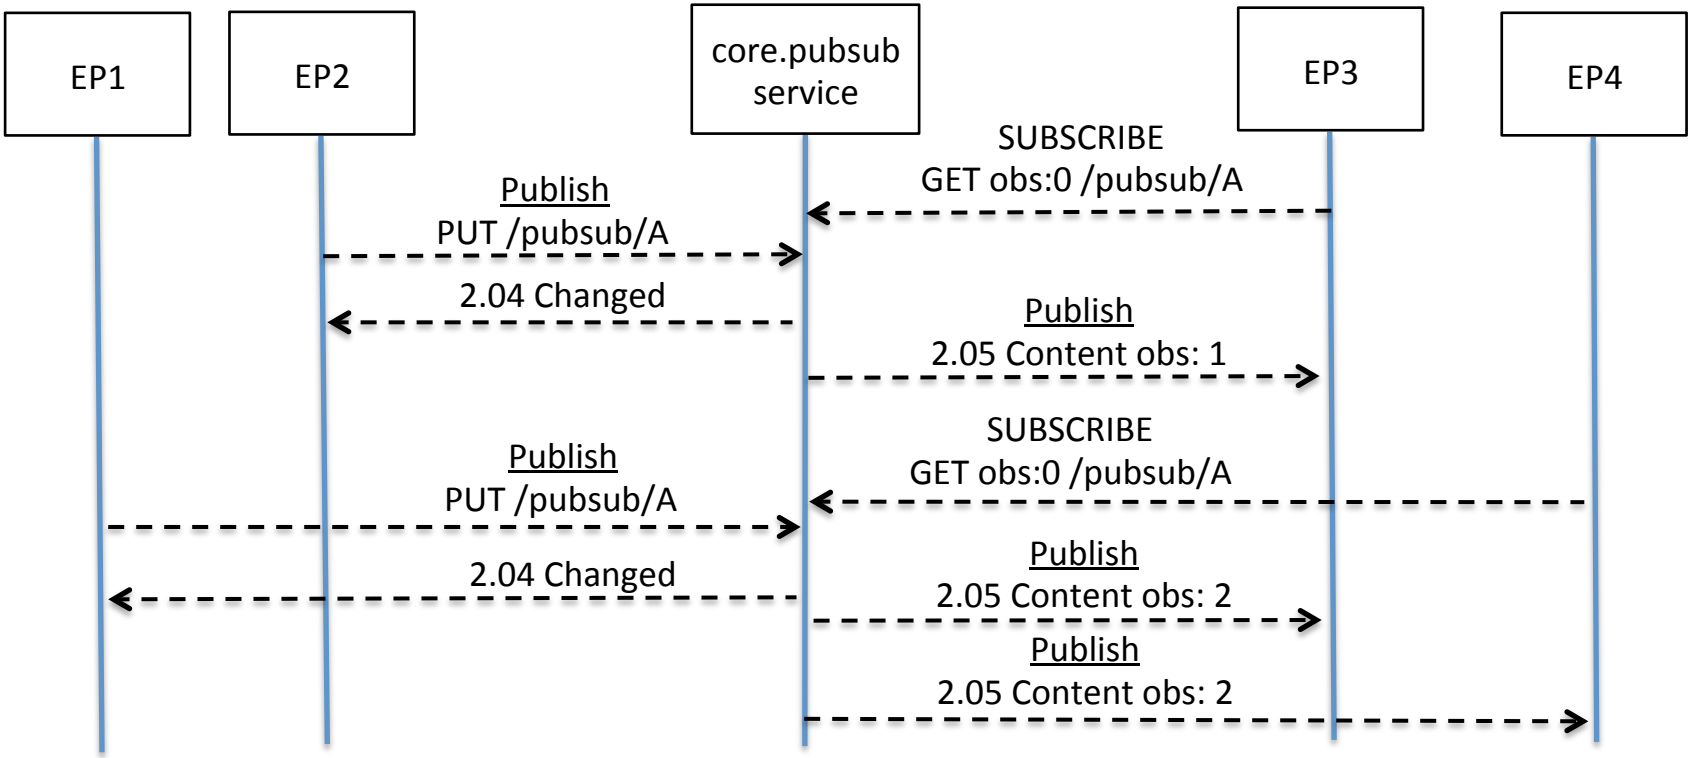

Design Overview and Issues

- Topic creation is separate from endpoint registration
- Topics are created upon first publication
- Subscriptions to non-existent topics may be rejected
- Unused, unsubscribed topics may be removed
- Rules for topic construction and wildcards needed
- CoAP-PubSub can act as a Pub-Sub to REST bridge

Next Steps

- Update the draft to reflect the consensus on the changes
- Define security model and considerations

Wednesday

- **09:00–09:03 Intro** All times are in time-warped HST
- **09:03–09:45 CoRE Management (PV)**
- **09:45–10:15 alt trans, continued (KL)**
- **10:15–10:25 No-Response (AB)**
- **10:25–10:45 endpoint IDs (OK, KL, ...)**
- **10:45–10:55 patience (KL)**
- **10:55–11:15 Congestion Control (CG)**
- **11:15–11:30 Flextime**

Flextime

http://trac.tools.ietf.org/wg/core/trac/wiki/CoreBacklog

core

Search

Login | About Trac | Preferences | Help/Guide

Wiki | Timeline | Roadmap | Browse Source | View Tickets | Search

wiki: CoreBacklog

Start Page | Index | History

This page maintains a backlog of work that the WG has identified as highest priority to work on next.

Work Item	Priority	Status	Related Work
Observe	High	IESG	draft-ietf-core-observe
Block	High	WG Document	draft-ietf-core-block
Resource Director	High	WG Document	draft-ietf-core-resource-directory
CoAP over TCP	High		draft-bormann-core-coap-tcp, draft-tschofenig-core-coap-tcp-tls, draft-silverajan-core-coap-alternative-transports
JSON Links	Normal	WG Document	draft-ietf-core-links-json
HTTP Mapping	Normal	WG Document	draft-ietf-core-http-mapping
SenML	Normal		draft-jennings-senml
CoRE Interfaces	Normal	WG Document	draft-ietf-core-interfaces
CoAP Management	Normal		draft-vanderstok-core-comi
CoAP Timeout Estimation	Low		draft-bormann-core-cocoa
CoAP Pub Sub	Low		draft-koster-core-coap-pubsub
CBOR Links	Low		draft-li-core-links-cbor

The following priority levels are used:

- High: This work item is high priority, and should be the next to try and progress through the WG
- Normal: This work item is normal priority
- Low: This work item is low priority, and would be nice to have, but should wait until higher priority work is complete