IETF93 - Prague

IPv6 over the TSCH mode of IEEE 802.15.4e

Chairs:
Pascal Thubert
Thomas Watteyne
Mailing list:
6tisch@ietf.org
Jabber:
6tisch@jabber.ietf.org
Etherpad for minutes:
http://etherpad.tools.ietf.org:9000/p/notes-ietf-93-6tisch

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.

Reminder:

Minutes are taken *
This meeting is recorded **
Presence is logged ***

* Scribe: please contribute online to the minutes at
http://etherpad.tools.ietf.org:9000/p/notes-ietf-91-6tisch
** Recordings and Minutes are public and may be subject to discovery in the event of litigation.
*** Please make sure you sign the blue sheets

Administrivia

- Blue Sheets
- Scribes
- Jabber

Agenda

Intro and Status (Chairs) [5min]
* Note-Well, Blue Sheets, Scribes, Agenda Bashing
* Drafts progression vs. plan

[13.05]
6TiSCH Plugtests report [15min]
* Organization report (Miguel Angel Reina Ortega)
* Tools:
  * <draft-munoz-6tisch-minimal-examples> (Dominique Barthel)
  * Golden Device (Tengfei Chang)
* Summary test cases, anonymized success highlights (Maria Rita Palattella)
* Lesson learned (Thomas Watteyne)

[13.20]
Hackathon report [5min]

[13.25]
<draft-ietf-6tisch-architecture-08> [15min]
* INT-DIR review and resolutions (Pascal Thubert)

Agenda

[13.40]
<draft-ietf-6tisch-minimal-10>
[30min]
* RPL artifacts and 6LoRH at 6lo (Gabriel Montenegro, 5min)
* Example format and security section (Xavi Vilajosana, 20min)
* Shepherd status and IESG submission (Pascal Thubert, 5min)

[14.10]
<draft-ietf-6tisch-6top-interface-04>
[10min]
* Update on changes (Qin Wang, Xavi alt.)
* Readiness assessment (all)

[14.20]
CoMI News [15min]
* <draft-vanderstok-core-comi-07> and
<draft-vanderstok-core-patch-01> 6TiSCH context (Peter van der Stok, 10min)
* Michel's Proposal to avoid hash collisions (Michel Veillette, 5min)

Agenda

[14.35]
Distributed scheduling [10min]
* <draft-dujovne-6tisch-on-the-fly-06> (Diego Dujovne)
* <draft-wang-6tisch-6top-coapie-01> (Qin Wang)

[14.45]
DetNet and dependencies [10min]
* BoF news (Lou Berger)
* <draft-wang-6tisch-track-use-cases-01> (Chonggang Wang)
* <draft-thubert-6tisch-4detnet-01> (Pascal Thubert)

[14.55]
Re-chartering kickstart and AOB (chairs) [5min]

Draft Progress

Draft	Status	Exp Milestone
Minimal	IESG after plugtest	IESG by Nov 2014
TSCH	RFC 7554	Pub Q by June 2015
Architecture	IESG review	IESG by Dec 2014
Terminology	Stable	Dec 2014
COAP IE	Dependency on CoMi	IESG by Nov 2014
6top Interface	Completeness feedback	IESG by Nov 2014

Plugtest/hackathon report

Agenda (20mn)

- 1st 6TiSCH Plugtests report
- draft-munoz-6tisch-minimal-examples-00
- OpenWSN/6TiSCH Hackathon

1st 6TiSCH Plugtests report

Miguel Angel Reina Ortega
Centre for Testing and Interoperability (CTI)
European Telecommunications Standards Institute (ETSI)

Agenda

- Overview of the event
- Plugtests agenda
- Test Plan and Tool
- Result reporting (ETSI Test Session Report tool)
- Test Cases
- High-Level Test Results
- Global Results for Test Cases
- Conclusions
- Roadmap

Overview of the Event

Event organized by:
- ETSI (European Telecommunications Standards Institute)

With support from:
- OpenMote

Participants:
- 21 participants
- 12 participating companies
- 3 observer company (4 people)
- 4 independent implementations

Tests:
- 23 total test pairings, each 1:30 hours in duration
- Single-hop configuration (12 test cases)
- Multi-hop configuration (8 teste cases), optional

Preparation calls (2)
- Organized and lead by ETSI and Experts group
- Including all Vendor Participants

Plugtests Agenda

	6TiSCH Plugtests Agenda - JULY 2015		
Time	Friday 17	Saturday 18	Sunday 19
08:30	Registration & Set up		
08:30 09:00	6TiSCH event Presentations	Room Opening	Room Opening
09:00 13:00	TEST SESSIONS	TEST SESSIONS	HACKATHON
13:00 14:00	NETWORKING LUNCH	NETWORKING LUNCH	NETWORKING LUNCH
14:00 18:00	TEST SESSIONS	TEST SESSIONS	HACKATHON
18:00 18:30		WRAP UP SESSION	Tear Down

Test Plan and Tools

- Developed by “Expert Group”: Xavier Vilajosana, Maria Rita Palattella, Tengfei Chang, Thomas Watteyne

- Test Plan
  - Covers IEEE802.15.4e/6TiSCH in detail
    - Synchronization, Header formats, later-2 security
  - Touches upon 6LoWPAN, RPL, ICMPv6
  - 18 test cases, including 10 Single Hop + 8 Multi Hop.
  - *(details in next slide)*

- Tools
  - Hardware: each participant has received an OpenMote kit
  - Firmware: “Golden Image” written using OpenWSN
  - Wireshark dissector *(Orange Labs)*

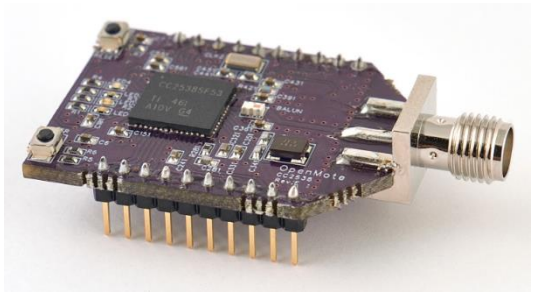

Result Reporting

- The results of each interoperability test session have been recorded in a dedicated web application software: the ETSI Test Report Tool (TRT)
  - After each test execution the interoperability result is agreed among all participants and then recorded
  - After each test session the report is submitted to ETSI

Single Hop Test List

• TD_6TiSCH_SYN_01	Check that a 6N can synchronize to the EB sent by the DR.
• TD_6TiSCH_SYN_02	Check that a 6N can synchronize to DR using KA messages.
• TD_6TiSCH_SYN_03	Check that a 6N's clock drifts if there is no re-synchronization.
• TD_6TiSCH_SYN_04	Check that the 6N can recover synchronization after de-synchronization.
• TD_6TiSCH_FORMAT_01	Check the format of the IEEE802.15.4e EB packet is correctly assembled.
• TD_6TiSCH_FORMAT_02	Check the timing template of TSCH time slot defined in draft-ietf-6tisch-minimal-11is correctly implemented.
• TD_6TiSCH_FORMAT_03	Check channel hopping is correctly implemented according to draft-ietf-6tisch-minimal-11.
• TD_6TiSCH_FORMAT_04	Check the number of retransmissions is implemented following draft-ietf-6tisch-minimal-11.
• TD_6TiSCH_FORMAT_05	Check the minimal schedule is implemented according to draft-ietf-6tisch-minimal-11.
• TD_6TiSCH_FORMAT_06	Check the 6N sets its slotframe size correctly when joining the network.
• TD_6TiSCH_SEC_01	Check the 6N is correctly authenticated with K1, when it synchronizes to DR with EB.
• TD_6TiSCH_SEC_02	Check the data packet sent by 6N is correctly encrypted with K2.

Multi Hop Test List

TD_6TiSCH_RPL_01	Check the value of EB join priority of a child 6N and a parent DR.
TD_6TiSCH_RPL_02	Check the rank of 6N is computed correctly according to draft-ietf-6tisch-minimal-11.
TD_6TiSCH_RPL_03	Check a 6N child changes its time source neighbor (parent) correctly.
TD_6TiSCH_RPL_04	Check the format of RPL DIO message.
TD_6TiSCH_RPL_05	Check the format of RPL DAO message.
TD_6TiSCH_RPL_06	Check IP extension header in 6LoWPAN.

High level test results

Interoperability			Not Executed		Totals	
OK	NO		NA/OT		Run	Results
207 (93.7%)	14 (6.3%)		58 (20.8%)		221 (79.2%)	279

On the 279 mandatory test cases planned, 207 have been executed and 58 not executed.

On the 221 Mandatory TC performed, 207 have been OK, which represents a success rate of 93.7%.`

Global Results per tests

	Interoperability		Not Executed	Totals	
	OK	NO	NA/OT	Run	Results
TD_6TiSCH_SYN_01	18 (85.7%)	3 (14.3%)	0 (0.0%)	21 (100.0%)	21
TD_6TiSCH_SYN_02	16 (94.1%)	1 (5.9%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_SYN_03	17 (100.0%)	0 (0.0%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_SYN_04	16 (94.1%)	1 (5.9%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_FORMAT_01	17 (100.0%)	0 (0.0%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_FORMAT_02	16 (94.1%)	1 (5.9%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_FORMAT_03	16 (100.0%)	0 (0.0%)	3 (15.8%)	16 (84.2%)	19
TD_6TiSCH_FORMAT_04	17 (100.0%)	0 (0.0%)	2 (10.5%)	17 (89.5%)	19
TD_6TiSCH_FORMAT_05	14 (100.0%)	0 (0.0%)	5 (26.3%)	14 (73.7%)	19
TD_6TiSCH_FORMAT_06	14 (100.0%)	0 (0.0%)	3 (17.6%)	14 (82.4%)	17
TD_6TiSCH_RPL_01	4 (50.0%)	4 (50.0%)	7 (46.7%)	8 (53.3%)	15
TD_6TiSCH_RPL_02	2 (100.0%)	0 (0.0%)	1 (33.3%)	2 (66.7%)	3
TD_6TiSCH_RPL_03	0 (0.0%)	0 (0.0%)	3 (100.0%)	0 (0.0%)	3
TD_6TiSCH_RPL_04	15 (100.0%)	0 (0.0%)	3 (16.7%)	15 (83.3%)	18
TD_6TiSCH_RPL_05	7 (100.0%)	0 (0.0%)	6 (46.2%)	7 (53.8%)	13
TD_6TiSCH_RPL_06	0 (0.0%)	0 (0.0%)	3 (100.0%)	0 (0.0%)	3
TD_6TiSCH_SEC_01	10 (83.3%)	2 (16.7%)	6 (33.3%)	12 (66.7%)	18
TD_6TiSCH_SEC_02	8 (80.0%)	2 (20.0%)	6 (37.5%)	10 (62.5%)	16

Conclusions

- The event has been successful and gave good interoperability results , especially for the first interop event on this technology.
  - 221 testcases performed with 93.7 % success rate
- Running code is the only way of working out all details
- Full IEEE802.15.4e TSCH synchronization between all implementations!!
- The participants were satisfied and gave very good feedback in the satisfaction survey
- The number of participants allowed that all the vendors met each other in test sessions during the event
- Dissector and golden image are essential tools

Roadmap for next 6TiSCH Plugtests

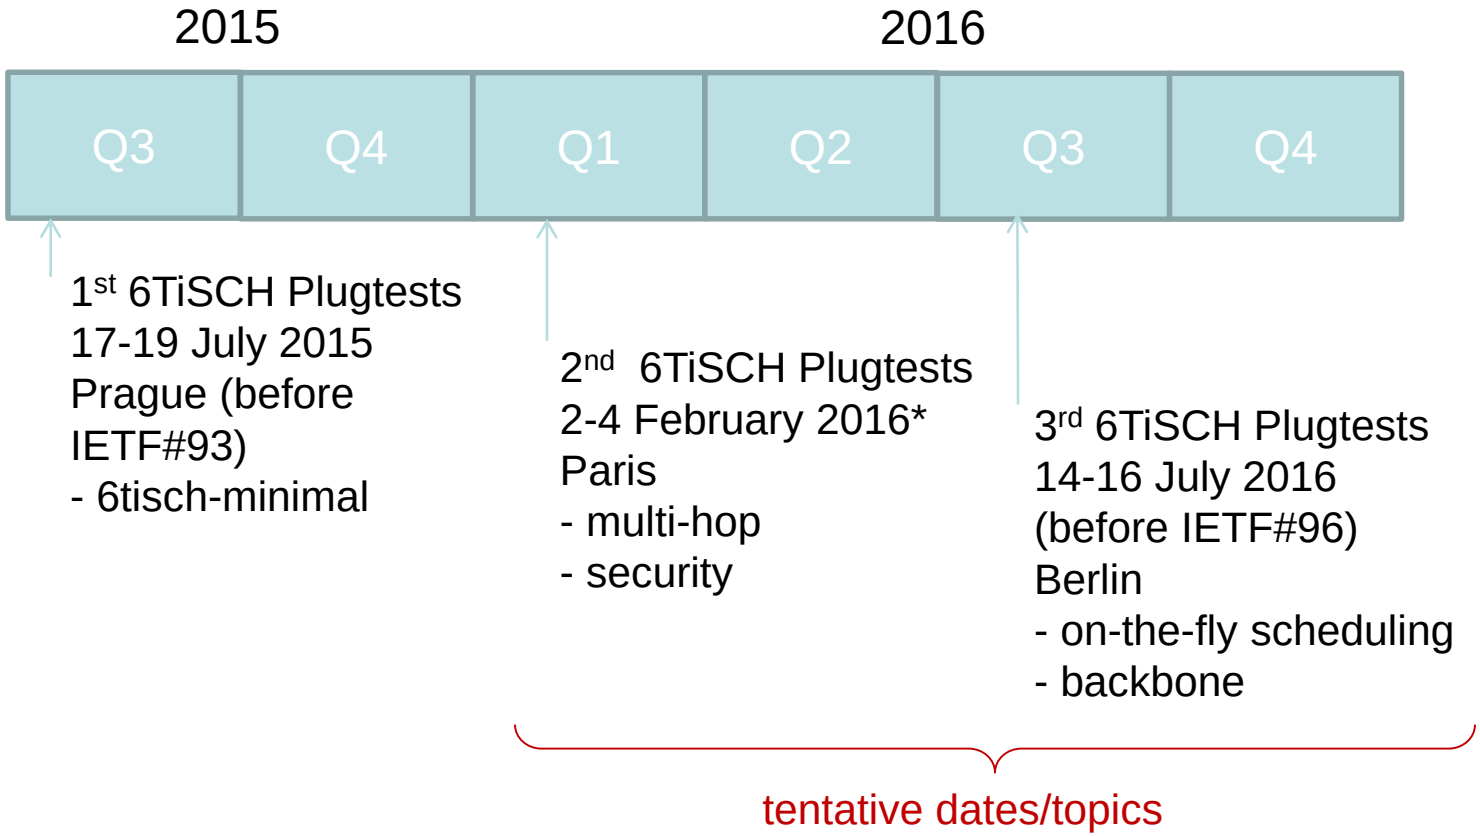

ETSI

World Class Standards

THANK YOU!

Miguel Angel Reina Ortega
Centre for Testing and Interoperability (CTI)
MiguelAngel.ReinaOrtega@etsi.org

Plugtest Feedback to WG

- Lots of clarifications during the preparation process:
  - 3 updates to `draft-ietf-6tisch-minimal` (-09, -10, -11)
  - publication of `draft-munoz-6tisch-minimal-examples`

Question

- Current text states:

> For downstream route maintenance, in a minimal configuration, RPL **SHOULD** be set to operate in the Non-Storing mode as described by [RFC6550] Section 9.7. Storing mode ([RFC6550] Section 9.8) **MAY** be supported in less constrained devices.

Should there be a MUST somewhere?

draft-munoz-6tisch-minimal-examples-00

Jonathan Munoz (Ed.)
Guillaume Gaillard
Dominique Barthel (dominique.barthel@orange.com)

Wireshark output for 6TiSCH-minimal

- we updated the Wireshark IEEE802.15.4 dissector with IEEE802.15.4e TSCH
- ran on OpenWSN implementation of 6TiSCH-minimal

```
Network prefix: bbbb::/64
MAC address:    14-15-92-cc-00-00-00-0x

        PDR=100%        PDR=100%
 +-----+         +-----+         +-----+
 | x=1 |---------| x=2 |---------| x=3 |
 +-----+         +-----+         +-----+
 DAGroot
```

  - EB, DIO, DAO, ACK, Echo requests/replies
- copy-pasted Wireshark output in draft

Next Steps / Questions

- Intended as educational/reference document
  - To see example frames with all protocols
  - For implementers of 6TiSCH-minimal
- Work-in-progress
  - Comments welcome! Many details, we expect lots of comments
- Questions
  - Useful?
  - Should be published? How?

OpenWSN/6TiSCH Hackathon

Prague, 19th July 2015

Thanks to:

OpenWSN/6TiSCH Technologies

1. **Distributed blacklisting for improving FHSS**
   - Pedro Henrique Gomes (Univ. Southern California)
2. **6top-to-6top 6TiSCH negotiation in OpenWSN**
   - Tengfei Chang, Qin Wang (Univ. Sc. Techno. Beijing)
3. **OpenWSN on the IoT-lab**
   - Nicola Accettura (UC Berkeley)
4. **Contiki 6TiSCH implementation and hardcell allocation**
   - Sedat Gormus, YiChao Jin
5. **IEEE802.15.4e/6TiSCH dissectors**
   - Jonathan Munoz, Guillaume Gaillard, Dominique Barthel (Orange Labs)
6. **Node Monitoring framework**
   - Dominique Barthel (Orange Labs)
7. **6TiSCH Scheduler-free prototype in Contiki**
   - Simon Duquennoy (SICS)
8. **Dust Networks/Linear Technology's SmartMesh IP**
   - Thomas Watteyne (Linear Technology/Dust)
9. **uPnP: Automatic recognition of connected sensors**
   - Prof. Danny Hughes (KU Leuven)
10. **Flexible HW/SW CCM* security implementations in OpenWSN**
    - Malisa Vucinic (ST Micro)
11. **6TiSCH layer-2 security implementation in OpenWSN**
    - Savio Sciancalepore, Giuseppe Piro (U. Bari)





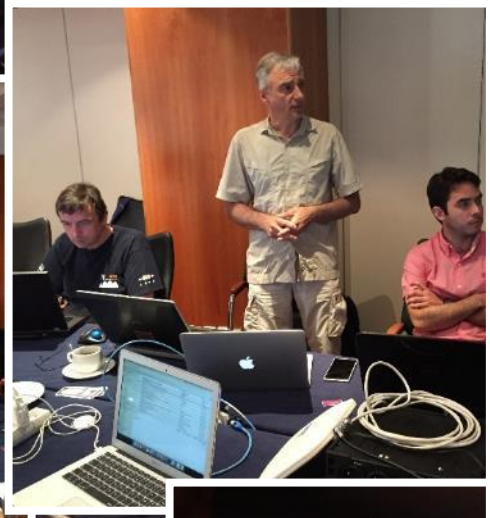



NSMN
OPENWSN
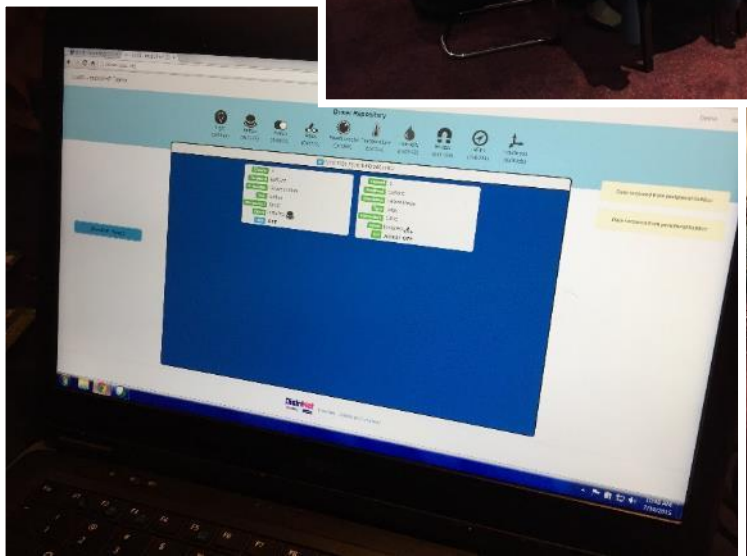


Prizes

- Best hackathon project
  - First place: 1 OpenMote Bronze kit + Raspberry Pi
    - Tengfei Chang
      "FreeRTOS/OpenWSN integration"
  - Second place: 1 OpenMote Bronze kit + sensor
    - Savio Sciancalepore
      "link-layer security OpenMote-CC2530 / TelosB interop"
- Best presented project
  - 1 OpenMote Bronze kit
    - Dominique Barthel, Quentin Lampin
      "*Node Monitoring framework"*

Special THANKs to

ETSI World Class Standards

openmote openhardwarefortheinternetofthings

Inria INVENTORS FOR THE DIGITAL WORLD

draft-ietf-6tisch-architecture-08

MANY Authors; Pascal Thubert, Cisco, Editor

Status

- INT AREA review by Ralph Droms
- Complete review published on the ML
- Main problem raised:
  - Mid-level Architecture incomplete, misses:
    - Security,
    - Dynamic scheduling (OTF..) and
    - DetNet applied to 6TiSCH networks
  - Deeper dive on chartered item
    - Inconsistent level in the document

Recommendations

- “In my opinion, it would be better to publish the complete mid-level architecture in one document, and the specific details of the components in subsequent documents as those additional details are developed. Those subsequent documents might be that actual protocol specifications or the system specifications that describe how the various components use IETF and other standards (something like the CableLabs DOCSIS spec).”

=> Reopen Archie till the whole mid-level architecture is complete

=> Publish separately the deeper dive work

  – Architecture elements?

Example: Section 6

6. 6LoWPAN (and RPL)
    - 6.1. RPL Leaf Support in 6LoWPAN ND
    - 6.2. registration Failures Due to Movement .
    - 6.3. Proxy registration
    - 6.4. Target Registration
    - 6.5. RPL root vs. 6LBR
    - 6.6. Securing the Registration

“. Section 6, on the other hand, gives specific design details that would be better expressed in a design or specification document. Similarly, section 10 specifies the current, preliminary design for the join process, rather than an architecture for security that describes all of the required security functions and how they relate to each other.”

Minimal 6TiSCH Configuration draft-ietf-6tisch-minimal-11

Xavier Vilajosana Universitat Oberta de Catalunya
Kris Pister University of California Berkeley

Status

Status:

- Adopted at IETF88 Vancouver
- Latest version (11) published on 7 July 2015 https://datatracker.ietf.org/doc/draft-ietf-6tisch-minimal/

News

- Security section wrapped up
- Packet examples section
- IEEE.802.15.4 Specific Header Fields and Considerations added

Next

- RPL in minimal, MUST? Should we not restrict 1hop networks?
  - MUST non-storing?
  - SHOULD storing?
  - Both might be needed for multihop settings (Outcome from Plugtest)
- Start topologies + RPL?

Security Section

Consensus on:

As this document refers to the interaction between Layer 3 and Layer 2 protocols, this interaction MUST be secured by L2 security mechanisms as defined by [IEEE802154]. Two security mechanisms are considered, authentication and encryption, authentication applies to all packet content while encryption applies to header IEs and MAC payload. Key distribution is out of scope of this document, but examples include pre-configured keys at the nodes, shared keys among peers or well-known keys. Refer to the 6TiSCH architecture document [I-D.ietf-6tisch-architecture] for further details on key distribution and advanced security aspects.

The present document assumes the existence of two cryptographic keys, which can be pre-configured. One of the keys (K1) is used to authenticate EBs. As defined in Section 4, EBs MUST be authenticated, with no payload encryption. This facilitates logical segregation of distinct networks. A second key (K2) is used to authenticate DATA, ACKNOWLEDGEMENT, MAC COMMAND frame types and respective header IEs, with payload encryption. Depending on security policy, these keys could be the same (i.e., K1=K2).

For early interoperability, K1 MAY be set to 36 54 69 53 43 48 20 6D 69 6E 69 6D 61 6C 31 35 ("6TiSCH minimal15").

Examples Section

4 Examples:

- Information Elements in EBs
- Information Elements in EBs with non-default TSCH Timeslot IE
- ACK/NACK Time Correction IE

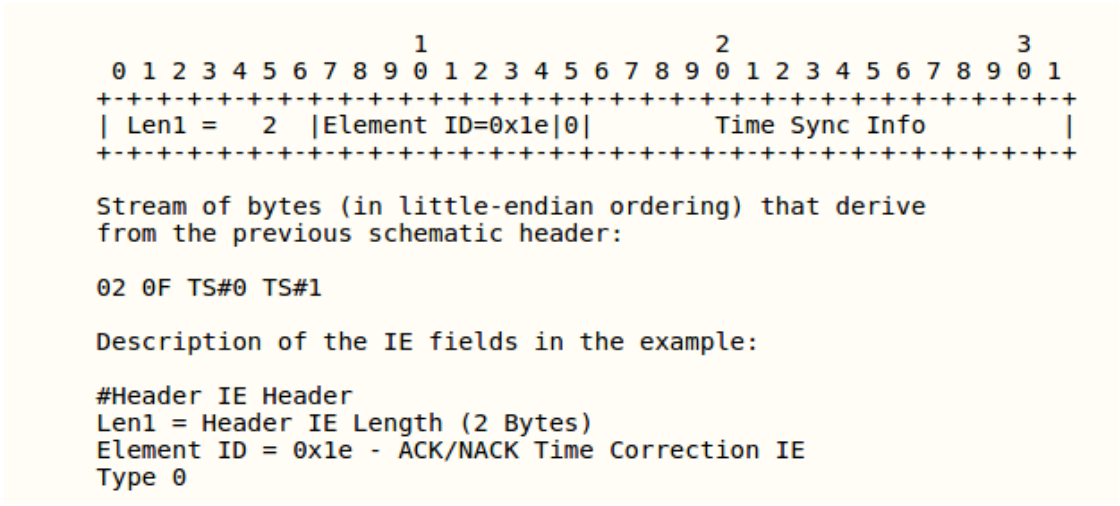

- Auxiliary Security Header

IEEE.802.15.4 Specific Header Fields

The IEEE802.15.4 header of all frames MUST include the Sequence Number field, the Source Address field and the Destination Address field. In the Frame Control Field, this translates to:

- The Frame Version field MUST be set to 0b10 (Frame Version 2)
- The Sequence Number Suppression bit MUST be set to 0b0
- The Source Addressing Mode MUST set to 0b11 (long address)
- The Destination Addressing Mode MUST set to 0b11 (long address) except for the broadcast address for which Destination Addressing Mode SHOULD set to 0b10 (short address). The use of long addresses is a REQUIRED as no association procedure is defined in this document.
- The PAN ID Compression bit MUST be set to 0b0. According to the Table 2a in [IEEE802154-2012], this translates into the Destination PAN ID field being "Present" and the Source PAN ID field being "Not Present".

Discussion

RPL 0F0

- Step of rank computation is 2*ETX.
- Proposed change to 2*ETX + 3

RPL in minimal

“Nodes in the network MUST use the RPL routing protocol [RFC6550] and implement the RPL Objective Function Zero [RFC6552].”

- MUST? What about star networks?

10.2.1. Mode of Operation

For downstream route maintenance, in a minimal configuration, RPL SHOULD be set to operate in the Non-Storing mode as described by [RFC6550] Section 9.7. Storing mode ([RFC6550] Section 9.8) MAY be supported in less constrained devices.

- *SHOULD and MAY?*

Thanks

Xavier Vilajosana
Universitat Oberta de Catalunya

Shepherd status and IESG submission

- Highly successful plugtest
- Proven interoperability
- Will submit to IESG ASAP
  - Waiting for final update after plugtest
  - And IPR confirmation from al authors

draft-ietf-6tisch-6top-interface-04

Qin Wang (Ed.)
Xavier Vilajosana

Status

- Adopted at IETF89
- Latest version published on 2015-07-06
  http://tools.ietf.org/html/draft-ietf-6tisch-6top-interface

Changes(1)

- Modify Container SecurityAttributes

```
+--rw container  SecurityAttributes
   +--rw leaf-list  K1*                          uint8, min-elements 16
   +--rw leaf       EBSecurityLevel              enumeration
   +--rw list       K2List*   [NodeAddress]
      +--rw  leaf      NodeAddress               nodeaddresstype
      +--rw leaf-list  K2*                       uint8, min-elements 16
      +--rw leaf       SecurityLevel             enumeration
```

Change (2)

- Re-order the attributes of YANG model

- Version
- SlotframeList
- CellList
- MonitoringStatusList
- StatisticsMetricsList
- EBList
- TimeSource
- NeighborList
- QueueList
- LabelSwitchList
- TrackList
- ChunkList
- ChunkCellList

- TSCHSpecificPIBAttributes
- TSCHmacTimeslotTemplate
- TSCHHoppingSequence
- SecurityAttributes

Change (3)

- Remove section 5 “Commands”
- Re-write section 4 “Generic Data Model”
- Typo corrections

Next Step

- Refine the r/w feature of each attribute
- Clarify mandatory/optional attributes
- Define a generic method to expose 15.4 PIB and related primitives to 6top users.
- Coordinate with http://tools.ietf.org/html/draft-wang-6tisch-6top-coapie-01.txt, and merge softcell negotiation RPC into YANG model.
- WG members review the YANG model

CoMi News

Agenda (15mn)

- draft-vanderstok-core-comi and draft-vanderstok-core-patch (Peter, 10mn)
- Michel's Proposal to avoid hash collisions (Michel with Alex, 5mn)

draft-vanderstok-core-comi-07

P. Van der Stok Independent
A. Bierman Yuma works
A. Sehgal Independent
J. Schoenwalder Jacobs University

Status

- Provide access to management variables specified in YANG between “reduced resource” clients and servers.
- Discussion points
  - Hashing of names
  - Use of patch content format

Hashing, 1

There is a probability that a hash clash occurs inside a server.

The hash clash probability for a 32 bit hash is given by:

Number of names	30084	927	10
probability	0.1	10^{-4}	10^{-8}

Supposing 10 server types with 1000 names; probability of a clash: .001, independent of total number of servers.

Hashing, 2

To reduce impact of rehash handling,

One measure:

- Hash detection when clash is used
- Distinguish clashes by module name

Assumption: no clashes within a module
Needs table in client with hash to module_name, for all hashes

Get example.com/mg/hash_clash
Returns: (wild syntax)
Module1, new_hash_1
Module2, new_hash_2

Hashing, 3

Wish to introduce alternative name to number conversions.

The server resource /mg/num.type

returns name to numbering scheme

Default: “yanghash”

Patch format 1

Wish to modify a single entry in a YANG list, identified by its key value.

This is not supported by current JSON Patch formats: RFC 7386, and RFC 6902

Patch format, 2

Simplified example:

```
List CellList{
    Key “CellID”
    Leaf CellID {type uint16;}
    Leaf ChannelOffset { type uint16;}
    Leaf CelType{ type enumeration}
}
```

Use of hashes

```
    CellList          -> hash1
    CellID            -> hash2
    ChannelOffset     -> hash3
    CelType           -> hash4
}
```

Suppose CellID 512 identifies the list item, and Celtype is modified to HARD

Two solutions:

(1) Send whole item in fixed field order: hash1[512, 34, HARD]

(2) Modify the wanted field only:
hash1{ {hash2: 512} : {hash4: HARD} }

CoOL
COnstrained Objects Language

Michel Veillette, Alexander Pelov

What is CoOL

- COnstrained Objects Language
- CoMI extension based on structured IDs
  - Based on the concept of alternate numbering scheme introduced in CoMI version 7.
  - IDs assigned to data nodes are managed (registered) instead of unmanaged (YANG hash)
- No hash clashes
  - Avoid associated overhead and issues

Why CoOL (1/3)

Currently known YANG hash issues

- Access to an unimplemented data node may affect (get, put, delete) a different object. The CoMI client need to known or infer which objects are supported on each peer CoMI server

Why CoOL (2/3)

- Notification sent to a CoMI client may have a clash on the client side. The CoMI client need to known or infer the list of notifications supported by each peer CoMI server to resolve the notification ID.
- Rehash caused by the dynamic loading of a YANG module are not disseminated to CoMI client. Subsequent rehashed notifications will confuse CoMI clients.

Why CoOL (3/3)

- CoMI client footprint
  - CoMI clients need to retrieve and store the rehash table of each peer node (e.g. 100 nodes * 100 bytes)
  - CoMI clients need to retrieve and store information about the list data nodes and notifications of each peer node (e.g. list of modules implemented)
  - CoMI clients need known the data node path of each object accessed in order to use rehash tables (e.g. 100 objects * 100 bytes)
- Message payloads are in average 50% larger

Managed IDs

- Composed of two parts
  - 20 bits registered Module ID
  - 10 bits assigned YANG ID
- Long form vs short form
  - JSON qualified-member-name -> 30 bits Module ID | 10 bits YANG ID
  - JSON member-name -> 10 bits YANG ID
- IANA registration of Module ID
  - 1 048 576 Module IDs available to SDOs or manufacturers
  - 3/4 the IDs reserved for future use
- Assignment of YANG ID
  - Automatic based on the location in the YANG module or manually using a new YANG statement
  - ID 0 to 23 are encoded using a single byte, can be manually assigned to frequently used nodes

CoOL example

- Perform on a single resource (e.g. GET /cdat)
- *"select"* option contains the list of nodes selected, encoded using a CBOR array

Qualified

REQ: GET /mg?select([14337, 18, 19]) Token(0x324a)

CBOR Unqualified

RES: 2.05 Content Token(0x324a) (Content-Format: application/cbor)

```
{
  14337: 57,
  18 : 76,
  19 : 837
}
```

CBOR

CoOL also supports

- Update, Create, Delete
  Using CoAP PUT, POST, DELETE
- Patch
  Based on [I-D.ietf-netconf-yang-patch]
- Protocol operations (YANG rpc)
  Based on [I-D. ietf-netconf-restconf]
- Notification stream (YANG notification)
  Based on RFC 5277
- Reporting
  Based on [I-D.ietf-core-observe]
- Resource discovery based on YANG module(s)
  (e.g. ietf-yang-library, ietf-restconf-monitoring

For mode details

- Michel Veillette
  Michel.Veillette@trilliantinc.com
- Alexander Pelov
  alexander.pelov@telecom-bretagne.eu

Distributed Scheduling

Agenda (10mn)

- draft-dujovne-6tisch-on-the-fly
- draft-wang-6tisch-6top-coapie-01

draft-dujovne-6tisch-on-the-fly

6TiSCH On-the-Fly Scheduling

Diego Dujovne – Universidad Diego Portales
Luigi Alfredo Grieco – Politecnico di Bari
Maria Rita Palattella - University of Luxembourg
Nicola Accettura - University of California Berkeley

About OTF

- Distributed Scheduling with or without a PCE
- Event-triggered
- Parametrized Allocation Policy
- External interface using CoAP
- Currently on -06 version

When vs. How Many

- Establish the difference between the triggering time and the number of requested cells:
  - Bandwidth estimation algorithm defines WHEN (after triggered)
  - Allocation Policy defines HOW MANY
  - However, the Allocation Policy will not reserve/delete more cells than what it is requested by the Bandwidth estimation algorithm

Bundles

- On a Layer 3 link, there are two associated bundles: Incoming and Outgoing.
- OTF requests cells on the Outgoing bundle only (TX cells); 6top negotiates this request with the neighbor.

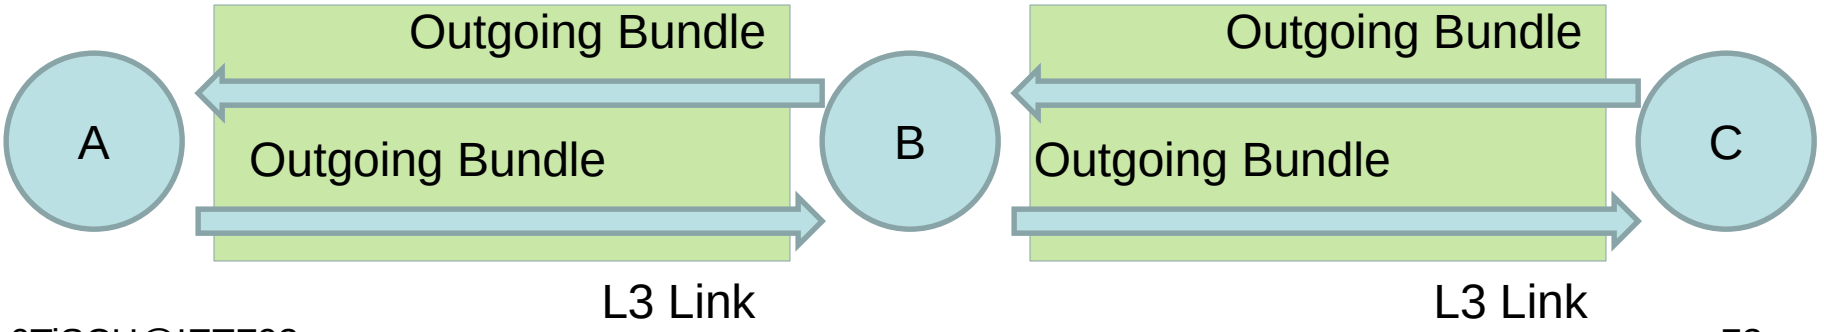

Default algorithm (I)

- **Step 1:**Collect the bandwidth requests from child nodes ***(incoming bundle soft cell allocation from 6top-to-6top negotiation).***
- **Step 2:**Collect the node bandwidth requirement from the application ***(self/local traffic, from the application soft cell pending requests).***
- **Step 3:**Collect the current outgoing scheduled bandwidth ***(outgoing traffic).***
- **Step 4**:If (outgoing < incoming + self) then SCHEDULE soft cells to satisfy bandwidth requirements.
- **Step 5:**If (outgoing > incoming + self) then DELETE the soft cells that are not used.
- **Step 6**:Return to step 1.

Default algorithm (II)

- *Defined sources of statistics*
- *Does not use the allocation policy:*The default bandwidth estimation algorithm adopts a “reactive allocation policy, i.e., it uses OTFTHRESHLOW = 0 and OTFTHRESHHIGH = 0;”
- *Defines the triggering:* “The algorithm is triggered either by Step 4 or Step 5.“

Next

- How to deal with L2 Tracks?
  - Not only best effort track.
- Chunk appropriation?
  - Restricts the number of available softcells with respect to the whole CDU
  - Use of shared softcells between chunks
- Change SCHEDULE to ADD on default algorithm

Questions?

draft-wang-6tisch-6top-coapie-01

Qin Wang (Ed.)
Xavi Vilajosana
Thomas Watteyne
Raghuram Sudhaakar
Pouria Zand

Status

- First version published at IETF90
- Latest version published on 2015-07-02 https://www.ietf.org/internet-drafts/draft-wang-6tisch-6top-coapie-01.txt

Changes

Format of a CoAP IE.

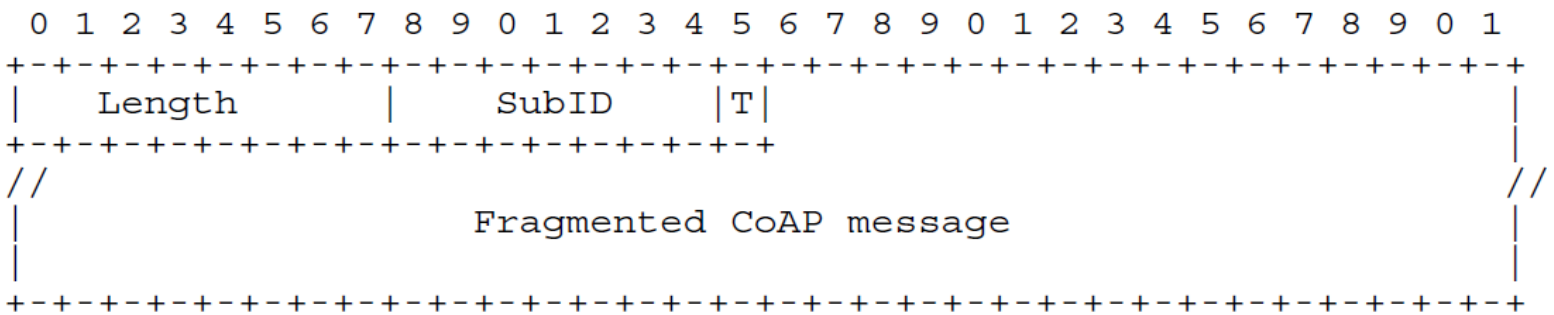

Octets: 1/2	0/1	0/2	0/1/2/8	0/2	0/1/2/8	0/1/5/6/1 0/14	variabl		variable	2
Frame Control	Sequence Number	Destination PAN Identifier	Destination Address	Source PAN Identifier	Source Address	Auxiliary Security Header	Informat Elemen		Frame Payload	FCS
		Addressing fields					Header IEs	Payload IEs		
MHR								MAC Payload		MFR

Change from "Header IE" to "Payload IE"

Next Step

- Merge RPC for softcell negotiation process into draft-ietf-6tisch-6top-interface, i.e. 6top YANG model

DetNet and dependencies

Agenda (10mn)

- BoF news (Pascal, Lou, 4min)
- draft-wang-6tisch-track-use-cases (Chonggang, 3min)
- draft-thubert-6tisch-4detnet (Pascal, 3min)

DetNet BoF news

Lou Berger

DetNet BoF Summary

- WG Forming BoF – Held Monday
- Very well attended
  - Filled Grand Hilton Ballroom
- Covered multiple use cases, problem statement, proposed WG scope
  - Professional audio, Electrical utilities, Industrial automation, Building automation, Radio/mobile access networks
- Normal polling
  - General support for problem statement and IETF work
  - Good support in room for contributing to a DetNet WG

Use Cases and Requirements for Using Track in 6TiSCH Networks

draft-wang-6tisch-track-use-cases-01

Zhuo Chen, Chonggang Wang

Status

- Status:
  - Latest version -01 published on 07.06.15
    available at: http://tools.ietf.org/html/draft-wang-6tisch-track-use-cases-01

Background

- Use Case – Industrial Networks
  - Industry Process Control and Automation Applications
  - Industrial Monitoring Applications

- Requirements for Track Reservation
  - Centralized Track Reservation
    - Need a protocol for LLN devices to report their topology and TSCH schedule information to the central controller.
    - Need a lightweight protocol for the central controller to configure hard cells of LLN Devices.
  - Distributed Track Reservation
    - Need a fast reaction protocol to reserve a Track.
    - Need a protocol which can quickly detect a Track reservation failure.
    - Need an efficient negotiation protocol between LLN Devices multi-hop away from each other.

Updates

- 1. The relationship between Track and DetNet.
  - Track is an instance of a deterministic path, so DetNet can be used as a reference design for Track managment.
  - Track in 6TiSCH is targeted to Low-power and Lossy Networks (LLNs) , solutions in DetNet must be customized/tailored for Track management in 6TiSCH considering
    - Low-power consumption
    - TSCH MAC
    - Constrained devices with limited buffer and computation strength
  - Track management should be studied in 6TiSCH, and the solutions can influence the design of DetNet.

Updates

- 2. Reliability support via path redundancy
  - Non-6TiSCH work
    - WirelessHart uses a redundant path to deliver a packet if a path is not available.
    - DetNet architecture proposes seamless redundancy, replicating packets and sending them along at least two different paths.
  - 6TiSCH
    - Limited number of redundant paths between source and destination.
      - RPL (DODAG) have limited redundant paths comparing with a true mesh network, e.g. WirlessHart
    - Power consumption concerns
      - Replicating packets will dramatically increase the energy consumption of the network. However, a critical emergency message MAY be replicated via multiple Tracks.

Updates

- 3. Latency-related metrics for industry networks in RFC 5673
  - In fast control, tens of milliseconds of latency is typical. In many of these systems, if a packet does not arrive within the specified interval, the system enters an emergency shutdown state, often with substantial financial repercussions.
  - Non-critical closed-loop applications have a latency requirement that can be as low as 100 milliseconds but many control loops are tolerant of latencies above 1 second.
  - Most monitoring latencies will be in seconds to minutes.

Next Step

- Add more use-cases
- Update the drafts based on comments

6TiSCH requirements for DetNet

draft-thubert-6tisch-4detnet-01

Pascal Thubert

Status

- Goal
  - Support work @ DetNet
  - Make sure 6TiSCH agrees on what's needed
  - Make sure 6TiSCH needs are addressed

- Status:
  - Latest draft-thubert-6tisch-4detnet-01
  - published on June 11

- Room for coauthors
  - Merge with Chonggang

What do you want to do differently in the future?

- Hour glass model to replace silos
  - E2e principle, with one network, one network management, many applications
  - allowing evolution and dropping costs
- Open Protocols, Open source implementations
  - leveraging IETF, IEEE and ETSI
- Mix of deterministic and stochastic traffic
  - using IPv6 to reach widespread non critical devices for Industrial Internet
- Virtualized networks
  - with perfect isolation of IP flows vs. each individual (deterministic) control flow
- Deterministic properties spanning beyond wireless
  - over backbone to fog running virtual appliances

What would you like the IETF to deliver?

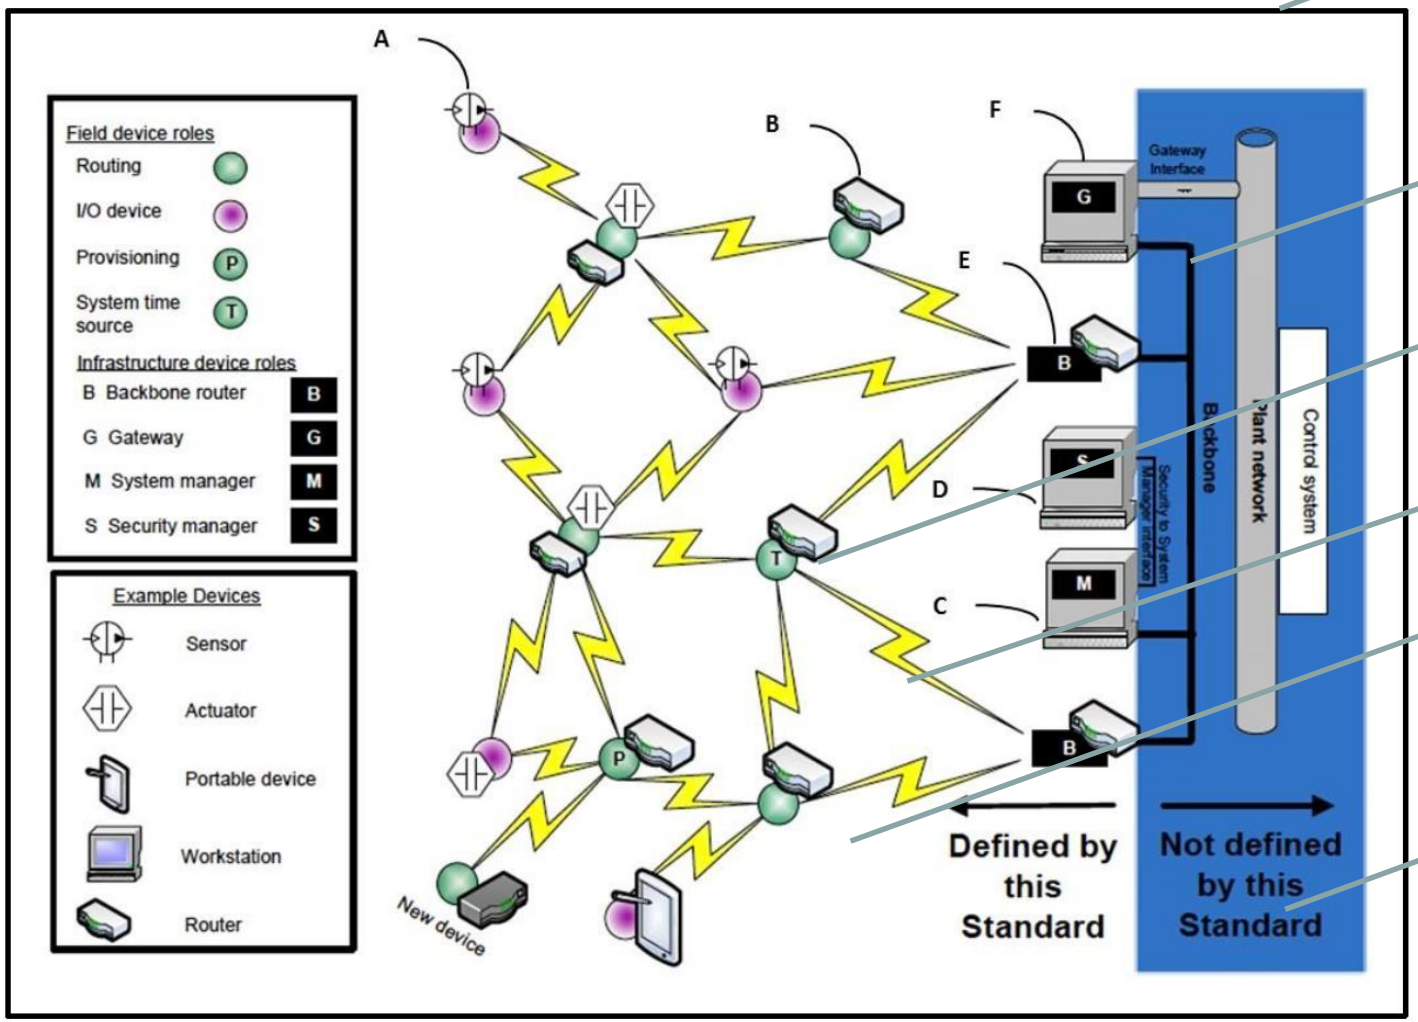

- ~~Silo~~ but open standards
- End to end tagging
- Replication & Elimination
- Per-Flow State
- IP routing (RPL) on same network
- End-to-end deterministic in standard

ML Questions

- What is a track?
  - Capability to do frame replication and elimination on top of ARQ
  - Multiple listeners in a time lot for replication
    - Use of 0xFFFF dest MAC in tracks
  - Multiple equivalent rcv slots for elimination?
  - One bundle per packet vs. sequence counter

AOB?

Thank You!