

# DOTS Telemetry

<https://datatracker.ietf.org/doc/draft-ietf-dots-telemetry/>

June 2020

M. Boucadair, T. Reddy, E. Doron, M. Chen,  
J. Shallow, K. Nishizuka

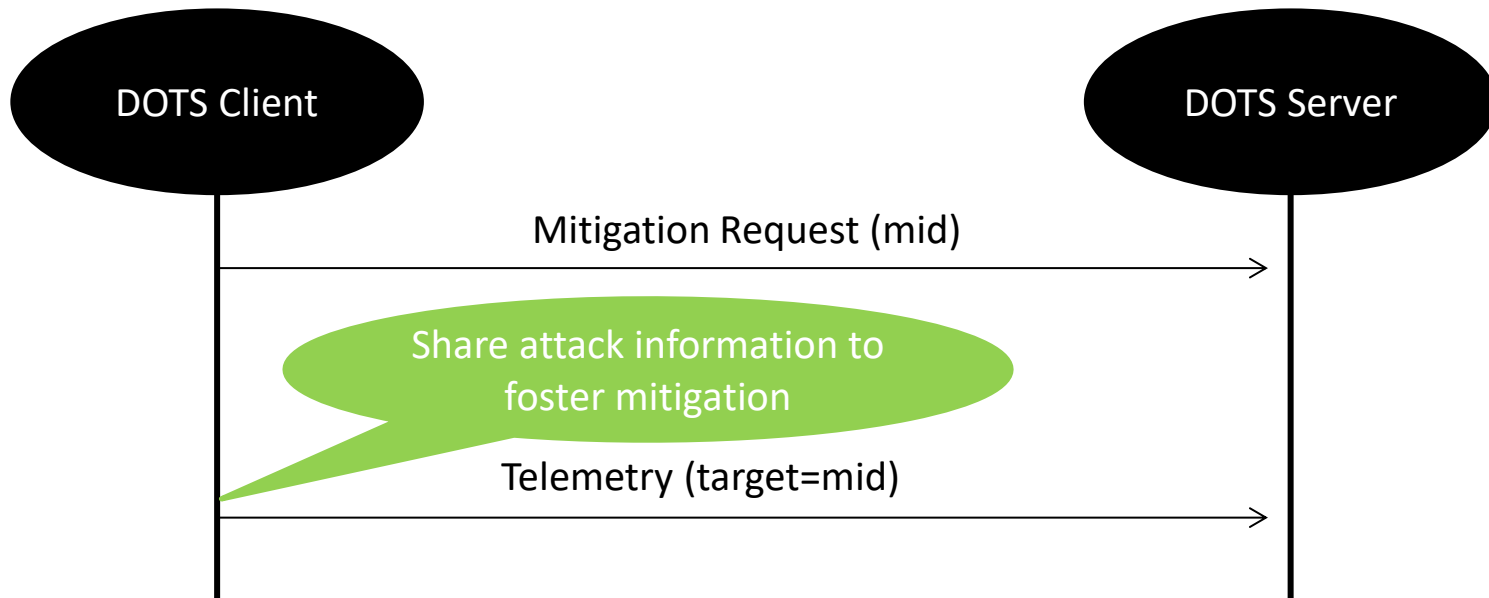
# Agenda

- Status
- Zoom on some key design points
- Next Steps

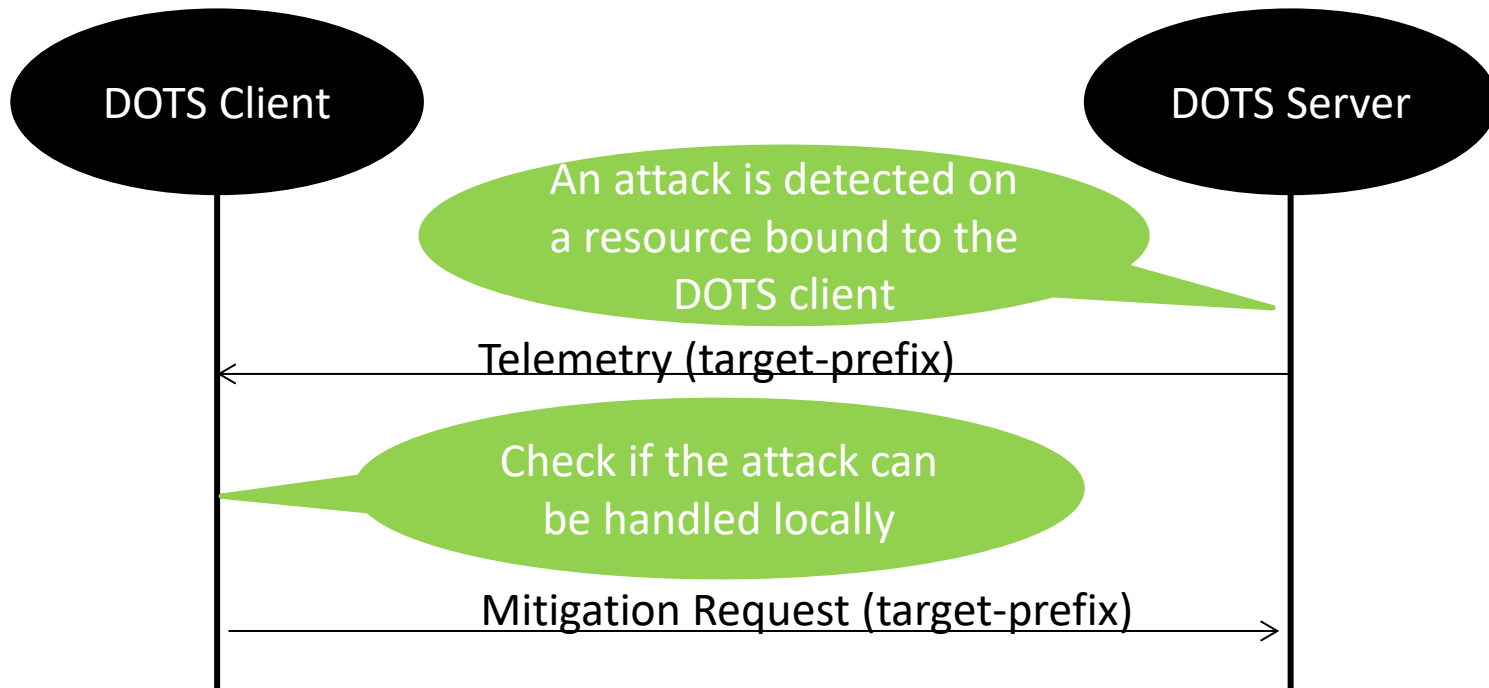
# Changes

- 12/19: WG Adoption
- 9 Revisions since then
  - 42 (00) vs. 105 pages (09)
  - 1 (00) vs. 2 YANG modules (09)
  - 51 (00) vs. 117 attributes (09)
  - +40 Examples/Figures in 09
  - Integrated implementation feedback
- We won't list all the changes but we will focus on some of them

# When to Share Telemetry Data: Some Samples



# When to Share Telemetry Data: Some Samples



# How to Share Telemetry Data

```
augment /ietf-signal:dots-signal/ietf-signal:message-type:
```

```
...
```

```
+---:(telemetry) {dots-telemetry}?
```

```
  +---rw pre-or-ongoing-mitigation* [cuid tmid]
```

```
    +---rw cuid                string
```

```
    +---rw cdid?               string
```

```
    +---rw tmid                uint32
```

```
    +---rw target
```

```
      | ...
```

```
    +---rw total-traffic* [unit]
```

```
      | ...
```

```
    +---rw total-traffic-protocol* [unit protocol]
```

```
      | ...
```

```
    +---rw total-traffic-port* [unit port]
```

```
      | ...
```

```
    +---rw total-attack-traffic* [unit]
```

```
      | ...
```

```
    +---rw total-attack-traffic-protocol* [unit protocol]
```

```
      | ...
```

```
    +---rw total-attack-traffic-port* [unit port]
```

```
      | ...
```

```
    +---rw total-attack-connection
```

```
      | ...
```

```
    +---rw total-attack-connection-port
```

```
      | ...
```

```
    +---rw attack-detail* [vendor-id attack-id]
```

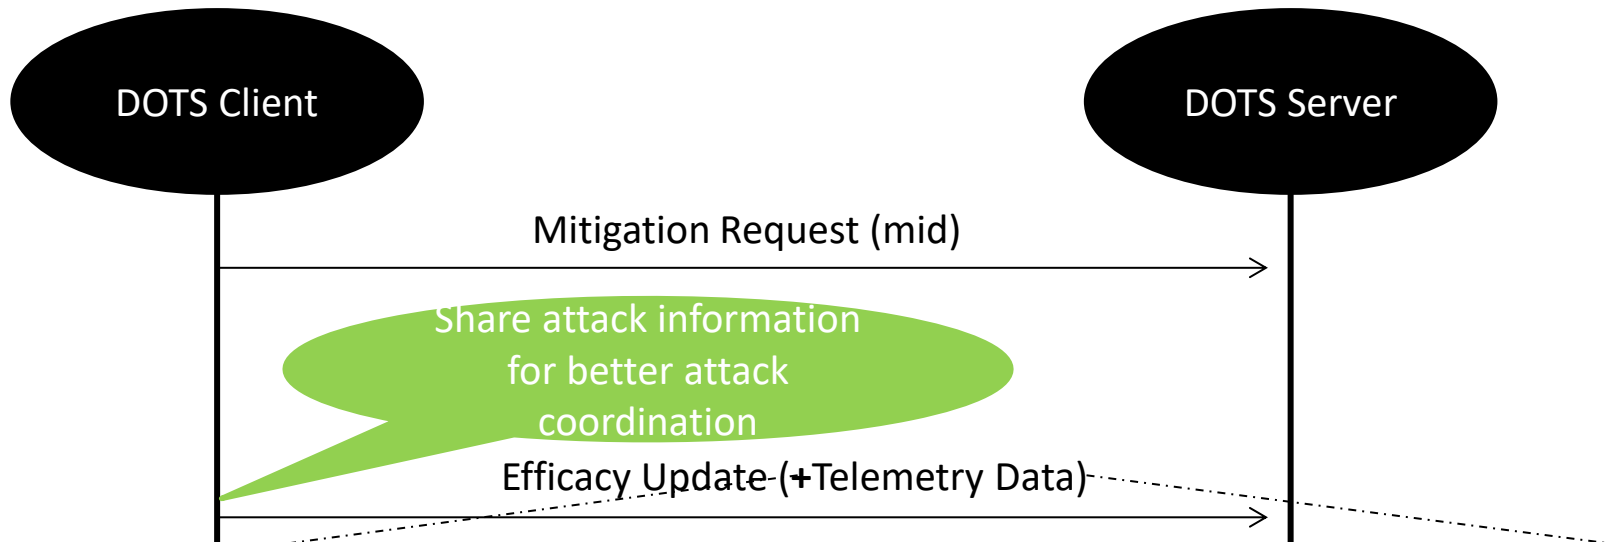
```
    ...
```

NEW Operation Path= /tm

Used to bind telemetry data with a mitigation request

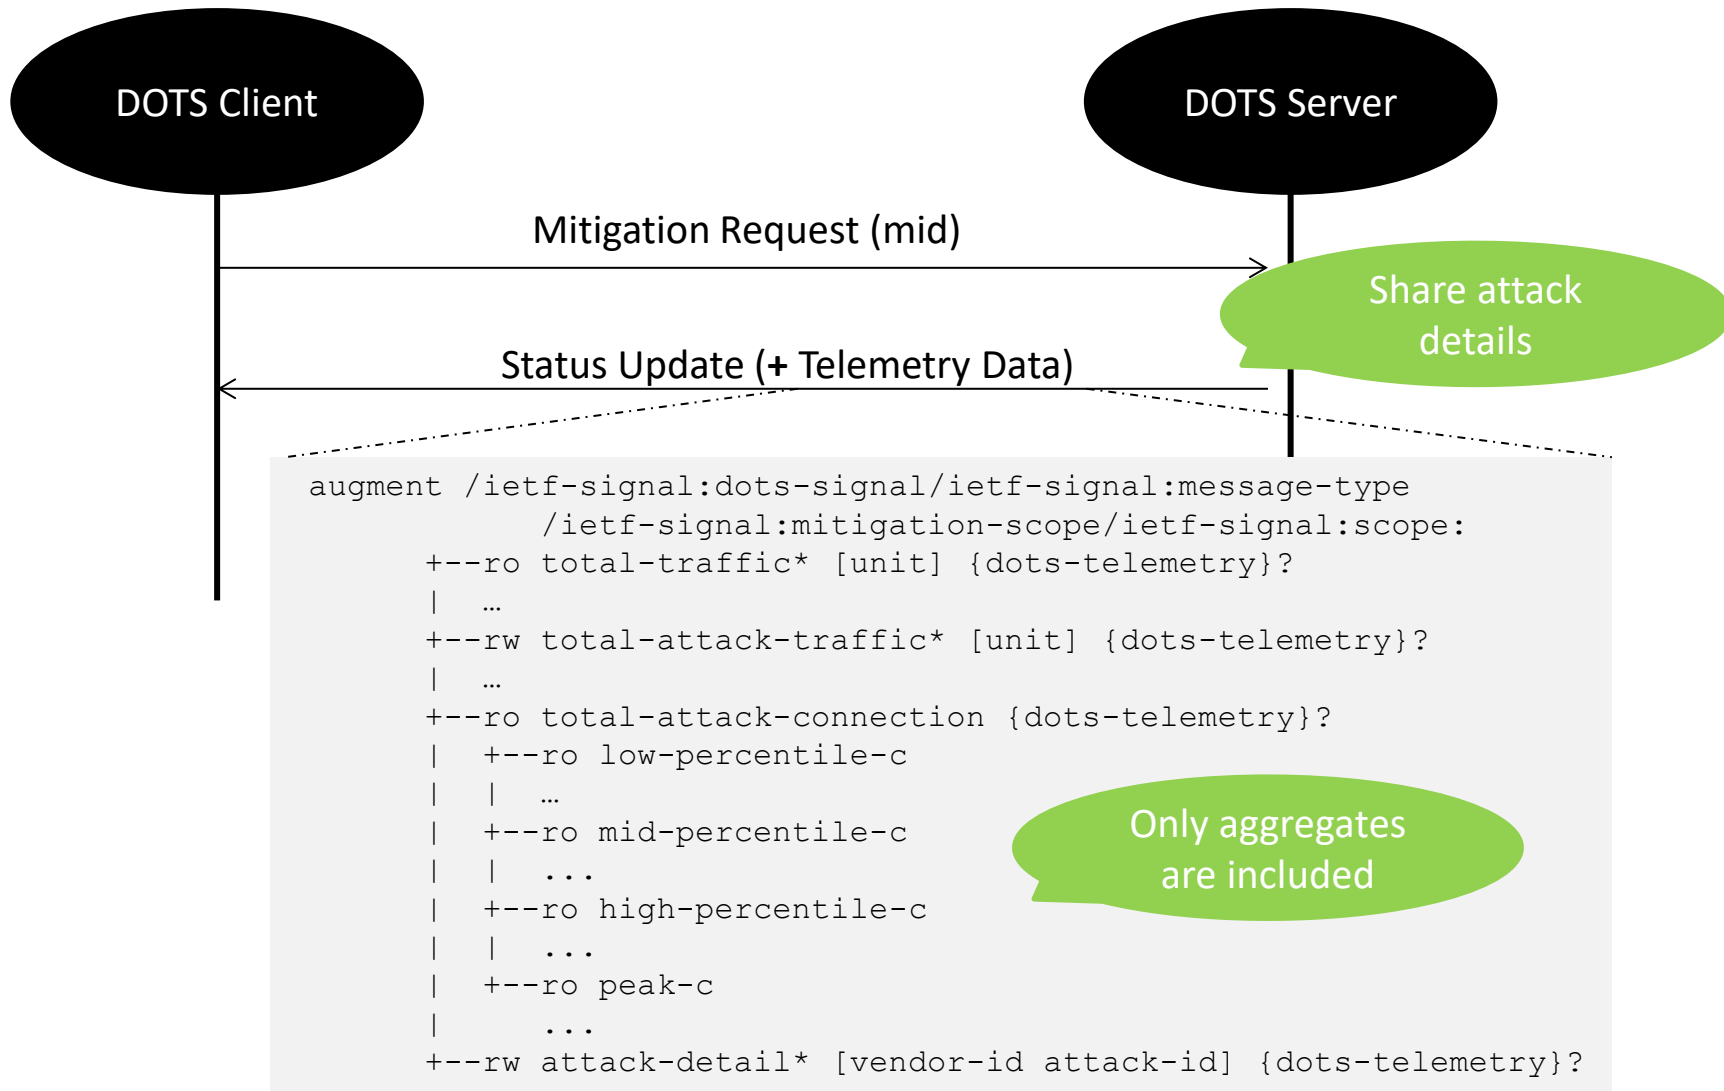
Comprehensive set of data. The granularity is controlled by DOTS agents

# When to Share Telemetry Data: Some Samples (Con'd)



```
augment /ietf-signal:dots-signal/ietf-signal:message-type
  /ietf-signal:mitigation-scope/ietf-signal:scope:
  +--rw total-attack-traffic* [unit] {dots-telemetry}?
  | ...
  +--rw attack-detail* [vendor-id attack-id] {dots-telemetry}?
    +--rw vendor-id          uint32
    +--rw attack-id          uint32
    +--rw attack-name?       string
    +--rw attack-severity?   attack-severity
    +--rw start-time?        uint64
    +--rw end-time?          uint64
    +--rw source-count
    | ...
    +--rw top-talker
    ...
```

# When to Share Telemetry Data: Some Samples (Con'd)





# When to Share Telemetry Data: Some Samples (Con'd)

DOTS Client

DOTS Server


Telemetry (Share Normal Traffic Baseline)

```
augment /ietf-signal:dots-signal/ietf-signal:message-type:
  +---:(telemetry-setup) {dots-telemetry}?
  |
  |   ...
  |   +---:(baseline)
  |     +---rw baseline* [id]
  |       +---rw id                               uint32
  |       +---rw target-prefix*                   inet:ip-prefix
  |       +---rw target-port-range* [lower-port]
  |         | ...
  |       +---rw target-protocol*                   uint8
  |       +---rw target-fqdn*                       inet:domain-name
  |       +---rw target-uri*                         inet:uri
  |       +---rw alias-name*                         string
  |       +---rw total-traffic-normal* [unit]
  |         | ...
  |       +---rw total-traffic-normal-per-protocol* [unit port]
  |         | ...
  |       +---rw total-traffic-normal-per-port* [unit port]
  |         | ...
  |       +---rw total-connection-capacity* [protocol]
  |         | ...
  |       +---rw total-connection-capacity-per-port* [protocol port]
  |         | ...
  |       ...
```

Can be used to  
detect abnormal  
traffic

Comprehensive set of data

# When to Share Telemetry Data: Some Samples (Con'd)



```
augment /ietf-signal:dots-signal/ietf-signal:message-type:
  +--:(telemetry-setup) {dots-telemetry}?
  |   ...
  |       +--:(pipe)
  |       |   +--rw total-pipe-capacity* [link-id unit]
  |       |   |   +--rw link-id      nt:link-id
  |       |   |   +--rw capacity     uint64
  |       |   |   +--rw unit         unit
  |       |   ...
  |       ...
```

Applies for the DOTS client domain

# Telemetry Configuration

```
augment /ietf-signal:dots-signal/ietf-signal:message-type:
  +--:(telemetry-setup) {dots-telemetry}?
  | +--ro max-config-values
  | | +--ro measurement-interval?          interval
  | | +--ro measurement-sample?           sample
  | | +--ro low-percentile?                percentile
  | | +--ro mid-percentile?                percentile
  | | +--ro high-percentile?               percentile
  | | +--ro server-originated-telemetry?   boolean
  | | +--ro telemetry-notify-interval?     uint32
  | +--ro min-config-values
  | | ...
  | +--ro supported-units
  | | +--ro unit-config* [unit]
  | |   +--ro unit          unit-type
  | |   +--ro unit-status   boolean
  | +--ro query-type*      query-type
  |   +--rw (setup-type)?
  |   +--:(telemetry-config)
  |   | ...
  |   +--:(pipe)
  |   | ...
  |   +--:(baseline)
  |   ...
+--:(telemetry) {dots-telemetry}?
...
```

Controls whether the server can include telemetry data in the status update

Controls the pace of telemetry notifications

# One or Two Key Values?

```
"ietf-dots-telemetry:total-attack-traffic": [  
  {  
    "ietf-dots-telemetry:unit": "megabit-ps",  
    "ietf-dots-telemetry:mid-percentile-g": "900"  
  }  
]
```

Efficacy Update

```
"total-attack-traffic": [  
  {  
    "unit": "megabit-ps",  
    "mid-percentile-g": "900"  
  }  
]
```

Telemetry

# ~~One or Two~~ Key Values?

```

"ietf-dots-telemetry:total-attack-traffic": [
  {
    "ietf-dots-telemetry:unit": "megabit-ps",
    "ietf-dots-telemetry:mid-percentile-g": "900"
  }
]

```

Efficacy Update

```

"total-attack-traffic": [
  {
    "unit": "megabit-ps",
    "mid-percentile-g": "900"
  }
]

```

Telemetry

Parameter Name	YANG Type	CBOR Key	CBOR Major Type & Information	JSON Type
total-attack-traffic	list	TBA17	4 array	Array
ietf-dots-telemetry:				
total-attack-traffic	list	TBA87	4 array	Array

# Telemetry Attributes

- Comprehension-Required or Comprehension-Optional?

# Telemetry Attributes

- ~~Comprehension Required or Comprehension-Optional?~~
  - Telemetry data are hints
  - Should not exacerbate message processing failures
  - Consistent with other specs (e.g., "source-prefix")
- This has an implication on the size as the key values will consume 3 bytes; hence this proposal:

Any objection?

Range	Registration Procedures	Note
1-127	IETF Review	comprehension-required
<b>128-255</b>	<b>IETF Review</b>	<b>comprehension-optional</b>
255-16383	IETF Review	comprehension-required

# Minimize the Size of Telemetry Data

- Filter out the asynchronous notifications (Uri-Query)
- Share the attack mapping details

```
module: ietf-dots-mapping
  augment /ietf-data:dots-data/ietf-data:dots-client:
    +--rw vendor-mapping {dots-telemetry}?
      +--rw vendor* [vendor-id]
        +--rw vendor-id          uint32
        +--rw attack-mapping* [attack-id]
          +--rw attack-id       uint32
          +--rw attack-name     string
  augment /ietf-data:dots-data/ietf-data:capabilities:
    +--ro vendor-mapping-enabled?  boolean {dots-telemetry}?
  augment /ietf-data:dots-data:
    +--ro vendor-mapping {dots-telemetry}?
      +--ro vendor* [vendor-id]
        +--ro vendor-id          uint32
        +--ro attack-mapping* [attack-id]
          +--ro attack-id       uint32
          +--ro attack-name     string
```

Avoids the need to include this attribute in the signal channel message



# Oversized Packet Handling

- Application break up data into Chunks
  - YANG <anydata> requires chunk to be full JSON as per RFC7951
  - How to break data down to minimize number of chunks?
- Use BLOCK1 and BLOCK2: Has limitations
  - Performance (symmetric traffic requires 'ACK' before next block is sent)
  - Handling lossy environments
- Use BLOCK3 and BLOCK4: Faster transmission without Block1/Block2 issues

# Block3/Block4 Options

## draft-bosh-core-new-block

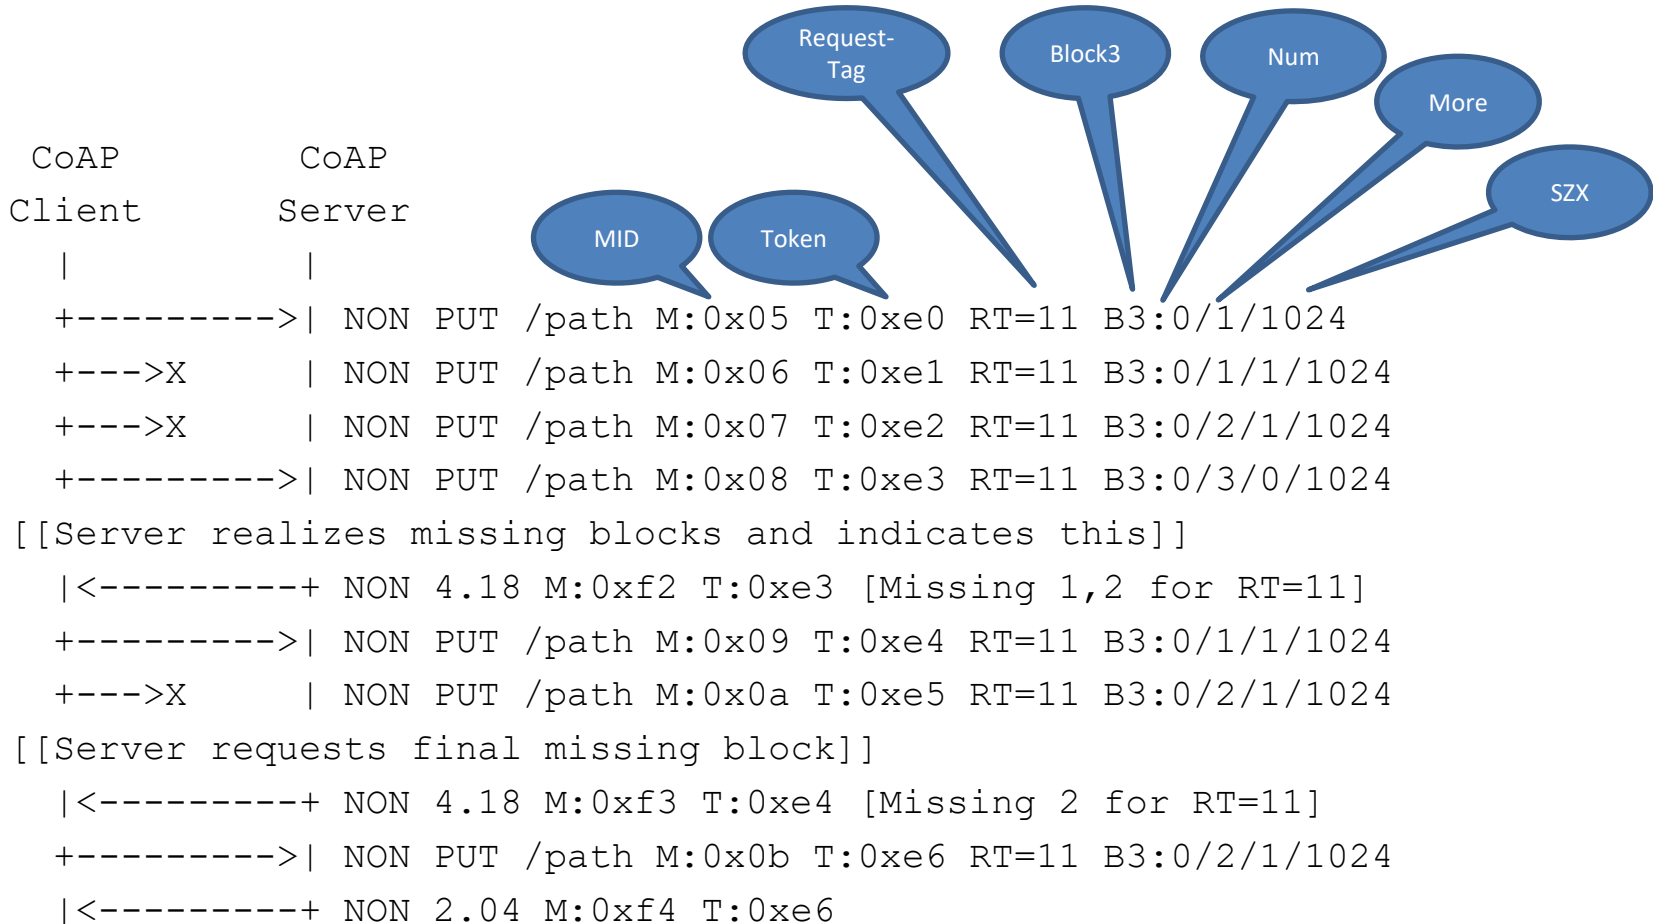
- Applicability Scope: DOTS-like
- Guards to prevent a CoAP agent from overloading the network
  - PROBING\_RATE clarification
  - MAX\_PAYLOADS defined, with a default value of 10

Congestion Control

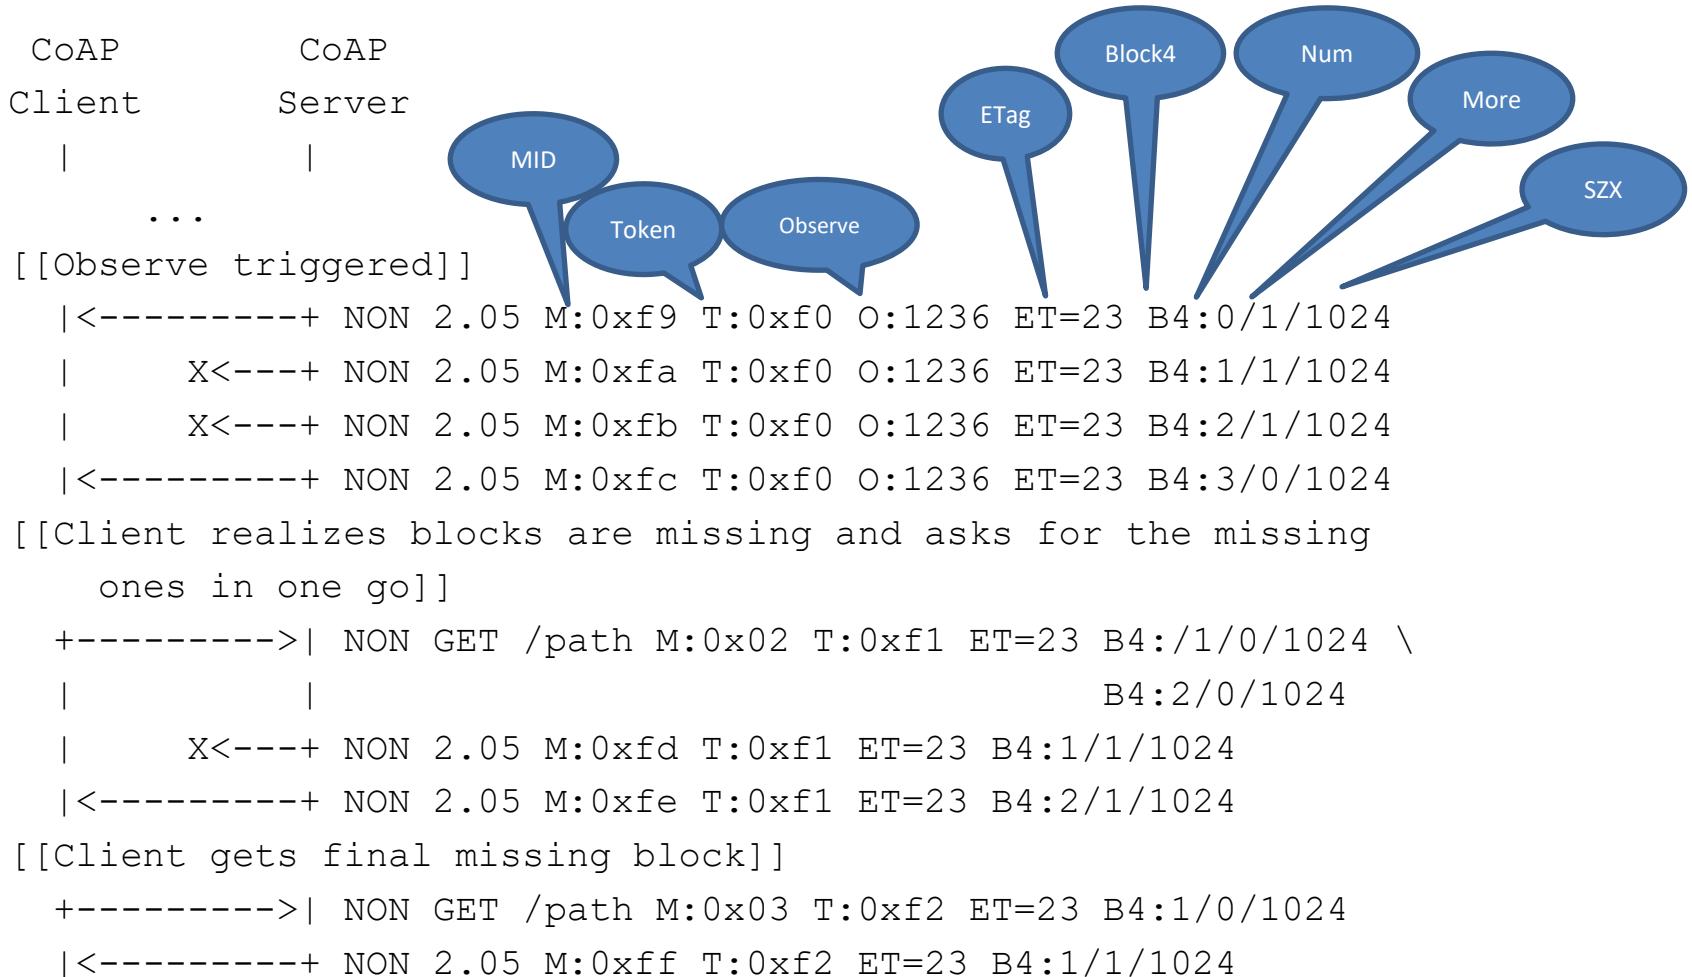
- Detailed description of Block3/Block4 Option
  - Including the use of Etag and Request-Tag
- New CoAP Response Code for missing blocks

Protocol Machinery

# Block3: An Example



# Block4 Example



# Current Discussion in the Draft

DOTS clients can use Block-wise transfer [RFC7959] with the recommendation detailed in Section 4.4.2 of [RFC8782] to control the size of a response when the data to be returned does not fit within a single datagram.

DOTS clients can also use CoAP Block1 Option in a PUT request (see Section 2.5 of [RFC7959]) to initiate large transfers, but these Block1 transfers will fail if the inbound "pipe" is running full, so consideration needs to be made to try to fit this PUT into a single transfer, or to separate out the PUT into several discrete PUTs where each of them fits into a single packet.

Block3 and Block 4 Options that are similar to the CoAP Block1 and Block2 Options, but enable faster transmissions of big blocks of data with less packet interchanges, are defined in [I-D.bosh-core-new-block]. DOTS implementations **can** consider the use of Block3 and Block 4 Options.

No normative language is used on purpose

# Implementation & Interop

# Implementation Status (Jon)

- RFC8782: DOTS Signal Channel Specification
- RFC8783: DOTS Data Channel Specification
- RFC8768: CoAP Hop-Limit Option
- draft-ietf-dots-signal-filter-control-06
- draft-ietf-dots-signal-call-home-08
- draft-ietf-dots-server-discovery-10
- **draft-ietf-dots-telemetry-07**
  - **Currently using Block2 if needed (8.1 may trigger Block1 usage)**
  - **Still in the process of implementing the vendor-id/attack-id changes to bring things up to the -09 spec**

# godots (Kaname)

- RFCs
  - RFC 8782 / RFC 8783
- drafts
  - draft-ietf-dots-signal-filter-control-06
  - draft-ietf-dots-signal-call-home-08
  - draft-ietf-dots-telemetry-09 <- Today's topic



# Interop (Kaname/Jon)

- Interoperability test of dots-telemetry has been done intensively on March/April
- Issues were fixed at both ends
- Essentially Sections 6, 7, and 8 were 'working'
  - Sections 6.5 and 8.1 still to be done
- General conclusion was that this was useful to have to pass back and forth additional information

# DOTS Telemetry Coverage

		godots	Interop tests	NCC
<b>6. DOTS Telemetry Setup Configuration</b>				
6.1	Telemetry Configuration	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6.2	Total Pipe Capacity	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6.3	Telemetry Baseline	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6.4	Reset Installed Telemetry Setup	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6.5	Conflict with Other DOTS Clients of the Same Domain	<input type="checkbox"/>	<input type="checkbox"/>	
<b>7. DOTS Pre-or-Ongoing Mitigation Telemetry</b>				
7.1	Pre-or-Ongoing-Mitigation DOTS Telemetry Attributes	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7.2	From DOTS Clients to DOTS Servers	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7.3	From DOTS Servers to DOTS Clients	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<b>8. DOTS Telemetry Mitigation Status Update</b>				
8.1	DOTS Clients to Servers Mitigation Efficacy			
8.2	DOTS Servers to Clients Mitigation Status			<input type="checkbox"/>

Not tested:

- Block-wise transfer Block3/Block4 (4.5)
- Vendor attack mapping(7.1.5)

# In summary

No significant issue is found but some functionality needs clarification

1. Current Uri-Path can't distinguish telemetry resources posted by the client or created by the server.
  - Those are totally different resources. For example, Uri-Query won't be applicable to the telemetry resources posted by the client.

**Proposal: Kaname to identify precisely which parts of the text need to be updated. Will then discuss whether/which changes are required.**

# In summary

No significant issue is found but some functionality needs clarification

1. Current Uri-Path can't distinguish telemetry resources posted by the client or created by the server.
2. Is the created vendor mapping shared among different DOTS clients?

**Clarification: No**

```
module: ietf-dots-data-channel
  +--rw dots-data
    +--rw dots-client* [cuid]
      | +--rw cuid string
```

```
module: ietf-dots-mapping
  augment /ietf-data:dots-data/ietf-data:dots-client:
    +--rw vendor-mapping {dots-telemetry}?
```

**DOTS clients of the same domain may interface with distinct DDoS protection technologies. Will add some text to better reflect this.**

# In summary

## No significant issue is found but some functionality needs clarification

1. Current Uri-Path can't distinguish telemetry resources posted by the client or created by the server.
2. Is the created vendor mapping shared among different DOTS client?
3. "DOTS agents MUST NOT include 'attack-name' attribute except if the corresponding attack mapping details were not shared with the peer DOTS agent". But how can the DOTS server know they're shared or not?

- **Clarification: This is implementation-specific. The DOTS server may record which version of the mapping table it shared with a DOTS client.**
- **Question: Do you think that we need to touch on these details?**

# Next Steps

- Release -10 with any required fixes to take into account in particular the implementation feedback and any comments raised today
- WGLC on -10