

# Applying COSE Signatures for YANG Data Provenance

draft-lopez-opsawg-yang-provenance-01

D. López, A. Pastor (*Telefónica*)

IETF#118, Prague (Czech Republic), November 2023

# As a Reminder: Provenance | 'prävən(ə)ns |

- The goal
  - Assurance of the origin and integrity of YANG datasets
  - Whenever the dataset is used beyond an original online flow
    - Use of data intermediaries, such as data lakes
    - AI/ML training and validation
    - Audit trails, including forensics evidence
- The means
  - An element containing a COSE signature
  - For any serialization method: XML, JSON, CBOR...
  - Detached payload

# The Changes in -01

- Some errata corrected
  - Including references for the provenance-signature typedef
- Comments addressed
  - Applying provenance in data pipelines
  - The recursion issue
- Still open comments
  - Choices for signature placement
  - Beyond the current proposal of a specific leaf element
  - Multiple signatures, what might bring considering attestation mechanisms

# The Recursion Issue

- The draft allows a provenance-signature leaf to appear anywhere in the enclosing element
  - But only once
- This is applicable to other non-leaf elements
  - Below or above
- The rules for (detached) signature generation and validation applies
  - Consistently dealing with any enclosed signature
- Support for recursive provenance verification
  - Data aggregation
  - Specific verification of relevant children

# The Signature Placement Issue

- Several proposals in addition to the one in the draft
  - Not necessarily exclusive among them, and with the original one
  - Though the implications for recursion need to be explored
- Annotations
  - Serialization neutrality
  - Are they defined for CBOR?
- YANG-push notifications
  - As RPC parameters they would require a change in the schema
  - As YANG elements, it would become a particular case of the current proposal
- Adding a provenance leaf to YANG-based files
  - Would it limit transparent combination of files?
  - Another particular case as leaf in the metadata schema
- How feasible would be an update of notification and metadata schemas?

# What Comes Next

- Sort the signature placement issue(s)
  - Considering multiple choices and their compatibility
- Refining and detailing use cases
  - Proposals welcome
- Consider implications of multiple signatures
  - And even beyond
- Practical evaluation
  - (Finally) hired a developer
  - And involved a COSE expert
- And continue seeking for WG comments and support