A YANG Data Model for Challenge-Response-based Remote Attestation Procedures using TPMs

draft-ietf-rats-yang-tpm-charra-18

IETF 113, March 2022, RATS WG

H. Birkholz
M. Eckel
Fraunhofer SIT

S. Bhandari
ThoughtSpot

E. Voit
B. Sulzen
Cisco

L. Xia
Huawei

T. Laffey
HPE

G. Fedorkow
Juniper

Relationship between drafts

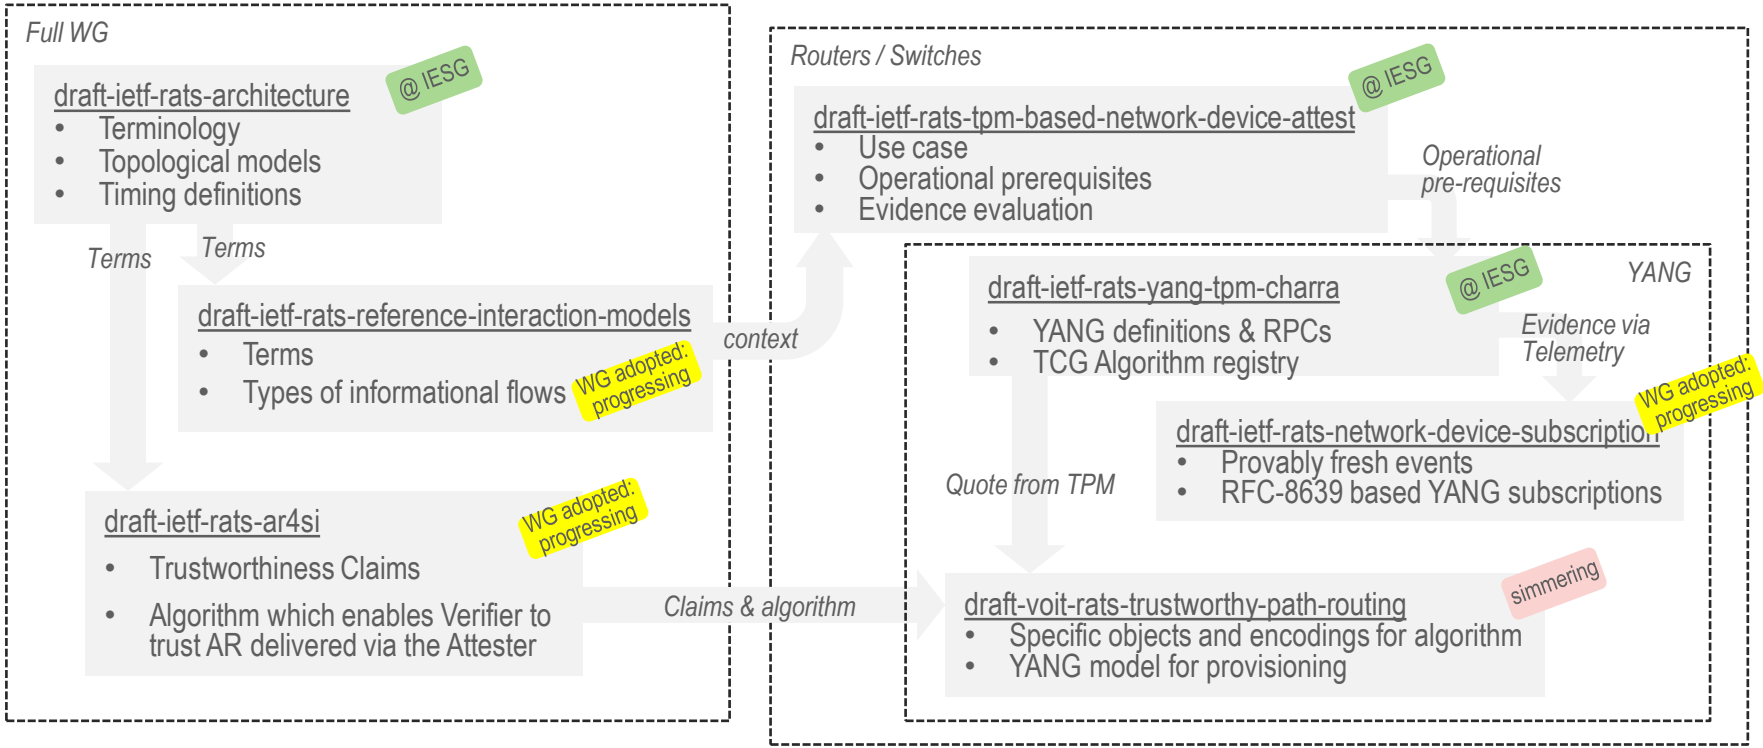

Status
One last IESG “Yes” or “No Objection” to pass

Discuss
Robert Wilton
Warren Kumari

Yes
Roman Danyliw

No Objection
Alvaro Retana
Benjamin Kaduk
Erik Kline
Francesca Palombini
John Scudder
Lars Eggert
Murray Kucherawy
Zaheduzzaman Sarker

No Record
Martin Duke
Martin Vigoureux
Éric Vyncke

- Tweaks made during ongoing IESG review
  - Appendix describing IMA, as Linux Kernel could not be referred to as Normative.
  - YANG model references included
  - XPATH syntax tweaks suggested by requested XPATH experts. Proposal included in new v18.
- No scope / functionality changes
- Nothing seen at this time expected to block Ballot closure and document acceptance

Attestation Event Stream Subscription

draft-ietf-rats-network-device-subscription-01

IETF 113, March 2022, RATS WG

Henk Birkholz
Fraunhofer SIT
henk.birkholz@sit.fraunhofer.de

Eric Voit
Cisco
evoit@cisco.com

Wei Pan
Huawei
william.panwei@huawei.com

Purpose & Scope

- Defines how to subscribe to a stream of attestation related Evidence on TPM-based network devices.
  - When subscribed, a Telemetry stream of verifiably fresh YANG notifications are pushed to the subscriber.
  - Notifications are generated for the Evidence going into TPM PCRs, and when the PCRs are extended.
- Result
  - Verifier is pushed new verifiably fresh Evidence whenever PCRs change.

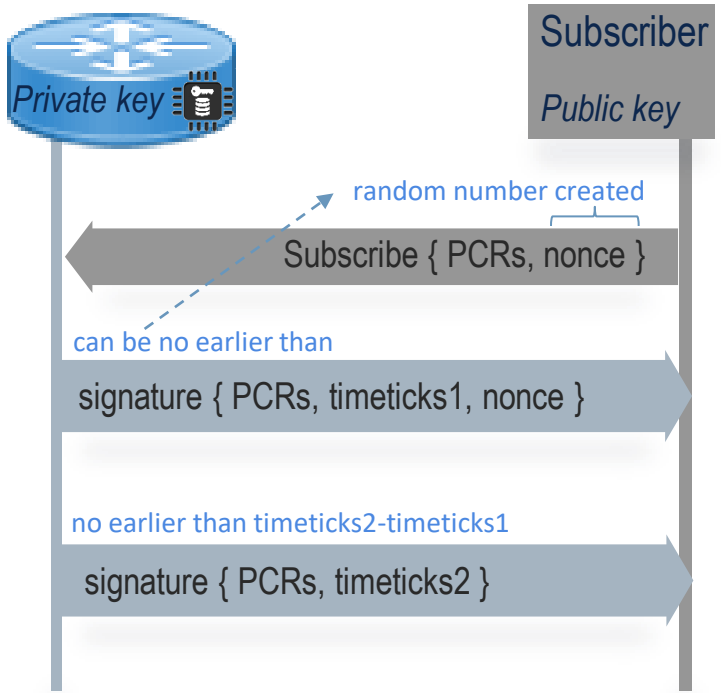

Status

- Stable as a direct combination of RFC-8639 & Charra
- Socialize Security Considerations section text (to be written)
- Then request WGLC

Attestation Results for Secure Interactions

draft-ietf-rats-ar4si-02

IETF 113, March 2022, RATS WG

Eric Voit
Cisco
evoit@cisco.com

Henk Birkholz
Fraunhofer SIT
henk.birkholz@sit.fraunhofer.de

Thomas Hardjono
MIT
hardjono@mit.edu

Thomas Fossati
Arm Limited
Thomas.Fossati@arm.com

Vincent Scarlata
Intel
vincent.r.scarlata@intel.com

Contents

- Part 1: Information Element definitions for Attestation Results (AR) generated by Verifier to support Secure Interactions between Attester and Relying Party
- Part 2: End-to-end implementation options: (a) Background check, (b) AR Augmented Evidence
- Implementations:
  - Trusted Path Routing (Proprietary – Cisco)
  - Veraison (Open Source, aspiration = Confidential Compute Consortium adoption)

Changes since IETF112

- WG Adoption
- Text clarifications on values of specific Trustworthiness Claims
- Mailing list comparison with EAT ‘security-level’
- Mailing list comparison with EAT ‘swresults’
- Continued alignment of instance draft:

Awaiting meaningful market uptake before requesting WG adoption

Slide from IETF 111, new Yellow Highlighting

Trustworthiness Claim Delivery

Based on draft-ietf-rats-architecture: Passport Model

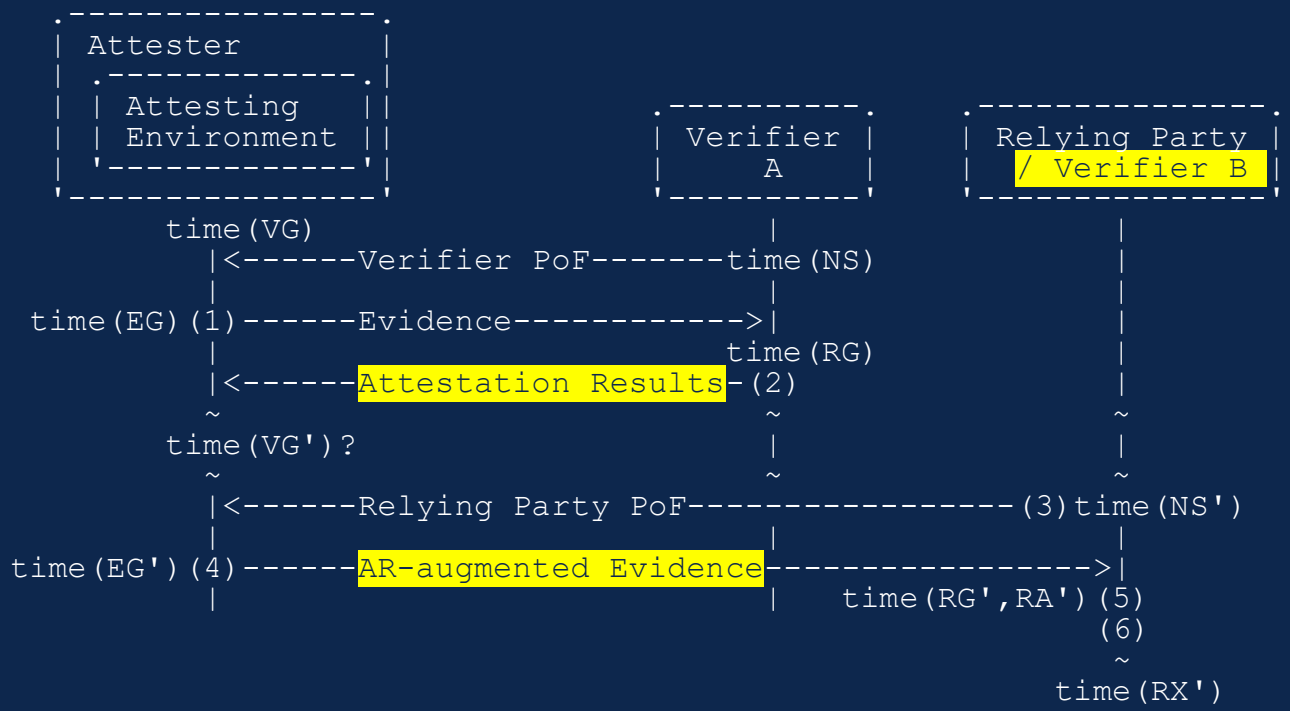

Section 2.3.1: AR Design Principles for Trustworthiness Claims

Design Principle	Reason
(1) Expose a small number of Trustworthiness Claims	A plethora of similar Trustworthiness Claims will result in divergent choices made on which to support between different Verifiers. This would place a lot of complexity in the Relying Party as it would be up to the Relying Party (and its policy language) to enable normalization across rich but incompatible Verifier object definitions.
(2) Each Trustworthiness Claim enumerates only the specific states that could viably result in a different outcome after the Policy for Attestation Results has been applied	By explicitly disallowing the standardization of enumerated states which cannot easily be connected to a use case, we avoid forcing implementers from making incompatible guesses on what these states might mean.
(3) Verifier and RP developers need explicit definitions of each state	Without such guidance, the Verifier will append plenty of raw supporting info. This relieves the Verifier of making the hard decisions. Of course, this raw info will be mostly non-interpretable and therefore non-actionable by the Relying Party.
(4) Support standards and non-standard extensibility	Standard types of Verifier generated Trustworthiness Claims should be vetted by the full RATS working group, rather than being maintained in a repository which doesn't follow the RFC process. This will keep a tight lid on extensions which must be considered by the Relying Party's policy language. Because this process takes time, non-standard extensions will be needed for implementation speed and flexibility

Comparing Trustworthiness Claims & swresults
(undergoing tweaks in EAT)

	Trustworthiness Claim (AR4SI)		EAT Claim
	'executables'	'file-system'	'swresults'
Attestation target	All runtime software/object loaded into Attester memory	A Verifier specified set of directories within the Attester file system	A Verifier specified set of software and/or multiple sets of software modules
Encodable states	Seven	Five	Six. Might need to encode more than one (e.g., Firmware & Kernel)
Vendor extensible	Yes		No
Claim consistency	Common claim generalizations across Verifier generated AR: (Affirming, Warning, Contraindicated, None)		No generalized claim abstractions across generated AR claims
RP claim interpretation	Claim always references the full attestation target		Claim references either attestation target or submodule(s). An RP parser must understand context within structured message.
Purpose	Only encodes information likely to be actioned by RP		Can encode both actionable information as well as supplementary information for debug logs
Encodings/serialization	Transport independent, example serialization in draft-voit-rats-trustworthy-path-routing		JSON, CBOR, could add more
Information Model	English prose		English prose & CDDL