

Workgroup: RATS Working Group
Internet-Draft:
draft-ietf-rats-reference-interaction-
models-09
Published: 4 March 2024
Intended Status: Informational
Expires: 5 September 2024
Authors: H. Birkholz M. Eckel
 Fraunhofer SIT Fraunhofer SIT
 W. Pan E. Voit
 Huawei Technologies Cisco

Reference Interaction Models for Remote Attestation Procedures

Abstract

This document describes interaction models for remote attestation procedures (RATS). Three conveying mechanisms -- Challenge/Response, Uni-Directional, and Streaming Remote Attestation -- are illustrated and defined. Analogously, a general overview about the information elements typically used by corresponding conveyance protocols are highlighted.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 September 2024.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with

respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
 - [2.1. Disambiguation](#)
- [3. Scope and Intent](#)
- [4. Essential Requirements](#)
 - [4.1. Endorsement of Attesting Environments](#)
- [5. Normative Prerequisites](#)
- [6. Generic Information Elements](#)
- [7. Interaction Models](#)
 - [7.1. Challenge/Response Remote Attestation](#)
 - [7.1.1. Models and Example Sequences of Challenge/Response Remote Attestation](#)
 - [7.2. Uni-Directional Remote Attestation](#)
 - [7.3. Streaming Remote Attestation](#)
 - [7.3.1. Streaming Remote Attestation without a Broker](#)
 - [7.3.2. Streaming Remote Attestation with a Broker](#)
- [8. Additional Application-Specific Requirements](#)
 - [8.1. Confidentiality](#)
 - [8.2. Mutual Authentication](#)
 - [8.3. Hardware-Enforcement/Support](#)
- [9. Implementation Status](#)
 - [9.1. Implementer](#)
 - [9.2. Implementation Name](#)
 - [9.3. Implementation URL](#)
 - [9.4. Maturity](#)
 - [9.5. Coverage and Version Compatibility](#)
 - [9.6. License](#)
 - [9.7. Implementation Dependencies](#)
 - [9.8. Contact](#)
- [10. Security and Privacy Considerations](#)
- [11. Acknowledgments](#)
- [12. References](#)
 - [12.1. Normative References](#)
 - [12.2. Informative References](#)
- [Appendix A. CDDL Specification for a simple CoAP Challenge/Response Interaction](#)
- [Authors' Addresses](#)

1. Introduction

Remote Attestation procedures (RATS, [[RFC9334](#)]) are workflows composed of roles and interactions, in which Verifiers create

Attestation Results about the trustworthiness of an Attester's system component characteristics. The Verifier's assessment in the form of Attestation Results is created based on Attestation Policies and Evidence -- trustable and tamper-evident Claims Sets about an Attester's system component characteristics -- generated by an Attester. The roles *Attester* and *Verifier*, as well as the Conceptual Messages *Evidence* and *Attestation Results* are concepts defined by the RATS Architecture [[RFC9334](#)]. This document defines interaction models that can be used in specific RATS-related solution documents. The primary focus of this document is the conveyance of attestation Evidence. The reference models defined can also be applied to the conveyance of other Conceptual Messages in RATS. Specific goals of this document are to:

1. prevent inconsistencies in descriptions of interaction models in other documents (due to text cloning and evolution over time), and to
2. enable to highlight an exact delta/divergence between the core set of characteristics captured here in this document and variants of these interaction models used in other specifications or solutions.

In summary, this document enables the specification and design of trustworthy and privacy preserving conveyance methods for attestation Evidence from an Attester to a Verifier. While the conveyance of other Conceptual Messages is out-of-scope the methods described can also be applied to the conveyance of, for example, Endorsements or Attestation Results.

2. Terminology

This document uses the following set of terms, roles, and concepts as defined in [[RFC9334](#)]: Attester, Verifier, Relying Party, Conceptual Message, Evidence, Endorsement, Attestation Result, Appraisal Policy, Attesting Environment, Target Environment

A PKIX Certificate is an X.509v3 format certificate as specified by [[RFC5280](#)].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2.1. Disambiguation

The term "Remote Attestation" is a common expression and often associated or connoted with certain properties. The term "Remote" in

this context does not necessarily refer to a remote entity in the scope of network topologies or the Internet. It rather refers to decoupled systems or entities that exchange the payload of the Conceptual Message type called Evidence [[RFC9334](#)]. This conveyance can also be "Local", if the Verifier role is part of the same entity as the Attester role, e.g., separate system components of the same Composite Device (a single RATS entity). Even if an entity takes on two or more different roles, the functions they provide typically reside in isolated environments that are components of the same entity. Examples of such isolated environments include: a Trusted Execution Environment (TEE), Baseboard Management Controllers (BMCs), as well as other physical or logical protected/isolated/shielded Computing Environments (e.g. embedded Secure Elements (eSE) or Trusted Platform Modules (TPM)). Readers of this document should be familiar with the concept of Layered Attestation as described in Section 3.1 Two Types of Environments of an Attester in [[RFC9334](#)] and the definition of Attestation as described in [[I-D.ietf-rats-tpm-based-network-device-attest](#)].

3. Scope and Intent

This document focuses on generic interaction models between Attesters and Verifiers in order to convey Evidence. Complementary procedures, functions, or services that are required for a complete semantic binding of the concepts defined in [[RFC9334](#)] are out-of-scope of this document. Examples include: identity establishment, key distribution and enrollment, time synchronization, as well as certificate revocation.

Furthermore, any processes and duties that go beyond carrying out remote attestation procedures are out-of-scope.

For instance, using the results of a remote attestation procedure that are created by the Verifier, e.g., how to triggering remediation actions or recovery processes, as well as such remediation actions and recovery processes themselves, are also out-of-scope.

The interaction models illustrated in this document are intended to provide a stable basis and reference for other solutions documents inside or outside the IETF. Solution documents of any kind can reference the interaction models in order to avoid text clones and to avoid the danger of subtle discrepancies. Analogously, deviations from the generic model descriptions in this document can be illustrated in solutions documents to highlight distinct contributions.

4. Essential Requirements

In order to ensure appropriate conveyance of Evidence, there exist essential requirements which MUST be fulfilled:

Integrity: Information provided by an Attester MUST be integral. This may be achieved by means of a digital signature over Attestation Evidence. The signature may be symmetric, such as an HMAC, or asymmetric, such as ECDSA.

Authentication: The information provided by the Attester MUST be authentic. For that purpose, the Attester should authenticate itself to the Verifier. This may be an implicit authentication by means of a digital signature over the Attestation Evidence, which does not require additional protocol steps, or may be achieved by using a confidential channel by means of encryption.

4.1. Endorsement of Attesting Environments

Via its Attesting Environments, an Attester only generates Evidence about its Target Environments. After being appraised to be trustworthy, a Target Environment may become a new Attesting Environment in charge of generating Evidence for further Target Environments. [\[RFC9334\]](#) explains this as Layered Attestation. Layered Attestation has to start with an initial Attesting Environment. In essence, there cannot be turtles all the way down [\[turtles\]](#). At this rock bottom of Layered Attestation, the Attesting Environments are always called Roots of Trust (RoT). An Attester cannot generate Evidence about its own RoTs by design. As a consequence, a Verifier requires trustable statements about this subset of Attesting Environments from a different source than the Attester itself. The corresponding trustable statements are called Endorsements and originate from external, trustable entities that take on the role of an Endorser (e.g., supply chain entities).

5. Normative Prerequisites

In order to ensure an appropriate conveyance of Evidence via interaction models in general, the following set of prerequisites MUST be in place to support the implementation of interaction models:

Authentication Secret: An Authentication Secret MUST be available exclusively to an Attesting Environment of an Attester.

The Attester MUST protect Claims with that Authentication Secret, thereby proving the authenticity of the Claims included in Evidence. The Authentication Secret MUST be established before RATS can take place.

Attester Identity:

A statement about a distinguishable Attester made by an Endorser.

The provenance of Evidence with respect to a distinguishable Attesting Environment MUST be correct and unambiguous.

An Attester Identity MAY be an Authentication Secret which is available exclusively to one of the Attesting Environments of an Attester. It MAY be a unique identity, MAY be included in a zero-knowledge proof (ZKP), MAY be part of a group signature, or it MAY be a randomized DAA credential [[DAA](#)].

Attestation Evidence Authenticity: Attestation Evidence MUST be authentic.

In order to provide proofs of authenticity, Attestation Evidence SHOULD be cryptographically associated with an identity document (e.g., a PKIX certificate or trusted key material, or a randomized DAA credential [[DAA](#)]), or SHOULD include a correct, unambiguous and stable reference to an accessible identity document.

Evidence Freshness: Evidence MUST include an indicator about its freshness that can be understood by a Verifier. Analogously, interaction models MUST support the conveyance of proofs of freshness in a way that is useful to Verifiers and their appraisal procedures.

Evidence Protection: Evidence MUST be a set of well-formatted and well-protected Claims that an Attester can create and convey to a Verifier in a tamper-evident manner.

6. Generic Information Elements

This section defines the information elements that are vital to all kinds interaction models. Varying from solution to solution, generic information elements can be either included in the scope of protocol messages (instantiating Conceptual Messages) or can be included in additional protocol parameters or payload. Ultimately, the following information elements are required by any kind of scalable remote attestation procedure using one or more of the interaction models provided.

Authentication Secret IDs ('authSecIDs'):

mandatory

A statement representing an identifier list that MUST be associated with corresponding Authentication Secrets used to protect Claims included in Evidence.

Each distinguishable Attesting Environment has access to a protected capability that provides an Authentication Secret associated with that Attesting Environment. Consequently, an Authentication Secret ID can also identify an Attesting Environment.

Handle ('handle'): *mandatory*

A statement that is intended to uniquely distinguish received Evidence and/or determine the freshness of Evidence.

A Verifier can also use a Handle as an indicator for authenticity or attestation provenance, as only Attesters and Verifiers that are intended to exchange Evidence should have knowledge of the corresponding Handles. Examples include Nonces or signed timestamps.

Claims ('claims'): *mandatory*

Claims are assertions that represent characteristics of an Attester's Target Environment.

Claims are part of a Conceptual Message and are, for example, used to appraise the integrity of Attesters via Verifiers. The other information elements in this section can be expressed as Claims in any type of Conceptual Messages.

Event Logs ('eventLogs'): *optional*

Event Logs accompany Claims by providing event trails of security-critical events in a system. The primary purpose of Event Logs is to support Claim reproducibility by providing information on how Claims originated.

Reference Values ('refValues') *mandatory*

Reference Values as defined in [\[RFC9334\]](#). This specific type of Claims is used to appraise Claims incorporated in Evidence. For example, Reference Values MAY be Reference Integrity Measurements (RIM) or assertions that are implicitly trusted because they are signed by a trusted authority (see Endorsements in [\[RFC9334\]](#)). Reference Values typically represent (trusted) Claim sets about an Attester's intended platform operational state.

Claim Selection ('claimSelection'): *optional*

A (sub-)set of Claims which can be created by an Attester.

Claim Selections act as filters to specify the exact set of Claims to be included in Evidence. In a remote attestation process, a Verifier sends a Claim Selection, among other elements, to an Attester. An Attester MAY decide whether or not to provide all requested Claims from a Claim Selection to the Verifier.

Collected Claims ('collectedClaims'): *mandatory*

Collected Claims represent a (sub-)set of Claims created by an Attester.

Collected Claims are gathered based on the Claims selected in the Claim Selection. If a Verifier does not provide a Claim Selection, then all available Claims on the Attester are part of the Collected Claims.

Evidence ('evidence'): *mandatory*

A set of Claims that consists of a list of Authentication Secret IDs that each identifies an Authentication Secret in a single Attesting Environment, the Attester Identity, Claims, and a Handle. Attestation Evidence MUST cryptographically bind all of these information elements. Evidence MUST be protected via an Authentication Secret. The Authentication Secret MUST be trusted by the Verifier as authoritative.

Attestation Result ('attestationResult'): *mandatory*

An Attestation Result is produced by the Verifier as the output of the appraisal of Evidence. Attestation Results include condensed assertions about integrity or other characteristics of the corresponding Attester that are processible by Relying Parties.

7. Interaction Models

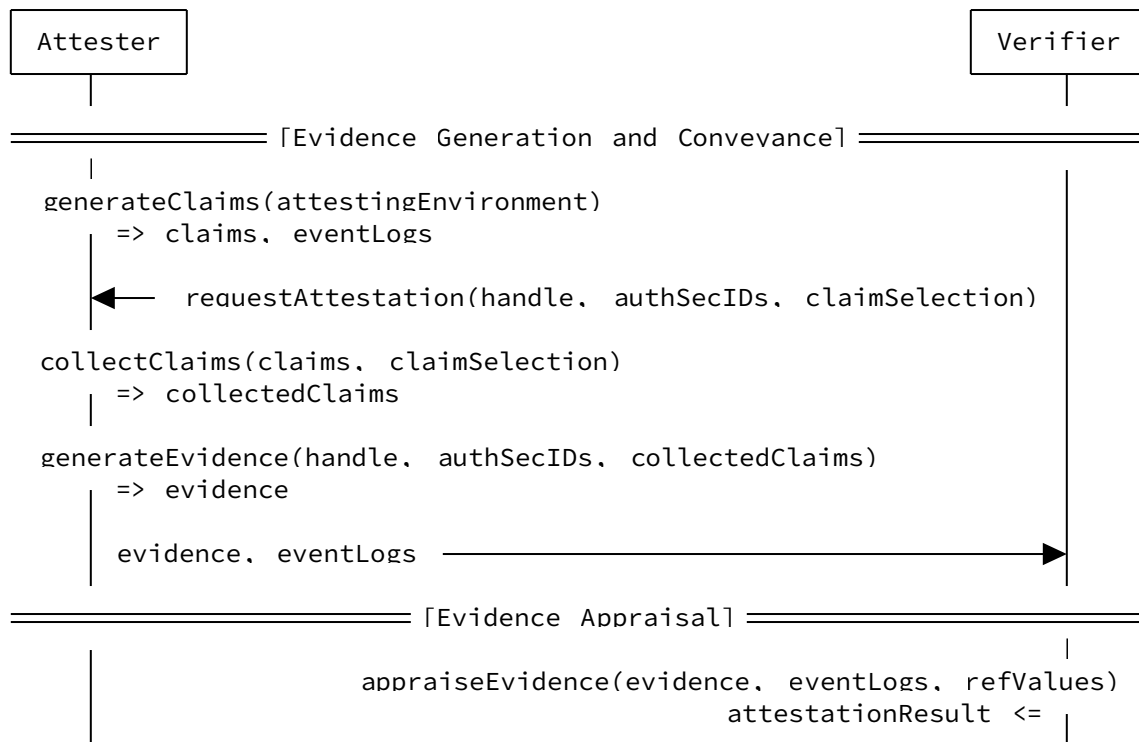
The following subsections introduce and illustrate the interaction models:

1. Challenge/Response Remote Attestation
2. Uni-Directional Remote Attestation
3. Streaming Remote Attestation

Each section starts with a sequence diagram illustrating the interactions between Attester and Verifier. While the presented interaction models focus on the conveyance of Evidence, the intention of this document is in support of future work that applies the presented models to the conveyance of other Conceptual Messages, namely Attestation Results, Endorsements, Reference Values, or Appraisal Policies.

All interaction models have a strong focus on the use of a handle to incorporate a type of proof of freshness and to prevent replay attacks. The way these handles are processed is the most prominent difference between the three interaction models.

7.1. Challenge/Response Remote Attestation



The Attester boots up and thereby produces claims about its boot state and its operational state. Event Logs accompany the produced claims by providing an event trail of security-critical events in a system. Claims are produced by all attesting Environments of an Attester system.

The Challenge/Response remote attestation procedure is initiated by the Verifier by sending a remote attestation request to the Attester. A request includes a Handle, a list of Authentication Secret IDs, and a Claim Selection.

In the Challenge/Response model, the handle is composed of qualifying data in the form of a practically infeasible to guess nonce, such as a cryptographically strong random number. The Verifier-generated nonce is intended to guarantee Evidence freshness and to prevent replay attacks.

The list of Authentication Secret IDs selects the attestation keys with which the Attester is requested to sign the Attestation Evidence. Each selected key is uniquely associated with an Attesting Environment of the Attester. As a result, a single Authentication Secret ID identifies a single Attesting Environment. Correspondingly, a particular set of Evidence originating from a particular Attesting Environment in a composite device can be requested via multiple Authentication Secret IDs. Methods to acquire Authentication Secret IDs or mappings between Attesting Environments to Authentication Secret IDs are out-of-scope of this document.

The Attester collects Claims based on the Claim Selection. With the Claim Selection the Verifier defines the set of Claims it requires. Correspondingly, collected Claims can be a subset of the produced Claims. This could be all available Claims, depending on the Claim Selection. If the Claim Selection is omitted, then by default all Claims that are known and available on the Attester MUST be used to create corresponding Evidence. For example, when performing a boot integrity evaluation, a Verifier may only be requesting a particular subset of claims about the Attester, such as Evidence about BIOS/UEFI and firmware that the Attester booted up, and not include information about all currently running software.

With the Handle, the Authentication Secret IDs, and the collected Claims, the Attester produces signed Evidence. That is, it digitally signs the Handle and the collected Claims with a cryptographic secret identified by the Authentication Secret ID. This is done once per Attesting Environment which is identified by the particular Authentication Secret ID. The Attester communicates the signed Evidence as well as all accompanying Event Logs back to the Verifier.

While it is crucial that Claims, the Handle, and the Attester Identity information (i.e., the Authentication Secret) MUST be cryptographically bound to the signature of Evidence, they MAY be presented obfuscated, encrypted, or cryptographically blinded. For further reference see section [Section 10](#).

As soon as the Verifier receives the Evidence and the Event Logs, it appraises the Evidence. For this purpose, it validates the signature, the Attester Identity, and the Handle, and then appraises the Claims. Appraisal procedures are application-specific and can be conducted via comparison of the Claims with corresponding Reference

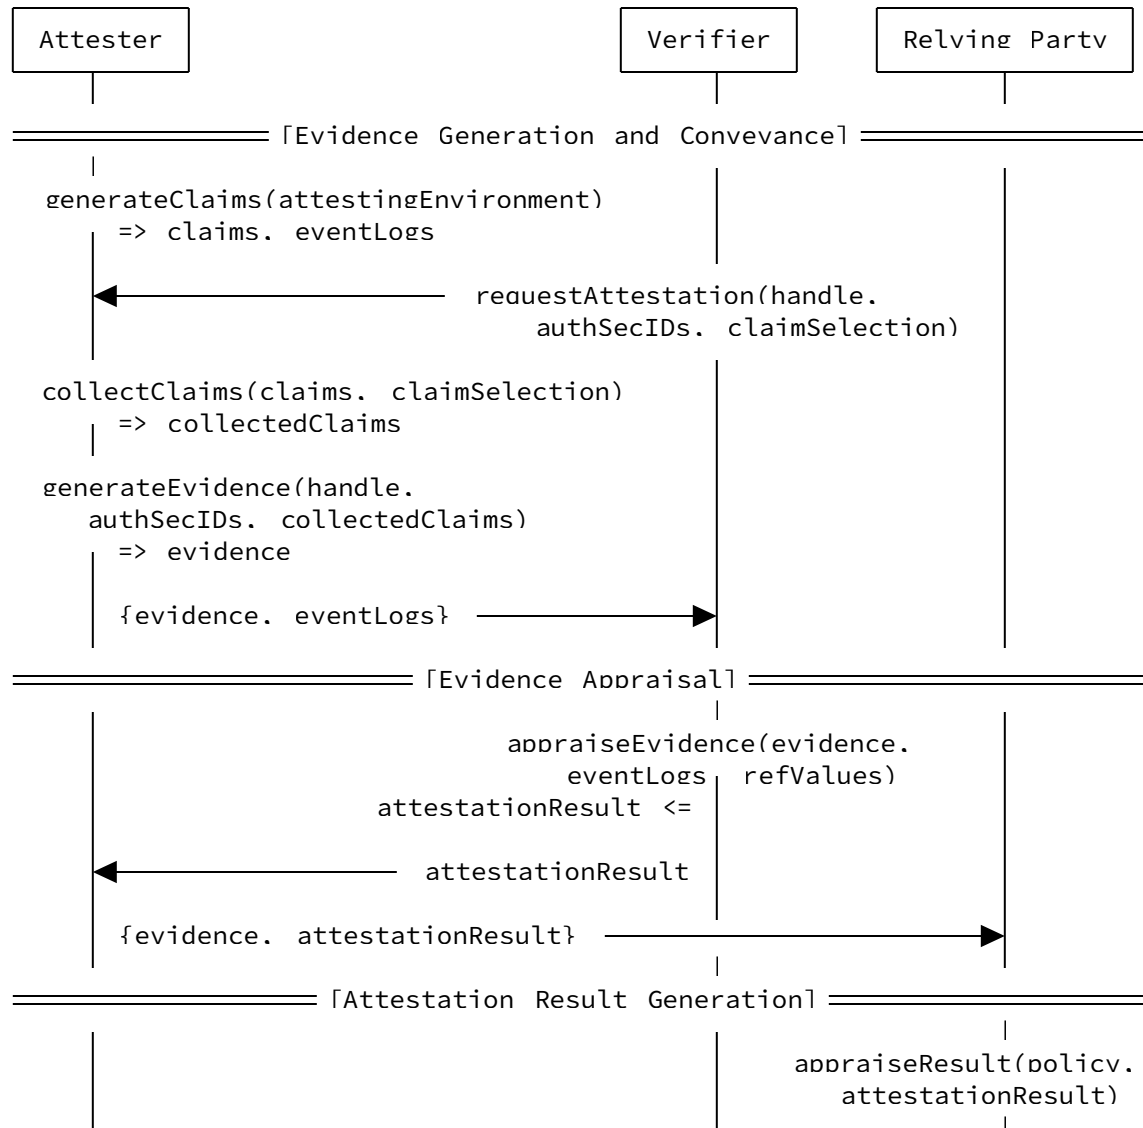
Values, such as Reference Integrity Measurements. The final output of the Verifier are Attestation Results. Attestation Results constitute new Claim Sets about the properties and characteristics of an Attester, which enables Relying Parties, for example, to assess an Attester's trustworthiness.

7.1.1. Models and Example Sequences of Challenge/Response Remote Attestation

According to the RATS Architecture, two reference models for Challenge/Response Attestation have been proposed. This section highlights the information flows between the Attester, Verifier, and Relying Party undergoing Remote Attestation Procedure, using these models.

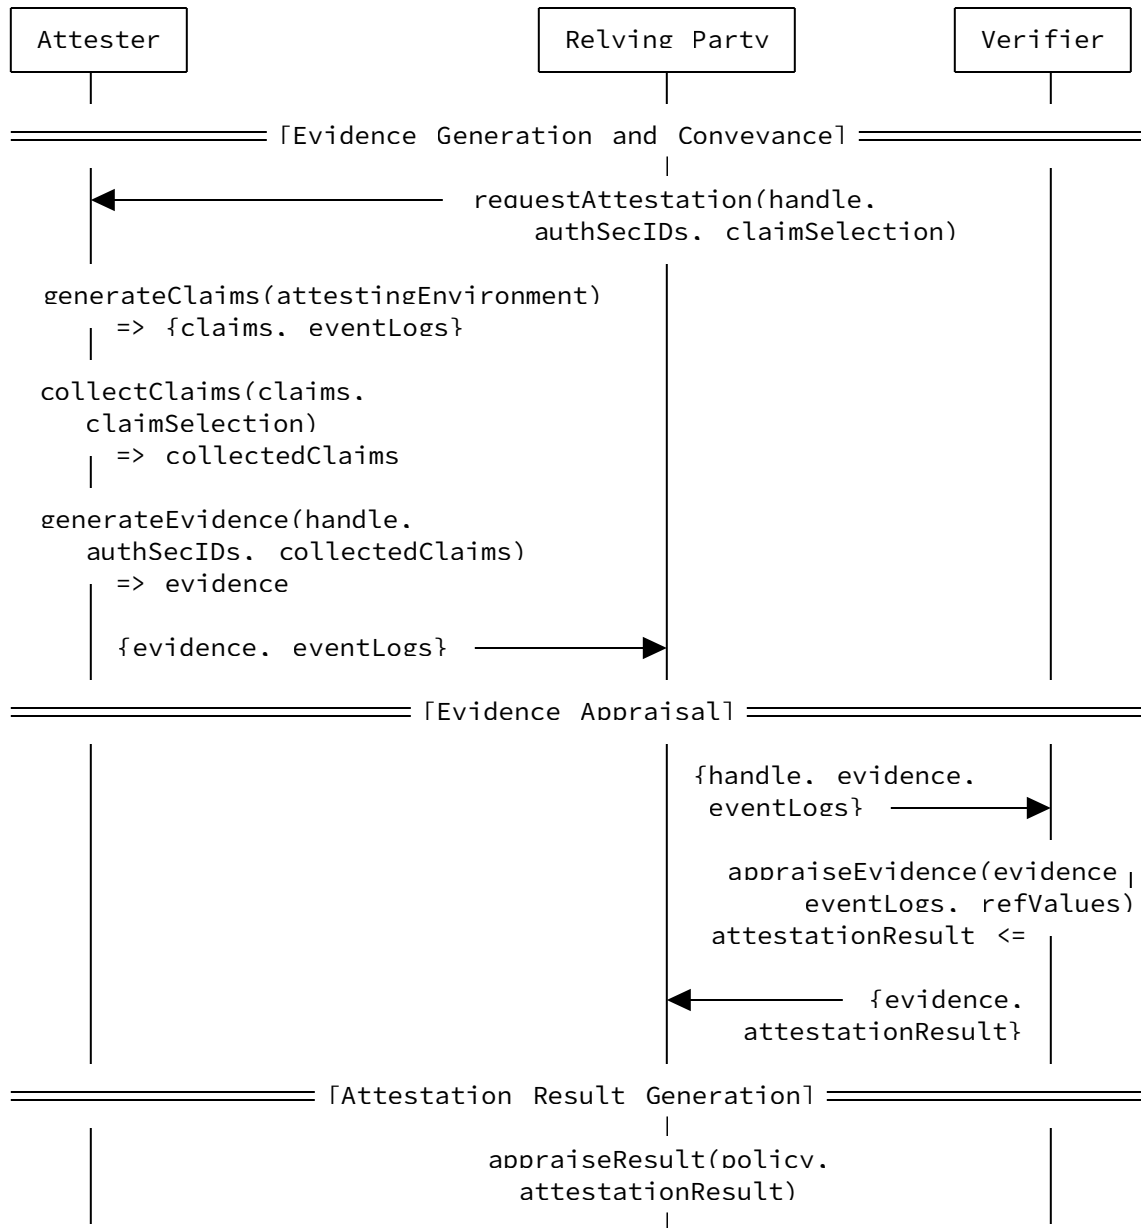
7.1.1.1. Passport Model

The passport model is so named because of its resemblance to how nations issue passports to their citizens. In this model, the attestation sequence is a two-step procedure. In the first step, an Attester conveys Evidence to a Verifier, which compares the Evidence against its appraisal policy. The Verifier then gives back an Attestation Result to the Attester, which simply caches it. In the second step, the Attester presents the Attestation Result (and possibly additional Claims/Evidence) to a Relying Party, which then compares this information against its own appraisal policy to establish the trustworthiness of the Attester.

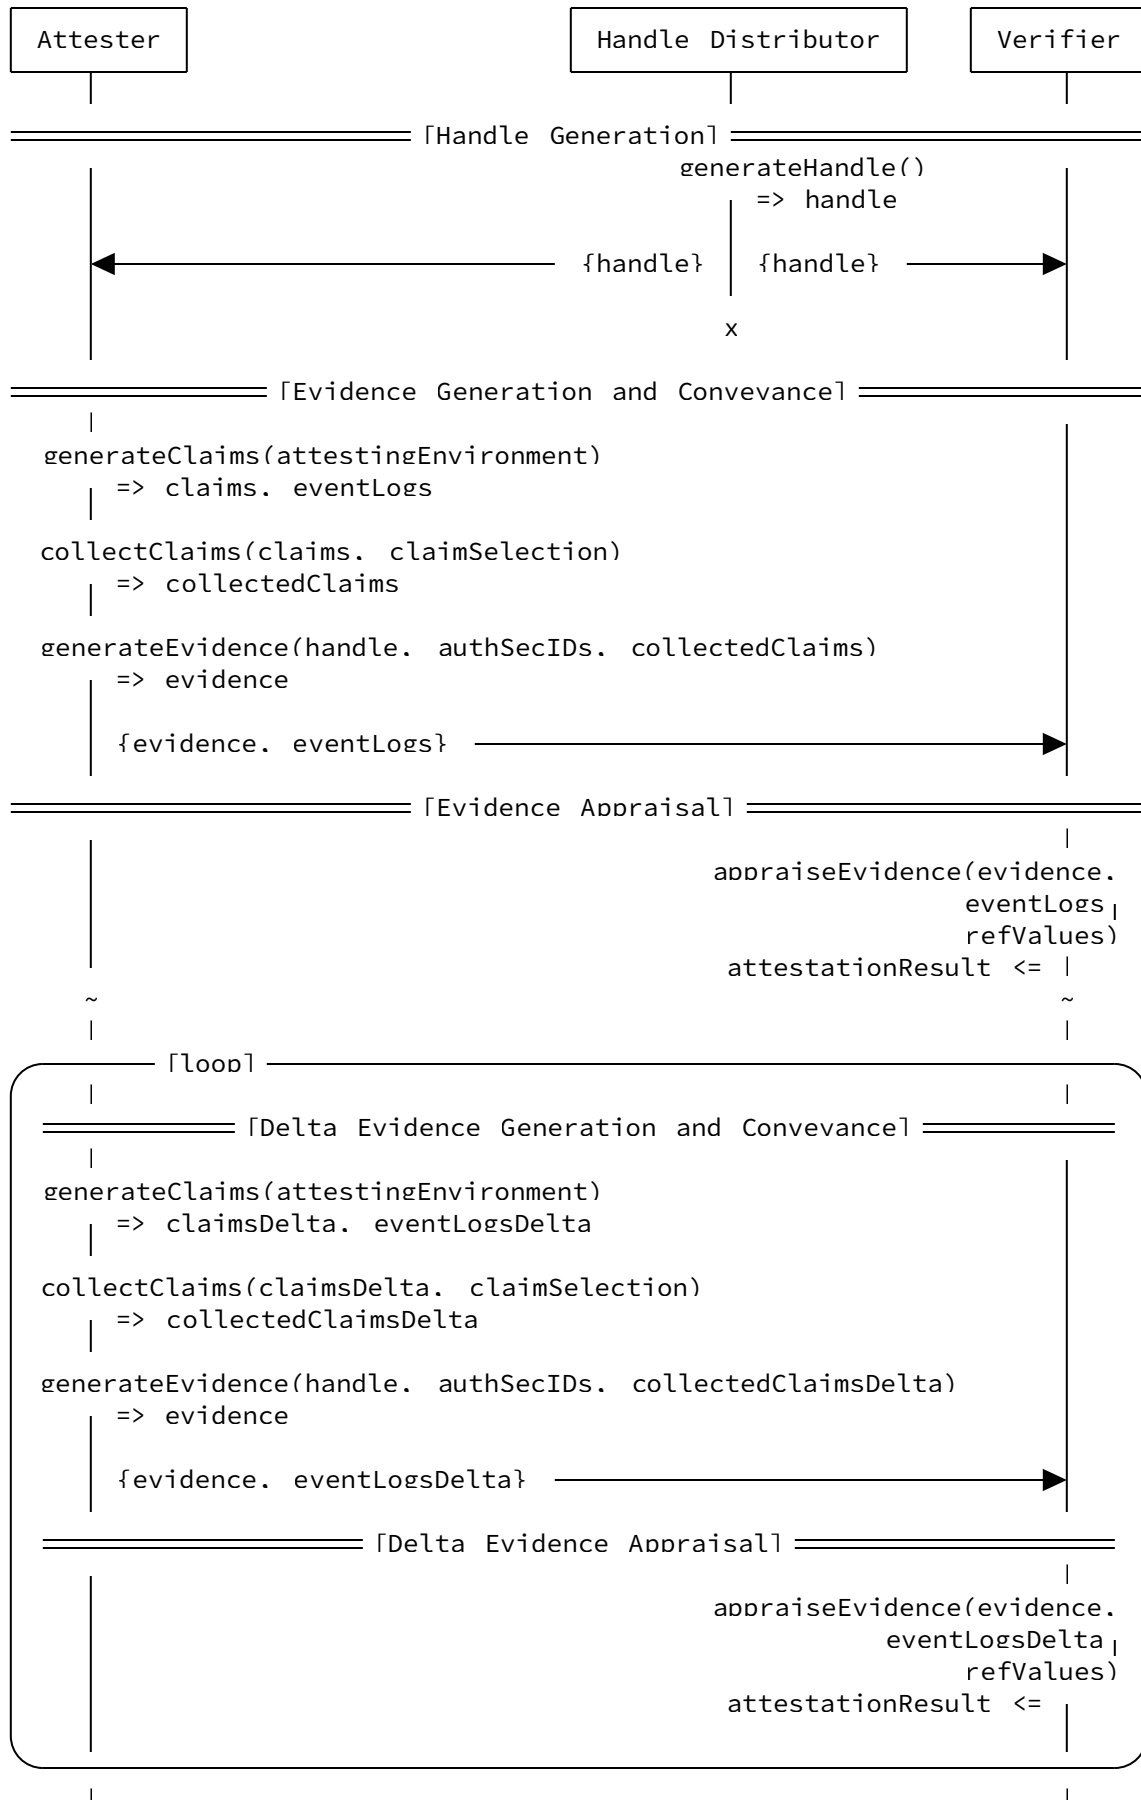


7.1.1.2. Background-Check Model

The background-check model is so named because of the resemblance of how employers and volunteer organizations perform background checks. In this model, the attestation sequence is initiated by a Relying Party. The Attester conveys Evidence to the Relying Party, which does not process its payload, but relays the message and optionally checks its signature against a policed trust anchor store. Upon receiving the Evidence, the Relying Party initiates a session with the Verifier. Once the session is established, it forwards the received Evidence to the Verifier. The Verifier appraises the received Evidence according to its appraisal policy for Evidence and returns a corresponding Attestation Result to the Relying Party. The Relying Party then checks the Attestation Result against its own appraisal policy to conclude attestation.



7.2. Uni-Directional Remote Attestation



Uni-Directional Remote Attestation procedures can be initiated both by the Attester and by the Verifier. Initiation by the Attester can result in unsolicited pushes of Evidence to the Verifier. Initiation by the Verifier always results in solicited pushes to the Verifier.

The Uni-Directional model uses the same information elements as the Challenge/Response model. In the sequence diagram above, the Attester initiates the conveyance of Evidence (comparable with a RESTful POST operation or the emission of a beacon). While a request of Evidence from the Verifier would result in a sequence diagram more similar to the Challenge/Response model (comparable with a RESTful GET operation). The specific manner how Handles are created and used always remains as the distinguishing quality of this model.

In the Uni-Directional model, handles are composed of cryptographically signed trusted timestamps as shown in [\[I-D.birkholz-rats-tuda\]](#), potentially including other qualifying data. The Handles are created by an external 3rd entity -- the Handle Distributor -- which includes a trustworthy source of time, and takes on the role of a Time Stamping Authority (TSA, as initially defined in [\[RFC3161\]](#)). Timestamps created from local clocks (absolute clocks using a global timescale, as well as relative clocks, such as tick-counters) of Attesters and Verifiers MUST be cryptographically bound to fresh Handles received from the Handle Distributor. This binding provides a proof of synchronization that MUST be included in all produced Evidence. Correspondingly, conveyed Evidence in this model provides a proof that it was fresh at a certain point in time.

While periodically pushing Evidence to the Verifier, the Attester only needs to generate and convey evidence generated from Claim values that have changed and new Event Log entries since the previous conveyance. These updates reflecting the differences are called "delta" in the sequence diagram above.

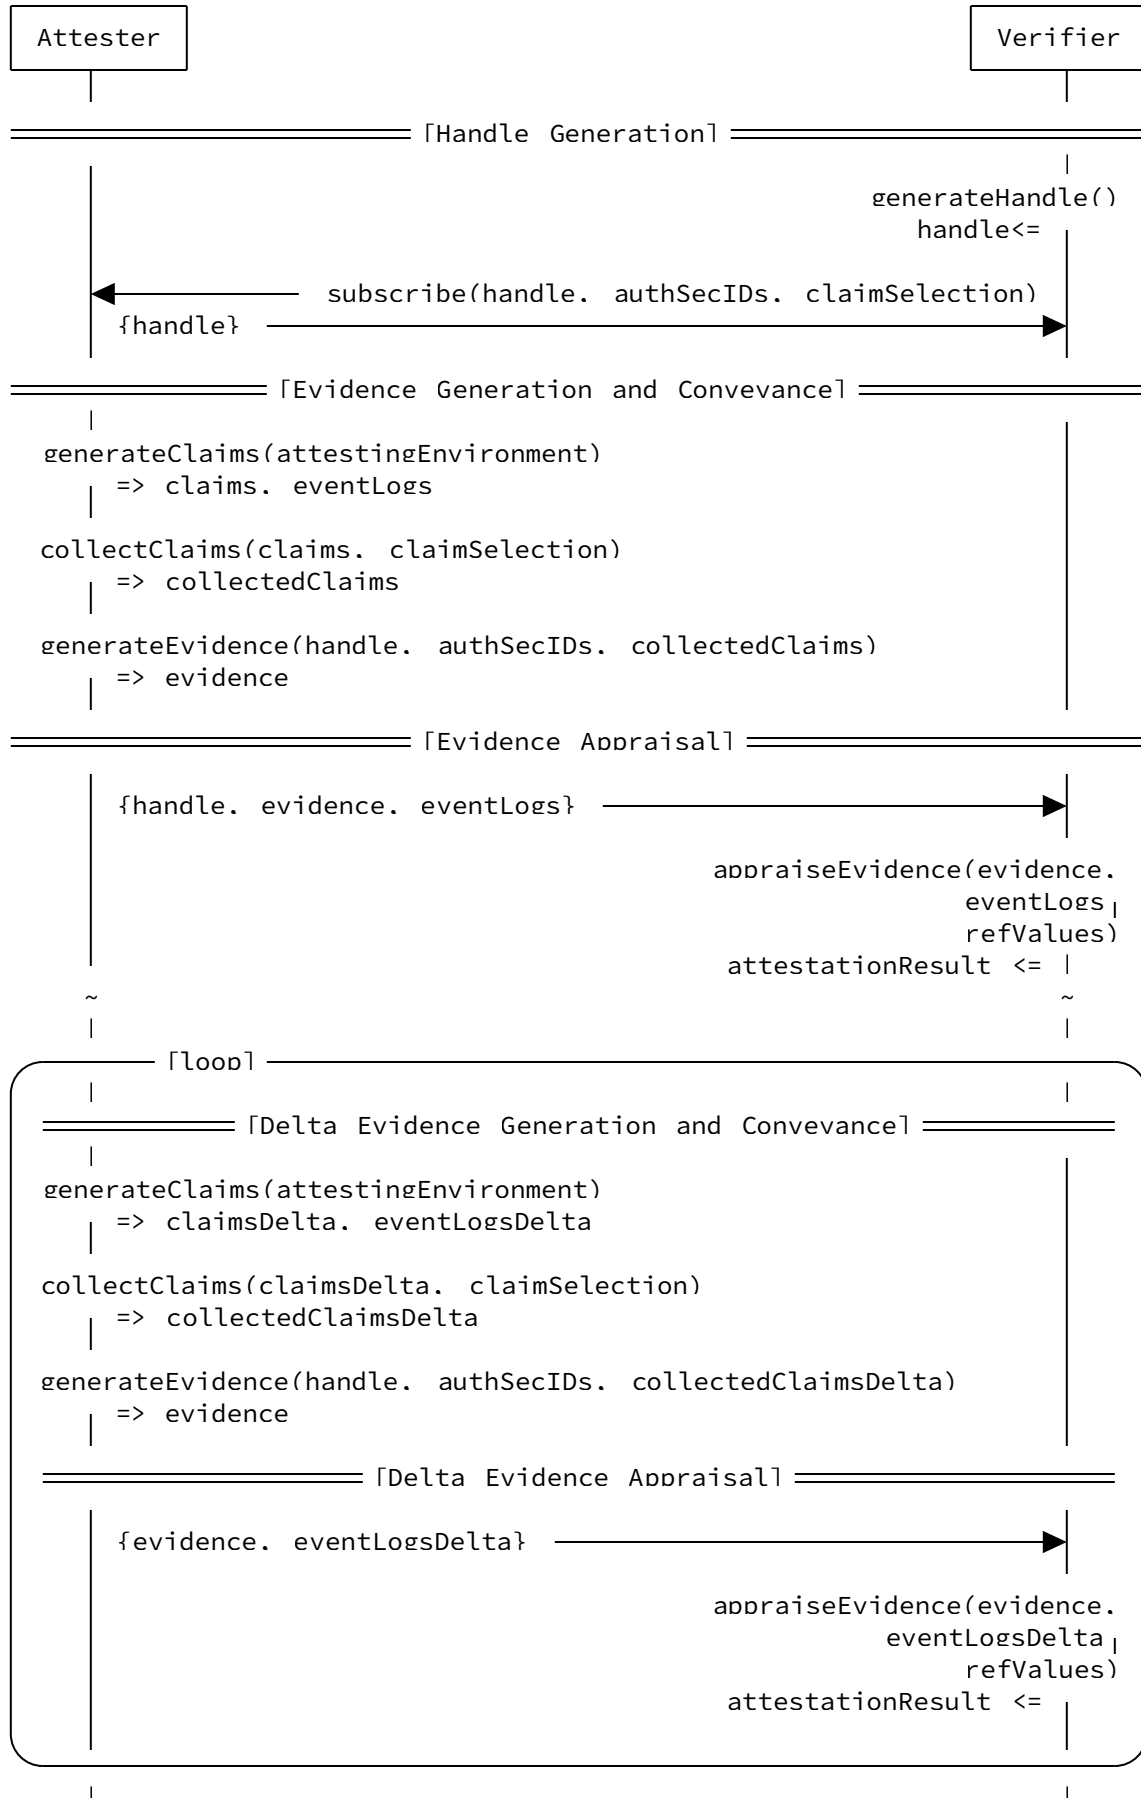
Effectively, the Uni-Directional model allows for a series of Evidence to be pushed to multiple Verifiers simultaneously. Methods to detect excessive time drift that would mandate a fresh Handle to be received by the Handle Distributor as well as timing of Handle distribution are out-of-scope of this document.

7.3. Streaming Remote Attestation

Streaming Remote Attestation serves as the foundational concept for both the observer pattern ([\[ISIS\]](#)) and the publish-subscribe pattern ([\[DesignPatterns\]](#)). It entails establishing subscription states to enable continuous remote attestation. The observer pattern directly connects observers to subjects without a broker, while the publish-subscribe pattern involves a central broker for message

distribution. In the following Subsections, streaming remote attestation without a broker (observer pattern) as well as with a broker (publish-subscribe pattern) are illustrated.

7.3.1. Streaming Remote Attestation without a Broker



The observer pattern is employed in scenarios where message delivery does not involve a central broker. Instead, an observer directly subscribes to observed resources via a dedicated mechanism. Consequently, these dedicated mechanisms contain information about the observer and are responsible for maintaining subscription state. Setting up subscription state between a Verifier and an Attester is conducted via a subscribe operation. The subscribe operation is used to convey Handles required for Evidence generation. Effectively, this allows for a series of Evidence to be pushed to a Verifier, similar to the Uni-Directional model. While a Handle Distributor is not mandatory in this model, the model is also limited to bi-lateral subscription relationships, in which each Verifier has to create and provide Handles individually. Handles provided by a specific subscribing Verifier MUST be used in Evidence generation for that specific Verifier. The streaming model without a broker uses the same information elements as the Challenge/Response and the Uni-Directional model. Methods to detect excessive time drift that would render Handles stale and mandate a fresh Handles to be conveyed via another subscribe operation are out-of-scope of this document.

7.3.2. Streaming Remote Attestation with a Broker

The publish-subscribe messaging pattern is widely used for communication in different areas. Unlike the *Streaming Remote Attestation without a Broker* interaction model, Attesters do not (need to) be aware of corresponding Verifiers. In scenarios with large numbers of Attesters and Verifiers, the publish-subscribe pattern may reduce interdependencies and improve scalability.

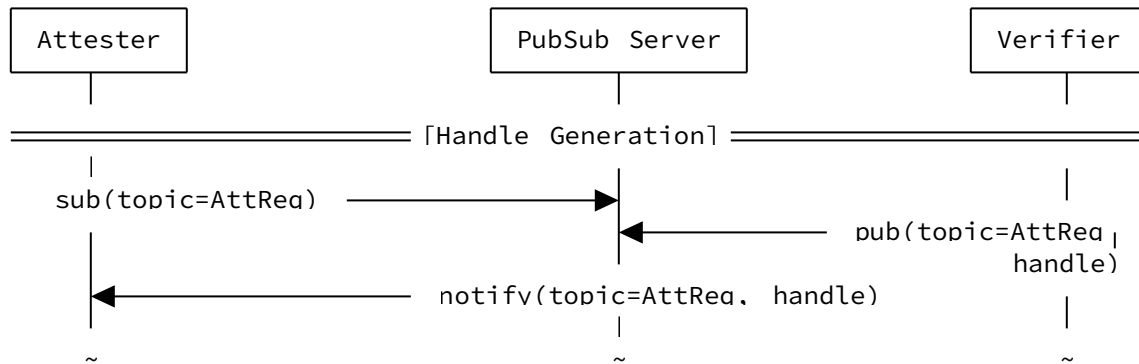
With publish-subscribe, clients typically *connect* to (or *register* with) a publish-subscribe server (PubSub server or Broker). Clients may *publish* data in the form of a *message* under a certain *topic*. *Subscribers* to that topic get *notified* whenever a message arrives under a topic, and the appropriate message is forwarded to them. Depending on the particular publish-subscribe model and implementation, clients can be either publishers or subscribers or both.

In the following sections, the interaction models *Challenge/Response Remote Attestation over Publish-Subscribe* and *Uni-Directional Remote Attestation over Publish-Subscribe* are described. There are different phases that both models go through:

1. Handle Generation
2. Evidence Generation and Conveyance
3. Evidence Appraisal
4. Attestation Result Generation

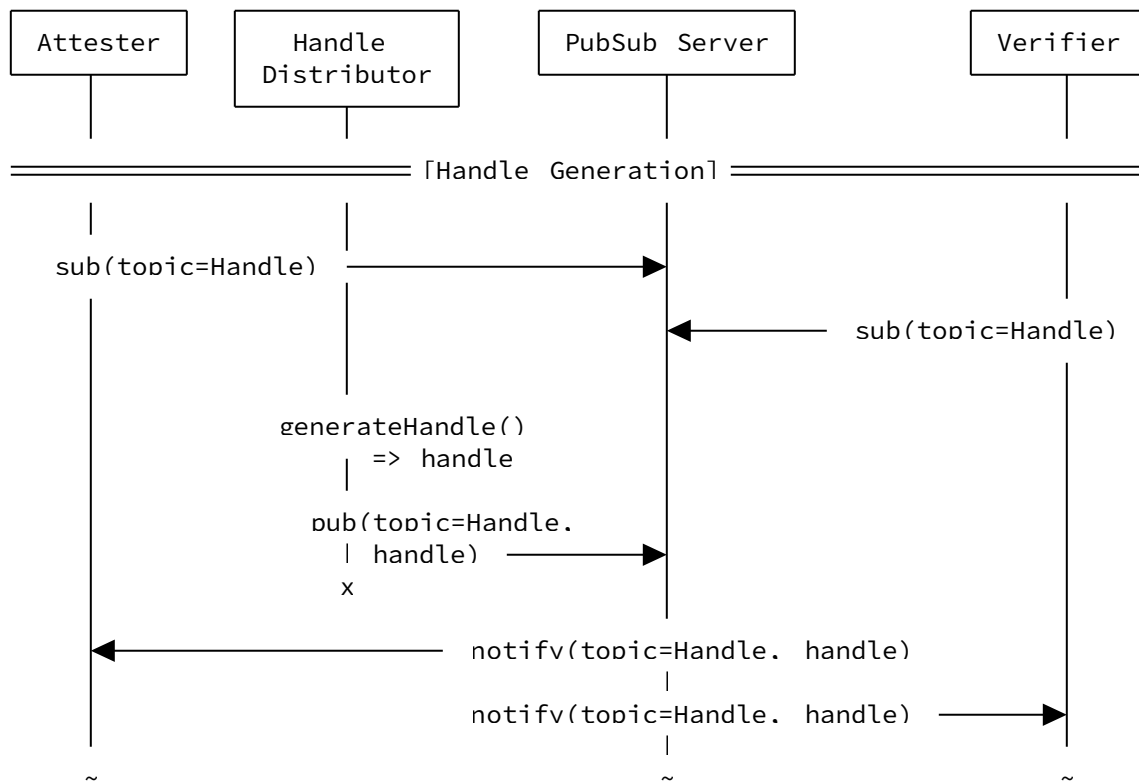
The models only differ in the handle generation phase. From a remote attestations procedure's point of view Evidence Generation, Conveyance, and Appraisal, as well as Attestation Result Generation are identical in both models.

7.3.2.1. Handle Generation for Challenge/Response Remote Attestation over Publish-Subscribe



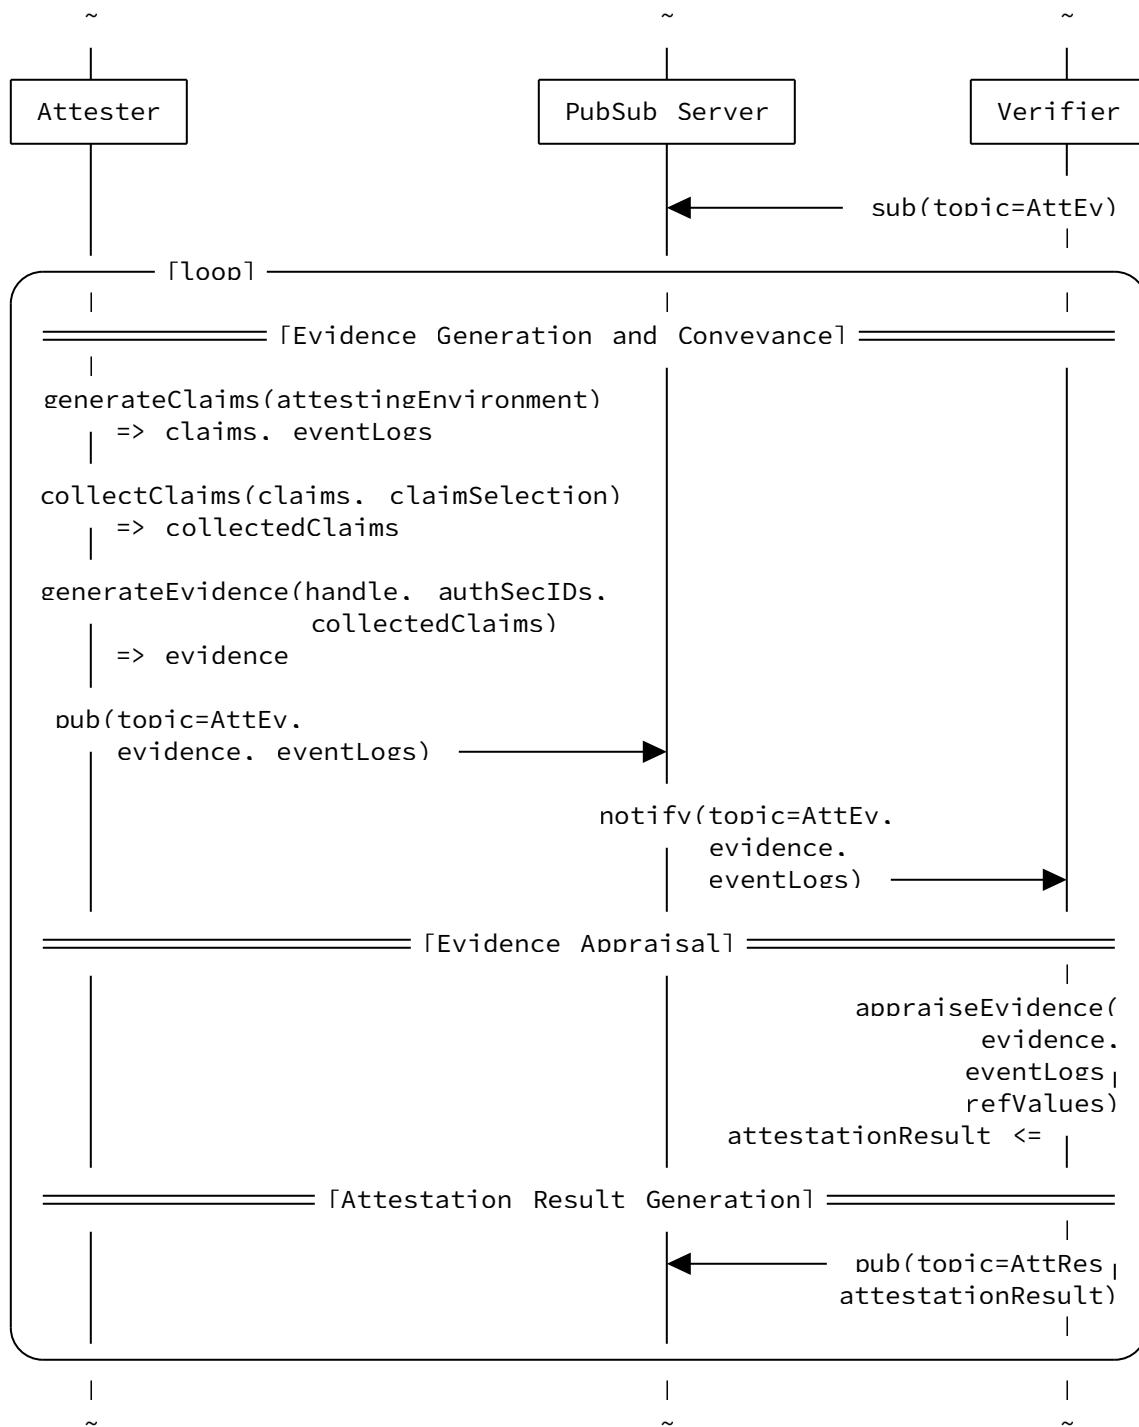
The *Challenge/Response Remote Attestation over Publish-Subscribe* interaction model uses the same information elements as the *Challenge/Response Remote Attestation* interaction model. Handles are provided by a Verifier on a per-request basis. In the sequence diagram above, an Attester subscribes to the "AttReq" (= Attestation Request) topic on the PubSub server. The Verifier publishes a Handle to the "AttReq" topic, which the PubSub server forwards to the Attester by notifying it.

7.3.2.2. Handle Generation for Uni-Directional Remote Attestation over Publish-Subscribe



The *Uni-Directional Remote Attestation over Publish-Subscribe* model uses the same information elements as the Uni-Directional Remote Attestation model. Accordingly, Handles are created by a 3rd party, the Handle Distributor. In the sequence diagram above, both an Attester and a Verifier subscribe to the topic "Handle" on the PubSub server. When the Handle Distributor generates and publishes a Handle to the "Handle" topic on the PubSub server, the PubSub server notifies the subscribers, Attester and Verifier, and forwards ("notify") the Handle to them during Handle Generation.

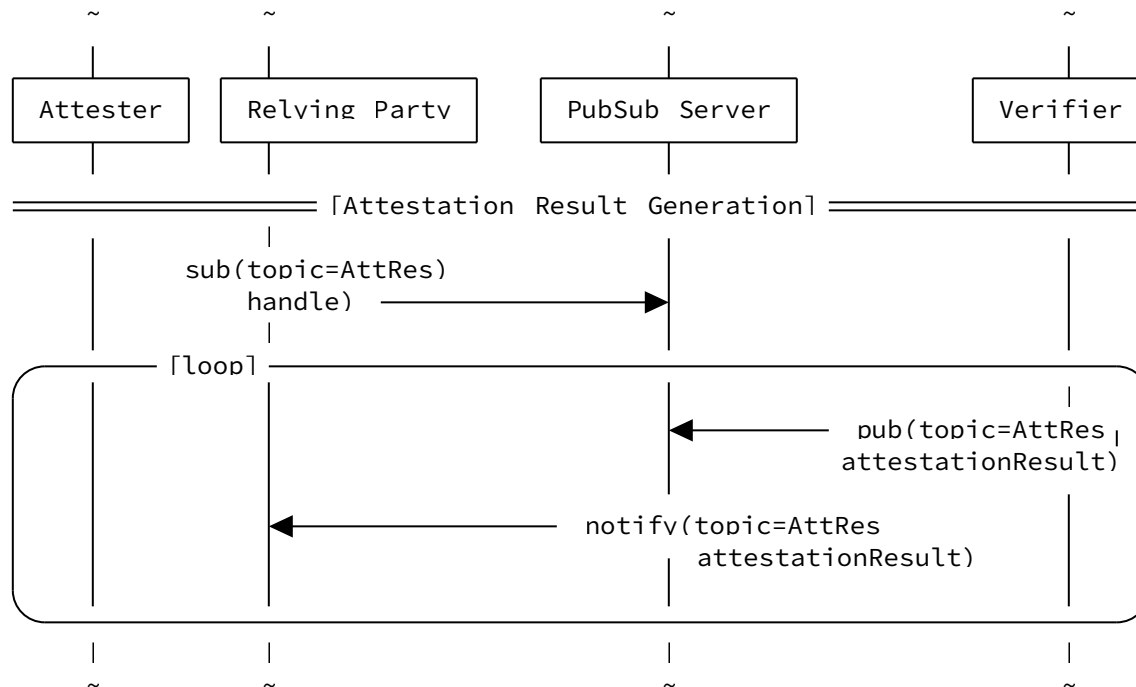
7.3.2.3. Evidence Generation and Appraisal



Exactly as in the Challenge/Response and Uni-Directional interaction models, there is an Evidence Generation-Appraisal loop, in which the Attester generates Evidence and the Verifier appraises it. In the Publish-Subscribe model above, the Attester publishes Evidence to the topic "AttEv" (= Attestation Evidence) on the PubSub server, to

which a Verifier subscribed before. The PubSub server notifies Verifiers, accordingly, by forwarding the attestation Evidence. Although the above diagram depicts only full attestation Evidence and Event Logs, later attestations may use 'deltas' for Evidence and Event Logs. Verifiers appraise the Evidence and publish the Attestation Result to topic "AttRes" (= Attestation Result) on the PubSub server.

7.3.2.4. Attestation Result Generation



Attestation Result Generation is the same for both publish-subscribe models, *Challenge/Response Remote Attestation over Publish-Subscribe* and *Uni-Directional Remote Attestation over Publish-Subscribe*. Relying Parties subscribe to topic AttRes (= Attestation Result) on the PubSub server. The PubSub server forwards Attestation Results to the Relying Parties as soon as they are published to topic AttRes.

7.3.2.5. Publish/Subscribe Topics

Many publish-subscribe models provide hierarchical organization of topics. This way, subscribers can subscribe to either all attestations (topic AttRes), or, for example, to topic AttRes/DbServers/Germany to receive only attestations from database servers in Germany. Further, it may be required to distinguish between uni-directional and challenge-response attestation evidence.

8. Additional Application-Specific Requirements

Depending on the use cases covered, there can be additional requirements. An exemplary subset is illustrated in this section.

8.1. Confidentiality

Confidentiality of exchanged attestation information may be desirable. This requirement usually is present when communication takes place over insecure channels, such as the public Internet. In such cases, TLS may be used as a suitable communication protocol which provides confidentiality protection. In private networks, such as carrier management networks, it must be evaluated whether or not the transport medium is considered confidential.

8.2. Mutual Authentication

In particular use cases, mutual authentication may be desirable in such a way that a Verifier also needs to prove its identity to the Attester, instead of only the Attester proving its identity to the Verifier.

8.3. Hardware-Enforcement/Support

Depending on given usage scenarios, hardware support for secure storage of cryptographic keys, crypto accelerators, as well as protected or isolated execution environments can be mandatory requirements. Well-known technologies in support of these requirements are roots of trusts, such as Hardware Security Modules (HSM), Physically Unclonable Functions (PUFs), Shielded Secrets, or Trusted Executions Environments (TEEs).

9. Implementation Status

Note to RFC Editor: Please remove this section as well as references to [\[BCP205\]](#) before AUTH48.

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [\[BCP205\]](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [BCP205], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

9.1. Implementer

The open-source implementation was initiated and is maintained by the Fraunhofer Institute for Secure Information Technology SIT.

9.2. Implementation Name

The open-source implementation is named "CHALLENGE-Response based Remote Attestation" or in short: CHARRA.

9.3. Implementation URL

The open-source implementation project resource can be located via: <https://github.com/fraunhofer-sit/charra>

9.4. Maturity

The code's level of maturity is considered to be "prototype".

9.5. Coverage and Version Compatibility

The current version ('6194b3b') implements a challenge/response interaction model and is aligned with the exemplary specification of the CoAP FETCH bodies defined in Section [Appendix A](#) of this document.

9.6. License

The CHARRA project and all corresponding code and data maintained on GitHub are provided under the BSD 3-Clause "New" or "Revised" license.

9.7. Implementation Dependencies

The implementation requires the use of the official Trusted Computing Group (TCG) open-source Trusted Software Stack (TSS) for the Trusted Platform Module (TPM) 2.0. The corresponding project resources (code and data) for Linux-based operating systems are maintained on GitHub at <https://github.com/tpm2-software/tpm2-tss/>.

The implementation uses the Constrained Application Protocol [RFC7252] (<http://coap.technology/>) and the Concise Binary Object Representation [RFC7049] (<https://cbor.io/>).

9.8. Contact

Michael Eckel (michael.eckel@sit.fraunhofer.de)

10. Security and Privacy Considerations

In a remote attestation procedure the Verifier or the Attester MAY want to cryptographically blind several attributes. For instance, information can be part of the signature after applying a one-way function (e. g., a hash function).

There is also a possibility to scramble the Nonce or Attester Identity with other information that is known to both the Verifier and Attester. A prominent example is the IP address of the Attester that usually is known by the Attester itself as well as the Verifier. This extra information can be used to scramble the Nonce in order to counter certain types of relay attacks.

11. Acknowledgments

Olaf Bergmann, Michael Richardson, and Ned Smith

12. References

12.1. Normative References

- [BCP205] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://doi.org/10.17487/RFC7942>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://doi.org/10.17487/RFC2119>>.
- [RFC3161] Adams, C., Cain, P., Pinkas, D., and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", RFC 3161, DOI 10.17487/RFC3161, August 2001, <<https://doi.org/10.17487/RFC3161>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://doi.org/10.17487/RFC5280>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://doi.org/10.17487/RFC7049>>.

[RFC7252]

Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://doi.org/10.17487/RFC7252>>.

[RFC8174]

Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://doi.org/10.17487/RFC8174>>.

[RFC8610]

Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://doi.org/10.17487/RFC8610>>.

[RFC9334]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedures (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <<https://doi.org/10.17487/RFC9334>>.

12.2. Informative References

[DAA]

Brickell, E., Camenisch, J., and L. Chen, "Direct Anonymous Attestation", page 132-145, ACM Proceedings of the 11th ACM conference on Computer and Communications Security, 2004.

[DesignPatterns]

Gamma, E., Helm, R., Johnson, R., and J. Vlissides, "Design Patterns - Elements of Reusable Object-Oriented Software", Publisher Addison-Wesley, 1994.

[I-D.birkholz-rats-tuda]

Fuchs, A., Birkholz, H., McDonald, I., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-07, 10 July 2022, <<https://datatracker.ietf.org/doc/html/draft-birkholz-rats-tuda-07>>.

[I-D.ietf-rats-tpm-based-network-device-attest]

Fedorkow, G., Voit, E., and J. Fitzgerald-McKay, "TPM-based Network Device Remote Integrity Verification", Work in Progress, Internet-Draft, draft-ietf-rats-tpm-based-network-device-attest-14, 22 March 2022, <<https://datatracker.ietf.org/>

[doc/html/draft-ietf-rats-tpm-based-network-device-attest-14](https://doi.org/10.1145/41457.37515)>.

- [ISIS] Birman, K. and T. Joseph, "Exploiting Virtual Synchrony in Distributed Systems", DOI 10.1145/41457.37515, 1987, <<https://doi.org/10.1145/41457.37515>>.
- [MQTT] OASIS, "Message Queuing Telemetry Transport (MQTT) Version 5.0 Committee Specification 02", Specification Version 5.0, 2018.
- [TNC] TCG, "TCG Trusted Network Communications TNC Architecture for Interoperability", Specification Version 2.0 Revision 13, 2017.
- [turtles] Rudnicki, R., "Turtles All the Way Down: Foundation, Edifice, and Ruin in Faulkner and McCarthy", DOI 10.1353/fau.2010.0002, The Faulkner Journal 25.2, 2010, <<https://doi.org/10.1353/fau.2010.0002>>.

Appendix A. CDDL Specification for a simple CoAP Challenge/Response Interaction

The following CDDL specification is an exemplary proof-of-concept to illustrate a potential implementation of the Challenge/Response Interaction Model. The communication protocol used is CoAP. Both the request message and the response message are exchanged via the FETCH operation and corresponding FETCH request and FETCH response body.

In this example, Evidence is created via the root-of-trust for reporting primitive operation "quote" that is provided by a TPM 2.0.

charra-bodies = charra-attestation-request / charra-attestation-response

```
charra-attestation-request = [  
    hello: bool,      ; if true, the TPM 2.0 AK Cert shall be conveyed  
    key-id: bytes,    ; the key ID to use for signing  
    nonce: bytes,     ; a (random) nonce, providing freshness and/or recen  
    pcr-selections: [ * pcr-selection ]  
]
```

```
pcr-selection = [  
    tcg-hash-alg-id: uint .size 2, ; TPM2_ALG_ID  
    pcrcs: [  
        pcr: uint .size 2  
    ]  
]
```

```
charra-attestation-response = [  
    attestation-data: bytes, ; TPMS_ATTEST.quoted  
    tpm2-signature: bytes,  
    ? ak-cert: bytes,        ; TPM2 attestation key certificate (AK Cer  
]
```

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@ietf.contact

Michael Eckel
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: michael.eckel@sit.fraunhofer.de

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com

Eric Voit
Cisco Systems

Email: evoit@cisco.com