

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 13 January 2022

H. Birkholz
Fraunhofer SIT
T. Fossati
Y. Deshpande
Arm Limited
N. Smith
Intel
W. Pan
Huawei Technologies
12 July 2021

Concise Reference Integrity Manifest
draft-birkholz-rats-corim-00

Abstract

Abstract

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the RATS Working Group mailing list (rats@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/rats/>.

Source for this draft and an issue tracker can be found at <https://github.com/ietf-rats/draft-birkholz-rats-corim>.

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 13 January 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Notation	4
2. Concise Reference Integrity Manifests	4
2.1. Typographical Conventions	5
2.2. Prefixes and Namespaces	6
2.3. Extensibility	6
2.4. Concise RIM Extension Points	6
2.5. CDDL Generic Types	7
2.5.1. Non-Empty	7
2.5.2. One-Or-More	7
3. Concise RIM Data Definition	8
3.1. The signed-corim Container	8
3.1.1. The corim-meta-map Container	8
3.1.2. The corim-entity-map Container	9
3.1.3. The validity-map Container	9
3.2. The unsigned-corim-map Container	10
3.2.1. The corim-locator-map Container	10
3.3. The concise-mid-tag Container	11
3.4. The tag-identity-map Container	11
3.5. The entity-map Container	12
3.6. The linked-tag-map Container	13
3.7. The triples-map Container	13
3.8. The environment-map Container	14
3.9. The class-map Container	15
3.10. The measurement-map and measurement-values-map Containers	16
3.10.1. The version-map Container	18
3.10.2. The svn-type-choice Enumeration	18
3.10.3. The raw-value-group Container	19
3.10.4. The ip-addr-type-choice Enumeration	19
3.10.5. The mac-addr-type-choice Enumeration	19
3.11. The verification-key-map Container	20

4. Full CDDL Definition	20
5. Privacy Considerations	34
6. Security Considerations	34
7. IANA Considerations	34
8. References	34
8.1. Normative References	34
8.2. Informative References	35
Authors' Addresses	35

1. Introduction

The Remote Attestation Procedures (RATS) architecture [I-D.ietf-rats-architecture] describes appraisal procedures for attestation Evidence and Attestation Results. Appraisal procedures for Evidence are conducted by Verifiers and are intended to assess the trustworthiness of a remote peer. Appraisal procedures for Attestation Results are conducted by Relying Parties and are intended to operationalize the assessment about a remote peer and to act appropriately based on the assessment. In order to enable their intent, appraisal procedures consume Appraisal Policies, Reference Values, and Endorsements.

This documents specifies a binary encoding for Reference Values using the Concise Binary Object Representation (CBOR). The encoding is based on three parts that are defined using the Concise Data Definition Language (CDDL):

- * Concise Reference Integrity Manifests (CoRIM),
- * Concise Module Identifiers (CoMID), and
- * Concise Software Identifier (CoSWID).

CoRIM and CoMID tags are defined in this document, CoSWID tags are defined in [I-D.ietf-sacm-coswid]. CoRIM provide a wrapper structure, in which CoMID tags, CoSWID tags, as well as corresponding metadata can be bundled and signed as a whole. CoMID tags represent hardware components and provide a counterpart to CoSWID tags, which represent software components.

In accordance to [RFC4949], software components that are stored in hardware modules are referred to as firmware. While firmware can be represented as a software component, it is also very hardware-specific and often resides directly on block devices instead of a file system. In this specification, firmware and their Reference Values are represented via CoMID tags. Reference Values for any other software components stored on a file system are represented via CoSWID tags.

In addition to CoRIM - and respective CoMID tags - this specification defines a Concise Manifest Revocation that represents a list of reference to CoRIM that are actively marked as invalid before their expiration time.

1.1. Requirements Notation

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. Concise Reference Integrity Manifests

This section specifies the Concise RIM (CoRIM) format, the Concise MID format (CoMID), and the extension to the CoSWID specification that augments CoSWID tags to express specific relationships to CoMID tags.

While each specification defines its own start rule, only CoMID and CoSWID are stand-alone specifications. The CoRIM specification - as the bundling format - has a dependency on CoMID and CoSWID and is not a stand-alone specification.

While stand-alone CoSWID tags may be signed [I-D.ietf-sacm-coswid], CoMID tags are not intended to be stand-alone and are always part of a CoRIM that must be signed. [I-D.ietf-sacm-coswid] specifies the use of COSE [RFC7231] for signing. This specification defines how to generate signed CoRIM tags with COSE to enable proof of authenticity and temper-evidence.

This document uses the Concise Data Definition Language (CDDL [RFC8610]) to define the data structure of CoRIM and CoMID tags, as well as the extensions to CoSWID. The CDDL definitions provided define nested containers. Typically, the CDDL types used for nested containers are maps. Every key used in the maps is a named type that is associated with an corresponding uint via a block of rules appended at the end of the CDDL definition.

Every set of uint keys that is used in the context of the "collision domain" of map is intended to be collision-free (each key is intended to be unique in the scope of a map, not a multimap). To accomplish that, for each map there is an IANA registry for the map members of maps.

2.1. Typographical Conventions

Type names in the following CDDL definitions follow the naming convention illustrated in table Table 1.

type trait	example	typo convention
extensible type choice	int / text / ...	"\$NAME"-type-choice"
closed type choice	int / text	NAME"-type-choice"
group choice	(1 => int // 2 => text)	"\$\$NAME"-group-choice"
group	(1 => int, 2 => text)	NAME"-group"
type	int	NAME"-type"
tagged type	#6.123(int)	"tagged-NAME"-type"
map	{ 1 => int, 2 => text }	NAME-"map"
flags	&(a: 1, b: 2)	NAME-"flags"

Table 1: Type Traits & Typographical Convention

2.2. Prefixes and Namespaces

The semantics of the information elements (attributes) defined for CoRIM, CoMID tags, and CoSWID tags are sometimes very similar, but often do not share the same scope or are actually quite different. In order to not overload the already existing semantics of the software-centric IANA registries of CoSWID tags with, for example, hardware-centric semantics of CoMID tags, new type names are introduced. For example: both CoSWID tags and CoMID tags define a tag-id. As CoSWID already specifies "tag-id", the tag-id in CoMID tags is prefixed with "comid." to disambiguate the context, resulting in "comid.tag-id". This prefixing provides a well-defined scope for the use of the types defined in this document and guarantees interoperability (no type name collisions) with the CoSWID CDDL definition. Effectively, the prefixes used in this specification enable simple hierarchical namespaces. The prefixing introduced is also based on the anticipated namespace features for CDDL.

2.3. Extensibility

Both the CoRIM and the CoMID tag specification include extension points using CDDL sockets (see [RFC8610] Section 3.9). The use of CDDL sockets allows for well-formed extensions to be defined in supplementary CDDL definitions that support additional uses of CoRIM and CoMID tags.

There are two types of extensibility supported via the extension points defined in this document. Both types allow for the addition of keys in the scope of a map.

Custom Keys: The CDDL definition allows for the use of negative integers as keys. These keys cannot take on a well-defined global semantic. They can take on custom-defined semantics in a limited or local scope, e.g. vendor-defined scope.

Registered Keys: Additional keys can be registered at IANA via separate specifications.

Both types of extensibility also allow for the definition of new nested maps that again can include additional defined keys.

2.4. Concise RIM Extension Points

The following CDDL sockets (extension points) are defined in the CoRIM specification, which allow the addition of new information structures to their respective CDDL groups.

Map Name	CDDL Socket	Defined in
corim-entity-map	\$\$corim-entity-map-extension	Section 3.1.2
unsigned-corim-map	\$\$unsigned-corim-map-extension	Section 3.2
concise-mid-tag	\$\$comid-extension	Section 3.3
tag-identity-map	\$\$tag-identity-map-extension	Section 3.4
entity-map	\$\$entity-map-extension	Section 3.5
triples-map	\$\$triples-map-extension	Section 3.7
measurement-values-map	\$\$measurement-values-map-extension	Section 3.10

Table 2: CoMID CDDL Group Extension Points

2.5. CDDL Generic Types

The CDDL definitions for CoRIM and CoMID tags use the two following generic types.

2.5.1. Non-Empty

The non-empty generic type is used to express that a map with only optional members MUST at least include one of the optional members.

```
non-empty<M> = (M) .within ({ + any => any })
```

2.5.2. One-Or-More

The one-or-more generic type allows to omit an encapsulating array, if only one member would be present.

```
one-or-more<T> = T / [ 2* T ] ; 2*
```

3. Concise RIM Data Definition

A CoRIM is a bundle of CoMID tags and/or CoSWID tags that can reference each other and that includes additional metadata about that bundle.

The root of the CDDL specification provided for CoRIM is the rule "corim" :

```
start = corim
```

3.1. The signed-corim Container

A CoRIM is signed using [RFC7231]. The additional CoRIM-specific COSE header member label corim-meta is defined as well as the corresponding type corim-meta-map as its value. This rule and its constraints MUST be followed when generating or validating a signed CoMID tag.

```
signed-corim = #6.18(COSE-Sign1-corim)
```

```
protected-signed-corim-header-map = {  
  corim.alg-id => int  
  corim.content-type => "application/rim+cbor"  
  corim.issuer-key-id => bstr  
  corim.meta => corim-meta-map  
  * cose-label => cose-values  
}
```

```
unprotected-signed-corim-header-map = {  
  * cose-label => cose-values  
}
```

```
COSE-Sign1-corim = [  
  protected: bstr .cbor protected-signed-corim-header-map  
  unprotected: unprotected-signed-corim-header-map  
  payload: bstr .cbor unsigned-corim-map  
  signature: bstr  
]
```

3.1.1. The corim-meta-map Container

This map contains the two additionally defined attributes "corim-entity-map" and "validity-map" that are used to annotate a CoRIM with metadata.


```
corim-meta-map = {  
  corim.signer => one-or-more<corim-entity-map>  
  ? corim.validity => validity-map  
}
```

corim.signer: One or more entities that created and/or signed the issued CoRIM.

corim.validity: A time period defining the validity span of a CoRIM.

3.1.2. The corim-entity-map Container

This map is used to identify the signer of a CoRIM via a dedicated entity name, a corresponding role and an optional identifying URI.

```
corim-entity-map = {  
  corim.entity-name => $entity-name-type-choice  
  ? corim.reg-id => uri  
  corim.role => $corim-role-type-choice  
  * $$corim-entity-map-extension  
}
```

\$corim-role-type-choice /= corim.manifest-creator

\$corim-role-type-choice /= corim.manifest-signer

corim.entity-name: The name of the organization that takes on the role expressed in "corim.role"

corim.reg-id: The registration identifier of the organization that has authority over the namespace for "corim.entity-name".

corim.role: The role type that is associated with the entity, e.g. the creator of the CoRIM or the signer of the CoRIM.

\$\$corim-entity-map-extension: This CDDL socket is used to add new information elements to the corim-entity-map container. See FIXME.

3.1.3. The validity-map Container

The members of this map indicate the life-span or period of validity of a CoRIM that is baked into the protected header at the time of signing.

```
validity-map = {  
  ? corim.not-before => time  
  corim.not-after => time  
}
```

corim.not-before: The timestamp indicating the CoRIM's begin of its validity period.

corim.not-after: The timestamp indicating the CoRIM's end of its validity period.

3.2. The unsigned-corim-map Container

This map contains the payload of the COSE envelope that is used to sign the CoRIM. This rule and its constraints MUST be followed when generating or validating an unsigned Concise RIM.

```
unsigned-corim-map = {  
  corim.id => $corim-id-type-choice  
  corim.tags => one-or-more<$concise-tag-type-choice>  
  ? corim.dependent-rims => one-or-more<corim-locator-map>  
  * $$unsigned-corim-map-extension  
}  
  
$corim-id-type-choice /= tstr  
$corim-id-type-choice /= uuid-type  
  
$concise-tag-type-choice /= #6.TBD-SWID(bytes .cbor concise-swid-tag)  
$concise-tag-type-choice /= #6.TBD-CoMID(bytes .cbor concise-mid-tag)  
  
corim.id: Typically a UUID or a text string that MUST uniquely  
  identify a CoRIM in a given scope.  
  
corim.tags: A collection of one or more CoMID tags and/or CoSWID  
  tags.  
  
corim.dependent-rims: One or more services available via the  
  Internet that can supply additional, possibly dependent manifests  
  (or other associated resources).  
  
$$unsigned-corim-map-extension: This CDDL socket is used to add new  
  information elements to the unsigned-corim-map container. See  
  FIXME.
```

3.2.1. The corim-locator-map Container

This map is used to locate and verify the integrity of resources provided by external services, e.g. the CoRIM provider.

```
corim-locator-map = {  
  corim.href => uri  
  ? corim.thumbprint => hash-entry  
}
```

corim.href: A pointer to a services that supplies dependent files or records.

corim.thumbprint: A digest of the reference resource.

3.3. The concise-mid-tag Container

The CDDL specification for the root concise-mid-tag map is as follows. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
concise-mid-tag = {  
  ? comid.language => language-type  
  comid.tag-identity => tag-identity-map  
  ? comid.entity => one-or-more<entity-map>  
  ? comid.linked-tags => one-or-more<linked-tag-map>  
  comid.triples => triples-map  
  * $$concise-mid-tag-extension  
}
```

The following describes each member of the concise-mid-tag root map.

comid.language: A textual language tag that conforms with the IANA Language Subtag Registry [IANA.language-subtag-registry].

comid.tag-identity: A composite identifier containing identifying attributes that enable global unique identification of a CoMID tag across versions.

comid.entity: A list of entities that contributed to the CoMID tag.

comid.linked-tags: A list of tags that are linked to this CoMID tag.

comid.triples: A set of relationships in the form of triples, representing a graph-like and semantic reference structure between tags.

\$\$comid-mid-tag-extension: This CDDL socket is used to add new information elements to the concise-mid-tag root container. See FIXME.

3.4. The tag-identity-map Container

The CDDL specification for the tag-identity-map includes all identifying attributes that enable a consumer of information to anticipate required capabilities to process the corresponding tag that map is included in. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
tag-identity-map = {  
  comid.tag-id => $tag-id-type-choice  
  comid.tag-version => tag-version-type  
  * $$tag-identity-map-extension  
}
```

```
$tag-id-type-choice /= tstr  
$tag-id-type-choice /= uuid-type
```

```
tag-version-type = uint .default 0
```

The following describes each member of the tag-identity-map container.

comid.tag-id: An identifier for a CoMID that MUST be globally unique.

comid.tag-version: An unsigned integer used as a version identifier.

\$\$tag-identity-map-extension: This CDDL socket is used to add new information elements to the tag-identity-map container. See FIXME.

3.5. The entity-map Container

This Container provides qualifying attributes that provide more context information describing the module as well its origin and purpose. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
entity-map = {  
  comid.entity-name => $entity-name-type-choice  
  ? comid.reg-id => uri  
  comid.role => one-or-more<$comid-role-type-choice>  
  * $$entity-map-extension  
}
```

```
$comid-role-type-choice /= comid.tag-creator  
$comid-role-type-choice /= comid.creator  
$comid-role-type-choice /= comid.maintainer
```

The following describes each member of the tag-identity-map container.

comid.entity-name: The name of an organization that performs the roles as indicated by comid.role.

comid.reg-id: The registration identifier of the organization that

has authority over the namespace for "comid.entity-name".

comid.role: The list of roles a CoMID entity is associated with. The entity that generates the concise-mid-tag SHOULD include a \$comid-role-type-choice value of comid.tag-creator.

\$entity-map-extension: This CDDL socket is used to add new information elements to the entity-map container. See FIXME.

3.6. The linked-tag-map Container

A list of tags that are linked to this CoMID tag.

```
linked-tag-map = {  
  comid.linked-tag-id => $tag-id-type-choice  
  comid.tag-rel => $tag-rel-type-choice  
}
```

```
$tag-rel-type-choice /= comid.supplements  
$tag-rel-type-choice /= comid.replaces
```

The following describes each member of the linked-tag-map container.

comid.linked-tag-id: The tag-id of the linked tag. A linked tag MAY be a CoMID tag or a CoSWID tag.

comid.tag-rel: The relationship type with the linked tag. The relationship type MAY be "supplements" or "replaces", as well as other types well-defined by additional specifications.

3.7. The triples-map Container

A set of directed properties that associate sets of data to provide reference values, endorsed values, verification key material or identifying key material for a specific hardware module that is a component of a composite device. The map provides the core element of CoMID tags that associate remote attestation relevant data with a distinct hardware component that runs an execution environment (a module that is either a Target Environment and/or an Attesting Environment). This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
triples-map = non-empty<{  
  ? comid.reference-triples => one-or-more<reference-triple-record>  
  ? comid.endorsed-triples => one-or-more<endorsed-triple-record>  
  ? comid.attest-key-triples => one-or-more<attest-key-triple-record>  
  ? comid.identity-triples => one-or-more<identity-triple-record>  
  * $$triples-map-extension  

```

The following describes each member of the triple-map container.

comid.reference-triples: A directed property that associates reference measurements with a module that is a Target Environment.

comid.endorsed-triples: A directed property that associates endorsed measurements with a module that is a Target Environment or Attesting Environment.

comid.attest-key-triples: A directed property that associates key material used to verify evidence generated from a module that is an attesting environment.

comid.identity-triples: A directed property that associates key material used to identify a module instance or a module class that is an identifying part of a device(-set).

\$\$triples-map-extension: This CDDL socket is used to add new information elements to the triples-map container. See FIXME.

3.8. The environment-map Container

This map represents the module(s) that a triple-map can point directed properties (relationships) from in order to associate them with external information for remote attestation, such as reference values, endorsement and endorsed values, verification key material for evidence, or identifying key material for module (re-)identification. This map can identify a single module instance via "comid.instance" or groups of modules via "comid.group". Referencing classes of modules requires the use of the more complex "class-map" container. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
environment-map = non-empty<{  
  ? comid.class => class-map  
  ? comid.instance => $instance-id-type-choice  
  ? comid.group => $group-id-type-choice  

```

```
$instance-id-type-choice /= tagged-ueid-type  
$instance-id-type-choice /= tagged-uuid-type
```

```
$group-id-type-choice /= tagged-uuid-type
```

The following describes each member of the environment-map container.

comid.class: A composite identifier for classes of environments/modules.

comid.instance: An identifier for distinct instances of environments/modules that is either a UEID or a UUID.

comid.group: An identifier for a group of environments/modules that is a UUID.

3.9. The class-map Container

This map enables a composite identifier intended to uniquely identify modules that are of a distinct class of devices. Effectively, all provided members in combination are a composite module class identifier. This rule and its constraints MUST be followed when generating or validating a CoMID tag. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
class-map = non-empty<{  
  ? comid.class-id => $class-id-type-choice  
  ? comid.vendor => tstr  
  ? comid.model => tstr  
  ? comid.layer => uint  
  ? comid.index => uint  

```

```
$class-id-type-choice /= tagged-oid-type  
$class-id-type-choice /= tagged-impl-id-type  
$class-id-type-choice /= tagged-uuid-type
```

The following describes each member of the class-map container.

3.10. The measurement-map and measurement-values-map Containers

One of the targets (range) that a triple-map can point to in order to associate it with a module (domain) is the measurement-map. This map is used to provide reference measurements values that can be compared with Evidence Claim values or Endorsements and endorsed values from other sources than the corresponding CoRIM. "measurement-map" comes with a measurement key that identifies the measured element with via a OID reference or a UUID. "measurement-values-map" contains the actual measurements associated with the module(s). Byte strings with corresponding bit masks that highlights which bits in the byte string are used as reference measurements or endorsement are located in "raw-value-group". The members of "measurement-values-map" provide well-defined and well-scoped semantics for reference measurement or endorsements with respect to a given module instance, class, or group. This rule and its constraints MUST be followed when generating or validating a CoMID tag.


```
measurement-map = {  
  ? comid.mkey => $measured-element-type-choice  
  comid.mval => measurement-values-map  
}
```

```
$measured-element-type-choice /= tagged-oid-type  
$measured-element-type-choice /= tagged-uuid-type
```

```
measurement-values-map = non-empty<{  
  ? comid.ver => version-map  
  ? comid.svn => svn-type-choice  
  ? comid.digests => digests-type  
  ? comid.flags => flags-type  
  ? raw-value-group  
  ? comid.mac-addr => mac-addr-type-choice  
  ? comid.ip-addr => ip-addr-type-choice  
  ? comid.serial-number => serial-number-type  
  ? comid.ueid => ueid-type  
  ? comid.uuid => uuid-type  
  * $$measurement-values-map-extension  

```

```
flags-type = bytes .bits operational-flags
```

```
operational-flags = &(  
  not-configured: 0  
  not-secure: 1  
  recovery: 2  
  debug: 3  
)
```

```
serial-number-type = text
```

```
digests-type = one-or-more<hash-entry>
```

The following describes each member of the measurement-map and the measurement-values-map container.

comid.mkey: An identifier for the set of measurements expressed in measurement-values-map that is either an OID or a UUID.

comid.ver: A version number measurement.

comid.svn: A security related version number measurement.

comid.digests: A digest (typically a hash value) measurement.

comid.flags: Measurements that reflect operational modes that are

made permanent at manufacturing time such that they are not expected to change during normal operation of the Attester.

raw-value-group: A measurement in the form of a byte string that can come with a corresponding bit mask defining the relevance of each bit in the byte string as a measurement.

comid.mac-addr: An EUI-48 or EUI-64 MAC address measurement.

comid.ip-addr: An Ipv4 or Ipv6 address measurement.

comid.serial-number: A measurement of a serial number in text.

comid.ueid: A measurement of a Unique Enough Identifier (UEID).

comid.uuid: A measurement of a Universally Unique Identifier (UUID).

\$\$measurement-values-map-extension: This CDDL socket is used to add new information elements to the measurement-values-map container. See FIXME.

3.10.1. The version-map Container

This map expresses reference values about version information.

```
version-map = {  
  comid.version => version-type  
  ? comid.version-scheme => $version-scheme  
}
```

```
version-type = text .default '0.0.0'
```

The following describes each member of the version-map container.

comid.version: The version in the form of a text string.

comid-version-scheme: The version-scheme of the text string value as defined in [I-D.ietf-sacm-coswid]

3.10.2. The svn-type-choice Enumeration

This choice defines the CBOR tagged Security Version Numbers (SVN) that can be used as reference values for Evidence and Endorsements.

```
svn = int
min-svn = int
tagged-svn = #6.TBD-SVN(svn)
tagged-min-svn = #6.TBD-minSVN(min-svn)
svn-type-choice = tagged-svn / tagged-min-svn
```

The following describes the types in the svn-type-choice enumeration.

tagged-svn: A specific SVN.

tagged-min-svn: A lower bound for allowed SVN.

3.10.3. The raw-value-group Container

FIXME This group can express a single raw byte value and can come with an optional bit mask that defines which bits in the byte string is used as a reference value, by setting corresponding position in the bit mask to 1.

```
raw-value-group = (
  comid.raw-value => raw-value-type
  ? comid.raw-value-mask => raw-value-mask-type
)
```

```
raw-value-type = bytes
raw-value-mask-type = bytes
```

The following describes the types in the raw-value-group Container.

comid.raw-value: FIXME Bit positions in raw-value-type that correspond to bit positions in raw-value-mask-type.

comid.raw-value-mask: A raw-value-mask-type bit corresponding to a bit in raw-value-type MUST be 1 to evaluate the corresponding raw-value-type bit.

3.10.4. The ip-addr-type-choice Enumeration

This type choice expresses IP addresses as reference values.

```
ip-addr-type-choice = ip4-addr-type / ip6-addr-type
ip4-addr-type = bytes .size 4
ip6-addr-type = bytes .size 16
```

3.10.5. The mac-addr-type-choice Enumeration

This type choice expresses MAC addresses as reference values.

```
mac-addr-type-choice = eui48-addr-type / eui64-addr-type
eui48-addr-type = bytes .size 6
eui64-addr-type = bytes .size 8
```

3.11. The verification-key-map Container

One of the targets (range) that a triple-map can point to in order to associate it with a module (domain). This map is used to provide the key material for evidence verification (effectively signature checking or a lightweight proof-of-possession of private signing key material) or for identity assertion/check (where a proof-of-possession implies a certain device identity). In support of informed trust decisions, an optional trust anchor in the form a PKIX certification path that is associated with the provided key material can be included. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
verification-key-map = {
  comid.key => pkix-base64-key-type
  ? comid.keychain => [ + pkix-base64-cert-type ]
}
```

```
pkix-base64-key-type = tstr
pkix-base64-cert-type = tstr
```

The following describes each member of the verification-key-map container.

comid.key: Verification key material in DER format base64 encoded. Typically, but not necessarily, a public key.

comid.keychain: One or more base64 encoded PKIX certificates. The certificate containing the public key in comid.key MUST be the first certificate. Additional certificates MAY follow. Each subsequent certificate SHOULD certify the previous certificate.

4. Full CDDL Definition

This section aggregates the CDDL definitions specified in this document in a full CDDL definitions including:

- * the COSE envelope for CoRIM: signed-corim
- * the CoRIM document: unsigned-corim
- * the CoMID document: concise-mid-tag

Not included in the full CDDL definition are CDDL dependencies to CoSWID. The following CDDL definitions can be found in [I-D.ietf-sacm-coswid]:

```
* the COSE envelope for CoRIM: signed-coswid

* the CoSWID document: concise-swid-tag

<CODE BEGINS>
corim = #6.500($concise-reference-integrity-manifest-type-choice)

$concise-reference-integrity-manifest-type-choice /= #6.501(unsigned-corim-map
)
$concise-reference-integrity-manifest-type-choice /= #6.502(signed-corim)

signed-corim = #6.18(COSE-Sign1-corim)

protected-signed-corim-header-map = {
  corim.alg-id => int
  corim.content-type => "application/rim+cbor"
  corim.issuer-key-id => bstr
  corim.meta => corim-meta-map
  * cose-label => cose-values
}

corim-meta-map = {
  corim.signer => one-or-more<corim-entity-map>
  ? corim.validity => validity-map
}

corim-entity-map = {
  corim.entity-name => $entity-name-type-choice
  ? corim.reg-id => uri
  corim.role => $corim-role-type-choice
  * $$corim-entity-map-extension
}

$corim-role-type-choice /= corim.manifest-creator
$corim-role-type-choice /= corim.manifest-signer

validity-map = {
  ? corim.not-before => time
  corim.not-after => time
}

unprotected-signed-corim-header-map = {
  * cose-label => cose-values
}
```

```
COSE-Sign1-corim = [  
  protected: bstr .cbor protected-signed-corim-header-map  
  unprotected: unprotected-signed-corim-header-map  
  payload: bstr .cbor unsigned-corim-map  
  signature: bstr  
]  
  
unsigned-corim-map = {  
  corim.id => $corim-id-type-choice  
  corim.tags => one-or-more<$concise-tag-type-choice>  
  ? corim.dependent-rims => one-or-more<corim-locator-map>  
  * $$unsigned-corim-map-extension  
}  
  
corim-locator-map = {  
  corim.href => uri  
  ? corim.thumbprint => hash-entry  
}  
  
$concise-tag-type-choice /= #6.505(bytes .cbor concise-swid-tag)  
$concise-tag-type-choice /= #6.506(bytes .cbor concise-mid-tag)  
  
concise-mid-tag = {  
  ? comid.language => language-type  
  comid.tag-identity => tag-identity-map  
  ? comid.entity => one-or-more<entity-map>  
  ? comid.linked-tags => one-or-more<linked-tag-map>  
  comid.triples => triples-map  
  * $$concise-mid-tag-extension  
}  
  
language-type = text  
  
tag-identity-map = {  
  comid.tag-id => $tag-id-type-choice  
  ? comid.tag-version => tag-version-type  
}  
  
$tag-id-type-choice /= tstr  
$tag-id-type-choice /= uuid-type  
  
tag-version-type = uint .default 0  
  
entity-map = {  
  comid.entity-name => $entity-name-type-choice  
  ? comid.reg-id => uri
```

```
    comid.role => one-or-more<$comid-role-type-choice>
    * $$entity-map-extension
}

$comid-role-type-choice /= comid.tag-creator
$comid-role-type-choice /= comid.creator
$comid-role-type-choice /= comid.maintainer

linked-tag-map = {
    comid.linked-tag-id => $tag-id-type-choice
    comid.tag-rel => $tag-rel-type-choice
}

$tag-rel-type-choice /= comid.supplements
$tag-rel-type-choice /= comid.replaces

triples-map = non-empty<{
    ? comid.reference-triples => one-or-more<reference-triple-record>
    ? comid.endorsed-triples => one-or-more<endorsed-triple-record>
    ? comid.attest-key-triples => one-or-more<attest-key-triple-record>
    ? comid.identity-triples => one-or-more<identity-triple-record>
    * $$triples-map-extension
}>

reference-triple-record = [
    environment-map ; target environment
    one-or-more<measurement-map> ; reference measurements
]

endorsed-triple-record = [
    environment-map ; (target or attesting) environment
    one-or-more<measurement-map> ; endorsed measurements
]

attest-key-triple-record = [
    environment-map ; attesting environment
    one-or-more<verification-key-map> ; attestation verification key(s)
]

identity-triple-record = [
    environment-map ; device identifier (instance or class)
    one-or-more<verification-key-map> ; DevID, or semantically equivalent
]

pkix-base64-key-type = tstr
pkix-base64-cert-type = tstr

verification-key-map = {
```

```

    ; Verification key in DER format base64-encoded.
    ; Typically, but not necessarily a public key.
    comid.key => pkix-base64-key-type
    ; Optional X.509 certificate chain corresponding to the public key
    ; in comid.key, encoded as an array of one or more base64-encoded
    ; DER PKIX certificates. The certificate containing the public key
    ; in comid.key MUST be the first certificate. This MAY be followed
    ; by additional certificates, with each subsequent certificate
    ; being the one used to certify the previous one.
    ? comid.keychain => [ + pkix-base64-cert-type ]
}

environment-map = non-empty<{
    ? comid.class => class-map
    ? comid.instance => $instance-id-type-choice
    ? comid.group => $group-id-type-choice
}>

class-map = non-empty<{
    ; TODO(tho) add text "class-map SHOULD include class-id"
    ? comid.class-id => $class-id-type-choice
    ? comid.vendor => tstr
    ? comid.model => tstr
    ? comid.layer => uint
    ? comid.index => uint
}>

$class-id-type-choice /= tagged-oid-type
$class-id-type-choice /= tagged-impl-id-type
$class-id-type-choice /= tagged-uuid-type

$instance-id-type-choice /= tagged-ueid-type
$instance-id-type-choice /= tagged-uuid-type

$group-id-type-choice /= tagged-uuid-type

oid-type = bytes
tagged-oid-type = #6.111(oid-type)

tagged-uuid-type = #6.37(uuid-type)

ueid-type = bytes .size 33
tagged-ueid-type = #6.550(ueid-type)

impl-id-type = bytes .size 32
tagged-impl-id-type = #6.551(impl-id-type)

$measured-element-type-choice /= tagged-oid-type

```



```
$measured-element-type-choice /= tagged-uuid-type

measurement-map = {
  ? comid.mkey => $measured-element-type-choice
  comid.mval => measurement-values-map
}

measurement-values-map = non-empty<{
  ? comid.ver => version-map
  ? comid.svn => svn-type-choice
  ? comid.digests => digests-type
  ? comid.flags => flags-type
  ? raw-value-group
  ? comid.mac-addr => mac-addr-type-choice
  ? comid.ip-addr => ip-addr-type-choice
  ? comid.serial-number => serial-number-type
  ? comid.ueid => ueid-type
  ? comid.uuid => uuid-type
  * $$measurement-values-map-extension
}>

version-map = {
  comid.version => version-type
  ? comid.version-scheme => $version-scheme
}
version-type = text .default '0.0.0'

svn = int
min-svn = int
tagged-svn = #6.552(svn)
tagged-min-svn = #6.553(min-svn)
svn-type-choice = tagged-svn / tagged-min-svn

flags-type = bytes .bits operational-flags

operational-flags = &(
  not-configured: 0
  not-secure: 1
  recovery: 2
  debug: 3
)

raw-value-group = (
  comid.raw-value => raw-value-type
  ? comid.raw-value-mask => raw-value-mask-type
)

raw-value-type = bytes
```

```
raw-value-mask-type = bytes

ip-addr-type-choice = ip4-addr-type / ip6-addr-type
ip4-addr-type = bytes .size 4
ip6-addr-type = bytes .size 16

mac-addr-type-choice = eui48-addr-type / eui64-addr-type
eui48-addr-type = bytes .size 6
eui64-addr-type = bytes .size 8

serial-number-type = text

digests-type = one-or-more<hash-entry>

concise-swid-tag = {
  tag-id => text / bstr .size 16,
  tag-version => integer,
  ? corpus => bool,
  ? patch => bool,
  ? supplemental => bool,
  software-name => text,
  ? software-version => text,
  ? version-scheme => $version-scheme,
  ? media => text,
  ? software-meta => one-or-more<software-meta-entry>,
  entity => one-or-more<entity-entry>,
  ? link => one-or-more<link-entry>,
  ? payload-or-evidence,
  * $$coswid-extension,
  global-attributes,
}

payload-or-evidence /= ( payload => payload-entry )
payload-or-evidence /= ( evidence => evidence-entry )

any-uri = uri
label = text / int

$version-scheme /= multipartnumeric
$version-scheme /= multipartnumeric-suffix
$version-scheme /= alphanumeric
$version-scheme /= decimal
$version-scheme /= semver
$version-scheme /= int / text

any-attribute = (
  label => one-or-more<text> / one-or-more<int>
```

```
)

one-or-more<T> = T / [ 2* T ]

global-attributes = (
  ? lang => text,
  * any-attribute,
)

hash-entry = [
  hash-alg-id: int,
  hash-value: bytes,
]

entity-entry = {
  entity-name => text,
  ? reg-id => any-uri,
  role => one-or-more<$role>,
  ? thumbprint => hash-entry,
  * $$entity-extension,
  global-attributes,
}

$role /= tag-creator
$role /= software-creator
$role /= aggregator
$role /= distributor
$role /= licensor
$role /= maintainer
$role /= int / text

link-entry = {
  ? artifact => text,
  href => any-uri,
  ? media => text,
  ? ownership => $ownership,
  rel => $rel,
  ? media-type => text,
  ? use => $use,
  * $$link-extension,
  global-attributes,
}

$ownership /= shared
$ownership /= private
$ownership /= abandon
$ownership /= int / text
```

```
$rel /= ancestor
$rel /= component
$rel /= feature
$rel /= installationmedia
$rel /= packageinstaller
$rel /= parent
$rel /= patches
$rel /= requires
$rel /= see-also
$rel /= supersedes
$rel /= supplemental
$rel /= -256..64436 / text

$use /= optional
$use /= required
$use /= recommended
$use /= int / text

software-meta-entry = {
  ? activation-status => text,
  ? channel-type => text,
  ? colloquial-version => text,
  ? description => text,
  ? edition => text,
  ? entitlement-data-required => bool,
  ? entitlement-key => text,
  ? generator => text,
  ? persistent-id => text,
  ? product => text,
  ? product-family => text,
  ? revision => text,
  ? summary => text,
  ? unspsc-code => text,
  ? unspsc-version => text,
  * $$software-meta-extension,
  global-attributes,
}

path-elements-group = ( ? directory => one-or-more<directory-entry>,
                        ? file => one-or-more<file-entry>,
                        )

resource-collection = (
  path-elements-group,
  ? process => one-or-more<process-entry>,
  ? resource => one-or-more<resource-entry>,
  * $$resource-collection-extension,
)
```

```
file-entry = {
  filesystem-item,
  ? size => uint,
  ? file-version => text,
  ? hash => hash-entry,
  * $$file-extension,
  global-attributes,
}

directory-entry = {
  filesystem-item,
  ? path-elements => { path-elements-group },
  * $$directory-extension,
  global-attributes,
}

process-entry = {
  process-name => text,
  ? pid => integer,
  * $$process-extension,
  global-attributes,
}

resource-entry = {
  type => text,
  * $$resource-extension,
  global-attributes,
}

filesystem-item = (
  ? key => bool,
  ? location => text,
  fs-name => text,
  ? root => text,
)

payload-entry = {
  resource-collection,
  * $$payload-extension,
  global-attributes,
}

evidence-entry = {
  resource-collection,
  ? date => integer-time,
  ? device-id => text,
  * $$evidence-extension,
  global-attributes,
```

```
}  
  
integer-time = #6.1(int)  
  
tag-id = 0  
software-name = 1  
entity = 2  
evidence = 3  
link = 4  
software-meta = 5  
payload = 6  
hash = 7  
corpus = 8  
patch = 9  
media = 10  
supplemental = 11  
tag-version = 12  
software-version = 13  
version-scheme = 14  
lang = 15  
directory = 16  
file = 17  
process = 18  
resource = 19  
size = 20  
file-version = 21  
key = 22  
location = 23  
fs-name = 24  
root = 25  
path-elements = 26  
process-name = 27  
pid = 28  
type = 29  
entity-name = 31  
reg-id = 32  
role = 33  
thumbprint = 34  
date = 35  
device-id = 36  
artifact = 37  
href = 38  
ownership = 39  
rel = 40  
media-type = 41  
use = 42  
activation-status = 43  
channel-type = 44
```

colloquial-version = 45
description = 46
edition = 47
entitlement-data-required = 48
entitlement-key = 49
generator = 50
persistent-id = 51
product = 52
product-family = 53
revision = 54
summary = 55
unspsc-code = 56
unspsc-version = 57

multipartnumeric = 1
multipartnumeric-suffix = 2
alphanumeric = 3
decimal = 4
semver = 16384

tag-creator=1
software-creator=2
aggregator=3
distributor=4
licensor=5
maintainer=6

shared=1
private=2
abandon=3

ancestor=1
component=2
feature=3
installationmedia=4
packageinstaller=5
parent=6
patches=7
requires=8
see-also=9
supersedes=10

optional=1
required=2
recommended=3

comid.language = 0
comid.tag-identity = 1

```
comid.entity = 2
comid.linked-tags = 3
comid.triples = 4

comid.tag-id = 0
comid.tag-version = 1

comid.entity-name = 0
comid.reg-id = 1
comid.role = 2

comid.linked-tag-id = 0
comid.tag-rel = 1

comid.reference-triples = 0
comid.endorsed-triples = 1
comid.identity-triples = 2
comid.attest-key-triples = 3

comid.class = 0
comid.instance = 1
comid.group = 2

comid.class-id = 0
comid.vendor = 1
comid.model = 2
comid.layer = 3
comid.index = 4

comid.mkey = 0
comid.mval = 1

comid.ver = 0
comid.svn = 1
comid.digests = 2
comid.flags = 3
comid.raw-value = 4
comid.raw-value-mask = 5
comid.mac-addr = 6
comid.ip-addr = 7
comid.serial-number = 8
comid.ueid = 9
comid.uuid = 10

comid.key = 0
comid.keychain = 1

comid.version = 0
```



```
comid.version-scheme = 1

comid.supplements = 0

comid.replaces = 1

comid.tag-creator = 0
comid.creator = 1
comid.maintainer = 2


corim.id = 0
corim.tags = 1
corim.dependent-rims = 2

corim.href = 0
corim.thumbprint = 1

corim.alg-id = 1
corim.content-type = 3
corim.issuer-key-id = 4
corim.meta = 8

corim.not-before = 0
corim.not-after = 1

corim.signer = 0
corim.validity = 1

corim.entity-name = 0
corim.reg-id = 1
corim.role = 2

corim.manifest-creator = 1

corim.manifest-signer = 2


non-empty<M> = (M) .within ({ + any => any })


cose-label = int / tstr
cose-values = any

$entity-name-type-choice /= text

$corim-id-type-choice /= tstr
```

```
$corim-id-type-choice /= uuid-type
```

```
uuid-type = bytes .size 16
```

```
<CODE ENDS>
```

5. Privacy Considerations

Privacy Considerations

6. Security Considerations

Security Considerations

7. IANA Considerations

See Body Section 2.

8. References

8.1. Normative References

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-rats-architecture-12>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-17, 22 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-sacm-coswid-17>>.

[IANA.language-subtag-registry]

IANA, "Language Subtag Registry", <<http://www.iana.org/assignments/language-subtag-registry>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.

- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/rfc/rfc7231>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/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://www.rfc-editor.org/rfc/rfc8610>>.

8.2. Informative References

- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/rfc/rfc4949>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Thomas Fossati
Arm Limited
United Kingdom

Email: Thomas.Fossati@arm.com

Yogesh Deshpande
Arm Limited
United Kingdom

Email: yogesh.deshpande@arm.com

Ned Smith
Intel Corporation
United States of America

Email: ned.smith@intel.com

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 30 July 2022

H. Birkholz
Fraunhofer SIT
T. Fossati
Y. Deshpande
Arm Limited
N. Smith
Intel
W. Pan
Huawei Technologies
26 January 2022

Concise Reference Integrity Manifest
draft-birkholz-rats-corim-02

Abstract

Remote Attestation Procedures (RATS) enable Relying Parties to put trust in the trustworthiness of a remote Attester and therefore to decide if to engage in secure interactions with it - or not. Evidence about trustworthiness can be rather complex, voluminous or Attester-specific. As it is deemed unrealistic that every Relying Party is capable of the appraisal of Evidence, that burden is taken on by a Verifier. In order to conduct Evidence appraisal procedures, a Verifier requires not only fresh Evidence from an Attester, but also trusted Endorsements and Reference Values from Endorsers, such as manufacturers, distributors, or owners. This document specifies Concise Reference Integrity Manifests (CoRIM) that represent Endorsements and Reference Values in CBOR format. Composite devices or systems are represented by a collection of Concise Module Identifiers (CoMID) and Concise Software Identifiers (CoSWID) bundled in a CoRIM document.

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 30 July 2022.

Copyright Notice

Copyright (c) 2022 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	3
1.1. Requirements Notation	4
2. Concise Reference Integrity Manifests	4
2.1. Typographical Conventions	5
2.2. Prefixes and Namespaces	6
2.3. Extensibility	7
2.4. Concise RIM Extension Points	7
2.5. CDDL Generic Types	8
2.5.1. Non-Empty	8
2.5.2. One-Or-More	8
3. Concise RIM Data Definition	9
3.1. The signed-corim Container	9
3.1.1. The corim-meta-map Container	9
3.1.2. The corim-signer-map Container	10
3.1.3. The validity-map Container	10
3.2. The corim-map Container	11
3.2.1. The corim-entity-map Container	12
3.2.2. The corim-locator-map Container	12
3.3. The concise-mid-tag Container	13
3.4. The tag-identity-map Container	13
3.5. The entity-map Container	14
3.6. The linked-tag-map Container	15
3.7. The triples-map Container	15
3.8. The environment-map Container	16
3.9. The class-map Container	17
3.10. The measurement-map and measurement-values-map Containers	17
3.10.1. The version-map Container	19
3.10.2. The svn-type-choice Enumeration	20
3.10.3. The raw-value-group Container	20

3.10.4. The ip-addr-type-choice Enumeration	21
3.10.5. The mac-addr-type-choice Enumeration	21
3.11. The verification-key-map Container	21
4. Full CDDL Definition	22
5. Privacy Considerations	35
6. Security Considerations	35
7. IANA Considerations	35
7.1. COSE Header Parameters Registry	35
7.2. CoRIM Map Items Registry	36
7.3. CoRIM Entity-Map Items Registry	36
7.4. CoRIM Entity-Types Registry	37
7.5. CoMID Map Items Registry	38
7.6. CoMID Entity-Map Items Registry	39
7.7. CoMID Triples-Map Items Registry	40
7.8. CoMID Measurement-Values-Map Items Registry	41
7.9. CoMID Tag-Relationship-Types Registry	42
7.10. CoMID Role-Types Registry	43
7.11. rim+cbor Media Type Registration	44
7.12. CoAP Content-Format Registration	45
7.13. CoRIM CBOR Tag Registration	46
7.14. CoMID CBOR Tag Registration	46
8. References	47
8.1. Normative References	47
8.2. Informative References	48
Authors' Addresses	48

1. Introduction

The Remote Attestation Procedures (RATS) architecture [I-D.ietf-rats-architecture] describes appraisal procedures for attestation Evidence and Attestation Results. Appraisal procedures for Evidence are conducted by Verifiers and are intended to assess the trustworthiness of a remote peer. Appraisal procedures for Attestation Results are conducted by Relying Parties and are intended to operationalize the assessment about a remote peer and to act appropriately based on the assessment. In order to enable their intent, appraisal procedures consume Appraisal Policies, Reference Values, and Endorsements.

This documents specifies a binary encoding for Reference Values using the Concise Binary Object Representation (CBOR). The encoding is based on three parts that are defined using the Concise Data Definition Language (CDDL):

- * Concise Reference Integrity Manifests (CoRIM),
- * Concise Module Identifiers (CoMID), and

* Concise Software Identifier (CoSWID).

CoRIM and CoMID tags are defined in this document, CoSWID tags are defined in [I-D.ietf-sacm-coswid]. CoRIM provide a wrapper structure, in which CoMID tags, CoSWID tags, as well as corresponding metadata can be bundled and signed as a whole. CoMID tags represent hardware components and provide a counterpart to CoSWID tags, which represent software components.

In accordance to [RFC4949], software components that are stored in hardware modules are referred to as firmware. While firmware can be represented as a software component, it is also very hardware-specific and often resides directly on block devices instead of a file system. In this specification, firmware and their Reference Values are represented via CoMID tags. Reference Values for any other software components stored on a file system are represented via CoSWID tags.

In addition to CoRIM - and respective CoMID tags - this specification defines a Concise Manifest Revocation that represents a list of reference to CoRIM that are actively marked as invalid before their expiration time.

1.1. Requirements Notation

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. Concise Reference Integrity Manifests

This section specifies the Concise RIM (CoRIM) format, the Concise MID format (CoMID), and the extension to the CoSWID specification that augments CoSWID tags to express specific relationships to CoMID tags.

While each specification defines its own start rule, only CoMID and CoSWID are stand-alone specifications. The CoRIM specification - as the bundling format - has a dependency on CoMID and CoSWID and is not a stand-alone specification.

While stand-alone CoSWID tags may be signed [I-D.ietf-sacm-coswid], CoMID tags are not intended to be stand-alone and are always part of a CoRIM that must be signed. [I-D.ietf-sacm-coswid] specifies the use of COSE [RFC7231] for signing. This specification defines how to generate signed CoRIM tags with COSE to enable proof of authenticity and temper-evidence.

This document uses the Concise Data Definition Language (CDDL [RFC8610]) to define the data structure of CoRIM and CoMID tags, as well as the extensions to CoSWID. The CDDL definitions provided define nested containers. Typically, the CDDL types used for nested containers are maps. Every key used in the maps is a named type that is associated with an corresponding uint via a block of rules appended at the end of the CDDL definition.

Every set of uint keys that is used in the context of the "collision domain" of map is intended to be collision-free (each key is intended to be unique in the scope of a map, not a multimap). To accomplish that, for each map there is an IANA registry for the map members of maps.

2.1. Typographical Conventions

Type names in the following CDDL definitions follow the naming convention illustrated in table Table 1.

type trait	example	typo convention
extensible type choice	int / text / ...	\$NAME-type-choice
closed type choice	int / text	NAME-type-choice
group choice	(1 => int // 2 => text)	\$\$NAME-group-choice
group	(1 => int, 2 => text)	NAME-group
type	int	NAME-type
tagged type	#6.123(int)	tagged-NAME-type
map	{ 1 => int, 2 => text }	NAME-map
flags	&(a: 1, b: 2)	NAME-flags

Table 1: Type Traits & Typographical Convention

2.2. Prefixes and Namespaces

The semantics of the information elements (attributes) defined for CoRIM, CoMID tags, and CoSWID tags are sometimes very similar, but often do not share the same scope or are actually quite different. In order to not overload the already existing semantics of the software-centric IANA registries of CoSWID tags with, for example, hardware-centric semantics of CoMID tags, new type names are introduced. For example: both CoSWID tags and CoMID tags define a tag-id. As CoSWID already specifies tag-id, the tag-id in CoMID tags is prefixed with `comid.` to disambiguate the context, resulting in `comid.tag-id`. This prefixing provides a well-defined scope for the use of the types defined in this document and guarantees interoperability (no type name collisions) with the CoSWID CDDL definition. Effectively, the prefixes used in this specification enable simple hierarchical namespaces. The prefixing introduced is also based on the anticipated namespace features for CDDL.

2.3. Extensibility

Both the CoRIM and the CoMID tag specification include extension points using CDDL sockets (see [RFC8610] Section 3.9). The use of CDDL sockets allows for well-formed extensions to be defined in supplementary CDDL definitions that support additional uses of CoRIM and CoMID tags.

There are two types of extensibility supported via the extension points defined in this document. Both types allow for the addition of keys in the scope of a map.

Custom Keys: The CDDL definition allows for the use of negative integers as keys. These keys cannot take on a well-defined global semantic. They can take on custom-defined semantics in a limited or local scope, e.g. vendor-defined scope.

Registered Keys: Additional keys can be registered at IANA via separate specifications.

Both types of extensibility also allow for the definition of new nested maps that again can include additional defined keys.

2.4. Concise RIM Extension Points

The following CDDL sockets (extension points) are defined in the CoRIM specification, which allow the addition of new information structures to their respective CDDL groups.

Map Name	CDDL Socket	Defined in
corim-entity-map	\$\$corim-entity-map-extension	Section 3.2.1
unsigned-corim-map	\$\$unsigned-corim-map-extension	Section 3.2
concise-mid-tag	\$\$comid-extension	Section 3.3
tag-identity-map	\$\$tag-identity-map-extension	Section 3.4
entity-map	\$\$entity-map-extension	Section 3.5
triples-map	\$\$triples-map-extension	Section 3.7
measurement-values-map	\$\$measurement-values-map-extension	Section 3.10

Table 2: CoMID CDDL Group Extension Points

2.5. CDDL Generic Types

The CDDL definitions for CoRIM and CoMID tags use the two following generic types.

2.5.1. Non-Empty

The non-empty generic type is used to express that a map with only optional members MUST at least include one of the optional members.

```
non-empty<M> = (M) .within ({ + any => any })
```

2.5.2. One-Or-More

The one-or-more generic type allows to omit an encapsulating array, if only one member would be present.

```
one-or-more<T> = T / [ 2* T ] ; 2*
```

3. Concise RIM Data Definition

A CoRIM is a bundle of CoMID tags and/or CoSWID tags that can reference each other and that includes additional metadata about that bundle.

The root of the CDDL specification provided for CoRIM is the rule `corim` :

```
start = corim
```

3.1. The signed-corim Container

A CoRIM is signed using [RFC7231]. The additional CoRIM-specific COSE header member label `corim-meta` is defined as well as the corresponding type `corim-meta-map` as its value. This rule and its constraints MUST be followed when generating or validating a signed CoRIM tag.

```
signed-corim = #6.18(COSE-Sign1-corim)
```

```
protected-corim-header-map = {  
  corim.alg-id => int  
  corim.content-type => "application/corim-unsigned+cbor"  
  corim.issuer-key-id => bstr  
  corim.meta => bstr .cbor corim-meta-map  
  * cose-label => cose-values  
}
```

```
unprotected-corim-header-map = {  
  * cose-label => cose-values  
}
```

```
COSE-Sign1-corim = [  
  protected: bstr .cbor protected-corim-header-map  
  unprotected: unprotected-corim-header-map  
  payload: bstr .cbor tagged-corim-map  
  signature: bstr  
]
```

3.1.1. The corim-meta-map Container

This map contains the two additionally defined attributes `corim-signer-map` and `validity-map` that are used to annotate a CoRIM with metadata.

```
corim-meta-map = {  
  corim.signer => corim-signer-map  
  ? corim.signature-validity => validity-map  
}
```

corim.signer: One or more entities that created and/or signed the issued CoRIM.

corim.signature-validity: A time period defining the validity span of the signature over the CoRIM.

3.1.2. The corim-signer-map Container

This map is used to identify the signer of a CoRIM via a name and an optional URI.

```
corim-signer-map = {  
  corim.signer-name => $entity-name-type-choice  
  ? corim.signer-uri => uri  
  * $$corim-signer-map-extension  
}
```

\$entity-name-type-choice /= text

corim.signer-name: The name of the organization that signs this CoRIM

corim.signer-uri: An URI uniquely linked to the organization that signs this CoRIM

\$\$corim-signer-map-extension: This CDDL socket is used to add new information elements to the corim-signer-map container. See FIXME.

3.1.3. The validity-map Container

The members of this map indicate the life-span or period of validity of a CoRIM that is baked into the protected header at the time of signing.

```
validity-map = {  
  ? corim.not-before => time  
  corim.not-after => time  
}
```

corim.not-before: The timestamp indicating the CoRIM's begin of its validity period.

corim.not-after: The timestamp indicating the CoRIM's end of its validity period.

3.2. The corim-map Container

This map contains the payload of the COSE envelope that is used to sign the CoRIM. This rule and its constraints MUST be followed when generating or validating an unsigned Concise RIM.

```
corim-map = {  
  corim.id => $corim-id-type-choice  
  corim.tags => [ + $concise-tag-type-choice ]  
  ? corim.dependent-rims => [ + corim-locator-map ]  
  ? corim.profile => [ + profile-type-choice ]  
  ? corim.rim-validity => validity-map  
  ? corim.entities => [ + corim-entity-map ]  
  * $$corim-map-extension  
}  
  
$corim-id-type-choice /= tstr  
$corim-id-type-choice /= uuid-type  
  
profile-type-choice = uri / tagged-oid-type  
  
$concise-tag-type-choice /= #6.505(bytes .cbor concise-swid-tag)  
$concise-tag-type-choice /= #6.506(bytes .cbor concise-mid-tag)  
  
corim.id: Typically a UUID or a text string that MUST uniquely  
  identify a CoRIM in a given scope.  
  
corim.tags: A collection of one or more CoMID tags and/or CoSWID  
  tags.  
  
corim.dependent-rims: One or more services available via the  
  Internet that can supply additional, possibly dependent manifests  
  (or other associated resources).  
  
corim.profile: One or more profiles that define the domain of  
  interpretation of the CoMID and/or CoSWID tags.  
  
corim.rim-validity: The validity of the CoRIM expressed as a  
  validity-map.  
  
corim.entities: One or more entities involved in the creation of  
  this CoRIM.  
  
$$corim-map-extension: This CDDL socket is used to add new  
  information elements to the corim-map container. See FIXME.
```

3.2.1. The corim-entity-map Container

This Container contains qualifying attributes that provide more context information about the RIM as well its origin and purpose. This rule and its constraints MUST be followed when generating or validating a CoRIM tag

```
corim-entity-map = {  
  corim.entity-name => $entity-name-type-choice  
  ? corim.reg-id => uri  
  corim.role => $corim-role-type-choice  
  * $$corim-entity-map-extension  
}
```

\$corim-role-type-choice /= corim.manifest-creator

corim.entity-name: The name of an organization that performs the roles as indicated by comid.role.

corim.reg-id: The registration identifier of the organization that has authority over the namespace for comid.entity-name.

corim.role: The list of roles the entity is associated with. The entity that generates the CoRIM SHOULD include a \$comid-role-type-choice value of corim.manifest-creator.

\$\$corim-entity-map-extension: This CDDL socket is used to add new information elements to the corim-entity-map container. See FIXME.

3.2.2. The corim-locator-map Container

This map is used to locate and verify the integrity of resources provided by external services, e.g. the CoRIM provider.

```
corim-locator-map = {  
  corim.href => uri  
  ? corim.thumbprint => hash-entry  
}
```

corim.href: A pointer to a services that supplies dependent files or records.

corim.thumbprint: A digest of the reference resource.

3.3. The concise-mid-tag Container

The CDDL specification for the root concise-mid-tag map is as follows. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
concise-mid-tag = {  
  ? comid.language => language-type  
  comid.tag-identity => tag-identity-map  
  ? comid.entities => [ + entity-map ]  
  ? comid.linked-tags => [ + linked-tag-map ]  
  comid.triples => triples-map  
  * $$concise-mid-tag-extension  
}
```

The following describes each member of the concise-mid-tag root map.

comid.language: A textual language tag that conforms with the IANA Language Subtag Registry [IANA.language-subtag-registry].

comid.tag-identity: A composite identifier containing identifying attributes that enable global unique identification of a CoMID tag across versions.

comid.entity: A list of entities that contributed to the CoMID tag.

comid.linked-tags: A list of tags that are linked to this CoMID tag.

comid.triples: A set of relationships in the form of triples, representing a graph-like and semantic reference structure between tags.

\$\$comid-mid-tag-extension: This CDDL socket is used to add new information elements to the concise-mid-tag root container. See FIXME.

3.4. The tag-identity-map Container

The CDDL specification for the tag-identity-map includes all identifying attributes that enable a consumer of information to anticipate required capabilities to process the corresponding tag that map is included in. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
tag-identity-map = {  
    comid.tag-id => $tag-id-type-choice  
    comid.tag-version => tag-version-type  
}
```

```
$tag-id-type-choice /= tstr  
$tag-id-type-choice /= uuid-type
```

```
tag-version-type = uint .default 0
```

The following describes each member of the tag-identity-map container.

comid.tag-id: An identifier for a CoMID that MUST be globally unique.

comid.tag-version: An unsigned integer used as a version identifier.

\$\$tag-identity-map-extension: This CDDL socket is used to add new information elements to the tag-identity-map container. See FIXME.

3.5. The entity-map Container

This Container provides qualifying attributes that provide more context information describing the module as well its origin and purpose. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
entity-map = {  
    comid.entity-name => $entity-name-type-choice  
    ? comid.reg-id => uri  
    comid.role => one-or-more<$comid-role-type-choice>  
    * $$entity-map-extension  
}
```

```
$comid-role-type-choice /= comid.tag-creator  
$comid-role-type-choice /= comid.creator  
$comid-role-type-choice /= comid.maintainer
```

The following describes each member of the tag-identity-map container.

comid.entity-name: The name of an organization that performs the roles as indicated by comid.role.

comid.reg-id: The registration identifier of the organization that has authority over the namespace for comid.entity-name.

comid.role: The list of roles a CoMID entity is associated with. The entity that generates the concise-mid-tag SHOULD include a \$comid-role-type-choice value of comid.tag-creator.

\$\$entity-map-extension: This CDDL socket is used to add new information elements to the entity-map container. See FIXME.

3.6. The linked-tag-map Container

A list of tags that are linked to this CoMID tag.

```
linked-tag-map = {  
  comid.linked-tag-id => $tag-id-type-choice  
  comid.tag-rel => $tag-rel-type-choice  
}
```

```
$tag-rel-type-choice /= comid.supplements  
$tag-rel-type-choice /= comid.replaces
```

The following describes each member of the linked-tag-map container.

comid.linked-tag-id: The tag-id of the linked tag. A linked tag MAY be a CoMID tag or a CoSWID tag.

comid.tag-rel: The relationship type with the linked tag. The relationship type MAY be supplements or replaces, as well as other types well-defined by additional specifications.

3.7. The triples-map Container

A set of directed properties that associate sets of data to provide reference values, endorsed values, verification key material or identifying key material for a specific hardware module that is a component of a composite device. The map provides the core element of CoMID tags that associate remote attestation relevant data with a distinct hardware component that runs an execution environment (a module that is either a Target Environment and/or an Attesting Environment). This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
triples-map = non-empty<{  
  ? comid.reference-triples => one-or-more<reference-triple-record>  
  ? comid.endorsed-triples => one-or-more<endorsed-triple-record>  
  ? comid.attest-key-triples => one-or-more<attest-key-triple-record>  
  ? comid.identity-triples => one-or-more<identity-triple-record>  
  * $$triples-map-extension  

```

The following describes each member of the triple-map container.

`comid.reference-triples`: A directed property that associates reference measurements with a module that is a Target Environment.

`comid.endorsed-triples`: A directed property that associates endorsed measurements with a module that is a Target Environment or Attesting Environment.

`comid.attest-key-triples`: A directed property that associates key material used to verify evidence generated from a module that is an attesting environment.

`comid.identity-triples`: A directed property that associates key material used to identify a module instance or a module class that is an identifying part of a device(-set).

`$$triples-map-extension`: This CDDL socket is used to add new information elements to the triples-map container. See FIXME.

3.8. The environment-map Container

This map represents the module(s) that a triple-map can point directed properties (relationships) from in order to associate them with external information for remote attestation, such as reference values, endorsement and endorsed values, verification key material for evidence, or identifying key material for module (re-)identification. This map can identify a single module instance via `comid.instance` or groups of modules via `comid.group`. Referencing classes of modules requires the use of the more complex class-map container. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
environment-map = non-empty<{
  ? comid.class => class-map
  ? comid.instance => $instance-id-type-choice
  ? comid.group => $group-id-type-choice
}>
```

```
$instance-id-type-choice /= tagged-ueid-type
$instance-id-type-choice /= tagged-uuid-type
```

```
$group-id-type-choice /= tagged-uuid-type
```

The following describes each member of the environment-map container.

`comid-class`: A composite identifier for classes of environments/modules.

comid.instance: An identifier for distinct instances of environments/modules that is either a UEID or a UUID.

comid.group: An identifier for a group of environments/modules that is a UUID.

3.9. The class-map Container

This map enables a composite identifier intended to uniquely identify modules that are of a distinct class of devices. Effectively, all provided members in combination are a composite module class identifier. This rule and its constraints MUST be followed when generating or validating a CoMID tag. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
class-map = non-empty<{  
  ? comid.class-id => $class-id-type-choice  
  ? comid.vendor => tstr  
  ? comid.model => tstr  
  ? comid.layer => uint  
  ? comid.index => uint  
>
```

```
$class-id-type-choice /= tagged-oid-type  
$class-id-type-choice /= tagged-uuid-type  
$class-id-type-choice /= tagged-int-type
```

The following describes each member of the class-map container.

comid.class-id: TODO

comid.vendor TODO

comid.model TODO

comid.layer TODO

comid.index TODO

3.10. The measurement-map and measurement-values-map Containers

One of the targets (range) that a triple-map can point to in order to associate it with a module (domain) is the measurement-map. This map is used to provide reference measurements values that can be compared with Evidence Claim values or Endorsements and endorsed values from other sources than the corresponding CoRIM. measurement-map comes with a measurement key that identifies the measured element with via a OID reference or a UUID. measurement-values-map contains the actual

measurements associated with the module(s). Byte strings with corresponding bit masks that highlights which bits in the byte string are used as reference measurements or endorsement are located in raw-value-group. The members of measurement-values-map provide well-defined and well-scoped semantics for reference measurement or endorsements with respect to a given module instance, class, or group. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
measurement-map = {
  ? comid.mkey => $measured-element-type-choice
  comid.mval => measurement-values-map
}
```

```
$measured-element-type-choice /= tagged-oid-type
$measured-element-type-choice /= tagged-uuid-type
$measured-element-type-choice /= uint
```

```
measurement-values-map = non-empty<{
  ? comid.ver => version-map
  ? comid.svn => svn-type-choice
  ? comid.digests => digests-type
  ? comid.flags => flags-type
  ? raw-value-group
  ? comid.mac-addr => mac-addr-type-choice
  ? comid.ip-addr => ip-addr-type-choice
  ? comid.serial-number => serial-number-type
  ? comid.ueid => ueid-type
  ? comid.uuid => uuid-type
  ? comid.name => tstr
  * $$measurement-values-map-extension
}>
```

```
flags-type = bytes .bits operational-flags
```

```
$operational-flags /= &( not-configured: 0 )
$operational-flags /= &( not-secure: 1 )
$operational-flags /= &( recovery: 2 )
$operational-flags /= &( debug: 3 )
$operational-flags /= &( not-replay-protected: 4 )
$operational-flags /= &( not-integrity-protected: 5 )
```

```
serial-number-type = text
```

```
digests-type = [ + hash-entry ]
```

The following describes each member of the measurement-map and the measurement-values-map container.

comid.mkey: An identifier for the set of measurements expressed in measurement-values-map that is either an OID or a UUID.

comid.ver: A version number measurement.

comid.svn: A security related version number measurement.

comid.digests: A digest (typically a hash value) measurement.

comid.flags: Measurements that reflect operational modes that are made permanent at manufacturing time such that they are not expected to change during normal operation of the Attester.

raw-value-group: A measurement in the form of a byte string that can come with a corresponding bit mask defining the relevance of each bit in the byte string as a measurement.

comid.mac-addr: An EUI-48 or EUI-64 MAC address measurement.

comid.ip-addr: An Ipv4 or Ipv6 address measurement.

comid.serial-number: A measurement of a serial number in text.

comid.ueid: A measurement of a Unique Enough Identifier (UEID).

comid.uuid: A measurement of a Universally Unique Identifier (UUID).

comid.name: TODO

\$\$measurement-values-map-extension: This CDDL socket is used to add new information elements to the measurement-values-map container. See FIXME.

3.10.1. The version-map Container

This map expresses reference values about version information.

```
version-map = {  
  comid.version => version-type  
  ? comid.version-scheme => $version-scheme  
}
```

```
version-type = text .default '0.0.0'
```

The following describes each member of the version-map container.

comid.version: The version in the form of a text string.

comid-version-scheme: The version-scheme of the text string value as defined in [I-D.ietf-sacm-coswid]

3.10.2. The svn-type-choice Enumeration

This choice defines the CBOR tagged Security Version Numbers (SVN) that can be used as reference values for Evidence and Endorsements.

```
svn-type = uint
svn = svn-type
min-svn = svn-type
tagged-svn = #6.552(svn)
tagged-min-svn = #6.553(min-svn)
svn-type-choice = tagged-svn / tagged-min-svn
```

The following describes the types in the svn-type-choice enumeration.

tagged-svn: A specific SVN.

tagged-min-svn: A lower bound for allowed SVN.

3.10.3. The raw-value-group Container

FIXME This group can express a single raw byte value and can come with an optional bit mask that defines which bits in the byte string is used as a reference value, by setting corresponding position in the bit mask to 1.

```
raw-value-group = (
  comid.raw-value => $raw-value-type-choice
  ? comid.raw-value-mask => raw-value-mask-type
)
```

\$raw-value-type-choice /= #6.560(bytes)

raw-value-mask-type = bytes

The following describes the types in the raw-value-group Container.

comid.raw-value: FIXME Bit positions in raw-value-type that correspond to bit positions in raw-value-mask-type.

comid.raw-value-mask: A raw-value-mask-type bit corresponding to a bit in raw-value-type MUST be 1 to evaluate the corresponding raw-value-type bit.

3.10.4. The ip-addr-type-choice Enumeration

This type choice expresses IP addresses as reference values.

```
ip-addr-type-choice = ip4-addr-type / ip6-addr-type
ip4-addr-type = bytes .size 4
ip6-addr-type = bytes .size 16
```

3.10.5. The mac-addr-type-choice Enumeration

This type choice expresses MAC addresses as reference values.

```
mac-addr-type-choice = eui48-addr-type / eui64-addr-type
eui48-addr-type = bytes .size 6
eui64-addr-type = bytes .size 8
```

3.11. The verification-key-map Container

One of the targets (range) that a triple-map can point to in order to associate it with a module (domain). This map is used to provide the key material for evidence verification (effectively signature checking or a lightweight proof-of-possession of private signing key material) or for identity assertion/check (where a proof-of-possession implies a certain device identity). In support of informed trust decisions, an optional trust anchor in the form a PKIX certification path that is associated with the provided key material can be included. This rule and its constraints MUST be followed when generating or validating a CoMID tag.

```
verification-key-map = {
  comid.key => pkix-base64-key-type
  ? comid.keychain => [ + pkix-base64-cert-type ]
}
```

```
pkix-base64-key-type = tstr
pkix-base64-cert-type = tstr
```

The following describes each member of the verification-key-map container.

comid.key: Verification key material in DER format base64 encoded. Typically, but not necessarily, a public key.

comid.keychain: One or more base64 encoded PKIX certificates. The certificate containing the public key in comid.key MUST be the first certificate. Additional certificates MAY follow. Each subsequent certificate SHOULD certify the previous certificate.

4. Full CDDL Definition

This section aggregates the CDDL definitions specified in this document in a full CDDL definitions including:

- * the COSE envelope for CoRIM: signed-corim
- * the CoRIM document: unsigned-corim
- * the CoMID document: concise-mid-tag

Not included in the full CDDL definition are CDDL dependencies to CoSWID. The following CDDL definitions can be found in [I-D.ietf-sacm-coswid]:

- * the COSE envelope for CoRIM: signed-coswid
- * the CoSWID document: concise-swid-tag

<CODE BEGINS>

```
corim = #6.500($concise-reference-integrity-manifest-type-choice)
```

```
$concise-reference-integrity-manifest-type-choice /= #6.501(unsigned-corim-map
)
```

```
$concise-reference-integrity-manifest-type-choice /= #6.502(signed-corim)
```

```
signed-corim = #6.18(COSE-Sign1-corim)
```

```
protected-signed-corim-header-map = {
  corim.alg-id => int
  corim.content-type => "application/rim+cbor"
  corim.issuer-key-id => bstr
  corim.meta => corim-meta-map
  * cose-label => cose-values
}
```

```
corim-meta-map = {
  corim.signer => [ + corim-entity-map ]
  ? corim.validity => validity-map
}
```

```
corim-entity-map = {
  corim.entity-name => $entity-name-type-choice
  ? corim.reg-id => uri
  corim.role => $corim-role-type-choice
  * $$corim-entity-map-extension
}
```

```
$corim-role-type-choice /= corim.manifest-creator
$corim-role-type-choice /= corim.manifest-signer

validity-map = {
  ? corim.not-before => time
  corim.not-after => time
}

unprotected-signed-corim-header-map = {
  * cose-label => cose-values
}

COSE-Sign1-corim = [
  protected: bstr .cbor protected-signed-corim-header-map
  unprotected: unprotected-signed-corim-header-map
  payload: bstr .cbor unsigned-corim-map
  signature: bstr
]

unsigned-corim-map = {
  corim.id => $corim-id-type-choice
  corim.tags => [ + $concise-tag-type-choice ]
  ? corim.dependent-rims => [ + corim-locator-map ]
  ? corim.profile => [ + profile-type-choice ]
  * $$unsigned-corim-map-extension
}

profile-type-choice = uri / tagged-oid-type

corim-locator-map = {
  corim.href => uri
  ? corim.thumbprint => hash-entry
}

$concise-tag-type-choice /= #6.505(bytes .cbor concise-swid-tag)
$concise-tag-type-choice /= #6.506(bytes .cbor concise-mid-tag)

concise-mid-tag = {
  ? comid.language => language-type
  comid.tag-identity => tag-identity-map
  ? comid.entity => [ + entity-map ]
  ? comid.linked-tags => [ + linked-tag-map ]
  comid.triples => triples-map
  * $$concise-mid-tag-extension
}
```

```
language-type = text

tag-identity-map = {
  comid.tag-id => $tag-id-type-choice
  ? comid.tag-version => tag-version-type
}

$tag-id-type-choice /= tstr
$tag-id-type-choice /= uuid-type

tag-version-type = uint .default 0

entity-map = {
  comid.entity-name => $entity-name-type-choice
  ? comid.reg-id => uri
  comid.role => [ + $comid-role-type-choice ]
  * $$entity-map-extension
}

$comid-role-type-choice /= comid.tag-creator
$comid-role-type-choice /= comid.creator
$comid-role-type-choice /= comid.maintainer

linked-tag-map = {
  comid.linked-tag-id => $tag-id-type-choice
  comid.tag-rel => $tag-rel-type-choice
}

$tag-rel-type-choice /= comid.supplements
$tag-rel-type-choice /= comid.replaces

triples-map = non-empty<{
  ? comid.reference-triples => [ + reference-triple-record ]
  ? comid.endorsed-triples => [ + endorsed-triple-record ]
  ? comid.attest-key-triples => [ + attest-key-triple-record ]
  ? comid.identity-triples => [ + identity-triple-record ]
  * $$triples-map-extension
}>

reference-triple-record = [
  environment-map ; target environment
  [ + measurement-map ] ; reference measurements
]

endorsed-triple-record = [
  environment-map ; (target or attesting) environment
  [ + measurement-map ] ; endorsed measurements
]
```

```
attest-key-triple-record = [  
  environment-map ; attesting environment  
  [ + verification-key-map ] ; attestation verification key(s)  
]  
  
identity-triple-record = [  
  environment-map ; device identifier (instance or class)  
  [ + verification-key-map ] ; DevID, or semantically equivalent  
]  
  
pkix-base64-key-type = tstr  
pkix-base64-cert-type = tstr  
  
verification-key-map = {  
  ; Verification key in DER format base64-encoded.  
  ; Typically, but not necessarily a public key.  
  comid.key => pkix-base64-key-type  
  ; Optional X.509 certificate chain corresponding to the public key  
  ; in comid.key, encoded as an array of one or more base64-encoded  
  ; DER PKIX certificates. The certificate containing the public key  
  ; in comid.key MUST be the first certificate. This MAY be followed  
  ; by additional certificates, with each subsequent certificate  
  ; being the one used to certify the previous one.  
  ? comid.keychain => [ + pkix-base64-cert-type ]  
}  
  
environment-map = non-empty<{  
  ? comid.class => class-map  
  ? comid.instance => $instance-id-type-choice  
  ? comid.group => $group-id-type-choice  
  
class-map = non-empty<{  
  ? comid.class-id => $class-id-type-choice  
  ? comid.vendor => tstr  
  ? comid.model => tstr  
  ? comid.layer => uint  
  ? comid.index => uint  
  
$class-id-type-choice /= tagged-oid-type  
$class-id-type-choice /= tagged-uuid-type  
  
$instance-id-type-choice /= tagged-uuid-type  
$instance-id-type-choice /= tagged-uuid-type  
  
$group-id-type-choice /= tagged-uuid-type
```

```
oid-type = bytes
tagged-oid-type = #6.111(oid-type)

tagged-uuid-type = #6.37(uuid-type)

uuid-type = bytes .size 33
tagged-uuid-type = #6.550(uuid-type)

$measured-element-type-choice /= tagged-oid-type
$measured-element-type-choice /= tagged-uuid-type

measurement-map = {
  ? comid.mkey => $measured-element-type-choice
  comid.mval => measurement-values-map
}

measurement-values-map = non-empty<{
  ? comid.ver => version-map
  ? comid.svn => svn-type-choice
  ? comid.digests => digests-type
  ? comid.flags => flags-type
  ? raw-value-group
  ? comid.mac-addr => mac-addr-type-choice
  ? comid.ip-addr => ip-addr-type-choice
  ? comid.serial-number => serial-number-type
  ? comid.ueid => ueid-type
  ? comid.uuid => uuid-type
  * $$measurement-values-map-extension
}>

version-map = {
  comid.version => version-type
  ? comid.version-scheme => $version-scheme
}
version-type = text .default '0.0.0'

svn = int
min-svn = int
tagged-svn = #6.552(svn)
tagged-min-svn = #6.553(min-svn)
svn-type-choice = tagged-svn / tagged-min-svn

flags-type = bytes .bits operational-flags

operational-flags = &(
  not-configured: 0
  not-secure: 1
  recovery: 2
```

```
    debug: 3
  )

raw-value-group = (
  comid.raw-value => raw-value-type
  ? comid.raw-value-mask => raw-value-mask-type
)

raw-value-type = bytes
raw-value-mask-type = bytes

ip-addr-type-choice = ip4-addr-type / ip6-addr-type
ip4-addr-type = bytes .size 4
ip6-addr-type = bytes .size 16

mac-addr-type-choice = eui48-addr-type / eui64-addr-type
eui48-addr-type = bytes .size 6
eui64-addr-type = bytes .size 8

serial-number-type = text

digests-type = [ + hash-entry ]

concise-swid-tag = {
  tag-id => text / bstr .size 16,
  tag-version => integer,
  ? corpus => bool,
  ? patch => bool,
  ? supplemental => bool,
  software-name => text,
  ? software-version => text,
  ? version-scheme => $version-scheme,
  ? media => text,
  ? software-meta => one-or-more<software-meta-entry>,
  entity => one-or-more<entity-entry>,
  ? link => one-or-more<link-entry>,
  ? payload-or-evidence,
  * $$coswid-extension,
  global-attributes,
}

payload-or-evidence //= ( payload => payload-entry )
payload-or-evidence //= ( evidence => evidence-entry )

any-uri = uri
label = text / int
```

```
$version-scheme /= multipartnumeric
$version-scheme /= multipartnumeric-suffix
$version-scheme /= alphanumeric
$version-scheme /= decimal
$version-scheme /= semver
$version-scheme /= int / text

any-attribute = (
  label => one-or-more<text> / one-or-more<int>
)

one-or-more<T> = T / [ 2* T ]

global-attributes = (
  ? lang => text,
  * any-attribute,
)

hash-entry = [
  hash-alg-id: int,
  hash-value: bytes,
]

entity-entry = {
  entity-name => text,
  ? reg-id => any-uri,
  role => one-or-more<$role>,
  ? thumbprint => hash-entry,
  * $$entity-extension,
  global-attributes,
}

$role /= tag-creator
$role /= software-creator
$role /= aggregator
$role /= distributor
$role /= licensor
$role /= maintainer
$role /= int / text

link-entry = {
  ? artifact => text,
  href => any-uri,
  ? media => text,
  ? ownership => $ownership,
  rel => $rel,
  ? media-type => text,
  ? use => $use,
```



```
    * $$link-extension,
    global-attributes,
  }

  $ownership /= shared
  $ownership /= private
  $ownership /= abandon
  $ownership /= int / text

  $rel /= ancestor
  $rel /= component
  $rel /= feature
  $rel /= installationmedia
  $rel /= packageinstaller
  $rel /= parent
  $rel /= patches
  $rel /= requires
  $rel /= see-also
  $rel /= supersedes
  $rel /= supplemental
  $rel /= -256..64436 / text

  $use /= optional
  $use /= required
  $use /= recommended
  $use /= int / text

  software-meta-entry = {
    ? activation-status => text,
    ? channel-type => text,
    ? colloquial-version => text,
    ? description => text,
    ? edition => text,
    ? entitlement-data-required => bool,
    ? entitlement-key => text,
    ? generator => text,
    ? persistent-id => text,
    ? product => text,
    ? product-family => text,
    ? revision => text,
    ? summary => text,
    ? unspsc-code => text,
    ? unspsc-version => text,
    * $$software-meta-extension,
    global-attributes,
  }

  path-elements-group = ( ? directory => one-or-more<directory-entry>,
```

```
        ? file => one-or-more<file-entry>,
    )

resource-collection = (
    path-elements-group,
    ? process => one-or-more<process-entry>,
    ? resource => one-or-more<resource-entry>,
    * $$resource-collection-extension,
)

file-entry = {
    filesystem-item,
    ? size => uint,
    ? file-version => text,
    ? hash => hash-entry,
    * $$file-extension,
    global-attributes,
}

directory-entry = {
    filesystem-item,
    ? path-elements => { path-elements-group },
    * $$directory-extension,
    global-attributes,
}

process-entry = {
    process-name => text,
    ? pid => integer,
    * $$process-extension,
    global-attributes,
}

resource-entry = {
    type => text,
    * $$resource-extension,
    global-attributes,
}

filesystem-item = (
    ? key => bool,
    ? location => text,
    fs-name => text,
    ? root => text,
)

payload-entry = {
    resource-collection,
```

```
    * $$payload-extension,  
    global-attributes,  
  }  
  
  evidence-entry = {  
    resource-collection,  
    ? date => integer-time,  
    ? device-id => text,  
    * $$evidence-extension,  
    global-attributes,  
  }  
  
  integer-time = #6.1(int)  
  
  tag-id = 0  
  software-name = 1  
  entity = 2  
  evidence = 3  
  link = 4  
  software-meta = 5  
  payload = 6  
  hash = 7  
  corpus = 8  
  patch = 9  
  media = 10  
  supplemental = 11  
  tag-version = 12  
  software-version = 13  
  version-scheme = 14  
  lang = 15  
  directory = 16  
  file = 17  
  process = 18  
  resource = 19  
  size = 20  
  file-version = 21  
  key = 22  
  location = 23  
  fs-name = 24  
  root = 25  
  path-elements = 26  
  process-name = 27  
  pid = 28  
  type = 29  
  entity-name = 31  
  reg-id = 32  
  role = 33  
  thumbprint = 34
```

date = 35
device-id = 36
artifact = 37
href = 38
ownership = 39
rel = 40
media-type = 41
use = 42
activation-status = 43
channel-type = 44
colloquial-version = 45
description = 46
edition = 47
entitlement-data-required = 48
entitlement-key = 49
generator = 50
persistent-id = 51
product = 52
product-family = 53
revision = 54
summary = 55
unspsc-code = 56
unspsc-version = 57

multipartnumeric = 1
multipartnumeric-suffix = 2
alphanumeric = 3
decimal = 4
semver = 16384

tag-creator=1
software-creator=2
aggregator=3
distributor=4
licensor=5
maintainer=6

shared=1
private=2
abandon=3

ancestor=1
component=2
feature=3
installationmedia=4
packageinstaller=5
parent=6
patches=7

```
requires=8
see-also=9
supersedes=10

optional=1
required=2
recommended=3

comid.language = 0
comid.tag-identity = 1
comid.entity = 2
comid.linked-tags = 3
comid.triples = 4

comid.tag-id = 0
comid.tag-version = 1

comid.entity-name = 0
comid.reg-id = 1
comid.role = 2

comid.linked-tag-id = 0
comid.tag-rel = 1

comid.reference-triples = 0
comid.endorsed-triples = 1
comid.identity-triples = 2
comid.attest-key-triples = 3

comid.class = 0
comid.instance = 1
comid.group = 2

comid.class-id = 0
comid.vendor = 1
comid.model = 2
comid.layer = 3
comid.index = 4

comid.mkey = 0
comid.mval = 1

comid.ver = 0
comid.svn = 1
comid.digests = 2
comid.flags = 3
comid.raw-value = 4
comid.raw-value-mask = 5
```

```
comid.mac-addr = 6
comid.ip-addr = 7
comid.serial-number = 8
comid.ueid = 9
comid.uuid = 10

comid.key = 0
comid.keychain = 1

comid.version = 0
comid.version-scheme = 1

comid.supplements = 0

comid.replaces = 1

comid.tag-creator = 0
comid.creator = 1
comid.maintainer = 2

corim.id = 0
corim.tags = 1
corim.dependent-rims = 2
corim.profile = 3

corim.href = 0
corim.thumbprint = 1

corim.alg-id = 1
corim.content-type = 3
corim.issuer-key-id = 4
corim.meta = 8

corim.not-before = 0
corim.not-after = 1

corim.signer = 0
corim.validity = 1

corim.entity-name = 0
corim.reg-id = 1
corim.role = 2

corim.manifest-creator = 1

corim.manifest-signer = 2
```

```
non-empty<M> = (M) .within ({ + any => any })
```

```
cose-label = int / tstr  
cose-values = any
```

```
$entity-name-type-choice /= text
```

```
$corim-id-type-choice /= tstr  
$corim-id-type-choice /= uuid-type
```

```
uuid-type = bytes .size 16
```

```
<CODE ENDS>
```

5. Privacy Considerations

Privacy Considerations

6. Security Considerations

Security Considerations

7. IANA Considerations

This document has a number of IANA considerations, as described in the following subsections. In summary, 6 new registries are established with this request, with initial entries provided for each registry. New values for 5 other registries are also requested.

7.1. COSE Header Parameters Registry

The 'corim metadata' parameter has been added to the "COSE Header Parameters" registry:

- * Name: 'corim metadata'
- * Label: 11
- * Value: corim-meta-map
- * Description: Provides a map of additional metadata for a CoRIM payload composed of (1) one or more entities that created or signed the corresponding CoRIM and (2) its period of validity

- * Reference: 'corim-meta-map' in {model-corim-meta-map} of this document

7.2. CoRIM Map Items Registry

This document defines a new registry titled "CoRIM Map". The registry uses integer values as index values for items in 'unsigned-corim-map' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 3: CoRIM Map Items
Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoRIM Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	corim.id	RFC-AAAA
1	corim.tags	RFC-AAAA
2	corim.dependent-rims	RFC-AAAA
3-255	Unassigned	

Table 4: CoRIM Map Items Initial Registrations

7.3. CoRIM Entity-Map Items Registry

This document defines a new registry titled "CoRIM Entity Map". The registry uses integer values as index values for items in 'corim-enentity-map' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 5: CoRIM Entity Map Items
Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoRIM Entity Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	corim.entity-name	RFC-AAAA
1	corim.reg-id	RFC-AAAA
2	corim.role	RFC-AAAA
3-255	Unassigned	

Table 6: CoRIM Entity Map Items Initial
Registrations

7.4. CoRIM Entity-Types Registry

This document defines a new registry titled "CoRIM Entity Types". The registry maintains well-defined integer values as choices for '\$entity-name-type-choice' CBOR uints.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 7: CoRIM Entity Types
Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoRIM Entity Types" registry are provided below. Assignments consist of an integer value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	corim.manifest-creator	RFC-AAAA
1	corim.manifest-signer	RFC-AAAA
2-255	Unassigned	

Table 8: CoRIM Entity Types Initial Registrations

7.5. CoMID Map Items Registry

This document defines a new registry titled "CoMID Map". The registry uses integer values as index values for items in 'concise-mid-tag' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 9: CoMID Map Items
Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.language	RFC-AAAA
1	comid.tag-identity	RFC-AAAA
2	comid.entity	RFC-AAAA
3	comid.linked-tags	RFC-AAAA
4	comid.triples	RFC-AAAA
5-255	Unassigned	

Table 10: CoMID Map Items Initial Registrations

7.6. CoMID Entity-Map Items Registry

This document defines a new registry titled "CoMID Entity Map". The registry uses integer values as index values for items in 'comid-entity-map' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 11: CoMID Entity Map Items Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Entity Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.entity-name	RFC-AAAA
1	comid.reg-id	RFC-AAAA
2	comid.role	RFC-AAAA
3-255	Unassigned	

Table 12: CoMID Entity Map Items Initial Registrations

7.7. CoMID Triples-Map Items Registry

This document defines a new registry titled "CoMID Triples Map". The registry uses integer values as index values for items in 'comid-triples-map' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 13: CoMID triples Map Items Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Triples Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.reference-triples	RFC-AAAA
1	comid.endorsed-triples	RFC-AAAA
2	comid.identity-triples	RFC-AAAA
3	comid.attest-key-triples	RFC-AAAA
4-255	Unassigned	

Table 14: CoMID Triples Map Items Initial Registrations

7.8. CoMID Measurement-Values-Map Items Registry

This document defines a new registry titled "CoMID Measurement-Values Map". The registry uses integer values as index values for items in 'comid-measurement-values-map' CBOR maps.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 15: CoMID Measurement-Values Map Items Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Measurement-Values Map" registry are provided below. Assignments consist of an integer index value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.ver	RFC-AAAA
1	comid.svn	RFC-AAAA
2	comid.digests	RFC-AAAA
3	comid.flags	RFC-AAAA
4	comid.raw-value	RFC-AAAA
5	comid.raw-value-mask	RFC-AAAA
6	comid.mac-addr	RFC-AAAA
7	comid.ip-addr	RFC-AAAA
8	comid.serial-number	RFC-AAAA
9	comid.ueid	RFC-AAAA
10	comid.uuid	RFC-AAAA
11-255	Unassigned	

Table 16: CoMID Measurement-Values Map Items
Initial Registrations

7.9. CoMID Tag-Relationship-Types Registry

This document defines a new registry titled "CoMID Tag-Relationship Types". The registry maintains well-defined integer values as choices for '\$tag-rel-type-choice' CBOR uints.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 17: CoMID Tag-Relationship
Types Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Tag-Relationship Types" registry are provided below. Assignments consist of an integer value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.supplements	RFC-AAAA
1	comid.replaces	RFC-AAAA
2-255	Unassigned	

Table 18: CoMID Tag-Relationship Types
Initial Registrations

7.10. CoMID Role-Types Registry

This document defines a new registry titled "CoMID Role Types". The registry maintains well-defined integer values as choices for '\$comid-role-type-choice' CBOR uints.

Future registrations for this registry are to be made based on [RFC8126] as follows:

Range	Registration Procedures
0-127	Standards Action
128-255	Specification Required

Table 19: CoMID Role Types
Registration Procedures

All negative values are reserved for Private Use.

Initial registrations for the "CoMID Role Types" registry are provided below. Assignments consist of an integer value, the item name, and a reference to the defining specification.

Index	Item Name	Specification
0	comid.tag-creator	RFC-AAAA
1	comid.creator	RFC-AAAA
2	comid.maintainer	RFC-AAAA
3-255	Unassigned	

Table 20: CoMID Role Types Initial
Registrations

7.11. rim+cbor Media Type Registration

IANA is requested to add the following to the IANA "Media Types" registry [IANA.media-types].

Type name: application

Subtype name: rim+cbor

Required parameters: none

Optional parameters: none

Encoding considerations: Must be encoded as using [RFC8949]. See RFC-AAAA for details.

Security considerations: See Section 6 of RFC-AAAA.

Interoperability considerations: Applications MAY ignore any key value pairs that they do not understand. This allows backwards compatible extensions to this specification.

Published specification: RFC-AAAA

Applications that use this media type: The type is used by remote attestation procedures, supply chain integrity management systems, vulnerability assessment systems, and in applications that rely on trustworthy endorsements and reference values describing the intended operational state of a system.

Fragment identifier considerations: Fragment identification for application/rim+cbor is supported by using fragment identifiers as specified by Section 9.5 of [RFC8949].

Additional information:

Magic number(s): first five bytes in hex: 43 4f 52 49 4d

File extension(s): corim

Macintosh file type code(s): none

Macintosh Universal Type Identifier code: org.ietf.corim conforms to public.data

Person & email address to contact for further information: Henk Birkholz <henk.birkholz@sit.fraunhofer.de>

Intended usage: COMMON

Restrictions on usage: None

Author: Henk Birkholz <henk.birkholz@sit.fraunhofer.de>

Change controller: IESG

7.12. CoAP Content-Format Registration

IANA is requested to assign a CoAP Content-Format ID for the CoRIM media type in the "CoAP Content-Formats" sub-registry, from the "IETF Review or IESG Approval" space (256..999), within the "CoRE Parameters" registry [RFC7252] [IANA.core-parameters]:

Media type	Encoding	ID	Reference
application/rim+cbor	-	TBD1	RFC-AAAA

Table 21: CoAP Content-Format IDs

7.13. CoRIM CBOR Tag Registration

IANA is requested to allocate tags in the "CBOR Tags" registry [IANA.cbor-tags], preferably with the specific value requested:

Tag	Data Item	Semantics
500	tagged array or tagged map	Concise Reference Integrity Manifest (CoRIM) [RFC-AAAA]
501	map	unsigned CoRIM [RFC-AAAA]
502	array	signed CoRIM [RFC-AAAA]
505	bstr	byte string with CBOR-encoded Concise SWID tag [RFC-AAAA]
506	bstr	byte string with CBOR-encoded Concise MID tag [RFC-AAAA]

Table 22: CoRIM CBOR Tags

7.14. CoMID CBOR Tag Registration

IANA is requested to allocate tags in the "CBOR Tags" registry [IANA.cbor-tags], preferably with the specific value requested:

Tag	Data Item	Semantics
550	bstr	UEID with max size of 33 bytes [RFC-AAAA]
551	int	Security Version Number [RFC-AAAA]
552	int	lower bound of allowed Security Version Number [RFC-AAAA]

Table 23: CoMID CBOR Tags

8. References

8.1. Normative References

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-14, 9 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-14.txt>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-20, 26 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-sacm-coswid-20.txt>>.

[IANA.cbor-tags]

IANA, "Concise Binary Object Representation (CBOR) Tags", <<http://www.iana.org/assignments/cbor-tags>>.

[IANA.core-parameters]

IANA, "Constrained RESTful Environments (CoRE) Parameters", <<http://www.iana.org/assignments/core-parameters>>.

[IANA.language-subtag-registry]

IANA, "Language Subtag Registry", <<http://www.iana.org/assignments/language-subtag-registry>>.

[IANA.media-types]

IANA, "Media Types", <<http://www.iana.org/assignments/media-types>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/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://www.rfc-editor.org/info/rfc8610>>.
- [RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/RFC8949, December 2020, <<https://www.rfc-editor.org/info/rfc8949>>.

8.2. Informative References

- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Thomas Fossati
Arm Limited
United Kingdom

Email: Thomas.Fossati@arm.com

Yogesh Deshpande
Arm Limited
United Kingdom

Email: yogesh.deshpande@arm.com

Ned Smith
Intel Corporation
United States of America

Email: ned.smith@intel.com

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 13 January 2022

H. Birkholz
Fraunhofer SIT
C. Newton
L. Chen
University of Surrey
12 July 2021

Direct Anonymous Attestation for the Remote Attestation Procedures
Architecture
draft-birkholz-rats-daa-01

Abstract

This document maps the concept of Direct Anonymous Attestation (DAA) to the Remote Attestation Procedures (RATS) Architecture. The role DAA Issuer is introduced and its interactions with existing RATS roles is specified.

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 13 January 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Direct Anonymous Attestation	3
4. DAA changes to the RATS Architecture	5
5. Additions to Remote Attestation principles	6
6. Privacy Considerations	7
7. Security Considerations	7
8. Implementation Considerations	7
9. IANA Considerations	8
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Authors' Addresses	8

1. Introduction

Remote Attestation procedures (RATS, [I-D.ietf-rats-architecture]) describe interactions between well-defined architectural constituents in support of Relying Parties that require an understanding about the trustworthiness of a remote peer. The identity of an Attester and its corresponding Attesting Environments play a vital role in RATS. A common way to refer to such an identity is the Authentication Secret ID as defined in the Reference Interaction Models for RATS [I-D.ietf-rats-reference-interaction-models]. The fact that every Attesting Environment can be uniquely identified in the context of the RATS architecture is not suitable for every application of remote attestation. Additional issues may arise when Personally identifiable information (PII) -- whether obfuscated or in clear text -- are included in attestation Evidence or even corresponding Attestation Results. This document illustrates how Direct Anonymous Attestation (DAA) can mitigate the issue of uniquely (re-)identifiable Attesting Environments. To accomplish that goal, a new RATS role -- the DAA Issuer -- is introduced and its duties as well as its interactions with other RATS roles are specified.

2. Terminology

This document uses the following set of terms, roles, and concepts as defined in [I-D.ietf-rats-architecture]: Attester, Verifier, Relying Party, Conceptual Message, Evidence, Attestation Result, Attesting Environment. The role of Endorser, also defined in [I-D.ietf-rats-architecture], needs to be adapted and details are given below.

Additionally, this document uses and adapts, as necessary, the following concepts and information elements as defined in [I-D.ietf-rats-reference-interaction-models]: Attester Identity, Authentication Secret, Authentication Secret ID

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.

3. Direct Anonymous Attestation

Figure 1 shows the data flows between the different RATS roles involved in DAA.

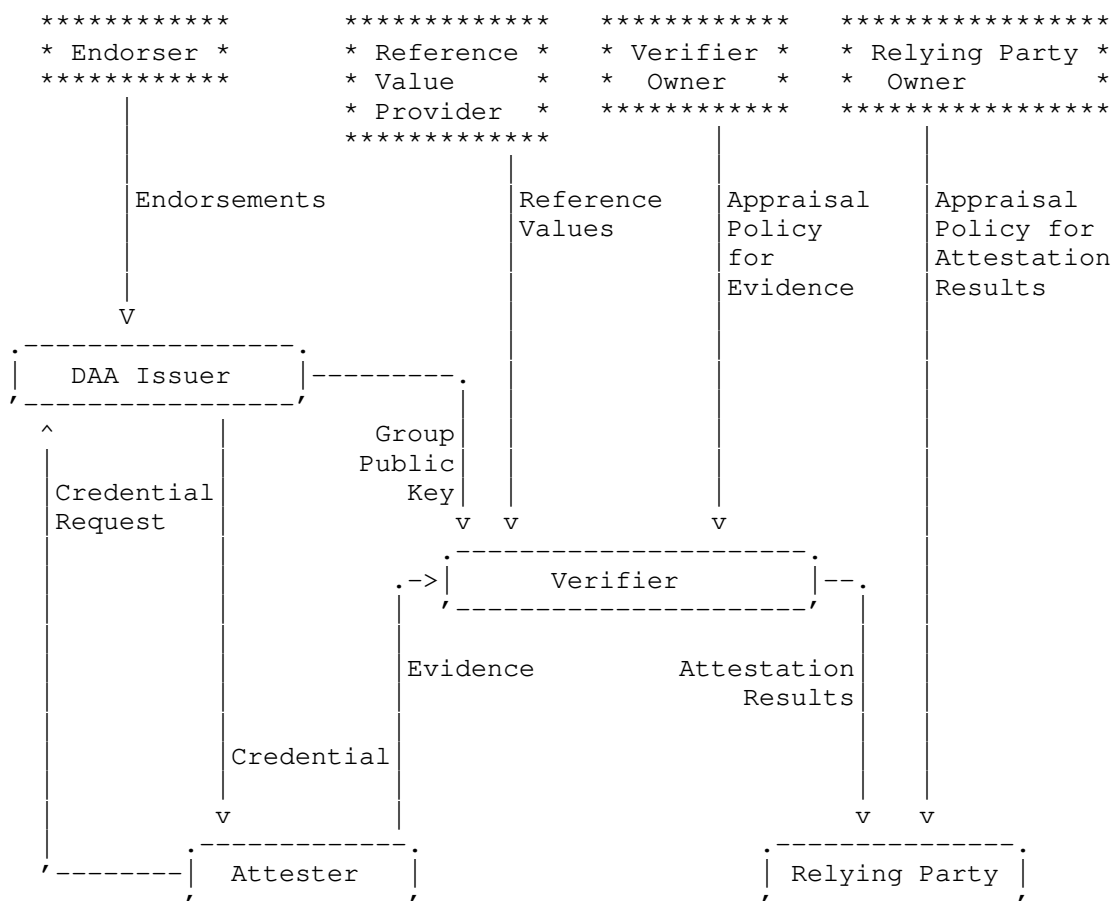


Figure 1: DAA data flows

DAA [DAA] is a signature scheme that allows the privacy of users that are associated with an Attester (e.g. its owner) to be maintained. Essentially, DAA can be seen as a group signature scheme with the feature that given a DAA signature no-one can find out who the signer is, i.e., the anonymity is not revocable. To be able to sign anonymously, an Attester has to obtain a credential from a DAA Issuer. The DAA Issuer uses a private/public key pair to generate credentials for a group of Attesters and makes the public key (in the form of a public key certificate) available to the verifier in order to enable them to validate the Evidence received.

In order to support these DAA signatures, the DAA Issuer MUST associate a single key pair with a group of Attesters and use the same key pair when creating the credentials for all of the Attesters

in this group. The DAA Issuer's group public key certificate replaces the individual Attester Identity documents during authenticity validation as a part of the appraisal of Evidence conducted by a verifier. This is in contrast to intuition that there has to be a unique Attester Identity per device.

For DAA, the role of the Endorser is essentially the same, but they now provide Attester endorsement documents to the DAA Issuer rather than directly to the verifier. These Attester endorsement documents enable the Attester to obtain a credential from the DAA Issuer.

4. DAA changes to the RATS Architecture

In order to enable the use of DAA, a new conceptual message, the Credential Request, is defined and a new role, the DAA Issuer role, is added to the roles defined in the RATS Architecture.

Credential Request: An Attester sends a Credential Request to the DAA Issuer to obtain a credential. This request contains information about the DAA key that the Attester will use to create evidence and together with Attester endorsement information that is provided by the Endorser to confirm that the request came from a valid Attester.

DAA Issuer: A RATS role that offers zero-knowledge proofs based on public-key certificates used for a group of Attesters (Group Public Keys) [DAA]. How this group of Attesters is defined is not specified here, but the group must be large enough for the necessary anonymity to be assured.

Effectively, these certificates share the semantics of Endorsements, with the following exceptions:

- * Upon receiving a Credential Request from an Attester, the associated group private key is used by the DAA Issuer to provide the Attester with a credential that it can use to convince the Verifier that its Evidence is valid. To keep their anonymity the Attester randomizes this credential each time that it is used. Although the DAA Issuer knows the Attester Identity and can associate this with the credential issued, randomisation ensures that the Attester's identity cannot be revealed to anyone, including the Issuer.
- * The Verifier can use the DAA Issuer's group public key certificate, together with the randomized credential from the Attester, to confirm that the Evidence comes from a valid Attester without revealing the Attester's identity.

- * A credential is conveyed from a DAA Issuer to an Attester in combination with the conveyance of the group public key certificate from DAA Issuer to Verifier.

5. Additions to Remote Attestation principles

In order to ensure an appropriate conveyance of Evidence via interaction models in general, the following prerequisite considering Attester Identity MUST be in place to support the implementation of interaction models.

Attestation Evidence Authenticity: Attestation Evidence MUST be correct and authentic.

In order to provide proofs of authenticity, Attestation Evidence SHOULD be cryptographically associated with an identity document that is a randomised DAA credential.

The following information elements define extensions for corresponding information elements defined in [I-D.ietf-rats-reference-interaction-models] and that are vital to all types of reference interaction models. Varying from solution to solution, generic information elements can be either included in the scope of protocol messages (instantiating Conceptual Messages defined by the RATS architecture) or can be included in additional protocol parameters of protocols that facilitate the conveyance of RATS Conceptual Messages. Ultimately, the following information elements are required by any kind of scalable remote attestation procedure using DAA with one of RATS's reference interaction models.

Attester Identity ('attesterIdentity'): _mandatory_

In DAA, the Attester's identity is not revealed to the verifier. The Attester is issued with a credential by the DAA Issuer that is randomised and then used to anonymously confirm the validity of their evidence. The evidence is verified using the DAA Issuer's group public key.

Authentication Secret IDs ('authSecID'): _mandatory_

In DAA, Authentication Secret IDs are represented by the DAA Issuer's group public key that MUST be used to create DAA credentials for the corresponding Authentication Secrets used to protect Evidence.

In DAA, an Authentication Secret ID does not identify a unique

Attesting Environment but is associated with a group of Attesting Environments. This is because an Attesting Environment should not be distinguishable and the DAA credential which represents the Attesting Environment is randomised each time it used.

6. Privacy Considerations

As outlined about for DAA to provide privacy for the Attester the DAA group must be large enough to stop the Verifier identifying the Attester.

Randomization of the DAA credential by the Attester means that collusion between the DAA Issuer and Verifier, will not give them any advantage when trying to identify the Attester.

For DAA, the Attestation Evidence conveyed to the Verifier MUST not uniquely identify the Attester. If the Attestation Evidence is unique to an Attester other cryptographic techniques can be used, for example, property based attestation. (Henk -- reference follows)

Chen L., Loehr H., Manulis M., Sadeghi AR. (2008) Property-Based Attestation without a Trusted Third Party. Information Security. ISC 2008. Lecture Notes in Computer Science, vol 5222. Springer. https://doi.org/10.1007/978-3-540-85886-7_3

7. Security Considerations

The anonymity property of DAA makes revocation difficult. Well known solutions include: 1. Rogue attester revocation -- if the an Attester's private key is compromised and known by the Verifier then any DAA signature from that Attester can be revoked. 2. EPID - Intel's Enhanced Privacy ID -- this requires the Attester to prove (as part of their Attestation) that their credential was not used to generate any signature in a signature revocation list.

There are no other special security conderations for DAA over and above those specifed in the RATS architecture document [I-D.ietf-rats-architecture].

8. Implementation Considerations

The new DAA Issuer role can be implemented in a number of ways, for example: 1. As a stand-alone service like a Certificate Authority, a Privacy CA. 2. As a part of the Attester's manufacture. The Endorser and the DAA Issuer could be the same entity and the manufacturer would then provide a certificate for the group public key to the Verifier.

9. IANA Considerations

We don't yet.

10. References

10.1. Normative References

- [DAA] Brickell, E., Camenisch, J., and L. Chen, "Direct Anonymous Attestation", page 132-145, ACM Proceedings of the 11rd ACM conference on Computer and Communications Security, 2004.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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://www.rfc-editor.org/info/rfc5280>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

10.2. Informative References

- [I-D.ietf-rats-architecture] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-12.txt>>.
- [I-D.ietf-rats-reference-interaction-models] Birkholz, H., Eckel, M., Pan, W., and E. Voit, "Reference Interaction Models for Remote Attestation Procedures", Work in Progress, Internet-Draft, draft-ietf-rats-reference-interaction-models-02, 25 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-reference-interaction-models-02.txt>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Christopher Newton
University of Surrey

Email: cn0016@surrey.ac.uk

Liqun Chen
University of Surrey

Email: liqun.chen@surrey.ac.uk

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 28 April 2022

H. Birkholz
Fraunhofer SIT
C. Newton
L. Chen
University of Surrey
D. Thaler
Microsoft
25 October 2021

Direct Anonymous Attestation for the Remote Attestation Procedures
Architecture
draft-birkholz-rats-daa-02

Abstract

This document maps the concept of Direct Anonymous Attestation (DAA) to the Remote Attestation Procedures (RATS) Architecture. The role DAA Issuer is introduced and its interactions with existing RATS roles is specified.

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 28 April 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Direct Anonymous Attestation	3
4. DAA changes to the RATS Architecture	5
5. Additions to Remote Attestation principles	6
6. Privacy Considerations	7
7. Security Considerations	7
8. Implementation Considerations	7
9. IANA Considerations	8
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Authors' Addresses	8

1. Introduction

Remote Attestation procedures (RATS, [I-D.ietf-rats-architecture]) describe interactions between well-defined architectural constituents in support of Relying Parties that require an understanding about the trustworthiness of a remote peer. The identity of an Attester and its corresponding Attesting Environments play a vital role in RATS. A common way to refer to such an identity is the Authentication Secret ID as defined in the Reference Interaction Models for RATS [I-D.ietf-rats-reference-interaction-models]. The fact that every Attesting Environment can be uniquely identified in the context of the RATS architecture is not suitable for every application of remote attestation. Additional issues may arise when Personally identifiable information (PII) -- whether obfuscated or in clear text -- are included in attestation Evidence or even corresponding Attestation Results. This document illustrates how Direct Anonymous Attestation (DAA) can mitigate the issue of uniquely (re-)identifiable Attesting Environments. To accomplish that goal, a new RATS role -- the DAA Issuer -- is introduced and its duties as well as its interactions with other RATS roles are specified.

2. Terminology

This document uses the following set of terms, roles, and concepts as defined in [I-D.ietf-rats-architecture]: Attester, Verifier, Relying Party, Conceptual Message, Evidence, Attestation Result, Attesting Environment. The role of Endorser, also defined in [I-D.ietf-rats-architecture], needs to be adapted and details are given below.

Additionally, this document uses and adapts, as necessary, the following concepts and information elements as defined in [I-D.ietf-rats-reference-interaction-models]: Attester Identity, Authentication Secret, Authentication Secret ID

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.

3. Direct Anonymous Attestation

Figure 1 shows the data flows between the different RATS roles involved in DAA.

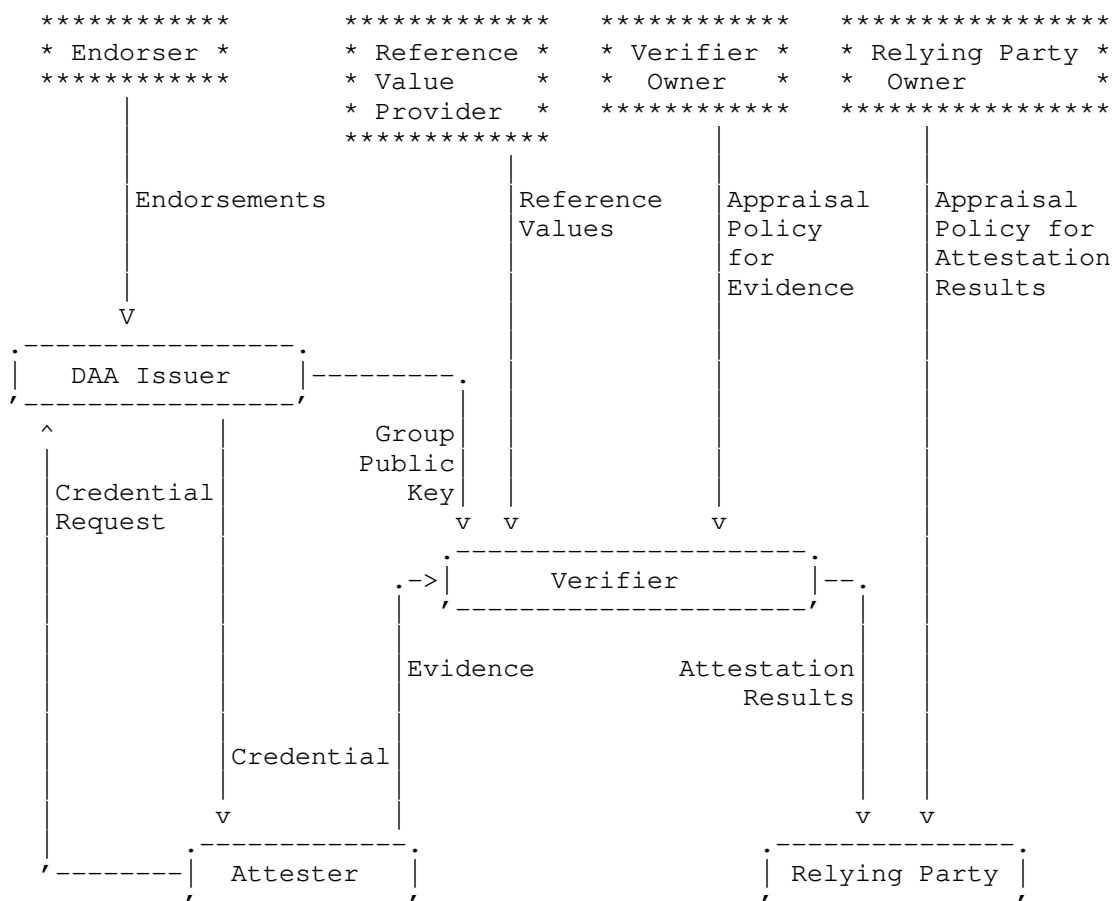


Figure 1: DAA data flows

DAA [DAA] is a signature scheme that allows the privacy of users that are associated with an Attester (e.g. its owner) to be maintained. Essentially, DAA can be seen as a group signature scheme with the feature that given a DAA signature no-one can find out who the signer is, i.e., the anonymity is not revocable. To be able to sign anonymously, an Attester has to obtain a credential from a DAA Issuer. The DAA Issuer uses a private/public key pair to generate credentials for a group of Attesters and makes the public key (in the form of a public key certificate) available to the verifier in order to enable them to validate the Evidence received.

In order to support these DAA signatures, the DAA Issuer MUST associate a single key pair with a group of Attesters and use the same key pair when creating the credentials for all of the Attesters

in this group. The DAA Issuer's group public key certificate replaces the individual Attester Identity documents during authenticity validation as a part of the appraisal of Evidence conducted by a verifier. This is in contrast to intuition that there has to be a unique Attester Identity per device.

For DAA, the role of the Endorser is essentially the same, but they now provide Attester endorsement documents to the DAA Issuer rather than directly to the verifier. These Attester endorsement documents enable the Attester to obtain a credential from the DAA Issuer.

4. DAA changes to the RATS Architecture

In order to enable the use of DAA, a new conceptual message, the Credential Request, is defined and a new role, the DAA Issuer role, is added to the roles defined in the RATS Architecture.

Credential Request: An Attester sends a Credential Request to the DAA Issuer to obtain a credential. This request contains information about the DAA key that the Attester will use to create evidence and together with Attester endorsement information that is provided by the Endorser to confirm that the request came from a valid Attester.

DAA Issuer: A RATS role that offers zero-knowledge proofs based on public-key certificates used for a group of Attesters (Group Public Keys) [DAA]. How this group of Attesters is defined is not specified here, but the group must be large enough for the necessary anonymity to be assured.

Effectively, these certificates share the semantics of Endorsements, with the following exceptions:

- * Upon receiving a Credential Request from an Attester, the associated group private key is used by the DAA Issuer to provide the Attester with a credential that it can use to convince the Verifier that its Evidence is valid. To keep their anonymity the Attester randomizes this credential each time that it is used. Although the DAA Issuer knows the Attester Identity and can associate this with the credential issued, randomisation ensures that the Attester's identity cannot be revealed to anyone, including the Issuer.
- * The Verifier can use the DAA Issuer's group public key certificate, together with the randomized credential from the Attester, to confirm that the Evidence comes from a valid Attester without revealing the Attester's identity.

- * A credential is conveyed from a DAA Issuer to an Attester in combination with the conveyance of the group public key certificate from DAA Issuer to Verifier.

5. Additions to Remote Attestation principles

In order to ensure an appropriate conveyance of Evidence via interaction models in general, the following prerequisite considering Attester Identity MUST be in place to support the implementation of interaction models.

Attestation Evidence Authenticity: Attestation Evidence MUST be correct and authentic.

In order to provide proofs of authenticity, Attestation Evidence SHOULD be cryptographically associated with an identity document that is a randomised DAA credential.

The following information elements define extensions for corresponding information elements defined in [I-D.ietf-rats-reference-interaction-models] and that are vital to all types of reference interaction models. Varying from solution to solution, generic information elements can be either included in the scope of protocol messages (instantiating Conceptual Messages defined by the RATS architecture) or can be included in additional protocol parameters of protocols that facilitate the conveyance of RATS Conceptual Messages. Ultimately, the following information elements are required by any kind of scalable remote attestation procedure using DAA with one of RATS's reference interaction models.

Attester Identity ('attesterIdentity'): _mandatory_

In DAA, the Attester's identity is not revealed to the verifier. The Attester is issued with a credential by the DAA Issuer that is randomised and then used to anonymously confirm the validity of their evidence. The evidence is verified using the DAA Issuer's group public key.

Authentication Secret IDs ('authSecID'): _mandatory_

In DAA, Authentication Secret IDs are represented by the DAA Issuer's group public key that MUST be used to create DAA credentials for the corresponding Authentication Secrets used to protect Evidence.

In DAA, an Authentication Secret ID does not identify a unique

Attesting Environment but is associated with a group of Attesting Environments. This is because an Attesting Environment should not be distinguishable and the DAA credential which represents the Attesting Environment is randomised each time it used.

6. Privacy Considerations

As outlined about for DAA to provide privacy for the Attester the DAA group must be large enough to stop the Verifier identifying the Attester.

Randomization of the DAA credential by the Attester means that collusion between the DAA Issuer and Verifier, will not give them any advantage when trying to identify the Attester.

For DAA, the Attestation Evidence conveyed to the Verifier MUST not uniquely identify the Attester. If the Attestation Evidence is unique to an Attester other cryptographic techniques can be used, for example, property based attestation. (Henk -- reference follows)

Chen L., Loehr H., Manulis M., Sadeghi AR. (2008) Property-Based Attestation without a Trusted Third Party. Information Security. ISC 2008. Lecture Notes in Computer Science, vol 5222. Springer. https://doi.org/10.1007/978-3-540-85886-7_3

7. Security Considerations

The anonymity property of DAA makes revocation difficult. Well known solutions include: 1. Rogue attester revocation -- if the an Attester's private key is compromised and known by the Verifier then any DAA signature from that Attester can be revoked. 2. EPID - Intel's Enhanced Privacy ID -- this requires the Attester to prove (as part of their Attestation) that their credential was not used to generate any signature in a signature revocation list.

There are no other special security conderations for DAA over and above those specifed in the RATS architecture document [I-D.ietf-rats-architecture].

8. Implementation Considerations

The new DAA Issuer role can be implemented in a number of ways, for example: 1. As a stand-alone service like a Certificate Authority, a Privacy CA. 2. As a part of the Attester's manufacture. The Endorser and the DAA Issuer could be the same entity and the manufacturer would then provide a certificate for the group public key to the Verifier.

9. IANA Considerations

We don't yet.

10. References

10.1. Normative References

- [DAA] Brickell, E., Camenisch, J., and L. Chen, "Direct Anonymous Attestation", page 132-145, ACM Proceedings of the 11rd ACM conference on Computer and Communications Security, 2004.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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://www.rfc-editor.org/info/rfc5280>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

10.2. Informative References

- [I-D.ietf-rats-architecture] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-12.txt>>.
- [I-D.ietf-rats-reference-interaction-models] Birkholz, H., Eckel, M., Pan, W., and E. Voit, "Reference Interaction Models for Remote Attestation Procedures", Work in Progress, Internet-Draft, draft-ietf-rats-reference-interaction-models-04, 26 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-reference-interaction-models-04.txt>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Christopher Newton
University of Surrey

Email: cn0016@surrey.ac.uk

Liqun Chen
University of Surrey

Email: liqun.chen@surrey.ac.uk

Dave Thaler
Microsoft
United States of America

Email: dthaler@microsoft.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 18 February 2022

H. Birkholz
Fraunhofer SIT
E. Voit
Cisco
W. Pan
Huawei
17 August 2021

Attestation Event Stream Subscription
draft-birkholz-rats-network-device-subscription-03

Abstract

This memo defines how to subscribe to YANG Event Streams for Remote Attestation Procedures (RATS). In RATS, Conceptual Messages, are defined. Analogously, the YANG module defined in this memo augments the YANG module for TPM-based Challenge-Response based Remote Attestation (CHARRA) to allow for subscription to remote attestation Evidence. Additionally, this memo provides the methods and means to define additional Event Streams for other Conceptual Message as illustrated in the RATS Architecture, e.g. Attestation Results, Endorsements, or Event Logs.

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 18 February 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	5
2.1. Requirements Notation	5
3. Operational Model	5
3.1. Sequence Diagram	5
3.2. Continuously Verifying Freshness	7
3.2.1. TPM 1.2 Quote	8
3.2.2. TPM 2 Quote	8
4. Remote Attestation Event Stream	9
4.1. Subscription to the <attestation> Event Stream	9
4.2. Replaying a history of previous TPM extend operations	10
4.2.1. TPM2 Heartbeat	11
4.3. YANG notifications placed on the <attestation> Event Stream	11
4.3.1. pcr-extend	11
4.3.2. tpm12-attestation	13
4.3.3. tpm20-attestation	13
4.4. Filtering Evidence at the Attester	14
4.5. Replaying previous PCR Extend events	14
4.6. Configuring the <attestation> Event Stream	14
5. YANG Module	15
6. Event Streams for Conceptual Messages	22
7. Security Considerations	22
8. IANA Considerations	22
9. References	22
9.1. Normative References	22
9.2. Informative References	23
Appendix A. Change Log	24
Acknowledgements	24
Authors' Addresses	24

1. Introduction

[I-D.ietf-rats-tpm-based-network-device-attest] and [I-D.ietf-rats-yang-tpm-charra] define the operational prerequisites and a YANG Model for the acquisition of Evidence and other Conceptual Messages from a TPM-based network device. However, there are limitations inherent in the challenge-response based remote attestation (CHARRA [I-D.ietf-rats-reference-interaction-models]) upon which these documents are based. One of these limitation is that it is a RATS role's duty to request Conceptual Messages, such as Evidence as provided by [I-D.ietf-rats-yang-tpm-charra], from another RATS entity. The result is that the interval between the occurrence of a security-relevant change event, and the event's visibility within the interested RATS entity, such as a Verifier or a Relying Party, can be unacceptably long. It is common to convey Conceptual Messages ad-hoc or periodically via requests. As new technologies emerge, some of these solutions require Conceptual Messages to be conveyed from one RATS entity to another without the need of continuous polling. Subscription to YANG Notifications [RFC8639] provides a set of standardized tools to facilitate these emerging requirements. This memo specifies a YANG augment to subscribe to YANG modeled remote attestation Evidence as defined in [I-D.ietf-rats-yang-tpm-charra]. Additionally, this memo provides the means to define further Event Streams to convey Conceptual Messages other than Evidence, such as Attestation Results, Endorsements, or Event Logs.

In essence, the limitation of poll-based interactions results in two adverse effects:

1. Conceptual Messages are not streamed to an interested consumer of information, e.g., Verifiers or Relying Parties, as soon as they are generated.
2. If they were to be streamed, Conceptual Messages are not appraisable for their freshness in every scenario. This becomes more important with Conceptual Messages that have a strong dependency on freshness, such as Evidence and corresponding Attestation Results.

This specification addresses the first adverse effect by enabling a consumer of Conceptual Messages (the subscriber) to request a continuous stream of new or updated Conceptual Messages via an [RFC8639] subscription to an <attestation> Event Stream. This new Event Stream is defined in this document and exists upon the producer of Conceptual Messages (the publisher). In the case of a Verifier's subscription to an Attester's Evidence, the Attester will continuously stream a requested set of freshly generated Evidence to the subscribing Verifier.

The second adverse effect results from the use of nonces in the challenge-response interaction model [I-D.ietf-rats-reference-interaction-models] realized in [I-D.ietf-rats-yang-tpm-charra]. In [I-D.ietf-rats-yang-tpm-charra], an Attester must wait for a new nonce from a Verifier before it generates a new TPM Quote. To address delays resulting from such a wait, this specification enables freshness to be asserted asynchronously via the streaming attestation interaction model [I-D.ietf-rats-reference-interaction-models]. To convey a RATS Conceptual Message, an initial nonce is provided during the subscription to an Event Stream.

There are several options to refresh a nonce provided by the initial subscription or its freshness characteristics. All of these methods are out-of-band of an established subscription to YANG Notifications. Two complementary methods are taken into account by this memo:

1. a central provider supplies new fresh nonces, e.g. via a Handle Provider that distributes Epoch IDs to all entities in a domain as described in [I-D.ietf-rats-architecture] and as facilitated by the Uni-Directional Remote Attestation described in [I-D.ietf-rats-reference-interaction-models] or
2. the freshness characteristics of a received nonce are updated by -- potentially periodic or ad-hoc -- out-of-band TPM Quote requests as facilitated by [I-D.ietf-rats-yang-tpm-charra].

Both approaches to update the freshness characteristics of the Conceptual Messages conveyed via subscription to YANG Notification that are taken into account by this memo assume that clock drift between involved entities can occur. In consequence, in some usage scenarios the timing considerations for freshness [I-D.ietf-rats-architecture] might have to be updated in some regular interval. Analogously, there can be additional methods that are not describe by but nevertheless supported by this memo.

This memo enables to remove the two adverse effects described by using the YANG augment specified. The YANG augment supports, for example, a RATS Verifier to maintain a continuous appraisal procedure of verifiably fresh Attester Evidence without relying on continuous polling.

2. Terminology

The following terms are imported from [I-D.ietf-rats-architecture]: Attester, Conceptual Message, Evidence, Relying Party, and Verifier. Also imported are the time definitions time(VG), time(NS), time(EG), time(RG), and time(RA) from that document's Appendix A. The following terms are imported from [RFC8639]: Event Stream, Subscription, Event Stream Filter, Dynamic Subscription.

2.1. Requirements Notation

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.

3. Operational Model

[I-D.ietf-rats-tpm-based-network-device-attest] describes the conveyance of TPM-based Evidence from a Verifier to an Attester using the CHARRA interaction model [I-D.ietf-rats-reference-interaction-models]. The operational model and corresponding sequence diagram described in this section is based on [I-D.ietf-rats-yang-tpm-charra]. The basis for interoperability required for additional types of Event Streams is covered in Section 6. The following sub-section focuses on subscription to YANG Notifications to the <attestation> Event Stream.

3.1. Sequence Diagram

Figure 1 below is a sequence diagram which updates Figure 5 of [I-D.ietf-rats-tpm-based-network-device-attest]. This sequence diagram replaces the [I-D.ietf-rats-tpm-based-network-device-attest] TPM-specific challenge-response interaction model with a [RFC8639] Dynamic Subscription to an <attestation> Event Stream. The contents of the <attestation> Event Stream are defined below within Section 4.

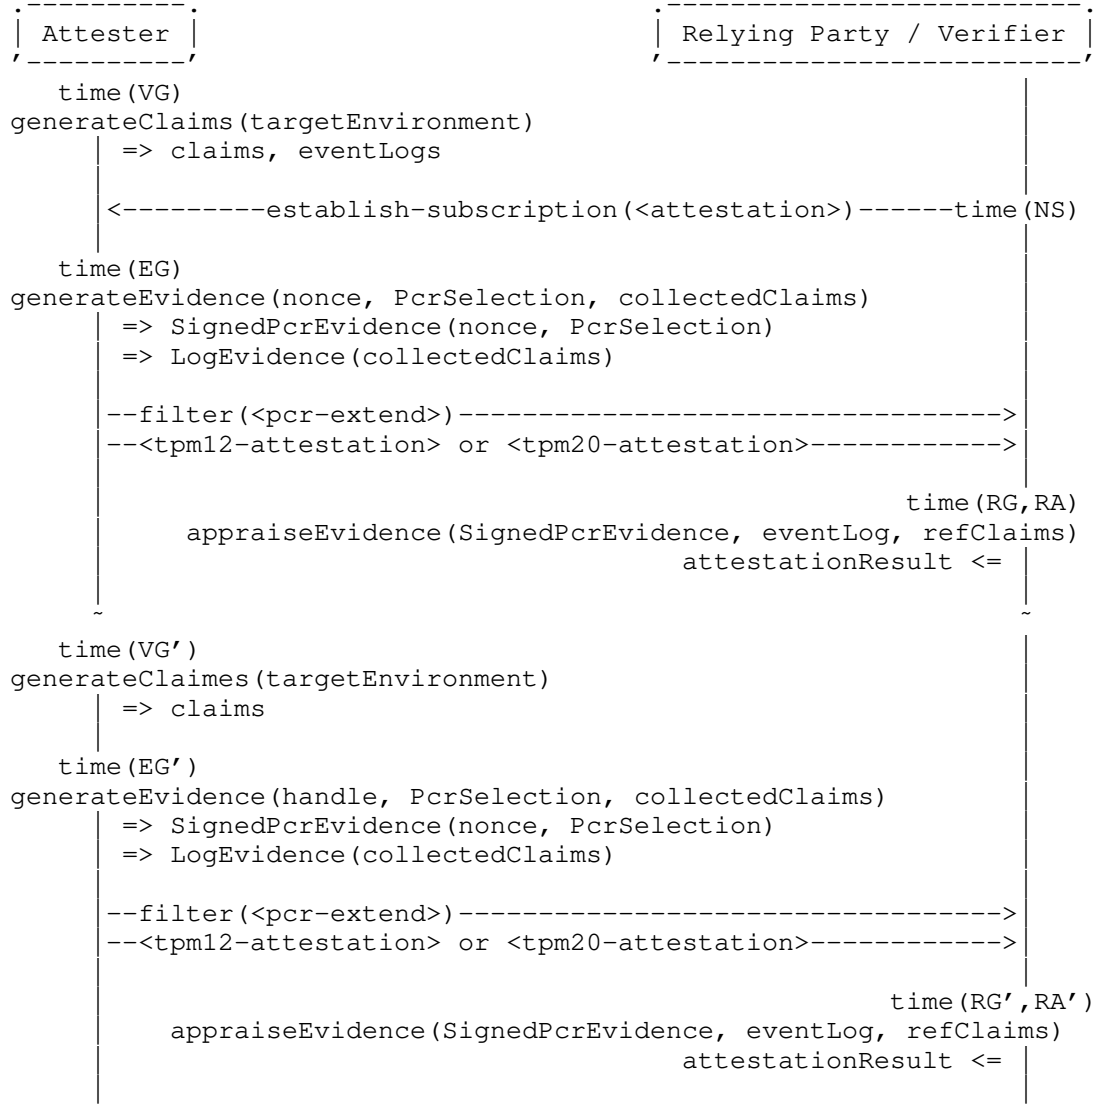


Figure 1: YANG Subscription Model for Remote Attestation

- * time(VG, RG, RA) are identical to the corresponding time definitions from [I-D.ietf-rats-tpm-based-network-device-attest].
- * time(VG', RG', RA') are subsequent instances of the corresponding times from Figure 5 in [I-D.ietf-rats-tpm-based-network-device-attest].

- * time(NS) - the subscriber generates a nonce and makes an [RFC8639] <establish-subscription> request based on a nonce. This request also includes the augmentations defined in this document's YANG model. Key subscription RPC parameters include:
 - the nonce,
 - a set of PCRs of interest which the wants to appraise, and
 - an optional filter which can reduce the logged events on the <attestation> stream pushed to the Verifier.
- * time(EG) - an initial response of Evidence is returned to the Verifier. This includes:
 - a replay of filtered log entries which have extended into a PCR of interest since boot are sent in the <pcr-extend> notification, and
 - a signed TPM quote that contains at least the PCRs from the <establish-subscription> RPC are included in a <tpm12-attestation> or <tpm20-attestation>). This quote must have included the nonce provided at time(NS).
- * time(VG',EG') - this occurs when a PCR is extended subsequent to time(EG). Immediately after the extension, the following information needs to be pushed to the Verifier:
 - any values extended into a PCR of interest,
 - a signed TPM Quote showing the result the PCR extension, and
 - and a handle (see Section 6. in [I-D.ietf-rats-reference-interaction-models], which is either the initially received nonce or a more recently received Epoch ID (see Section 10.3. in [I-D.ietf-rats-architecture] that contains a new nonce or equivalent qualified data.

One way to acquire a new time synchronization that allows for the reuse of the initially received nonce as a fresh handle is elaborated on in the follow section Section 3.2.

3.2. Continuously Verifying Freshness

As there is no new Verifier nonce provided at time(EG'), it is important to validate the freshness of TPM Quotes which are delivered at that time. The method of doing this verification will vary based on the capabilities of the TPM cryptoprocessor used.

3.2.1. TPM 1.2 Quote

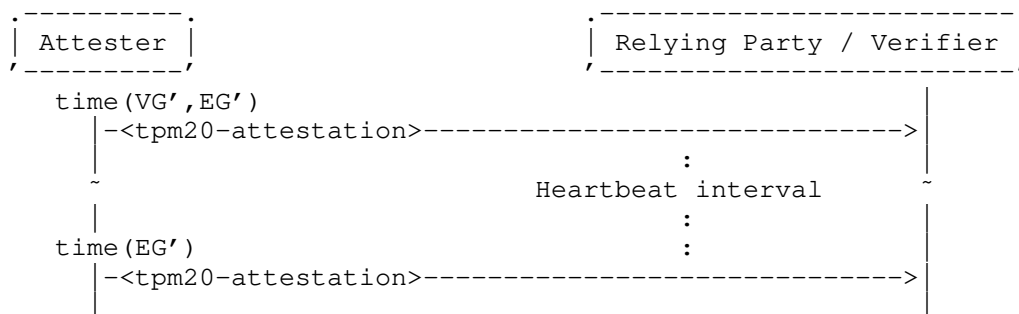
The [RFC8639] notification format includes the <eventTime> object. This can be used to determine the amount of time subsequent to the initial subscription each notification was sent. However this time is not part of the signed results which are returned from the Quote, and therefore is not trustworthy as objects returned in the Quote. Therefore a Verifier MUST periodically issue a new nonce, and receive this nonce within a TPM quote response in order to ensure the freshness of the results. This can be done using the <tpm12-challenge-response-attestation> RPC from [I-D.ietf-rats-yang-tpm-charra].

3.2.2. TPM 2 Quote

When the Attester includes a TPM2 compliant cryptoprocessor, internal time-related counters are included within the signed TPM Quote. By including a initial nonce in the [RFC8639] subscription request, fresh values for these counters are pushed as part of the first TPM Quote returned to the Verifier. And then as shown by [I-D.birkholz-rats-tuda], subsequent TPM Quotes delivered to the Verifier can the be appraised for freshness based on the predictable incrementing of these time-related counters.

The relevant internal time-related counters defined within [TPM2.0] can be seen within <tpms-clock-info>. These counters include the <clock>, <reset-counter>, and <restart-counter> objects. The rules for appraising these objects are as follows:

- * If the <clock> has incremented for no more than the same duration as both the <eventTime> and the Verifier's internal time since the initial time(EG) and any previous time(EG'), then the TPM Quote may be considered fresh. Note that [TPM2.0] allows for +/- 15% clock drift. However many chips significantly improve on this maximum drift. If available, chip specific maximum drifts SHOULD be considered during the appraisal process.
- * If the <reset-counter>, <restart-counter> has incremented. The existing subscription MUST be terminated, and a new <establish-subscription> SHOULD be generated.
- * If a TPM Quote on any subscribed PCR has not been pushed to the Verifier for a duration of an Attester defined heartbeat interval, then a new TPM Quote notification should be sent to the Verifier. This may often be the case, as certain PCRs might be infrequently updated.



4. Remote Attestation Event Stream

The <attestation> Event Stream is an [RFC8639] compliant Event Stream which is defined within this section and within the YANG Module of [I-D.ietf-rats-yang-tpm-charra]. This Event Stream contains YANG notifications which carry Evidence to assist a Verifier in appraising the Trustworthiness Level of an Attester. Data Nodes within Section 4.6 allow the configuration of this Event Stream's contents on an Attester.

This <attestation> Event Stream may only be exposed on Attesters supporting [I-D.ietf-rats-tpm-based-network-device-attest]. As with [I-D.ietf-rats-tpm-based-network-device-attest], it is up to the Verifier to understand which types of cryptoprocessors and keys are acceptable.

4.1. Subscription to the <attestation> Event Stream

To establish a subscription to an Attester in a way which provides provably fresh Evidence, initial randomness must be provided to the Attester. This is done via the augmentation of a <nonce-value> into [RFC8639] the <establish-subscription> RPC. Additionally, a Verifier must ask for PCRs of interest from a platform.

```

augment /sn:establish-subscription/sn:input:
  +---w nonce-value      binary
  +---w pcr-index*       tpm:pcr
  
```

The result of the subscription will be that passing of the following information:

1. <tpm12-attestation> and <tpm20-attestation> notifications which include the provided <nonce-value>. These attestation notifications MUST at least include all the <pcr-indicies> requested in the RPC.

2. a series of <pcr-extend> notifications which reference the requested PCRs on all TPM based cryptoprocessors on the Attester.
3. <tpm12-attestation> and <tpm20-attestation> notifications generated within a few seconds of the <pcr-extend> notifications. These attestation notifications MUST at least include any PCRs extended.

If the Verifier does not want to see the logged extend operations for all PCRs available from an Attester, an Event Stream Filter should be applied. This filter will remove Evidence from any PCRs which are not interesting to the Verifier.

4.2. Replaying a history of previous TPM extend operations

Unless it is relying on Known Good Values, a Verifier will need to acquire a history of PCR extensions since the Attester has been booted. This history may be requested from the Attester as part of the <establish-subscription> RPC. This request is accomplished by placing a very old <replay-start-time> within the original RPC request. As the very old <replay-start-time> will pre-date the time of Attester boot, a <replay-start-time-revision> will be returned in the <establish-subscription> RPC response, indicating when the Attester booted. Immediately following the response (and before the notifications above) one or more <pcr-extend> notifications which document all extend operations which have occurred for the requested PCRs since boot will be sent. Many extend operations to a single PCR index on a single TPM SHOULD be included within a single notification.

Note that if a Verifier has a partial history of extensions, the <replay-start-time> can be adjusted so that known extensions are not forwarded.

The end of this history replay will be indicated with the [RFC8639] <replay-completed> notification. For more on this sequence, see Section 2.4.2.1 of [RFC8639].

After the <replay-complete> notification is provided, a TPM Quote will be requested and the result passed to the Verifier via a <tpm12-attestation> and <tpm20-attestation> notification. If there have been any additional extend operations which have changed a subscribed PCR value in this quote, these MUST be pushed to the Verifier before the <tpm12-attestation> and <tpm20-attestation> notification.

At this point the Verifier has sufficient Evidence appraise the reported extend operations for each PCR, as well compare the expected value of the PCR value against that signed by the TPM.

4.2.1. TPM2 Heartbeat

For TPM2, make sure that every requested PCR is sent within an <tpm20-attestation> no less frequently than once per heartbeat interval. This MAY be done with a single <tpm20-attestation> notification that includes all requested PCRs every heartbeat interval. This MAY be done with several <tpm20-attestation> notifications at different times during that heartbeat interval.

4.3. YANG notifications placed on the <attestation> Event Stream

4.3.1. pcr-extend

This notification documents when a subscribed PCR is extended within a single TPM cryptoprocessor. It SHOULD be emitted no less than the <marshalling-period> after an the PCR is first extended. (The reason for the marshalling is that it is quite possible that multiple extensions to the same PCR have been made in quick succession, and these should be reflected in the same notification.) This notification MUST be emitted prior to a <tpm12-attestation> or <tpm20-attestation> notification which has included and signed the results of any specific PCR extension. If pcr extending events occur during the generation of the <tpm12-attestation> or <tpm20-attestation> notification, the marshalling period MUST be extended so that a new <pcr-extend> is not sent until the corresponding notifications have been sent.

```

+---n pcr-extend
+--ro certificate-name      certificate-name-ref
+--ro pcr-index-changed*   tpm:pcr
+--ro attested-event* []
+--ro attested-event
+--ro extended-with          binary
+--ro (event-details)?
+--:(bios-event-log)
+--ro bios-event-entry* [event-number]
+--ro event-number          uint32
+--ro event-type?           uint32
+--ro pcr-index?            pcr
+--ro digest-list* []
+--ro hash-algo?            identityref
+--ro digest*               binary
+--ro event-size?           uint32
+--ro event-data*           uint8
+--:(ima-event-log)
+--ro ima-event-entry* [event-number]
+--ro event-number          uint64
+--ro ima-template?         string
+--ro filename-hint?        string
+--ro filedata-hash?        binary
+--ro filedata-hash-algorithm? string
+--ro template-hash-algorithm? string
+--ro template-hash?        binary
+--ro pcr-index?            pcr
+--ro signature?            binary
+--:(netequip-boot-event-log)
+--ro boot-event-entry* [event-number]
+--ro event-number          uint64
+--ro filename-hint?        string
+--ro filedata-hash?        binary
+--ro filedata-hash-algorithm? string
+--ro file-version?         string
+--ro file-type?            string
+--ro pcr-index?            pcr

```

Each <pcr-extend> MUST include one or more values being extended into the PCR. These are passed within the <extended-with> object. For each extension, details of the event SHOULD be provided within the <event-details> object. The format of any included <event-details> is identified by the <event-type>. This document includes two YANG structures which may be inserted into the <event-details>. These two structures are: <ima-event-log> and <bios-event-log>. Implementations wanting to provide additional documentation of a type of PCR extension may choose to define additional YANG structures which can be placed into <event-details>.

4.3.2. tpm12-attestation

This notification contains an instance of a TPM1.2 style signed cryptoprocessor measurement. It is supplemented by Attester information which is not signed. This notification is generated and emitted from an Attester when at least one PCR identified within the subscribed <pcr-indices> has changed from the previous <tpm12-attestation> notification. This notification MUST NOT include the results of any PCR extensions not previously reported by a <pcr-extend>. This notification SHOULD be emitted as soon as a TPM Quote can extract the latest PCR hashed values. This notification MUST be emitted prior to a subsequent <pcr-extend>.

```
+---n tpm12-attestation {taa:TPM12}?
  +--ro certificate-name      tpm:certificate-name-ref
  +--ro up-time?              uint32
  +--ro TPM_QUOTE2?           binary
  +--ro TPM12-hash-algo?      identityref
  +--ro unsigned-pcr-values* []
    +--ro pcr-index*          tpm:pcr
    +--ro pcr-value*          binary
```

All YANG objects above are defined within [I-D.ietf-rats-yang-tpm-charra]. The <tpm12-attestation> is not replayable.

4.3.3. tpm20-attestation

This notification contains an instance of TPM2 style signed cryptoprocessor measurements. It is supplemented by Attester information which is not signed. This notification is generated at two points in time:

- * every time at least one PCR has changed from a previous tpm20-attestation. In this case, the notification SHOULD be emitted within 10 seconds of the corresponding <pcr-extend> being sent:
- * after a locally configurable minimum heartbeat period since a previous tpm20-attestation was sent.

```
+---n tpm20-attestation {taa:TPM20}?
  +--ro certificate-name      tpm:certificate-name-ref
  +--ro TPMS_QUOTE_INFO      binary
  +--ro quote-signature?     binary
  +--ro up-time?             uint32
  +--ro unsigned-pcr-values* []
    +--ro TPM20-hash-algo?   identityref
    +--ro pcr-values* [pcr-index]
      +--ro pcr-index       pcr
      +--ro pcr-value?     binary
```

All YANG objects above are defined within [I-D.ietf-rats-yang-tpm-charra]. The <tpm20-attestation> is not replayable.

4.4. Filtering Evidence at the Attester

It can be useful not to receive all Evidence related to a PCR. An example of this is would be a when a Verifier maintains known good values of a PCR. In this case, it is not necessary to send each extend operation.

To accomplish this reduction, when an RFC8639 <establish-subscription> RPC is sent, a <stream-filter> as per RFC8639, Section 2.2 can be set to discard a <pcr-extend> notification when the <pcr-index-changed> is uninteresting to the verifier.

4.5. Replaying previous PCR Extend events

To verify the value of a PCR, a Verifier must either know that the value is a known good value [KGV] or be able to reconstruct the hash value by viewing all the PCR-Extends since the Attester rebooted. Wherever a hash reconstruction might be needed, the <attestation> Event Stream MUST support the RFC8639 <replay> feature. Through the <replay> feature, it is possible for a Verifier to retrieve and sequentially hash all of the PCR extending events since an Attester booted. And thus, the Verifier has access to all the evidence needed to verify a PCR's current value.

4.6. Configuring the <attestation> Event Stream

Figure 2 is tree diagram which exposes the operator configurable elements of the <attestation> Event Stream. This allows an Attester to select what information should be available on the stream. A fetch operation also allows an external device such as a Verifier to understand the current configuration of stream.

Almost all YANG objects below are defined via reference from [I-D.ietf-rats-yang-tpm-charra]. There is one object which is new with this model however. <tpm2-heartbeat> defines the maximum amount of time which should pass before a subscriber to the Event Stream should get a <tpm20-attestation> notification from devices which contain a TPM2.

```
augment /tpm:rats-support-structures:
  +--rw marshalling-period?          uint8
  +--rw tpm12-subscribed-signature-scheme?
    |   -> ../tpm:attester-supported-algos/tpm12-asymmetric-signing
    |   {taa:TPM12}?
  +--rw tpm20-subscribed-signature-scheme?
    |   -> ../tpm:attester-supported-algos/tpm20-asymmetric-signing
    |   {taa:TPM20}?
  +--rw tpm20-subscription-heartbeat?    uint16
augment /tpm:rats-support-structures/tpm:tpms:
  +--rw subscription-aik?          tpm:certificate-name-ref
  +--rw (subscribable)?
    +--:(tpm12-stream) {taa:TPM12}?
    |   +--rw TPM12-hash-algo?    identityref
    |   +--rw tpm12-pcr-index*    tpm:pcr
    +--:(tpm20-stream) {taa:TPM20}?
    |   +--rw TPM20-hash-algo?    identityref
    |   +--rw tpm20-pcr-index*    tpm:pcr
```

Figure 2: Configuring the \<attestation\> Event Stream

5. YANG Module

This YANG module imports modules from [I-D.ietf-rats-yang-tpm-charra] and [RFC8639]. It is also work-in-progress.

```
<CODE BEGINS> ietf-rats-attestation-stream@2020-12-15.yang
module ietf-tpm-remote-attestation-stream {
  yang-version 1.1;
  namespace
    "urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation-stream";
  prefix tras;

  import ietf-subscribed-notifications {
    prefix sn;
    reference
      "RFC 8639: Subscription to YANG Notifications";
  }
  import ietf-tpm-remote-attestation {
    prefix tpm;
    reference
```

```
    "draft-ietf-rats-yang-tpm-charra";
  }
  import ietf-tcg-algs {
    prefix taa;
  }

  organization "IETF";
  contact
    "WG Web:  <http://tools.ietf.org/wg/rats/>
    WG List:  <mailto:rats@ietf.org>

    Editor:   Eric Voit
              <mailto:evoit@cisco.com>";

  description
    "This module contains conceptual YANG specifications for
    subscribing to attestation streams being generated from TPM chips.

    Copyright (c) 2021 IETF Trust and the persons identified
    as authors of the code. All rights reserved.

    Redistribution and use in source and binary forms, with
    or without modification, is permitted pursuant to, and
    subject to the license terms contained in, the Simplified
    BSD License set forth in Section 4.c of the IETF Trust's
    Legal Provisions Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); see the RFC
    itself for full legal notices.";

  revision 2021-05-11 {
    description
      "Initial version.";
    reference
      "draft-birkholz-rats-network-device-subscription";
  }

  /*
   * IDENTITIES
   */

  identity pcr-unsubscribable {
    base sn:establish-subscription-error;
    description
      "Requested PCR is subscribable by the Attester.";
  }
```

```
/*
 * Groupings
 */

grouping heartbeat {
  description
    "Allows an Attester to push verifiable, current TPM PCR values
    even when there have been no recent changes to PCRs.";
  leaf tpm20-subscription-heartbeat {
    type uint16;
    description
      "Number of seconds before the Attestation stream should send a
      new notification with a fresh quote. This allows confirmation
      that the PCR values haven't changed since the last
      tpm20-attestation.";
  }
}

/*
 * RPCs
 */

augment "/sn:establish-subscription/sn:input" {
  when 'derived-from-or-self(sn:stream, "attestation)';
  description
    "This augmentation adds a nonce to as a subscription parameters
    that apply specifically to datastore updates to RPC input.";
  uses tpm:nonce;
  leaf-list pcr-index {
    type tpm:pcr;
    min-elements 1;
    description
      "The numbers/indexes of the PCRs. This will act as a filter for
      the subscription so that 'tpm-extend' notifications related to
      non-requested PCRs will not be sent to a subscriber.";
  }
}

/*
 * NOTIFICATIONS
 */

notification pcr-extend {
  description
    "This notification indicates that one or more PCRs have been
    extended within a TPM based cryptoprocessor. In less than the
    'marshalling-period', it MUST be followed with either a
```



```
    corresponding tpm12-attestation or tpm20-attestation notification
    which exposes the result of the PCRs updated.";
uses tpm:certificate-name-ref;
leaf-list pcr-index-changed {
  type tpm:pcr;
  min-elements 1;
  description
    "The number of each PCR extended. This list MUST contain the
    set of PCRs described within the event log details. This leaf
    can be derived from the list of attested events, but exposing
    it here allows for easy filtering of the notifications of
    interest to a verifier.";
}
list attested-event {
  description
    "A set of events which extended an Attester PCR. The sequence
    of elements represented in list must match the sequence of
    events placed into the TPM's PCR.";
  container attested-event {
    description
      "An instance of an event which extended an Attester PCR";
    leaf extended-with {
      type binary;
      mandatory true;
      description
        "Information extending the PCR.";
    }
    choice event-details {
      description
        "Contains the event happened the Attester thought
        was worthy of recording in a PCR.

        choices are of types defined by the identityref
        base tpm:attested_event_log_type";
      case bios-event-log {
        if-feature "tpm:bios";
        description
          "BIOS/UEFI event log format";
        uses tpm:bios-event-log;
      }
      case ima-event-log {
        if-feature "tpm:ima";
        description
          "IMA event log format";
        uses tpm:ima-event-log;
      }
      case netequip-boot-event-log {
        if-feature "tpm:netequip_boot";
```

```
        description
          "IMA event log format";
        uses tpm:network-equipment-boot-event-log;
      }
    }
  }
}

notification tpm12-attestation {
  if-feature "taa:tpm12";
  description
    "Contains an instance of TPM1.2 style signed cryptoprocessor
    measurements. It is supplemented by unsigned Attester
    information.";
  leaf certificate-name {
    type tpm:certificate-name-ref;
    mandatory true;
    description
      "Allows a TPM quote to be associated with a certificate.";
  }
  uses tpm:tpm12-attestation;
  uses tpm:tpm12-hash-algo;
  list unsigned-pcr-values {
    description
      "Allows notifications to be filtered by PCR number or
      PCR value based on via YANG related mechanisms such as PATH.
      This is done without requiring the filtering structure to be
      applied against TCG structured data.";
    leaf-list pcr-index {
      type tpm:pcr;
      min-elements 1;
      description
        "PCR index number.";
    }
    leaf-list pcr-value {
      type binary;
      description
        "PCR value in a sequence which matches to the 'pcr-index'.";
    }
  }
}

notification tpm20-attestation {
  if-feature "taa:tpm20";
  description
    "Contains an instance of TPM2 style signed cryptoprocessor
    measurements. It is supplemented by unsigned Attester
```

```
        information.";
    leaf certificate-name {
        type tpm:certificate-name-ref;
        mandatory true;
        description
            "Allows a TPM quote to be associated with a certificate.";
    }
    uses tpm:tpm20-attestation {
        description
            "Provides the attestation info. Also ensures PCRs can be XPATH
            filtered by refining the unsigned data so that it appears.";
        refine unsigned-pcr-values {
            min-elements 1;
        }
        refine unsigned-pcr-values/pcr-values {
            min-elements 1;
        }
    }
}

/*
 * DATA NODES
 */

augment "/tpm:rats-support-structures" {
    description
        "Defines platform wide 'attestation' stream subscription
        parameters.";
    leaf marshallng-period {
        type uint8;
        default 5;
        description
            "The maximum number of seconds between the time an event
            extends a PCR, and the 'tpm-extend' notification which reports
            it to a subscribed Verifier. This period allows multiple
            extend operations bundled together and handled as a group.";
    }
    leaf tpm12-subscribed-signature-scheme {
        if-feature "taa:tpm12";
        type leafref {
            path "../tpm:attester-supported-algos" +
                "/tpm:tpm12-asymmetric-signing";
        }
        description
            "A single signature-scheme which will be used to sign the
            evidence from a TPM 1.2. which is then placed onto the
            'attestation' event stream.";
    }
}
```

```
}
leaf tpm20-subscribed-signature-scheme {
  if-feature "taa:tpm20";
  type leafref {
    path "../tpm:attester-supported-algos" +
        "/tpm:tpm20-asymmetric-signing";
  }
  description
    "A single signature-scheme which will be used to sign the
    evidence from a TPM 2.0. which is then placed onto the
    'attestation' event stream.";
}
uses heartbeat{
  if-feature "taa:tpm20";
}
}

augment "/tpm:rats-support-structures/tpm:tpms" {
  description
    "Allows the configuration 'attestation' stream parameters for a
    TPM.";
  leaf subscription-aik {
    type tpm:certificate-name-ref;
    description
      "Identifies the certificate-name associated with the
      notifications in the 'attestation' stream.";
  }
  choice subscribable {
    config true;
    description
      "Indicates that the set of notifications which comprise the
      'attestation' event stream can be modified or tuned by a
      network administrator.";
    case tpm12-stream {
      if-feature "taa:tpm12";
      description
        "Configuration elements for a TPM1.2 event stream.";
      uses tpm:tpm12-hash-algo;
      leaf-list tpm12-pcr-index {
        type tpm:pcr;
        description
          "The numbers/indexes of the PCRs which can be subscribed.";
      }
    }
    case tpm20-stream {
      if-feature "taa:tpm20";
      description
        "Configuration elements for a TPM2.0 event stream.";
    }
  }
}
```

```
    uses tpm:tpm20-hash-algo;
    leaf-list tpm20-pcr-index {
      type tpm:pcr;
      description
        "The numbers/indexes of the PCRs which can be subscribed.";
    }
  }
}
}
}
<CODE ENDS>
```

6. Event Streams for Conceptual Messages

Analogous to the [RFC8639] compliant <attestation> Event Stream for the conveyance of remote attestation Evidence as defined in Section Section 4, additional Event Streams can be defined for this YANG augment. Additional Event Streams require separate YANG augment specifications that provide the Event Stream definition and optionally a content format definition either via subscriptions to YANG datastores or dedicated YANG Notifications. It is possible to use either YANG subscription methods to other YANG modules for RATS Conceptual Messages or to define Event Streams for other none-YANG-modeled data. In the context of RATS Conceptual Messages, both options MUST be a specified via YANG augments to this specification.

7. Security Considerations

To be written.

8. IANA Considerations

To be written.

9. References

9.1. Normative References

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-12.txt>>.

[I-D.ietf-rats-reference-interaction-models]

Birkholz, H., Eckel, M., Pan, W., and E. Voit, "Reference Interaction Models for Remote Attestation Procedures",

Work in Progress, Internet-Draft, draft-ietf-rats-reference-interaction-models-04, 26 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-reference-interaction-models-04.txt>>.

[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-08, 26 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-tpm-based-network-device-attest-08.txt>>.

[I-D.ietf-rats-yang-tpm-charra]
Birkholz, H., Eckel, M., Bhandari, S., Voit, E., Sulzen, B., (Frank), L. X., Laffey, T., and G. C. Fedorkow, "A YANG Data Model for Challenge-Response-based Remote Attestation Procedures using TPMs", Work in Progress, Internet-Draft, draft-ietf-rats-yang-tpm-charra-10, 12 August 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-yang-tpm-charra-10.txt>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8639] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Subscription to YANG Notifications", RFC 8639, DOI 10.17487/RFC8639, September 2019, <<https://www.rfc-editor.org/info/rfc8639>>.

[TPM2.0] TCG, "TPM 2.0 Library Specification", n.d., <<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.

9.2. Informative References

[I-D.birkholz-rats-tuda]
Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-05, 12 July 2021, <<https://www.ietf.org/archive/id/draft-birkholz-rats-tuda-05.txt>>.

[KGV] TCG, "KGV", October 2003,
<https://trustedcomputinggroup.org/wp-content/uploads/TCG-NetEq-Attestation-Workflow-Outline_v1r9b_pubrev.pdf>.

Appendix A. Change Log

v01-v02

- * Match YANG changes/simplifications made to charra

v00-v01

- * rename notification: pcr-extended, which supports multiple PCRs
- * netequip boot added
- * YANG structure extension removed
- * Matched to structural changes made within charra

Acknowledgements

Thanks to ...

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Eric Voit
Cisco Systems, Inc.

Email: evoit@cisco.com

Wei Pan
Huawei Technologies
101 Software Avenue, Yuhuatai District
Nanjing, Jiangsu
210012
China

Email: william.panwei@huawei.com

RATS
Internet-Draft
Intended status: Standards Track
Expires: 16 July 2022

H. Birkholz
Fraunhofer SIT
B. Moran
Arm Limited
12 January 2022

Trustworthiness Vectors for the Software Updates of Internet of Things
(SUIT) Workflow Model
draft-birkholz-rats-suit-claims-03

Abstract

The IETF Remote Attestation Procedures (RATS) architecture defines Conceptual Messages as input and output of the appraisal process that assesses the trustworthiness of remote peers: Evidence and Attestation Results. Based on the Trustworthiness Vectors defined in Trusted Path Routing, this document defines a core set of Claims to be used in Evidence and Attestation Results for the Software Update for the Internet of Things (SUIT) Workflow Model. Consecutively, this document is in support of the Trusted Execution Environment Provisioning (TEEP) architecture, which defines the assessment of remote peers via RATS and uses SUIT for evidence generation as well as a remediation measure to improve trustworthiness of given remote peers.

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 16 July 2022.

Copyright Notice

Copyright (c) 2022 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
1.1. SUIT Workflow Model and Procedures	3
1.2. Terminology	4
2. Trustworthiness Vectors	4
3. SUIT Claims	5
3.1. System Properties Claims	5
3.1.1. vendor-identifier	6
3.1.2. class-identifier	6
3.1.3. device-identifier	6
3.1.4. image-digest	6
3.1.5. image-size	6
3.1.6. version	6
3.2. Interpreter Record Claims	7
3.2.1. record-success	7
3.2.2. component-index	7
3.2.3. dependency-index	7
3.2.4. command-index	7
3.2.5. nominal-parameters	7
3.3. Generic Record Conditions (TBD)	7
4. List of Commands (TBD)	8
5. References	9
5.1. Normative References	9
5.2. Informative References	9
Authors' Addresses	10

1. Introduction

Attestation Results are an essential output of Verifiers as defined in the Remote ATtestation procedureS (RATS) architecture [I-D.ietf-rats-architecture]. They are consumed by Relying Parties: the entities that intend to build future decisions on trustworthiness assessments of remote peers. Attestation Results must be easily appraised by Relying Parties -- in contrast to the rather complex or domain-specific Evidence appraised by Verifiers.

In order to create Attestation Results, a Verifier must consume Evidence generated by a given Attester (amongst other Conceptual Messages, such as Endorsements and Attestation Policies). Both Evidence and Attestation Results are composed of Claims. This document highlights and defines a set of Claims to be used in Evidence and Attestation Results that are based on the SUIT Workflow Model [I-D.ietf-suit-manifest]. In the scope of this document, an Attester takes on the role of a SUIT Recipient: the system that receives a SUIT Manifest.

1.1. SUIT Workflow Model and Procedures

This document focuses on Evidence and Attestation Results that can be generated based on the output of SUIT Procedures. The SUIT Workflow Model allows for two types of SUIT Procedures generating Reports on the Attester as defined in the SUIT Manifest specification [I-D.ietf-suit-manifest]:

Update Procedures: A procedure that updates a device by fetching dependencies, software images, and installing them.

An Update Procedure creates a Report about mutable software components that are installed or updated on hardware components.

Boot Procedures: A procedure that boots a device by checking dependencies and images, loading images, and invoking one or more image.

A Boot Procedure creates a Report on measured boot events (e.g. during Secure Boot).

The Records contained in each type of Report can be used as Claims in Evidence generation on the Attester for Remote Attestation Procedures as described in this document. Analogously, a corresponding Verifier appraising that Evidence can generate Attestation Results using the Claims defined in this document.

Both types of SUIT Procedures pass several stages (e.g. dependency-checking is one stage). The type and sequence of stages are defined by the Command Sequences included in a SUIT Manifest. For each stage in which a Command from the Command Sequence is executed a Record is created. All Records of a SUIT procedure contain binary results limited to "fail" or "pass". The aggregated sequence of all Records is composed into a Report.

This document specifies new Claims derived from Command Sequence Reports and relates them to Claims defined in Attestation Results for Secure Interactions [I-D.ietf-rats-ar4si] -- if applicable to the operational state of installed and updated software.

The Claims defined in this document are in support of the Trusted Execution Environment Provisioning (TEEP) architecture. During TEEP, the current operational state of an Attester is assessed via RATS. If the corresponding Attestation Results -- as covered in this document -- indicate insufficient Trustworthiness Tiers in a Trustworthiness Vector with respect to installed software, the SUIT Workflow Model is used for remediation.

1.2. Terminology

This document uses the terms and concepts defined in [I-D.ietf-rats-architecture], [I-D.ietf-suit-manifest], and [I-D.ietf-teep-architecture].

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. Trustworthiness Vectors

While there are usage scenarios where Attestation Results can be binary decisions, more often than not the assessment of trustworthiness is represented by a more fine-grained spectrum or based on multiple factors. These shades of Attestation Results are captured by the definition of Trustworthiness Vectors in Attestation Results for Secure Interaction [I-D.ietf-rats-ar4si]. Trustworthiness Vectors are sets of Trustworthiness Claims representing appraisal outputs produced by a Verifier (Attestation Results). Each of these Trustworthiness Claims has a Trustworthiness Tier ranging from Affirmed to None.

An Attester processing SUIT Manifests can manage three types of information about its Target Environments:

- * installed manifests including initial state (e.g. factory default),
- * hardware component identifiers that represent identifiable targets of updates, and

- * SUIT Interpreter results (e.g. test-failed) generated during updates.

Every SUIT Manifest maps to a certain intended state of a device. Every intended device composition of software components associated with hardware components can therefore be expressed based on a SUIT Manifest. The current operational state of a device can be represented in the same form, including the initial state.

As a result, the Claims defined in this document are bundled by the scope of the information represented in SUIT Manifests, i.e., dedicated blobs of software that are the payload of a SUIT Manifest. All Claims associated with an identifiable SUIT Manifest MUST always be bundled together in a Claims set that is limited to the Claims defined in this document.

3. SUIT Claims

The Claim description in this document uses CDDL as the formal modeling language for Claims. This approach is aligned with [I-D.ietf-rats-eat]. All Claims are based on information elements as used in the SUIT Manifest specification [I-D.ietf-suit-manifest]. For instance, a SUIT Class ID is represented as an UUID. Analogously, the corresponding class-identifier Claim found below is based on a UUID. SUIT Claims are differentiated in:

- * software and hardware characteristics (System Properties), and
- * reports about updates and their SUIT Commands (SUIT Records).
- * success/failure reports

Each type of Claims is always bundled in a dedicated Claim Set. Implementations can encode this information in various different ways (data models), e.g., sets, sequences, or nested structures.

The SUIT Report is defined in [I-D.ietf-suit-report]. It is used verbatim in this draft. The following subsections define the SUIT Report Claims for RATS.

3.1. System Properties Claims

System Properties Claims are composed of:

- * Hardware Component Claims and
- * Software Component Claims.

Correspondingly, the Claim definitions below highlight if a Claim is generic or hw/sw-component specific.

3.1.1. vendor-identifier

A RFC 4122 UUID representing the vendor of the Attester or one of its hardware and/or software components.

```
$$system-property-claim // = ( vendor-identifier =>
    (RFC4122_UUID / cbor-pen) )
cbor-pen = #6.112(bstr)
```

3.1.2. class-identifier

A RFC 4122 UUID representing the class of the Attester or one of its hardware and/or software components.

```
$$system-property-claim // = ( class-identifier => RFC4122_UUID )
```

3.1.3. device-identifier

A RFC 4122 UUID representing the Attester.

```
$$system-property-claim // = ( device-identifier => RFC4122_UUID )
```

3.1.4. image-digest

A fingerprint computed over a software component image on the Attester. This Claim is always bundled with a component-identifier or component-index.

```
$$system-property-claim // = ( image-digest => digest )
```

3.1.5. image-size

The size of a firmware image on the Attester.

```
$$system-property-claim // = ( image-size => size )
```

3.1.6. version

The Version of a hardware or software component of the Attester.

```
$$system-property-claim // = ( version => version-value )
```

3.2. Interpreter Record Claims

This class of Claims represents the content of SUIT Records generated by Interpreters running on Recipients. They are always bundled into Claim Sets representing SUIT Reports and are intended to be included in Evidence generated by an Attester. The Interpreter Record Claims appraised by a Verifier can steer a corresponding a Firmware Appraisal procedures that consumes this Evidence. Analogously, these Claims can be re-used in generated Attestation Results as Trustworthiness Vectors [I-D.ietf-rats-ar4si].

3.2.1. record-success

The result of a Command that was executed by the Interpreter on an Attester.

```
$$interpreter-record-claim //= ( record-success => bool )
```

3.2.2. component-index

A positive integer representing an entry in a flat list of indices mapped to software component identifiers to be updated.

```
$$system-property-claim //= ( component-index => uint )
```

3.2.3. dependency-index

A thumbprint of a software component that an update depends on.

```
$$interpreter-record-claim //= ( dependency-index => digest )
```

3.2.4. command-index

A positive integer representing an entry in a SUIT_Command_Sequence identifying a Command encoded as a SUIT Manifest Directive or SUIT Manifest Condition.

```
$$interpreter-record-claim //= ( command-index => uint )
```

3.2.5. nominal-parameters

A list of SUIT_Parameters associated with a specific Command that was executed by the Interpreter on an Attester.

```
$$interpreter-record-claim //= ( actual-parameters => parameter-list )
```

3.3. Generic Record Conditions (TBD)

- * test-failed
- * unsupported-command
- * unsupported-parameter
- * unsupported-component-id
- * payload-unavailable
- * dependency-unavailable
- * critical-application-failure
- * watchdog-timeout

4. List of Commands (TBD)

- * Check Vendor Identifier
- * Check Class Identifier
- * Verify Image
- * Set Component Index
- * Override Parameters
- * Set Dependency Index
- * Set Parameters
- * Process Dependency
- * Run
- * Fetch
- * Use Before
- * Check Component Offset
- * Check Device Identifier
- * Check Image Not Match
- * Check Minimum Battery

- * Check Update Authorized
- * Check Version
- * Abort
- * Try Each
- * Copy
- * Swap
- * Wait For Event
- * Run Sequence
- * Run with Arguments

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

5.2. Informative References

- [I-D.ietf-rats-ar4si] Voit, E., Birkholz, H., Hardjono, T., Fossati, T., and V. Scarlata, "Attestation Results for Secure Interactions", Work in Progress, Internet-Draft, draft-ietf-rats-ar4si-01, 2 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-ar4si-01.txt>>.
- [I-D.ietf-rats-architecture] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-14, 9 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-14.txt>>.

[I-D.ietf-rats-eat]

Lundblade, L., Mandyam, G., and J. O'Donoghue, "The Entity Attestation Token (EAT)", Work in Progress, Internet-Draft, draft-ietf-rats-eat-11, 24 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-eat-11.txt>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-19, 20 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-sacm-coswid-19.txt>>.

[I-D.ietf-suit-manifest]

Moran, B., Tschofenig, H., Birkholz, H., and K. Zandberg, "A Concise Binary Object Representation (CBOR)-based Serialization Format for the Software Updates for Internet of Things (SUIT) Manifest", Work in Progress, Internet-Draft, draft-ietf-suit-manifest-16, 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-suit-manifest-16.txt>>.

[I-D.ietf-suit-report]

Moran, B., "Secure Reporting of Update Status", Work in Progress, Internet-Draft, draft-ietf-suit-report-00, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-suit-report-00.txt>>.

[I-D.ietf-teep-architecture]

Pei, M., Tschofenig, H., Thaler, D., and D. Wheeler, "Trusted Execution Environment Provisioning (TEEP) Architecture", Work in Progress, Internet-Draft, draft-ietf-teep-architecture-15, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-teep-architecture-15.txt>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT

Email: henk.birkholz@sit.fraunhofer.de

Brendan Moran
Arm Limited

Internet-Draft

SUIT TV

January 2022

Email: Brendan.Moran@arm.com

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 25 October 2021

H. Birkholz
Fraunhofer SIT
D. Thaler
Microsoft
M. Richardson
Sandelman Software Works
N. Smith
Intel
W. Pan
Huawei Technologies
23 April 2021

Remote Attestation Procedures Architecture
draft-ietf-rats-architecture-12

Abstract

In network protocol exchanges it is often useful for one end of a communication to know whether the other end is in an intended operating state. This document provides an architectural overview of the entities involved that make such tests possible through the process of generating, conveying, and evaluating evidentiary claims. An attempt is made to provide for a model that is neutral toward processor architectures, the content of claims, and protocols.

Note to Readers

Discussion of this document takes place on the RATS Working Group mailing list (rats@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/rats/> (<https://mailarchive.ietf.org/arch/browse/rats/>).

Source for this draft and an issue tracker can be found at <https://github.com/ietf-rats-wg/architecture> (<https://github.com/ietf-rats-wg/architecture>).

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 25 October 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Reference Use Cases	5
2.1. Network Endpoint Assessment	5
2.2. Confidential Machine Learning Model Protection	5
2.3. Confidential Data Protection	6
2.4. Critical Infrastructure Control	6
2.5. Trusted Execution Environment Provisioning	7
2.6. Hardware Watchdog	7
2.7. FIDO Biometric Authentication	7
3. Architectural Overview	8
3.1. Layered Attestation Environments	12
3.2. Composite Device	14
3.3. Implementation Considerations	16
4. Terminology	17
4.1. Roles	17
4.2. Artifacts	18
5. Topological Patterns	19
5.1. Passport Model	20
5.2. Background-Check Model	21
5.3. Combinations	22
6. Roles and Entities	23
7. Trust Model	24
7.1. Relying Party	24
7.2. Attester	25
7.3. Relying Party Owner	25

7.4. Verifier	25
7.5. Endorser, Reference Value Provider, and Verifier Owner	27
8. Conceptual Messages	28
8.1. Evidence	28
8.2. Endorsements	28
8.3. Reference Values	29
8.4. Attestation Results	29
8.5. Appraisal Policies	30
9. Claims Encoding Formats	31
10. Freshness	32
10.1. Explicit Timekeeping using Synchronized Clocks	33
10.2. Implicit Timekeeping using Nonces	33
10.3. Implicit Timekeeping using Epoch IDs	33
10.4. Discussion	35
11. Privacy Considerations	35
12. Security Considerations	36
12.1. Attester and Attestation Key Protection	36
12.1.1. On-Device Attester and Key Protection	37
12.1.2. Attestation Key Provisioning Processes	37
12.2. Integrity Protection	39
12.3. Epoch ID-based Attestation	39
12.4. Trust Anchor Protection	40
13. IANA Considerations	40
14. Acknowledgments	41
15. Notable Contributions	41
16. Appendix A: Time Considerations	41
16.1. Example 1: Timestamp-based Passport Model Example	43
16.2. Example 2: Nonce-based Passport Model Example	44
16.3. Example 3: Epoch ID-based Passport Model Example	46
16.4. Example 4: Timestamp-based Background-Check Model Example	47
16.5. Example 5: Nonce-based Background-Check Model Example	48
17. References	49
17.1. Normative References	49
17.2. Informative References	49
Contributors	51
Authors' Addresses	53

1. Introduction

The question of how one system can know that another system can be trusted has found new interest and relevance in a world where trusted computing elements are maturing in processor architectures.

Systems that have been attested and verified to be in a good state (for some value of "good") can improve overall system posture. Conversely, systems that cannot be attested and verified to be in a good state can be taken out of service, or otherwise flagged for repair.

For example:

- * A bank back-end system might refuse to transact with another system that is not known to be in a good state.
- * A healthcare system might refuse to transmit electronic healthcare records to a system that is not known to be in a good state.

In Remote Attestation Procedures (RATS), one peer (the "Attester") produces believable information about itself - Evidence - to enable a remote peer (the "Relying Party") to decide whether to consider that Attester a trustworthy peer or not. RATS are facilitated by an additional vital party, the Verifier.

The Verifier appraises Evidence via appraisal policies and creates the Attestation Results to support Relying Parties in their decision process. This document defines a flexible architecture consisting of attestation roles and their interactions via conceptual messages. Additionally, this document defines a universal set of terms that can be mapped to various existing and emerging Remote Attestation Procedures. Common topological patterns and the sequence of data flows associated with them, such as the "Passport Model" and the "Background-Check Model", are illustrated. The purpose is to define useful terminology for remote attestation and enable readers to map their solution architecture to the canonical attestation architecture provided here. Having a common terminology that provides well-understood meanings for common themes such as roles, device composition, topological patterns, and appraisal procedures is vital for semantic interoperability across solutions and platforms involving multiple vendors and providers.

Amongst other things, this document is about trust and trustworthiness. Trust is a choice one makes about another system. Trustworthiness is a quality about the other system that can be used in making one's decision to trust it or not. This is subtle difference and being familiar with the difference is crucial for using this document. Additionally, the concepts of freshness and trust relationships with respect to RATS are elaborated on to enable implementers to choose appropriate solutions to compose their Remote Attestation Procedures.

2. Reference Use Cases

This section covers a number of representative and generic use cases for remote attestation, independent of specific solutions. The purpose is to provide motivation for various aspects of the architecture presented in this document. Many other use cases exist, and this document does not intend to have a complete list, only to illustrate a set of use cases that collectively cover all the functionality required in the architecture.

Each use case includes a description followed by an additional summary of the Attester and Relying Party roles derived from the use case.

2.1. Network Endpoint Assessment

Network operators want a trustworthy report that includes identity and version information about the hardware and software on the machines attached to their network, for purposes such as inventory, audit, anomaly detection, record maintenance and/or trending reports (logging). The network operator may also want a policy by which full access is only granted to devices that meet some definition of hygiene, and so wants to get Claims about such information and verify its validity. Remote attestation is desired to prevent vulnerable or compromised devices from getting access to the network and potentially harming others.

Typically, solutions start with a specific component (called a root of trust) that is intended to provide trustworthy device identity and protected storage for measurements. The system components perform a series of measurements that may be signed via functions provided by a root of trust, considered as Evidence about present system components, such as hardware, firmware, BIOS, software, etc.

Attester: A device desiring access to a network.

Relying Party: Network equipment such as a router, switch, or access point, responsible for admission of the device into the network.

2.2. Confidential Machine Learning Model Protection

A device manufacturer wants to protect its intellectual property. The intellectual property's scope primarily encompasses the machine learning (ML) model that is deployed in the devices purchased by its customers. The protection goals include preventing attackers, potentially the customer themselves, from seeing the details of the model.

This typically works by having some protected environment in the device go through a remote attestation with some manufacturer service that can assess its trustworthiness. If remote attestation succeeds, then the manufacturer service releases either the model, or a key to decrypt a model already deployed on the Attester in encrypted form, to the requester.

Attester: A device desiring to run an ML model.

Relying Party: A server or service holding ML models it desires to protect.

2.3. Confidential Data Protection

This is a generalization of the ML model use case above, where the data can be any highly confidential data, such as health data about customers, payroll data about employees, future business plans, etc. As part of the attestation procedure, an assessment is made against a set of policies to evaluate the state of the system that is requesting the confidential data. Attestation is desired to prevent leaking data via compromised devices.

Attester: An entity desiring to retrieve confidential data.

Relying Party: An entity that holds confidential data for release to authorized entities.

2.4. Critical Infrastructure Control

Potentially harmful physical equipment (e.g., power grid, traffic control, hazardous chemical processing, etc.) is connected to a network in support of critical infrastructure. The organization managing such infrastructure needs to ensure that only authorized code and users can control corresponding critical processes, and that these processes are protected from unauthorized manipulation or other threats. When a protocol operation can affect a critical system component of the infrastructure, devices attached to that critical component require some assurances depending on the security context, including that: a requesting device or application has not been compromised, and the requesters and actors act on applicable policies. As such, remote attestation can be used to only accept commands from requesters that are within policy.

Attester: A device or application wishing to control physical equipment.

Relying Party: A device or application connected to potentially

dangerous physical equipment (hazardous chemical processing, traffic control, power grid, etc.).

2.5. Trusted Execution Environment Provisioning

A Trusted Application Manager (TAM) server is responsible for managing the applications running in a Trusted Execution Environment (TEE) of a client device, as described in [I-D.ietf-tee-architecture]. To achieve its purpose, the TAM needs to assess the state of a TEE, or of applications in the TEE, of a client device. The TEE conducts Remote Attestation Procedures with the TAM, which can then decide whether the TEE is already in compliance with the TAM's latest policy. If not, the TAM has to uninstall, update, or install approved applications in the TEE to bring it back into compliance with the TAM's policy.

Attester: A device with a TEE capable of running trusted applications that can be updated.

Relying Party: A TAM.

2.6. Hardware Watchdog

There is a class of malware that holds a device hostage and does not allow it to reboot to prevent updates from being applied. This can be a significant problem, because it allows a fleet of devices to be held hostage for ransom.

A solution to this problem is a watchdog timer implemented in a protected environment such as a Trusted Platform Module (TPM), as described in [TCGarch] section 43.3. If the watchdog does not receive regular, and fresh, Attestation Results as to the system's health, then it forces a reboot.

Attester: The device that should be protected from being held hostage for a long period of time.

Relying Party: A watchdog capable of triggering a procedure that resets a device into a known, good operational state.

2.7. FIDO Biometric Authentication

In the Fast IDentity Online (FIDO) protocol [WebAuthN], [CTAP], the device in the user's hand authenticates the human user, whether by biometrics (such as fingerprints), or by PIN and password. FIDO authentication puts a large amount of trust in the device compared to typical password authentication because it is the device that verifies the biometric, PIN and password inputs from the user, not

the server. For the Relying Party to know that the authentication is trustworthy, the Relying Party needs to know that the Authenticator part of the device is trustworthy. The FIDO protocol employs remote attestation for this.

The FIDO protocol supports several remote attestation protocols and a mechanism by which new ones can be registered and added. Remote attestation defined by RATS is thus a candidate for use in the FIDO protocol.

Other biometric authentication protocols such as the Chinese IFAA standard and WeChat Pay as well as Google Pay make use of remote attestation in one form or another.

Attester: Every FIDO Authenticator contains an Attester.

Relying Party: Any web site, mobile application back-end, or service that relies on authentication data based on biometric information.

3. Architectural Overview

Figure 1 depicts the data that flows between different roles, independent of protocol or use case.

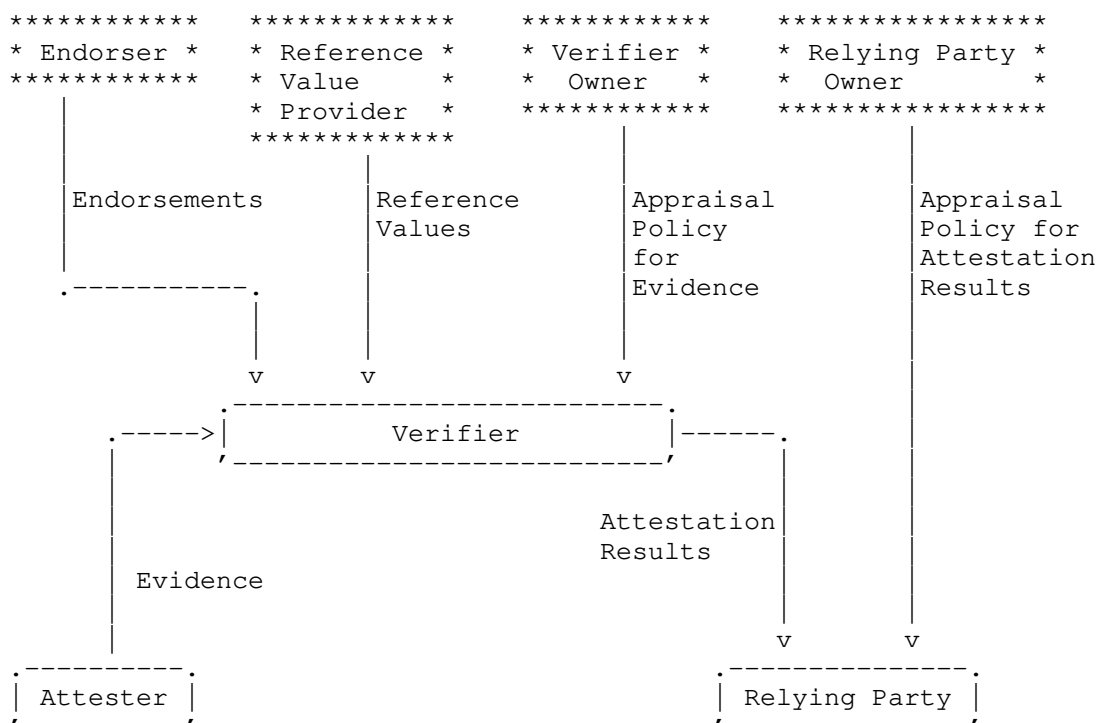


Figure 1: Conceptual Data Flow

The text below summarizes the activities conducted by the roles illustrated in Figure 1.

An Attester creates Evidence that is conveyed to a Verifier.

A Verifier uses the Evidence, any Reference Values from Reference Value Providers, and any Endorsements from Endorsers, by applying an Appraisal Policy for Evidence to assess the trustworthiness of the Attester. This procedure is called the appraisal of Evidence.

Subsequently, the Verifier generates Attestation Results for use by Relying Parties.

The Appraisal Policy for Evidence might be obtained from the Verifier Owner via some protocol mechanism, or might be configured into the Verifier by the Verifier Owner, or might be programmed into the Verifier, or might be obtained via some other mechanism.

A Relying Party uses Attestation Results by applying its own appraisal policy to make application-specific decisions, such as authorization decisions. This procedure is called the appraisal of Attestation Results.

The Appraisal Policy for Attestation Results might be obtained from the Relying Party Owner via some protocol mechanism, or might be configured into the Relying Party by the Relying Party Owner, or might be programmed into the Relying Party, or might be obtained via some other mechanism.

See Section 8 for further discussion of the conceptual messages shown in Figure 1. ## Two Types of Environments of an Attester

As shown in Figure 2, an Attester consists of at least one Attesting Environment and at least one Target Environment. In some implementations, the Attesting and Target Environments might be combined. Other implementations might have multiple Attesting and Target Environments, such as in the examples described in more detail in Section 3.1 and Section 3.2. Other examples may exist. All compositions of Attesting and Target Environments discussed in this architecture can be combined into more complex implementations.

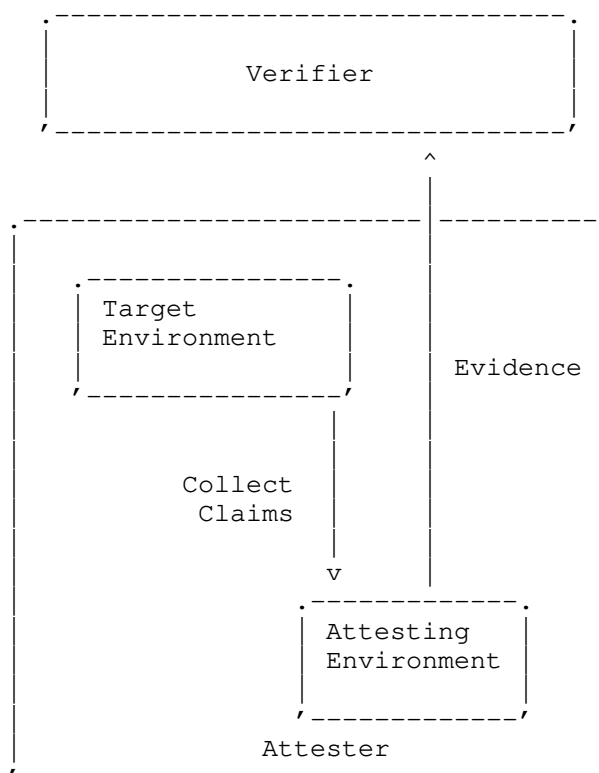


Figure 2: Two Types of Environments

Claims are collected from Target Environments. That is, Attesting Environments collect the values and the information to be represented in Claims, by reading system registers and variables, calling into subsystems, taking measurements on code, memory, or other security related assets of the Target Environment. Attesting Environments then format the Claims appropriately, and typically use key material and cryptographic functions, such as signing or cipher algorithms, to generate Evidence. There is no limit to or requirement on the types of hardware or software environments that can be used to implement an Attesting Environment, for example: Trusted Execution Environments (TEEs), embedded Secure Elements (eSEs), Trusted Platform Modules (TPMs) [TCGarch], or BIOS firmware.

An arbitrary execution environment may not, by default, be capable of Claims collection for a given Target Environment. Execution environments that are designed specifically to be capable of Claims collection are referred to in this document as Attesting Environments. For example, a TPM doesn't actively collect Claims

itself, it instead requires another component to feed various values to the TPM. Thus, an Attesting Environment in such a case would be the combination of the TPM together with whatever component is feeding it the measurements.

3.1. Layered Attestation Environments

By definition, the Attester role generates Evidence. An Attester may consist of one or more nested environments (layers). The root layer of an Attester includes at least one root of trust. In order to appraise Evidence generated by an Attester, the Verifier needs to trust the Attester's root of trust. Trust in the Attester's root of trust can be established in various ways as discussed in Section 7.4.

In layered attestation, a root of trust is the initial Attesting Environment. Claims can be collected from or about each layer. The corresponding Claims can be structured in a nested fashion that reflects the nesting of the Attester's layers. Normally, Claims are not self-asserted, rather a previous layer acts as the Attesting Environment for the next layer. Claims about a root of trust typically are asserted by an Endorser.

The example device illustrated in Figure 3 includes (A) a BIOS stored in read-only memory, (B) a bootloader, and (C) an operating system kernel.

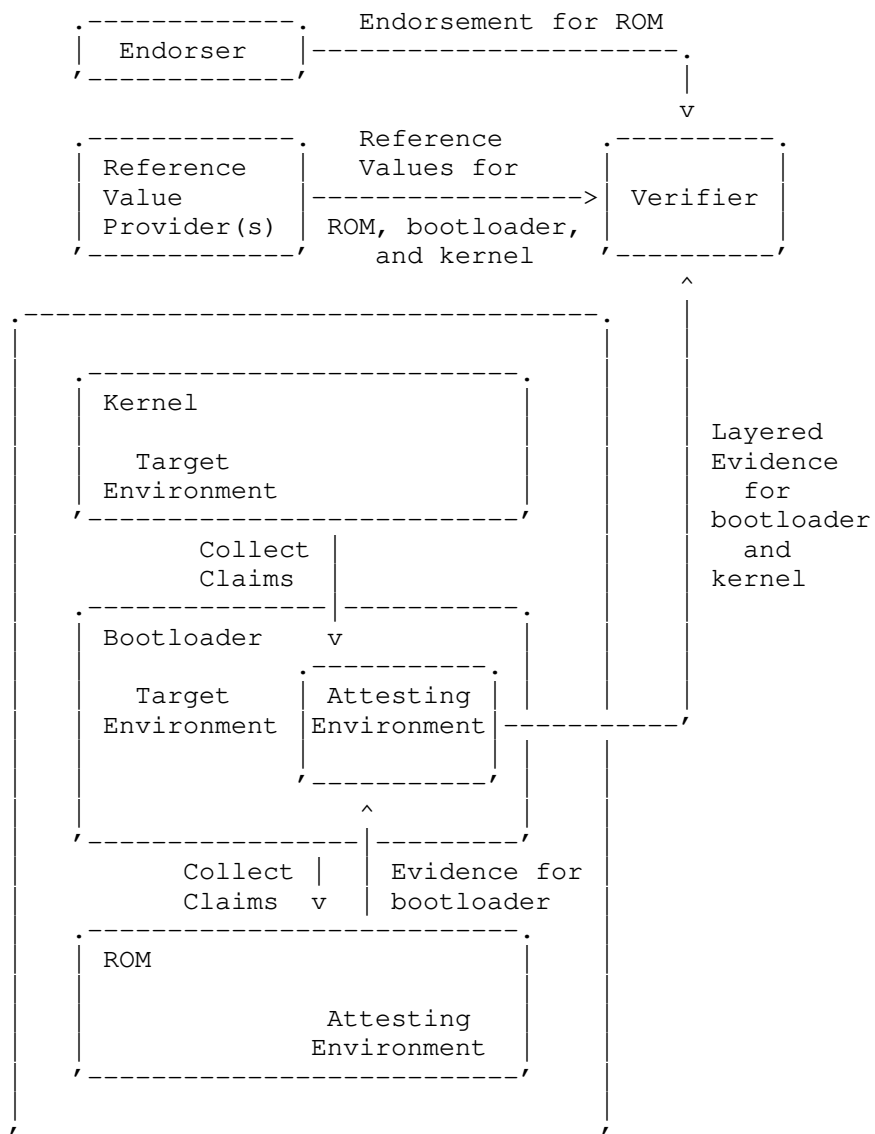


Figure 3: Layered Attester

The first Attesting Environment, the read-only BIOS in this example, has to ensure the integrity of the bootloader (the first Target Environment). There are potentially multiple kernels to boot, and the decision is up to the bootloader. Only a bootloader with intact integrity will make an appropriate decision. Therefore, the Claims relating to the integrity of the bootloader have to be measured securely. At this stage of the boot-cycle of the device, the Claims collected typically cannot be composed into Evidence.

After the boot sequence is started, the BIOS conducts the most important and defining feature of layered attestation, which is that the successfully measured bootloader now becomes (or contains) an Attesting Environment for the next layer. This procedure in layered attestation is sometimes called "staging". It is important that the bootloader not be able to alter any Claims about itself that were collected by the BIOS. This can be ensured having those Claims be either signed by the BIOS or stored in a tamper-proof manner by the BIOS.

Continuing with this example, the bootloader's Attesting Environment is now in charge of collecting Claims about the next Target Environment, which in this example is the kernel to be booted. The final Evidence thus contains two sets of Claims: one set about the bootloader as measured and signed by the BIOS, plus a set of Claims about the kernel as measured and signed by the bootloader.

This example could be extended further by making the kernel become another Attesting Environment for an application as another Target Environment. This would result in a third set of Claims in the Evidence pertaining to that application.

The essence of this example is a cascade of staged environments. Each environment has the responsibility of measuring the next environment before the next environment is started. In general, the number of layers may vary by device or implementation, and an Attesting Environment might even have multiple Target Environments that it measures, rather than only one as shown by example in Figure 3.

3.2. Composite Device

A composite device is an entity composed of multiple sub-entities such that its trustworthiness has to be determined by the appraisal of all these sub-entities.

Each sub-entity has at least one Attesting Environment collecting the Claims from at least one Target Environment, then this sub-entity generates Evidence about its trustworthiness. Therefore, each sub-

entity can be called an Attester. Among all the Attesters, there may be only some which have the ability to communicate with the Verifier while others do not.

For example, a carrier-grade router consists of a chassis and multiple slots. The trustworthiness of the router depends on all its slots' trustworthiness. Each slot has an Attesting Environment, such as a TEE, collecting the Claims of its boot process, after which it generates Evidence from the Claims.

Among these slots, only a "main" slot can communicate with the Verifier while other slots cannot. But other slots can communicate with the main slot by the links between them inside the router. So the main slot collects the Evidence of other slots, produces the final Evidence of the whole router and conveys the final Evidence to the Verifier. Therefore the router is a composite device, each slot is an Attester, and the main slot is the lead Attester.

Another example is a multi-chassis router composed of multiple single carrier-grade routers. Multi-chassis router setups create redundancy groups that provide higher throughput by interconnecting multiple routers in these groups, which can be treated as one logical router for simpler management. A multi-chassis router setup provides a management point that connects to the Verifier. Typically one router in the group is designated as the main router. Other routers in the multi-chassis setup are connected to the main router only via physical network links and are therefore managed and appraised via the main router's help. In consequence, a multi-chassis router setup is a composite device, each router is an Attester, and the main router is the lead Attester.

Figure 4 depicts the conceptual data flow for a composite device.

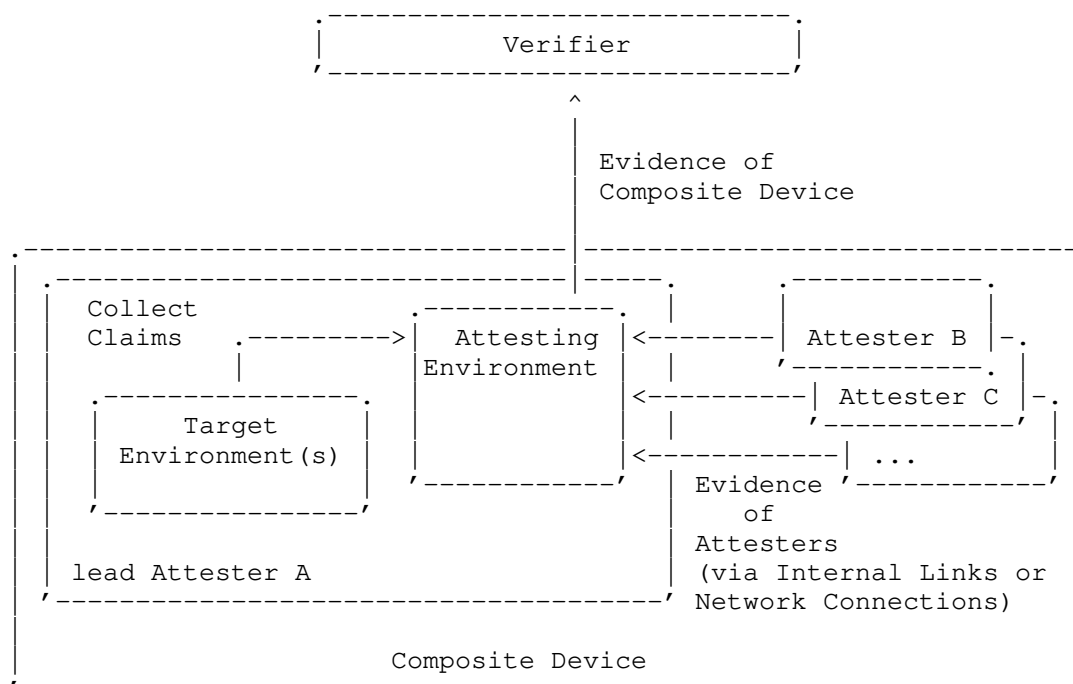


Figure 4: Composite Device

In a composite device, each Attester generates its own Evidence by its Attesting Environment(s) collecting the Claims from its Target Environment(s). The lead Attester collects Evidence from other Attesters and conveys it to a Verifier. Collection of Evidence from sub-entities may itself be a form of Claims collection that results in Evidence asserted by the lead Attester. The lead Attester generates Evidence about the layout of the whole composite device, while sub-Attesters generate Evidence about their respective (sub-)modules.

In this scenario, the trust model described in Section 7 can also be applied to an inside Verifier.

3.3. Implementation Considerations

An entity can take on multiple RATS roles (e.g., Attester, Verifier, Relying Party, etc.) at the same time. Multiple entities can cooperate to implement a single RATS role as well. In essence, the combination of roles and entities can be arbitrary. For example, in the composite device scenario, the entity inside the lead Attester can also take on the role of a Verifier, and the outer entity of

Verifier can take on the role of a Relying Party. After collecting the Evidence of other Attesters, this inside Verifier uses Endorsements and appraisal policies (obtained the same way as by any other Verifier) as part of the appraisal procedures that generate Attestation Results. The inside Verifier then conveys the Attestation Results of other Attesters to the outside Verifier, whether in the same conveyance protocol as part of the Evidence or not.

4. Terminology

This document uses the following terms.

4.1. Roles

Attester: A role performed by an entity (typically a device) whose Evidence must be appraised in order to infer the extent to which the Attester is considered trustworthy, such as when deciding whether it is authorized to perform some operation.

Produces: Evidence

Relying Party: A role performed by an entity that depends on the validity of information about an Attester, for purposes of reliably applying application specific actions. Compare /relying party/ in [RFC4949].

Consumes: Attestation Results

Verifier: A role performed by an entity that appraises the validity of Evidence about an Attester and produces Attestation Results to be used by a Relying Party.

Consumes: Evidence, Reference Values, Endorsements, Appraisal Policy for Evidence

Produces: Attestation Results

Relying Party Owner: A role performed by an entity (typically an administrator), that is authorized to configure Appraisal Policy for Attestation Results in a Relying Party.

Produces: Appraisal Policy for Attestation Results

Verifier Owner: A role performed by an entity (typically an administrator), that is authorized to configure Appraisal Policy for Evidence in a Verifier.

Produces: Appraisal Policy for Evidence

Endorser: A role performed by an entity (typically a manufacturer) whose Endorsements help Verifiers appraise the authenticity of Evidence.

Produces: Endorsements

Reference Value Provider: A role performed by an entity (typically a manufacturer) whose Reference Values help Verifiers appraise Evidence to determine if acceptable known Claims have been recorded by the Attester.

Produces: Reference Values

4.2. Artifacts

Claim: A piece of asserted information, often in the form of a name/value pair. Claims make up the usual structure of Evidence and other RATS artifacts. Compare /claim/ in [RFC7519].

Endorsement: A secure statement that an Endorser vouches for the integrity of an Attester's various capabilities such as Claims collection and Evidence signing.

Consumed By: Verifier

Produced By: Endorser

Evidence: A set of Claims generated by an Attester to be appraised by a Verifier. Evidence may include configuration data, measurements, telemetry, or inferences.

Consumed By: Verifier

Produced By: Attester

Attestation Result: The output generated by a Verifier, typically including information about an Attester, where the Verifier vouches for the validity of the results.

Consumed By: Relying Party

Produced By: Verifier

Appraisal Policy for Evidence: A set of rules that informs how a Verifier evaluates the validity of information about an Attester. Compare /security policy/ in [RFC4949].

Consumed By: Verifier

Produced By: Verifier Owner

Appraisal Policy for Attestation Results: A set of rules that direct how a Relying Party uses the Attestation Results regarding an Attester generated by the Verifiers. Compare /security policy/ in [RFC4949].

Consumed by: Relying Party

Produced by: Relying Party Owner

Reference Values: A set of values against which values of Claims can be compared as part of applying an Appraisal Policy for Evidence. Reference Values are sometimes referred to in other documents as known-good values, golden measurements, or nominal values, although those terms typically assume comparison for equality, whereas here Reference Values might be more general and be used in any sort of comparison.

Consumed By: Verifier

Produced By: Reference Value Provider

5. Topological Patterns

Figure 1 shows a data-flow diagram for communication between an Attester, a Verifier, and a Relying Party. The Attester conveys its Evidence to the Verifier for appraisal, and the Relying Party receives the Attestation Result from the Verifier. This section refines the data-flow diagram by describing two reference models, as well as one example composition thereof. The discussion that follows is for illustrative purposes only and does not constrain the interactions between RATS roles to the presented patterns.

5.1. Passport Model

The passport model is so named because of its resemblance to how nations issue passports to their citizens. The nature of the Evidence that an individual needs to provide to its local authority is specific to the country involved. The citizen retains control of the resulting passport document and presents it to other entities when it needs to assert a citizenship or identity Claim, such as an airport immigration desk. The passport is considered sufficient because it vouches for the citizenship and identity Claims, and it is issued by a trusted authority. Thus, in this immigration desk analogy, the passport issuing agency is a Verifier, the passport is an Attestation Result, and the immigration desk is a Relying Party.

In this model, an Attester conveys Evidence to a Verifier, which compares the Evidence against its appraisal policy. The Verifier then gives back an Attestation Result. If the Attestation Result was a successful one, the Attester can then present the Attestation Result (and possibly additional Claims) to a Relying Party, which then compares this information against its own appraisal policy.

Three ways in which the process may fail include:

- * First, the Verifier may not issue a positive Attestation Result due to the Evidence not passing the Appraisal Policy for Evidence.
- * The second way in which the process may fail is when the Attestation Result is examined by the Relying Party, and based upon the Appraisal Policy for Attestation Results, the result does not pass the policy.
- * The third way is when the Verifier is unreachable or unavailable.

Since the resource access protocol between the Attester and Relying Party includes an Attestation Result, in this model the details of that protocol constrain the serialization format of the Attestation Result. The format of the Evidence on the other hand is only constrained by the Attester-Verifier remote attestation protocol. This implies that interoperability and standardization is more relevant for Attestation Results than it is for Evidence.

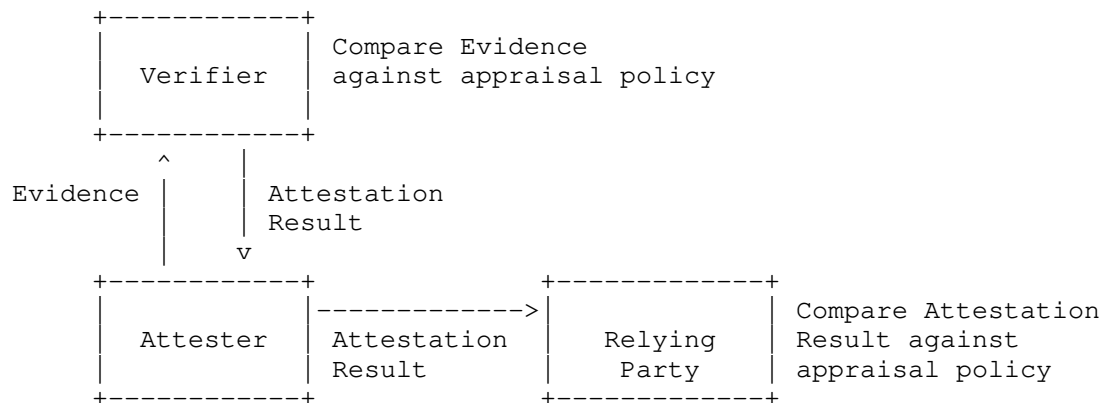


Figure 5: Passport Model

5.2. Background-Check Model

The background-check model is so named because of the resemblance of how employers and volunteer organizations perform background checks. When a prospective employee provides Claims about education or previous experience, the employer will contact the respective institutions or former employers to validate the Claim. Volunteer organizations often perform police background checks on volunteers in order to determine the volunteer's trustworthiness. Thus, in this analogy, a prospective volunteer is an Attester, the organization is the Relying Party, and the organization that issues a report is a Verifier.

In this model, an Attester conveys Evidence to a Relying Party, which simply passes it on to a Verifier. The Verifier then compares the Evidence against its appraisal policy, and returns an Attestation Result to the Relying Party. The Relying Party then compares the Attestation Result against its own appraisal policy.

The resource access protocol between the Attester and Relying Party includes Evidence rather than an Attestation Result, but that Evidence is not processed by the Relying Party. Since the Evidence is merely forwarded on to a trusted Verifier, any serialization format can be used for Evidence because the Relying Party does not need a parser for it. The only requirement is that the Evidence can be encapsulated in the format required by the resource access protocol between the Attester and Relying Party.

However, like in the Passport model, an Attestation Result is still consumed by the Relying Party. Code footprint and attack surface area can be minimized by using a serialization format for which the

Relying Party already needs a parser to support the protocol between the Attester and Relying Party, which may be an existing standard or widely deployed resource access protocol. Such minimization is especially important if the Relying Party is a constrained node.

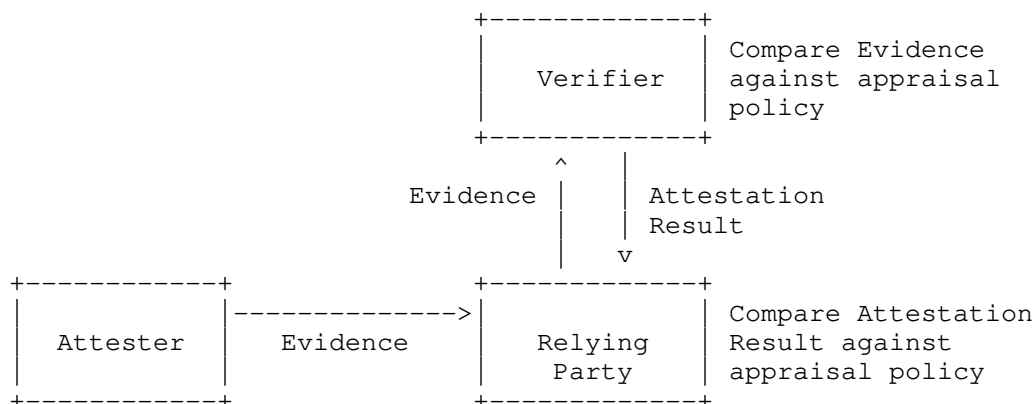


Figure 6: Background-Check Model

5.3. Combinations

One variation of the background-check model is where the Relying Party and the Verifier are on the same machine, performing both functions together. In this case, there is no need for a protocol between the two.

It is also worth pointing out that the choice of model depends on the use case, and that different Relying Parties may use different topological patterns.

The same device may need to create Evidence for different Relying Parties and/or different use cases. For instance, it would use one model to provide Evidence to a network infrastructure device to gain access to the network, and the other model to provide Evidence to a server holding confidential data to gain access to that data. As such, both models may simultaneously be in use by the same device.

Figure 7 shows another example of a combination where Relying Party 1 uses the passport model, whereas Relying Party 2 uses an extension of the background-check model. Specifically, in addition to the basic functionality shown in Figure 6, Relying Party 2 actually provides the Attestation Result back to the Attester, allowing the Attester to use it with other Relying Parties. This is the model that the Trusted Application Manager plans to support in the TEEP architecture [I-D.ietf-teep-architecture].

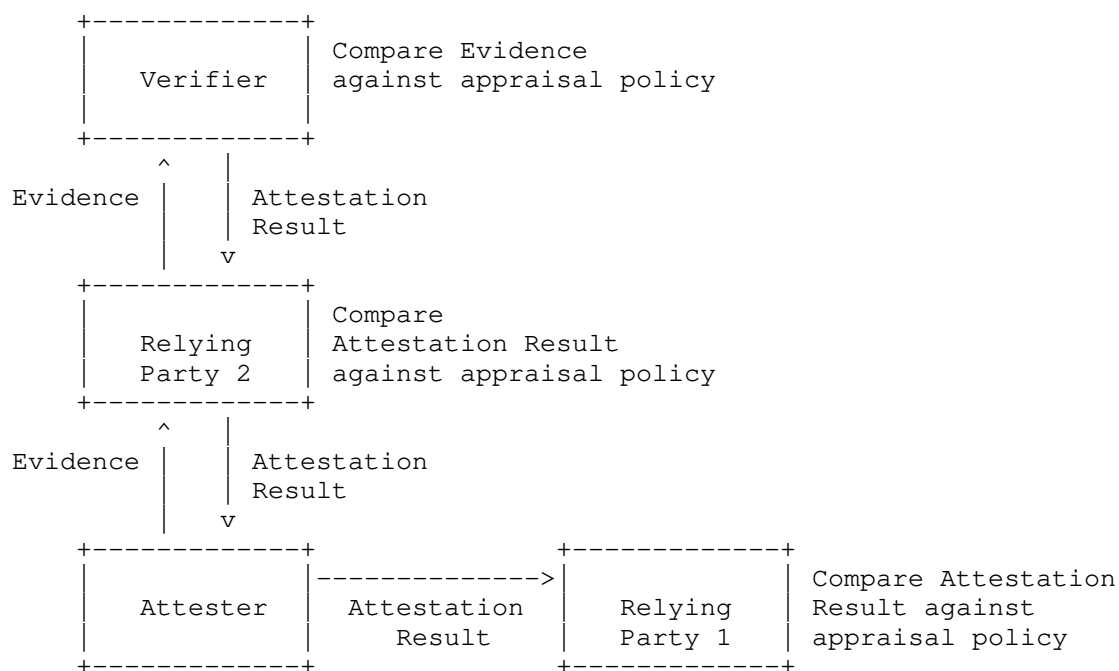


Figure 7: Example Combination

6. Roles and Entities

An entity in the RATS architecture includes at least one of the roles defined in this document.

An entity can aggregate more than one role into itself, such as being both a Verifier and a Relying Party, or being both a Reference Value Provider and an Endorser. As such, any conceptual messages (see Section 8 for more discussion) originating from such roles might also be combined. For example, Reference Values might be conveyed as part of an appraisal policy if the Verifier Owner and Reference Value Provider roles are combined. Similarly, Reference Values might be conveyed as part of an Endorsement if the Endorser and Reference Value Provider roles are combined.

Interactions between roles aggregated into the same entity do not necessarily use the Internet Protocol. Such interactions might use a loopback device or other IP-based communication between separate environments, but they do not have to. Alternative channels to convey conceptual messages include function calls, sockets, GPIO interfaces, local busses, or hypervisor calls. This type of conveyance is typically found in composite devices. Most

importantly, these conveyance methods are out-of-scope of RATS, but they are presumed to exist in order to convey conceptual messages appropriately between roles.

In essence, an entity that combines more than one role creates and consumes the corresponding conceptual messages as defined in this document.

7. Trust Model

7.1. Relying Party

This document covers scenarios for which a Relying Party trusts a Verifier that can appraise the trustworthiness of information about an Attester. Such trust might come by the Relying Party trusting the Verifier (or its public key) directly, or might come by trusting an entity (e.g., a Certificate Authority) that is in the Verifier's certificate path. Such trust is expressed by storing one or more "trust anchors" in a secure location known as a trust anchor store.

As defined in [RFC6024], "A trust anchor represents an authoritative entity via a public key and associated data. The public key is used to verify digital signatures, and the associated data is used to constrain the types of information for which the trust anchor is authoritative." The trust anchor may be a certificate or it may be a raw public key along with additional data if necessary such as its public key algorithm and parameters.

The Relying Party might implicitly trust a Verifier, such as in a Verifier/Relying Party combination where the Verifier and Relying Party roles are combined. Or, for a stronger level of security, the Relying Party might require that the Verifier first provide information about itself that the Relying Party can use to assess the trustworthiness of the Verifier before accepting its Attestation Results.

For example, one explicit way for a Relying Party "A" to establish such trust in a Verifier "B", would be for B to first act as an Attester where A acts as a combined Verifier/Relying Party. If A then accepts B as trustworthy, it can choose to accept B as a Verifier for other Attesters.

As another example, the Relying Party can establish trust in the Verifier by out of band establishment of key material, combined with a protocol like TLS to communicate. There is an assumption that between the establishment of the trusted key material and the creation of the Evidence, that the Verifier has not been compromised.

Similarly, the Relying Party also needs to trust the Relying Party Owner for providing its Appraisal Policy for Attestation Results, and in some scenarios the Relying Party might even require that the Relying Party Owner go through a remote attestation procedure with it before the Relying Party will accept an updated policy. This can be done similarly to how a Relying Party could establish trust in a Verifier as discussed above.

7.2. Attester

In some scenarios, Evidence might contain sensitive information such as Personally Identifiable Information (PII) or system identifiable information. Thus, an Attester must trust entities to which it conveys Evidence, to not reveal sensitive data to unauthorized parties. The Verifier might share this information with other authorized parties, according to a governing policy that address the handling of sensitive information (potentially included in Appraisal Policies for Evidence). In the background-check model, this Evidence may also be revealed to Relying Party(s).

When Evidence contains sensitive information, an Attester typically requires that a Verifier authenticates itself (e.g., at TLS session establishment) and might even request a remote attestation before the Attester sends the sensitive Evidence. This can be done by having the Attester first act as a Verifier/Relying Party, and the Verifier act as its own Attester, as discussed above.

7.3. Relying Party Owner

The Relying Party Owner might also require that the Relying Party first act as an Attester, providing Evidence that the Owner can appraise, before the Owner would give the Relying Party an updated policy that might contain sensitive information. In such a case, authentication or attestation in both directions might be needed, in which case typically one side's Evidence must be considered safe to share with an untrusted entity, in order to bootstrap the sequence. See Section 11 for more discussion.

7.4. Verifier

The Verifier trusts (or more specifically, the Verifier's security policy is written in a way that configures the Verifier to trust) a manufacturer, or the manufacturer's hardware, so as to be able to appraise the trustworthiness of that manufacturer's devices. Such trust is expressed by storing one or more trust anchors in the Verifier's trust anchor store.

In a typical solution, a Verifier comes to trust an Attester indirectly by having an Endorser (such as a manufacturer) vouch for the Attester's ability to securely generate Evidence, in which case the Endorser's key material is stored in the Verifier's trust anchor store.

In some solutions, a Verifier might be configured to directly trust an Attester by having the Verifier have the Attester's key material (rather than the Endorser's) in its trust anchor store.

Such direct trust must first be established at the time of trust anchor store configuration either by checking with an Endorser at that time, or by conducting a security analysis of the specific device. Having the Attester directly in the trust anchor store narrows the Verifier's trust to only specific devices rather than all devices the Endorser might vouch for, such as all devices manufactured by the same manufacturer in the case that the Endorser is a manufacturer.

Such narrowing is often important since physical possession of a device can also be used to conduct a number of attacks, and so a device in a physically secure environment (such as one's own premises) may be considered trusted whereas devices owned by others would not be. This often results in a desire to either have the owner run their own Endorser that would only endorse devices one owns, or to use Attesters directly in the trust anchor store. When there are many Attesters owned, the use of an Endorser enables better scalability.

That is, a Verifier might appraise the trustworthiness of an application component, operating system component, or service under the assumption that information provided about it by the lower-layer firmware or software is true. A stronger level of assurance of security comes when information can be vouched for by hardware or by ROM code, especially if such hardware is physically resistant to hardware tampering. In most cases, components that have to be vouched for via Endorsements because no Evidence is generated about them are referred to as roots of trust.

The manufacturer having arranged for an Attesting Environment to be provisioned with key material with which to sign Evidence, the Verifier is then provided with some way of verifying the signature on the Evidence. This may be in the form of an appropriate trust anchor, or the Verifier may be provided with a database of public keys (rather than certificates) or even carefully curated and secured lists of symmetric keys.

The nature of how the Verifier manages to validate the signatures produced by the Attester is critical to the secure operation of a remote attestation system, but is not the subject of standardization within this architecture.

A conveyance protocol that provides authentication and integrity protection can be used to convey Evidence that is otherwise unprotected (e.g., not signed). Appropriate conveyance of unprotected Evidence (e.g., [I-D.birkholz-rats-uccs]) relies on the following conveyance protocol's protection capabilities:

1. The key material used to authenticate and integrity protect the conveyance channel is trusted by the Verifier to speak for the Attesting Environment(s) that collected Claims about the Target Environment(s).
2. All unprotected Evidence that is conveyed is supplied exclusively by the Attesting Environment that has the key material that protects the conveyance channel
3. The root of trust protects both the conveyance channel key material and the Attesting Environment with equivalent strength protections.

As illustrated in [I-D.birkholz-rats-uccs], an entity that receives unprotected Evidence via a trusted conveyance channel always takes on the responsibility of vouching for the Evidence's authenticity and freshness. If protected Evidence is generated, the Attester's Attesting Environments take on that responsibility. In cases where unprotected Evidence is processed by a Verifier, Relying Parties have to trust that the Verifier is capable of handling Evidence in a manner that preserves the Evidence's authenticity and freshness. Generating and conveying unprotected Evidence always creates significant risk and the benefits of that approach have to be carefully weighed against potential drawbacks.

See Section 12 for discussion on security strength.

7.5. Endorser, Reference Value Provider, and Verifier Owner

In some scenarios, the Endorser, Reference Value Provider, and Verifier Owner may need to trust the Verifier before giving the Endorsement, Reference Values, or appraisal policy to it. This can be done similarly to how a Relying Party might establish trust in a Verifier.

As discussed in Section 7.3, authentication or attestation in both directions might be needed, in which case typically one side's identity or Evidence must be considered safe to share with an untrusted entity, in order to bootstrap the sequence. See Section 11 for more discussion.

8. Conceptual Messages

Figure 1 illustrates the flow of a conceptual messages between various roles. This section provides additional elaboration and implementation considerations. It is the responsibility of protocol specifications to define the actual data format and semantics of any relevant conceptual messages.

8.1. Evidence

Evidence is a set of Claims about the target environment that reveal operational status, health, configuration or construction that have security relevance. Evidence is appraised by a Verifier to establish its relevance, compliance, and timeliness. Claims need to be collected in a manner that is reliable. Evidence needs to be securely associated with the target environment so that the Verifier cannot be tricked into accepting Claims originating from a different environment (that may be more trustworthy). Evidence also must be protected from man-in-the-middle attackers who may observe, change or misdirect Evidence as it travels from Attester to Verifier. The timeliness of Evidence can be captured using Claims that pinpoint the time or interval when changes in operational status, health, and so forth occur.

8.2. Endorsements

An Endorsement is a secure statement that some entity (e.g., a manufacturer) vouches for the integrity of the device's signing capability. For example, if the signing capability is in hardware, then an Endorsement might be a manufacturer certificate that signs a public key whose corresponding private key is only known inside the device's hardware. Thus, when Evidence and such an Endorsement are used together, an appraisal procedure can be conducted based on appraisal policies that may not be specific to the device instance, but merely specific to the manufacturer providing the Endorsement. For example, an appraisal policy might simply check that devices from a given manufacturer have information matching a set of Reference Values, or an appraisal policy might have a set of more complex logic on how to appraise the validity of information.

However, while an appraisal policy that treats all devices from a given manufacturer the same may be appropriate for some use cases, it would be inappropriate to use such an appraisal policy as the sole means of authorization for use cases that wish to constrain which compliant devices are considered authorized for some purpose. For example, an enterprise using remote attestation for Network Endpoint Assessment [RFC5209] may not wish to let every healthy laptop from the same manufacturer onto the network, but instead only want to let devices that it legally owns onto the network. Thus, an Endorsement may be helpful information in authenticating information about a device, but is not necessarily sufficient to authorize access to resources which may need device-specific information such as a public key for the device or component or user on the device.

8.3. Reference Values

Reference Values used in appraisal procedures come from a Reference Value Provider and are then used by the Verifier to compare to Evidence. Reference Values with matching Evidence produces acceptable Claims. Additionally, appraisal policy may play a role in determining the acceptance of Claims.

8.4. Attestation Results

Attestation Results are the input used by the Relying Party to decide the extent to which it will trust a particular Attester, and allow it to access some data or perform some operation.

Attestation Results may carry a boolean value indicating compliance or non-compliance with a Verifier's appraisal policy, or may carry a richer set of Claims about the Attester, against which the Relying Party applies its Appraisal Policy for Attestation Results.

The quality of the Attestation Results depends upon the ability of the Verifier to evaluate the Attester. Different Attesters have a different Strength of Function [strengthoffunction], which results in the Attestation Results being qualitatively different in strength.

An Attestation Result that indicates non-compliance can be used by an Attester (in the passport model) or a Relying Party (in the background-check model) to indicate that the Attester should not be treated as authorized and may be in need of remediation. In some cases, it may even indicate that the Evidence itself cannot be authenticated as being correct.

By default, the Relying Party does not believe the Attester to be compliant. Upon receipt of an authentic Attestation Result and given the Appraisal Policy for Attestation Results is satisfied, the

Attester is allowed to perform the prescribed actions or access. The simplest such appraisal policy might authorize granting the Attester full access or control over the resources guarded by the Relying Party. A more complex appraisal policy might involve using the information provided in the Attestation Result to compare against expected values, or to apply complex analysis of other information contained in the Attestation Result.

Thus, Attestation Results often need to include detailed information about the Attester, for use by Relying Parties, much like physical passports and drivers licenses include personal information such as name and date of birth. Unlike Evidence, which is often very device- and vendor-specific, Attestation Results can be vendor-neutral, if the Verifier has a way to generate vendor-agnostic information based on the appraisal of vendor-specific information in Evidence. This allows a Relying Party's appraisal policy to be simpler, potentially based on standard ways of expressing the information, while still allowing interoperability with heterogeneous devices.

Finally, whereas Evidence is signed by the device (or indirectly by a manufacturer, if Endorsements are used), Attestation Results are signed by a Verifier, allowing a Relying Party to only need a trust relationship with one entity, rather than a larger set of entities, for purposes of its appraisal policy.

8.5. Appraisal Policies

The Verifier, when appraising Evidence, or the Relying Party, when appraising Attestation Results, checks the values of matched Claims against constraints specified in its appraisal policy. Examples of such constraints checking include:

- * comparison for equality against a Reference Value, or
- * a check for being in a range bounded by Reference Values, or
- * membership in a set of Reference Values, or
- * a check against values in other Claims.

Upon completing all appraisal policy constraints, the remaining Claims are accepted as input toward determining Attestation Results, when appraising Evidence, or as input to a Relying Party, when appraising Attestation Results.

9. Claims Encoding Formats

The following diagram illustrates a relationship to which remote attestation is desired to be added:

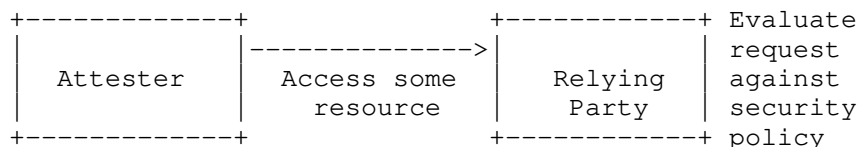


Figure 8: Typical Resource Access

In this diagram, the protocol between Attester and a Relying Party can be any new or existing protocol (e.g., HTTP(S), COAP(S), ROLIE [RFC8322], 802.1x, OPC UA [OPCUA], etc.), depending on the use case.

Typically, such protocols already have mechanisms for passing security information for authentication and authorization purposes. Common formats include JWTs [RFC7519], CWTs [RFC8392], and X.509 certificates.

Retrofitting already deployed protocols with remote attestation requires adding RATS conceptual messages to the existing data flows. This must be done in a way that does not degrade the security properties of the systems involved and should use native extension mechanisms provided by the underlying protocol. For example, if a TLS handshake is to be extended with remote attestation capabilities, attestation Evidence may be embedded in an ad-hoc X.509 certificate extension (e.g., [TCG-DICE]), or into a new TLS Certificate Type (e.g., [I-D.tschofenig-tls-cwt]).

Especially for constrained nodes there is a desire to minimize the amount of parsing code needed in a Relying Party, in order to both minimize footprint and to minimize the attack surface. While it would be possible to embed a CWT inside a JWT, or a JWT inside an X.509 extension, etc., there is a desire to encode the information natively in a format that is already supported by the Relying Party.

This motivates having a common "information model" that describes the set of remote attestation related information in an encoding-agnostic way, and allowing multiple encoding formats (CWT, JWT, X.509, etc.) that encode the same information into the Claims format needed by the Relying Party.

The following diagram illustrates that Evidence and Attestation Results might be expressed via multiple potential encoding formats, so that they can be conveyed by various existing protocols. It also

motivates why the Verifier might also be responsible for accepting Evidence that encodes Claims in one format, while issuing Attestation Results that encode Claims in a different format.

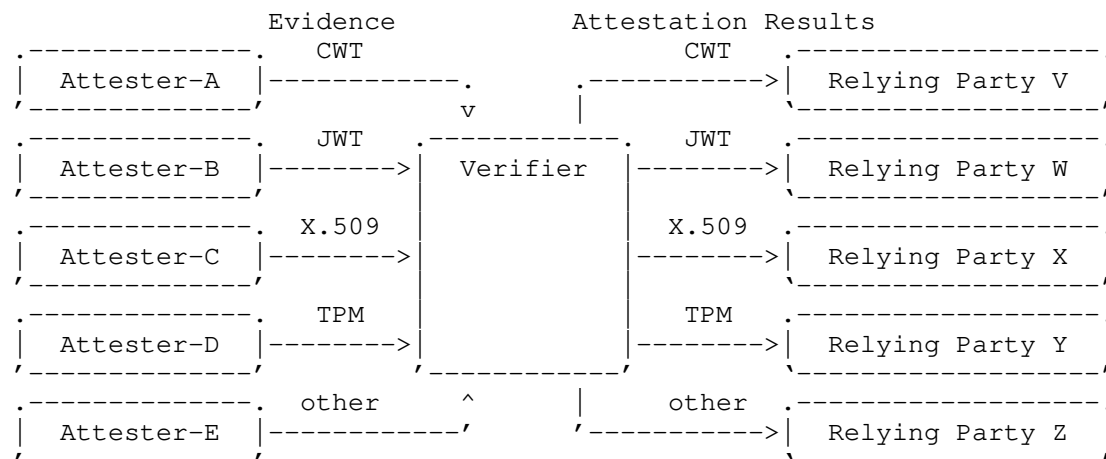


Figure 9: Multiple Attesters and Relying Parties with Different Formats

10. Freshness

A Verifier or Relying Party might need to learn the point in time (i.e., the "epoch") an Evidence or Attestation Result has been produced. This is essential in deciding whether the included Claims and their values can be considered fresh, meaning they still reflect the latest state of the Attester, and that any Attestation Result was generated using the latest Appraisal Policy for Evidence.

Freshness is assessed based on the Appraisal Policy for Evidence or Attestation Results that compares the estimated epoch against an "expiry" threshold defined locally to that policy. There is, however, always a race condition possible in that the state of the Attester, and the appraisal policies might change immediately after the Evidence or Attestation Result was generated. The goal is merely to narrow their recentness to something the Verifier (for Evidence) or Relying Party (for Attestation Result) is willing to accept. Some flexibility on the freshness requirement is a key component for enabling caching and reuse of both Evidence and Attestation Results, which is especially valuable in cases where their computation uses a substantial part of the resource budget (e.g., energy in constrained devices).

There are three common approaches for determining the epoch of Evidence or an Attestation Result.

10.1. Explicit Timekeeping using Synchronized Clocks

The first approach is to rely on synchronized and trustworthy clocks, and include a signed timestamp (see [I-D.birkholz-rats-tuda]) along with the Claims in the Evidence or Attestation Result. Timestamps can also be added on a per-Claim basis to distinguish the time of generation of Evidence or Attestation Result from the time that a specific Claim was generated. The clock's trustworthiness can generally be established via Endorsements and typically requires additional Claims about the signer's time synchronization mechanism.

In some use cases, however, a trustworthy clock might not be available. For example, in many Trusted Execution Environments (TEEs) today, a clock is only available outside the TEE and so cannot be trusted by the TEE.

10.2. Implicit Timekeeping using Nonces

A second approach places the onus of timekeeping solely on the Verifier (for Evidence) or the Relying Party (for Attestation Results), and might be suitable, for example, in case the Attester does not have a trustworthy clock or time synchronization is otherwise impaired. In this approach, a non-predictable nonce is sent by the appraising entity, and the nonce is then signed and included along with the Claims in the Evidence or Attestation Result. After checking that the sent and received nonces are the same, the appraising entity knows that the Claims were signed after the nonce was generated. This allows associating a "rough" epoch to the Evidence or Attestation Result. In this case the epoch is said to be rough because:

- * The epoch applies to the entire Claim set instead of a more granular association, and
- * The time between the creation of Claims and the collection of Claims is indistinguishable.

10.3. Implicit Timekeeping using Epoch IDs

A third approach relies on having epoch identifiers (or "IDs") periodically sent to both the sender and receiver of Evidence or Attestation Results by some "Epoch ID Distributor".

Epoch IDs are different from nonces as they can be used more than once and can even be used by more than one entity at the same time. Epoch IDs are different from timestamps as they do not have to convey information about a point in time, i.e., they are not necessarily monotonically increasing integers.

Like the nonce approach, this allows associating a "rough" epoch without requiring a trustworthy clock or time synchronization in order to generate or appraise the freshness of Evidence or Attestation Results. Only the Epoch ID Distributor requires access to a clock so it can periodically send new epoch IDs.

The most recent epoch ID is included in the produced Evidence or Attestation Results, and the appraising entity can compare the epoch ID in received Evidence or Attestation Results against the latest epoch ID it received from the Epoch ID Distributor to determine if it is within the current epoch. An actual solution also needs to take into account race conditions when transitioning to a new epoch, such as by using a counter signed by the Epoch ID Distributor as the epoch ID, or by including both the current and previous epoch IDs in messages and/or checks, by requiring retries in case of mismatching epoch IDs, or by buffering incoming messages that might be associated with a epoch ID that the receiver has not yet obtained.

More generally, in order to prevent an appraising entity from generating false negatives (e.g., discarding Evidence that is deemed stale even if it is not), the appraising entity should keep an "epoch window" consisting of the most recently received epoch IDs. The depth of such epoch window is directly proportional to the maximum network propagation delay between the first to receive the epoch ID and the last to receive the epoch ID, and it is inversely proportional to the epoch duration. The appraising entity shall compare the epoch ID carried in the received Evidence or Attestation Result with the epoch IDs in its epoch window to find a suitable match.

Whereas the nonce approach typically requires the appraising entity to keep state for each nonce generated, the epoch ID approach minimizes the state kept to be independent of the number of Attesters or Verifiers from which it expects to receive Evidence or Attestation Results, as long as all use the same Epoch ID Distributor.

10.4. Discussion

Implicit and explicit timekeeping can be combined into hybrid mechanisms. For example, if clocks exist and are considered trustworthy but are not synchronized, a nonce-based exchange may be used to determine the (relative) time offset between the involved peers, followed by any number of timestamp based exchanges.

It is important to note that the actual values in Claims might have been generated long before the Claims are signed. If so, it is the signer's responsibility to ensure that the values are still correct when they are signed. For example, values generated at boot time might have been saved to secure storage until network connectivity is established to the remote Verifier and a nonce is obtained.

A more detailed discussion with examples appears in Section 16.

For a discussion on the security of epoch IDs see Section 12.3.

11. Privacy Considerations

The conveyance of Evidence and the resulting Attestation Results reveal a great deal of information about the internal state of a device as well as potentially any users of the device. In many cases, the whole point of attestation procedures is to provide reliable information about the type of the device and the firmware/software that the device is running. This information might be particularly interesting to many attackers. For example, knowing that a device is running a weak version of firmware provides a way to aim attacks better.

Many Claims in Evidence and Attestation Results are potentially Personally Identifying Information (PII) depending on the end-to-end use case of the remote attestation procedure. Remote attestation that goes up to include containers and applications, e.g., a blood pressure monitor, may further reveal details about specific systems or users.

In some cases, an attacker may be able to make inferences about the contents of Evidence from the resulting effects or timing of the processing. For example, an attacker might be able to infer the value of specific Claims if it knew that only certain values were accepted by the Relying Party.

Evidence and Attestation Results are expected to be integrity protected (i.e., either via signing or a secure channel) and optionally might be confidentiality protected via encryption. If confidentiality protection via signing the conceptual messages is

omitted or unavailable, the protecting protocols that convey Evidence or Attestation Results are responsible for detailing what kinds of information are disclosed, and to whom they are exposed.

As Evidence might contain sensitive or confidential information, Attesters are responsible for only sending such Evidence to trusted Verifiers. Some Attesters might want a stronger level of assurance of the trustworthiness of a Verifier before sending Evidence to it. In such cases, an Attester can first act as a Relying Party and ask for the Verifier's own Attestation Result, and appraising it just as a Relying Party would appraise an Attestation Result for any other purpose.

Another approach to deal with Evidence is to remove PII from the Evidence while still being able to verify that the Attester is one of a large set. This approach is often called "Direct Anonymous Attestation". See [CCC-DeepDive] section 6.2 for more discussion.

12. Security Considerations

This document provides an architecture for doing remote attestation. No specific wire protocol is documented here. Without a specific proposal to compare against, it is impossible to know if the security threats listed below have been mitigated well. The security considerations below should be read as being essentially requirements against realizations of the RATS Architecture. Some threats apply to protocols, some are against implementations (code), and some threats are against physical infrastructure (such as factories).

12.1. Attester and Attestation Key Protection

Implementers need to pay close attention to the protection of the Attester and the manufacturing processes for provisioning attestation key material. If either of these are compromised, intended levels of assurance for RATS are compromised because attackers can forge Evidence or manipulate the Attesting Environment. For example, a Target Environment should not be able to tamper with the Attesting Environment that measures it, by isolating the two environments from each other in some way.

Remote attestation applies to use cases with a range of security requirements, so the protections discussed here range from low to high security where low security may be limited to application or process isolation by the device's operating system, and high security may involve specialized hardware to defend against physical attacks on a chip.

12.1.1. On-Device Attester and Key Protection

It is assumed that an Attesting Environment is sufficiently isolated from the Target Environment it collects Claims about and that it signs the resulting Claims set with an attestation key, so that the Target Environment cannot forge Evidence about itself. Such an isolated environment might be provided by a process, a dedicated chip, a TEE, a virtual machine, or another secure mode of operation. The Attesting Environment must be protected from unauthorized modification to ensure it behaves correctly. Confidentiality protection of the Attesting Environment's signing key is vital so it cannot be misused to forge Evidence.

In many cases the user or owner of a device that includes the role of Attester must not be able to modify or extract keys from the Attesting Environments, to prevent creating forged Evidence. Some common examples include the user of a mobile phone or FIDO authenticator. An essential value-add provided by RATS is for the Relying Party to be able to trust the Attester even if the user or owner is not trusted.

Measures for a minimally protected system might include process or application isolation provided by a high-level operating system, and restricted access to root or system privileges. In contrast, For really simple single-use devices that don't use a protected mode operating system, like a Bluetooth speaker, the only factual isolation might be the sturdy housing of the device.

Measures for a moderately protected system could include a special restricted operating environment, such as a TEE. In this case, only security-oriented software has access to the Attester and key material.

Measures for a highly protected system could include specialized hardware that is used to provide protection against chip decapping attacks, power supply and clock glitching, faulting injection and RF and power side channel attacks.

12.1.2. Attestation Key Provisioning Processes

Attestation key provisioning is the process that occurs in the factory or elsewhere to establish signing key material on the device and the validation key material off the device. Sometimes this is procedure is referred to as personalization or customization.

12.1.2.1. Off-Device Key Generation

One way to provision key material is to first generate it external to the device and then copy the key onto the device. In this case, confidentiality protection of the generator, as well as for the path over which the key is provisioned, is necessary. The manufacturer needs to take care to protect corresponding key material with measures appropriate for its value.

The degree of protection afforded to this key material can vary by device, based upon considerations as to a cost/benefit evaluation of the intended function of the device. The confidentiality protection is fundamentally based upon some amount of physical protection: while encryption is often used to provide confidentiality when a key is conveyed across a factory, where the attestation key is created or applied, it must be available in an unencrypted form. The physical protection can therefore vary from situations where the key is unencrypted only within carefully controlled secure enclaves within silicon, to situations where an entire facility is considered secure, by the simple means of locked doors and limited access.

The cryptography that is used to enable confidentiality protection of the attestation key comes with its own requirements to be secured. This results in recursive problems, as the key material used to provision attestation keys must again somehow have been provisioned securely beforehand (requiring an additional level of protection, and so on).

So, this is why, in general, a combination of some physical security measures and some cryptographic measures is used to establish confidentiality protection.

12.1.2.2. On-Device Key Generation

When key material is generated within a device and the secret part of it never leaves the device, then the problem may lessen. For public-key cryptography, it is, by definition, not necessary to maintain confidentiality of the public key: however integrity of the chain of custody of the public key is necessary in order to avoid attacks where an attacker is able get a key they control endorsed.

To summarize: attestation key provisioning must ensure that only valid attestation key material is established in Attesters.

12.2. Integrity Protection

Any solution that conveys information used for security purposes, whether such information is in the form of Evidence, Attestation Results, Endorsements, or appraisal policy must support end-to-end integrity protection and replay attack prevention, and often also needs to support additional security properties, including:

- * end-to-end encryption,
- * denial of service protection,
- * authentication,
- * auditing,
- * fine grained access controls, and
- * logging.

Section 10 discusses ways in which freshness can be used in this architecture to protect against replay attacks.

To assess the security provided by a particular appraisal policy, it is important to understand the strength of the root of trust, e.g., whether it is mutable software, or firmware that is read-only after boot, or immutable hardware/ROM.

It is also important that the appraisal policy was itself obtained securely. If an attacker can configure appraisal policies for a Relying Party or for a Verifier, then integrity of the process is compromised.

Security protections in RATS may be applied at different layers, whether by a conveyance protocol, or an information encoding format. This architecture expects conceptual messages (see Section 8) to be end-to-end protected based on the role interaction context. For example, if an Attester produces Evidence that is relayed through some other entity that doesn't implement the Attester or the intended Verifier roles, then the relaying entity should not expect to have access to the Evidence.

12.3. Epoch ID-based Attestation

Epoch IDs, described in Section 10.3, can be tampered with, replayed, dropped, delayed, and reordered by an attacker.

An attacker could be either external or belong to the distribution group, for example, if one of the Attester entities have been compromised.

An attacker who is able to tamper with epoch IDs can potentially lock all the participants in a certain epoch of choice for ever, effectively freezing time. This is problematic since it destroys the ability to ascertain freshness of Evidence and Attestation Results.

To mitigate this threat, the transport should be at least integrity protected and provide origin authentication.

Selective dropping of epoch IDs is equivalent to pinning the victim node to a past epoch. An attacker could drop epoch IDs to only some entities and not others, which will typically result in a denial of service due to the permanent staleness of the Attestation Result or Evidence.

Delaying or reordering epoch IDs is equivalent to manipulating the victim's timeline at will. This ability could be used by a malicious actor (e.g., a compromised router) to mount a confusion attack where, for example, a Verifier is tricked into accepting Evidence coming from a past epoch as fresh, while in the meantime the Attester has been compromised.

Reordering and dropping attacks are mitigated if the transport provides the ability to detect reordering and drop. However, the delay attack described above can't be thwarted in this manner.

12.4. Trust Anchor Protection

As noted in Section 7, Verifiers and Relying Parties have trust anchor stores that must be secured. Specifically, a trust anchor store must resist modification against unauthorized insertion, deletion, and modification.

If certificates are used as trust anchors, Verifiers and Relying Parties are also responsible for validating the entire certificate path up to the trust anchor, which includes checking for certificate revocation. See Section 6 of [RFC5280] for details.

13. IANA Considerations

This document does not require any actions by IANA.

14. Acknowledgments

Special thanks go to Joerg Borchert, Nancy Cam-Winget, Jessica Fitzgerald-McKay, Diego Lopez, Laurence Lundblade, Paul Rowe, Hannes Tschofenig, Frank Xia, and David Wooten.

15. Notable Contributions

Thomas Hardjono created initial versions of the terminology section in collaboration with Ned Smith. Eric Voit provided the conceptual separation between Attestation Provision Flows and Attestation Evidence Flows. Monty Wisemen created the content structure of the first three architecture drafts. Carsten Bormann provided many of the motivational building blocks with respect to the Internet Threat Model.

16. Appendix A: Time Considerations

Section 10 discussed various issues and requirements around freshness of evidence, and summarized three approaches that might be used by different solutions to address them. This appendix provides more details with examples to help illustrate potential approaches, to inform those creating specific solutions.

The table below defines a number of relevant events, with an ID that is used in subsequent diagrams. The times of said events might be defined in terms of an absolute clock time, such as the Coordinated Universal Time timescale, or might be defined relative to some other timestamp or timeticks counter, such as a clock resetting its epoch each time it is powered on.

ID	Event	Explanation of event
VG	Value generated	A value to appear in a Claim was created. In some cases, a value may have technically existed before an Attester became aware of it but the Attester might have no idea how long it has had that value. In such a case, the Value created time is the time at which the Claim containing the copy of the value was created.
NS	Nonce sent	A nonce not predictable to an Attester (recentness & uniqueness) is sent to an Attester.
NR	Nonce	A nonce is relayed to an Attester by

	relayed	another entity.
IR	Epoch ID received	An epoch ID is successfully received and processed by an entity.
EG	Evidence generation	An Attester creates Evidence from collected Claims.
ER	Evidence relayed	A Relying Party relays Evidence to a Verifier.
RG	Result generation	A Verifier appraises Evidence and generates an Attestation Result.
RR	Result relayed	A Relying Party relays an Attestation Result to a Relying Party.
RA	Result appraised	The Relying Party appraises Attestation Results.
OP	Operation performed	The Relying Party performs some operation requested by the Attester via a resource access protocol as depicted in Figure 8, e.g., across a session created earlier at time(RA).
RX	Result expiry	An Attestation Result should no longer be accepted, according to the Verifier that generated it.

Table 1

Using the table above, a number of hypothetical examples of how a solution might be built are illustrated below. This list is not intended to be complete, but is just representative enough to highlight various timing considerations.

All times are relative to the local clocks, indicated by an "_a" (Attester), "_v" (Verifier), or "_r" (Relying Party) suffix.

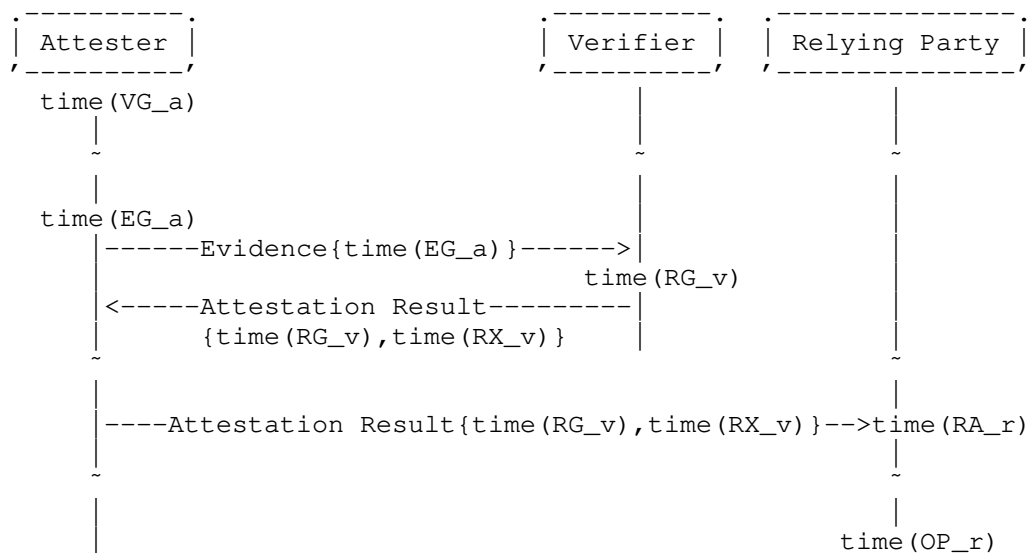
Times with an appended Prime (') indicate a second instance of the same event.

How and if clocks are synchronized depends upon the model.

In the figures below, curly braces indicate containment. For example, the notation `Evidence{foo}` indicates that 'foo' is contained in the Evidence and is thus covered by its signature.

16.1. Example 1: Timestamp-based Passport Model Example

The following example illustrates a hypothetical Passport Model solution that uses timestamps and requires roughly synchronized clocks between the Attester, Verifier, and Relying Party, which depends on using a secure clock synchronization mechanism. As a result, the receiver of a conceptual message containing a timestamp can directly compare it to its own clock and timestamps.



The Verifier can check whether the Evidence is fresh when appraising it at time(RG_v) by checking " $\text{time}(\text{RG}_v) - \text{time}(\text{EG}_a) < \text{Threshold}$ ", where the Verifier's threshold is large enough to account for the maximum permitted clock skew between the Verifier and the Attester.

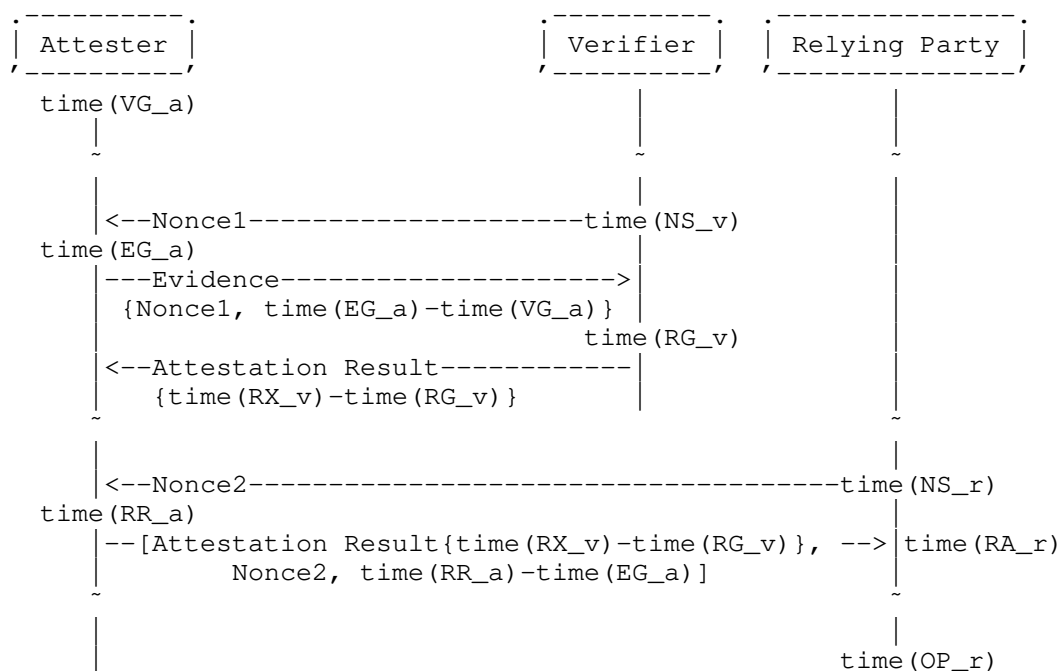
If time(VG_a) is also included in the Evidence along with the Claim value generated at that time, and the Verifier decides that it can trust the time(VG_a) value, the Verifier can also determine whether the Claim value is recent by checking " $\text{time}(\text{RG}_v) - \text{time}(\text{VG}_a) < \text{Threshold}$ ". The threshold is decided by the Appraisal Policy for Evidence, and again needs to take into account the maximum permitted clock skew between the Verifier and the Attester.

The Relying Party can check whether the Attestation Result is fresh when appraising it at $\text{time}(\text{RA}_r)$ by checking " $\text{time}(\text{RA}_r) - \text{time}(\text{RG}_v) < \text{Threshold}$ ", where the Relying Party's threshold is large enough to account for the maximum permitted clock skew between the Relying Party and the Verifier. The result might then be used for some time (e.g., throughout the lifetime of a connection established at $\text{time}(\text{RA}_r)$). The Relying Party must be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain fresh enough. Thus, it might allow use (at $\text{time}(\text{OP}_r)$) as long as " $\text{time}(\text{OP}_r) - \text{time}(\text{RG}_v) < \text{Threshold}$ ". However, if the Attestation Result contains an expiry time $\text{time}(\text{RX}_v)$ then it could explicitly check " $\text{time}(\text{OP}_r) < \text{time}(\text{RX}_v)$ ".

16.2. Example 2: Nonce-based Passport Model Example

The following example illustrates a hypothetical Passport Model solution that uses nonces instead of timestamps. Compared to the timestamp-based example, it requires an extra round trip to retrieve a nonce, and requires that the Verifier and Relying Party track state to remember the nonce for some period of time.

The advantage is that it does not require that any clocks are synchronized. As a result, the receiver of a conceptual message containing a timestamp cannot directly compare it to its own clock or timestamps. Thus we use a suffix ("a" for Attester, "v" for Verifier, and "r" for Relying Party) on the IDs below indicating which clock generated them, since times from different clocks cannot be compared. Only the delta between two events from the sender can be used by the receiver.



In this example solution, the Verifier can check whether the Evidence is fresh at "time(RG_v)" by verifying that "time(RG_v)-time(NS_v) < Threshold".

The Verifier cannot, however, simply rely on a Nonce to determine whether the value of a Claim is recent, since the Claim value might have been generated long before the nonce was sent by the Verifier. However, if the Verifier decides that the Attester can be trusted to correctly provide the delta "time(EG_a)-time(VG_a)", then it can determine recency by checking "time(RG_v)-time(NS_v) + time(EG_a)-time(VG_a) < Threshold".

Similarly if, based on an Attestation Result from a Verifier it trusts, the Relying Party decides that the Attester can be trusted to correctly provide time deltas, then it can determine whether the Attestation Result is fresh by checking "time(OP_r)-time(NS_r) + time(RR_a)-time(EG_a) < Threshold". Although the Nonce2 and "time(RR_a)-time(EG_a)" values cannot be inside the Attestation Result, they might be signed by the Attester such that the Attestation Result vouches for the Attester's signing capability.

The Relying Party must still be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain valid. Thus, if the Attestation Result sends a

validity lifetime in terms of `"time(RX_v)-time(RG_v)"`, then the Relying Party can check `"time(OP_r)-time(NS_r) < time(RX_v)-time(RG_v)"`.

16.3. Example 3: Epoch ID-based Passport Model Example

The example in Figure 10 illustrates a hypothetical Passport Model solution that uses epoch IDs instead of nonces or timestamps.

The Epoch ID Distributor broadcasts epoch ID "I" which starts a new epoch "E" for a protocol participant upon reception at `"time(IR)"`.

The Attester generates Evidence incorporating epoch ID "I" and conveys it to the Verifier.

The Verifier appraises that the received epoch ID "I" is "fresh" according to the definition provided in Section 10.3 whereby retries are required in the case of mismatching epoch IDs, and generates an Attestation Result. The Attestation Result is conveyed to the Attester.

After the transmission of epoch ID "I" a new epoch "E'" is established when "I'" is received by each protocol participant. The Attester relays the Attestation Result obtained during epoch "E" (associated with epoch ID "I") to the Relying Party using the epoch ID for the current epoch "I'". If the Relying Party had not yet received "I'", then the Attestation Result would be rejected, but in this example, it is received.

In the illustrated scenario, the epoch ID for relaying an Attestation Result to the Relying Party is current, while a previous epoch ID was used to generate Verifier evaluated evidence. This indicates that at least one epoch transition has occurred, and the Attestation Results may only be as fresh as the previous epoch. If the Relying Party remembers the previous epoch ID "I" during an epoch window as discussed in Section 10.3, and the message is received during that window, the Attestation Result is accepted as fresh, and otherwise it is rejected as stale.

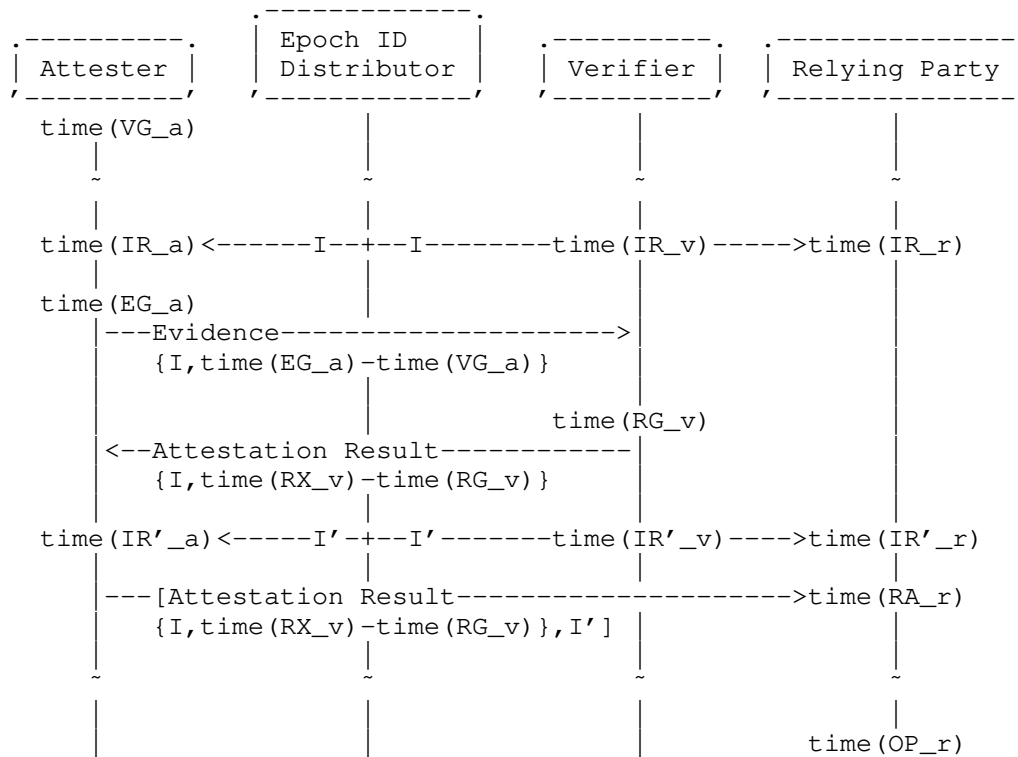
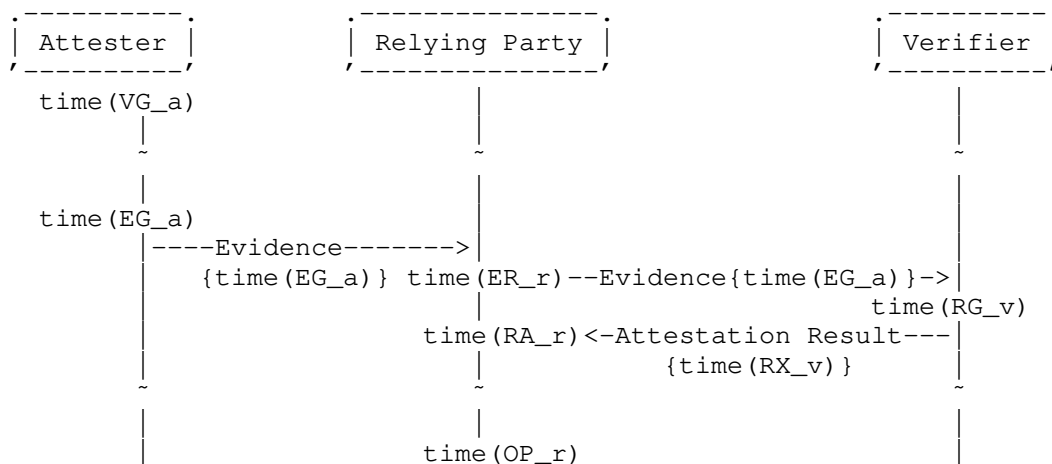


Figure 10: Epoch ID-based Passport Model

16.4. Example 4: Timestamp-based Background-Check Model Example

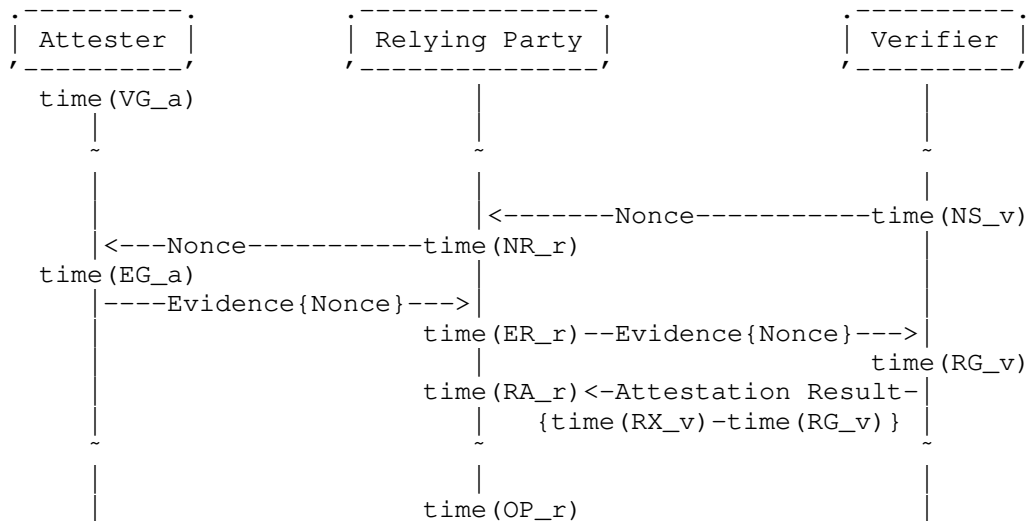
The following example illustrates a hypothetical Background-Check Model solution that uses timestamps and requires roughly synchronized clocks between the Attester, Verifier, and Relying Party.



The time considerations in this example are equivalent to those discussed under Example 1 above.

16.5. Example 5: Nonce-based Background-Check Model Example

The following example illustrates a hypothetical Background-Check Model solution that uses nonces and thus does not require that any clocks are synchronized. In this example solution, a nonce is generated by a Verifier at the request of a Relying Party, when the Relying Party needs to send one to an Attester.



The Verifier can check whether the Evidence is fresh, and whether a Claim value is recent, the same as in Example 2 above.

However, unlike in Example 2, the Relying Party can use the Nonce to determine whether the Attestation Result is fresh, by verifying that `"time(OP_r)-time(NR_r) < Threshold"`.

The Relying Party must still be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain valid. Thus, if the Attestation Result sends a validity lifetime in terms of `"time(RX_v)-time(RG_v)"`, then the Relying Party can check `"time(OP_r)-time(ER_r) < time(RX_v)-time(RG_v)"`.

17. References

17.1. Normative References

- [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://www.rfc-editor.org/rfc/rfc5280>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/rfc/rfc7519>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/rfc/rfc8392>>.

17.2. Informative References

- [CCC-DeepDive] Confidential Computing Consortium, "Confidential Computing Deep Dive", n.d., <<https://confidentialcomputing.io/whitepaper-02-latest>>.
- [CTAP] FIDO Alliance, "Client to Authenticator Protocol", n.d., <<https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-client-to-authenticator-protocol-v2.0-id-20180227.html>>.

- [I-D.birkholz-rats-tuda]
Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann,
"Time-Based Uni-Directional Attestation", Work in
Progress, Internet-Draft, draft-birkholz-rats-tuda-04, 13
January 2021,
<<https://tools.ietf.org/html/draft-birkholz-rats-tuda-04>>.
- [I-D.birkholz-rats-uccs]
Birkholz, H., O'Donoghue, J., Cam-Winget, N., and C.
Bormann, "A CBOR Tag for Unprotected CWT Claims Sets",
Work in Progress, Internet-Draft, draft-birkholz-rats-
uccs-03, 8 March 2021,
<<https://tools.ietf.org/html/draft-birkholz-rats-uccs-03>>.
- [I-D.ietf-teep-architecture]
Pei, M., Tschofenig, H., Thaler, D., and D. Wheeler,
"Trusted Execution Environment Provisioning (TEEP)
Architecture", Work in Progress, Internet-Draft, draft-
ietf-teep-architecture-14, 22 February 2021,
<<https://tools.ietf.org/html/draft-ietf-teep-architecture-14>>.
- [I-D.tschofenig-tls-cwt]
Tschofenig, H. and M. Brossard, "Using CBOR Web Tokens
(CWTs) in Transport Layer Security (TLS) and Datagram
Transport Layer Security (DTLS)", Work in Progress,
Internet-Draft, draft-tschofenig-tls-cwt-02, 13 July 2020,
<<https://tools.ietf.org/html/draft-tschofenig-tls-cwt-02>>.
- [OPCUA] OPC Foundation, "OPC Unified Architecture Specification,
Part 2: Security Model, Release 1.03", OPC 10000-2 , 25
November 2015, <[https://opcfoundation.org/developer-tools/
specifications-unified-architecture/part-2-security-
model/](https://opcfoundation.org/developer-tools/specifications-unified-architecture/part-2-security-model/)>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2",
FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007,
<<https://www.rfc-editor.org/rfc/rfc4949>>.
- [RFC5209] Sangster, P., Khosravi, H., Mani, M., Narayan, K., and J.
Tardo, "Network Endpoint Assessment (NEA): Overview and
Requirements", RFC 5209, DOI 10.17487/RFC5209, June 2008,
<<https://www.rfc-editor.org/rfc/rfc5209>>.
- [RFC6024] Reddy, R. and C. Wallace, "Trust Anchor Management
Requirements", RFC 6024, DOI 10.17487/RFC6024, October
2010, <<https://www.rfc-editor.org/rfc/rfc6024>>.

[RFC8322] Field, J., Banghart, S., and D. Waltermire, "Resource-Oriented Lightweight Information Exchange (ROLIE)", RFC 8322, DOI 10.17487/RFC8322, February 2018, <<https://www.rfc-editor.org/rfc/rfc8322>>.

[strengthoffunction] NISC, "Strength of Function", n.d., <https://csrc.nist.gov/glossary/term/strength_of_function>.

[TCG-DICE] Trusted Computing Group, "DICE Certificate Profiles", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/DICE-Certificate-Profiles-r01_3june2020-1.pdf>.

[TCGArch] Trusted Computing Group, "Trusted Platform Module Library - Part 1: Architecture", 8 November 2019, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_TPM2_r1p59_Part1_Architecture_pub.pdf>.

[WebAuthN] W3C, "Web Authentication: An API for accessing Public Key Credentials", n.d., <<https://www.w3.org/TR/webauthn-1/>>.

Contributors

Monty Wiseman

Email: montywiseman32@gmail.com

Liang Xia

Email: frank.xialiang@huawei.com

Laurence Lundblade

Email: lgl@island-resort.com

Eliot Lear

Email: ellear@cisco.com

Jessica Fitzgerald-McKay

Sarah C. Helbe

Andrew Guinn

Peter Loscocco

Email: pete.loscocco@gmail.com

Eric Voit

Thomas Fossati

Email: thomas.fossati@arm.com

Paul Rowe

Carsten Bormann

Email: cabo@tzi.org

Giri Mandyam

Email: mandyam@qti.qualcomm.com

Kathleen Moriarty

Email: kathleen.moriarty.ietf@gmail.com

Guy Fedorkow

Email: gfedorkow@juniper.net

Simon Frost

Email: Simon.Frost@arm.com

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Dave Thaler
Microsoft
United States of America

Email: dthaler@microsoft.com

Michael Richardson
Sandelman Software Works
Canada

Email: mcr+ietf@sandelman.ca

Ned Smith
Intel Corporation
United States of America

Email: ned.smith@intel.com

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 12 August 2022

H. Birkholz
Fraunhofer SIT
D. Thaler
Microsoft
M. Richardson
Sandelman Software Works
N. Smith
Intel
W. Pan
Huawei Technologies
8 February 2022

Remote Attestation Procedures Architecture
draft-ietf-rats-architecture-15

Abstract

In network protocol exchanges it is often useful for one end of a communication to know whether the other end is in an intended operating state. This document provides an architectural overview of the entities involved that make such tests possible through the process of generating, conveying, and evaluating evidentiary claims. An attempt is made to provide for a model that is neutral toward processor architectures, the content of claims, and protocols.

Note to Readers

Discussion of this document takes place on the RATS Working Group mailing list (rats@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/rats/> (<https://mailarchive.ietf.org/arch/browse/rats/>).

Source for this draft and an issue tracker can be found at <https://github.com/ietf-rats-wg/architecture> (<https://github.com/ietf-rats-wg/architecture>).

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 12 August 2022.

Copyright Notice

Copyright (c) 2022 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	3
2. Reference Use Cases	5
2.1. Network Endpoint Assessment	5
2.2. Confidential Machine Learning Model Protection	6
2.3. Confidential Data Protection	6
2.4. Critical Infrastructure Control	6
2.5. Trusted Execution Environment Provisioning	7
2.6. Hardware Watchdog	7
2.7. FIDO Biometric Authentication	8
3. Architectural Overview	8
3.1. Two Types of Environments of an Attester	10
3.2. Layered Attestation Environments	12
3.3. Composite Device	14
3.4. Implementation Considerations	16
4. Terminology	17
4.1. Roles	17
4.2. Artifacts	18
5. Topological Patterns	19
5.1. Passport Model	19
5.2. Background-Check Model	21
5.3. Combinations	22
6. Roles and Entities	23
7. Trust Model	24
7.1. Relying Party	24
7.2. Attester	25

7.3.	Relying Party Owner	25
7.4.	Verifier	26
7.5.	Endorser, Reference Value Provider, and Verifier Owner .	28
8.	Conceptual Messages	28
8.1.	Evidence	28
8.2.	Endorsements	29
8.3.	Reference Values	29
8.4.	Attestation Results	29
8.5.	Appraisal Policies	31
9.	Claims Encoding Formats	31
10.	Freshness	33
10.1.	Explicit Timekeeping using Synchronized Clocks	33
10.2.	Implicit Timekeeping using Nonces	34
10.3.	Implicit Timekeeping using Epoch IDs	34
10.4.	Discussion	35
11.	Privacy Considerations	36
12.	Security Considerations	37
12.1.	Attester and Attestation Key Protection	37
12.1.1.	On-Device Attester and Key Protection	37
12.1.2.	Attestation Key Provisioning Processes	38
12.2.	Integrity Protection	39
12.3.	Epoch ID-based Attestation	40
12.4.	Trust Anchor Protection	41
13.	IANA Considerations	41
14.	Acknowledgments	41
15.	Notable Contributions	42
16.	Appendix A: Time Considerations	42
16.1.	Example 1: Timestamp-based Passport Model Example . . .	44
16.2.	Example 2: Nonce-based Passport Model Example	45
16.3.	Example 3: Epoch ID-based Passport Model Example	46
16.4.	Example 4: Timestamp-based Background-Check Model Example	47
16.5.	Example 5: Nonce-based Background-Check Model Example .	48
17.	References	49
17.1.	Normative References	49
17.2.	Informative References	49
	Contributors	52
	Authors' Addresses	53

1. Introduction

The question of how one system can know that another system can be trusted has found new interest and relevance in a world where trusted computing elements are maturing in processor architectures.

Systems that have been attested and verified to be in a good state (for some value of "good") can improve overall system posture. Conversely, systems that cannot be attested and verified to be in a good state can be given reduced access or privileges, taken out of service, or otherwise flagged for repair.

For example:

- * A bank back-end system might refuse to transact with another system that is not known to be in a good state.
- * A healthcare system might refuse to transmit electronic healthcare records to a system that is not known to be in a good state.

In Remote Attestation Procedures (RATS), one peer (the "Attester") produces believable information about itself - Evidence - to enable a remote peer (the "Relying Party") to decide whether to consider that Attester a trustworthy peer or not. RATS are facilitated by an additional vital party, the Verifier.

The Verifier appraises Evidence via appraisal policies and creates the Attestation Results to support Relying Parties in their decision process. This document defines a flexible architecture consisting of attestation roles and their interactions via conceptual messages. Additionally, this document defines a universal set of terms that can be mapped to various existing and emerging Remote Attestation Procedures. Common topological patterns and the sequence of data flows associated with them, such as the "Passport Model" and the "Background-Check Model", are illustrated. The purpose is to define useful terminology for remote attestation and enable readers to map their solution architecture to the canonical attestation architecture provided here. Having a common terminology that provides well-understood meanings for common themes such as roles, device composition, topological patterns, and appraisal procedures is vital for semantic interoperability across solutions and platforms involving multiple vendors and providers.

Amongst other things, this document is about trust and trustworthiness. Trust is a choice one makes about another system. Trustworthiness is a quality about the other system that can be used in making one's decision to trust it or not. This is subtle difference and being familiar with the difference is crucial for using this document. Additionally, the concepts of freshness and trust relationships with respect to RATS are elaborated on to enable implementers to choose appropriate solutions to compose their Remote Attestation Procedures.

2. Reference Use Cases

This section covers a number of representative and generic use cases for remote attestation, independent of specific solutions. The purpose is to provide motivation for various aspects of the architecture presented in this document. Many other use cases exist, and this document does not intend to have a complete list, only to illustrate a set of use cases that collectively cover all the functionality required in the architecture.

Each use case includes a description followed by an additional summary of the Attester and Relying Party roles derived from the use case.

2.1. Network Endpoint Assessment

Network operators want trustworthy reports that include identity and version information about the hardware and software on the machines attached to their network. Examples of reports include purposes, such as inventory summaries, audit results, anomaly notifications, typically including the maintenance of log records or trend reports. The network operator may also want a policy by which full access is only granted to devices that meet some definition of hygiene, and so wants to get Claims about such information and verify its validity. Remote attestation is desired to prevent vulnerable or compromised devices from getting access to the network and potentially harming others.

Typically, a trustworthy solution starts with a specific component (sometimes referred to as a root of trust) that often provides trustworthy device identity, and performs a series of operations that enables trustworthiness appraisals for other components. Such components perform operations that help determine the trustworthiness of yet other components, by collecting, protecting or signing measurements. Measurements that have been signed by such components comprise Evidence that when evaluated either supports or refutes a claim of trustworthiness. Measurements can describe a variety of attributes of system components, such as hardware, firmware, BIOS, software, etc.

Attester: A device desiring access to a network.

Relying Party: Network equipment such as a router, switch, or access point, responsible for admission of the device into the network.

2.2. Confidential Machine Learning Model Protection

A device manufacturer wants to protect its intellectual property. The intellectual property's scope primarily encompasses the machine learning (ML) model that is deployed in the devices purchased by its customers. The protection goals include preventing attackers, potentially the customer themselves, from seeing the details of the model.

This typically works by having some protected environment in the device go through a remote attestation with some manufacturer service that can assess its trustworthiness. If remote attestation succeeds, then the manufacturer service releases either the model, or a key to decrypt a model already deployed on the Attester in encrypted form, to the requester.

Attester: A device desiring to run an ML model.

Relying Party: A server or service holding ML models it desires to protect.

2.3. Confidential Data Protection

This is a generalization of the ML model use case above, where the data can be any highly confidential data, such as health data about customers, payroll data about employees, future business plans, etc. As part of the attestation procedure, an assessment is made against a set of policies to evaluate the state of the system that is requesting the confidential data. Attestation is desired to prevent leaking data via compromised devices.

Attester: An entity desiring to retrieve confidential data.

Relying Party: An entity that holds confidential data for release to authorized entities.

2.4. Critical Infrastructure Control

Potentially harmful physical equipment (e.g., power grid, traffic control, hazardous chemical processing, etc.) is connected to a network in support of critical infrastructure. The organization managing such infrastructure needs to ensure that only authorized code and users can control corresponding critical processes, and that these processes are protected from unauthorized manipulation or other threats. When a protocol operation can affect a critical system component of the infrastructure, devices attached to that critical component require some assurances depending on the security context, including that: a requesting device or application has not been

compromised, and the requesters and actors act on applicable policies. As such, remote attestation can be used to only accept commands from requesters that are within policy.

Attester: A device or application wishing to control physical equipment.

Relying Party: A device or application connected to potentially dangerous physical equipment (hazardous chemical processing, traffic control, power grid, etc.).

2.5. Trusted Execution Environment Provisioning

A Trusted Application Manager (TAM) server is responsible for managing the applications running in a Trusted Execution Environment (TEE) of a client device, as described in [I-D.ietf-tee-architecture]. To achieve its purpose, the TAM needs to assess the state of a TEE, or of applications in the TEE, of a client device. The TEE conducts Remote Attestation Procedures with the TAM, which can then decide whether the TEE is already in compliance with the TAM's latest policy. If not, the TAM has to uninstall, update, or install approved applications in the TEE to bring it back into compliance with the TAM's policy.

Attester: A device with a TEE capable of running trusted applications that can be updated.

Relying Party: A TAM.

2.6. Hardware Watchdog

There is a class of malware that holds a device hostage and does not allow it to reboot to prevent updates from being applied. This can be a significant problem, because it allows a fleet of devices to be held hostage for ransom.

A solution to this problem is a watchdog timer implemented in a protected environment such as a Trusted Platform Module (TPM), as described in [TCGarch] section 43.3. If the watchdog does not receive regular, and fresh, Attestation Results as to the system's health, then it forces a reboot.

Attester: The device that should be protected from being held hostage for a long period of time.

Relying Party: A watchdog capable of triggering a procedure that resets a device into a known, good operational state.

2.7. FIDO Biometric Authentication

In the Fast IDentity Online (FIDO) protocol [WebAuthN], [CTAP], the device in the user's hand authenticates the human user, whether by biometrics (such as fingerprints), or by PIN and password. FIDO authentication puts a large amount of trust in the device compared to typical password authentication because it is the device that verifies the biometric, PIN and password inputs from the user, not the server. For the Relying Party to know that the authentication is trustworthy, the Relying Party needs to know that the Authenticator part of the device is trustworthy. The FIDO protocol employs remote attestation for this.

The FIDO protocol supports several remote attestation protocols and a mechanism by which new ones can be registered and added. Remote attestation defined by RATS is thus a candidate for use in the FIDO protocol.

Attester: FIDO Authenticator.

Relying Party: Any web site, mobile application back-end, or service that relies on authentication data based on biometric information.

3. Architectural Overview

Figure 1 depicts the data that flows between different roles, independent of protocol or use case.

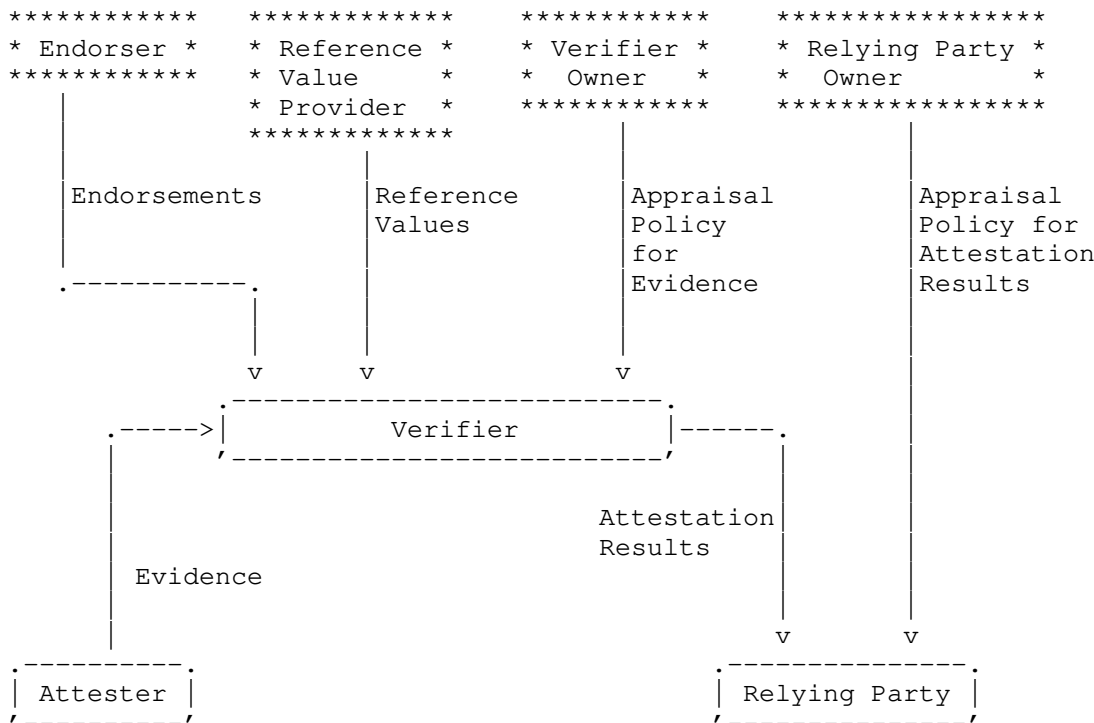


Figure 1: Conceptual Data Flow

The text below summarizes the activities conducted by the roles illustrated in Figure 1. Roles are assigned to entities. Entities are often system components [RFC4949], such as devices. As the term device is typically more intuitive than the term entity or system component, device is often used as a illustrative synonym throughout this document.

The Attester role is assigned to entities that create Evidence that is conveyed to a Verifier.

The Verifier role is assigned to entities that use the Evidence, any Reference Values from Reference Value Providers, and any Endorsements from Endorsers, by applying an Appraisal Policy for Evidence to assess the trustworthiness of the Attester. This procedure is called the appraisal of Evidence.

Subsequently, the Verifier role generates Attestation Results for use by Relying Parties.

The Appraisal Policy for Evidence might be obtained from the Verifier Owner via some protocol mechanism, or might be configured into the Verifier by the Verifier Owner, or might be programmed into the Verifier, or might be obtained via some other mechanism.

The Relying Party role is assigned to entities that uses Attestation Results by applying its own appraisal policy to make application-specific decisions, such as authorization decisions. This procedure is called the appraisal of Attestation Results.

The Appraisal Policy for Attestation Results might be obtained from the Relying Party Owner via some protocol mechanism, or might be configured into the Relying Party by the Relying Party Owner, or might be programmed into the Relying Party, or might be obtained via some other mechanism.

See Section 8 for further discussion of the conceptual messages shown in Figure 1. Section Section 4 provides a more complete definition of all RATS roles.

3.1. Two Types of Environments of an Attester

As shown in Figure 2, an Attester consists of at least one Attesting Environment and at least one Target Environment co-located in one entity. In some implementations, the Attesting and Target Environments might be combined into one environment. Other implementations might have multiple Attesting and Target Environments, such as in the examples described in more detail in Section 3.2 and Section 3.3. Other examples may exist. All compositions of Attesting and Target Environments discussed in this architecture can be combined into more complex implementations.

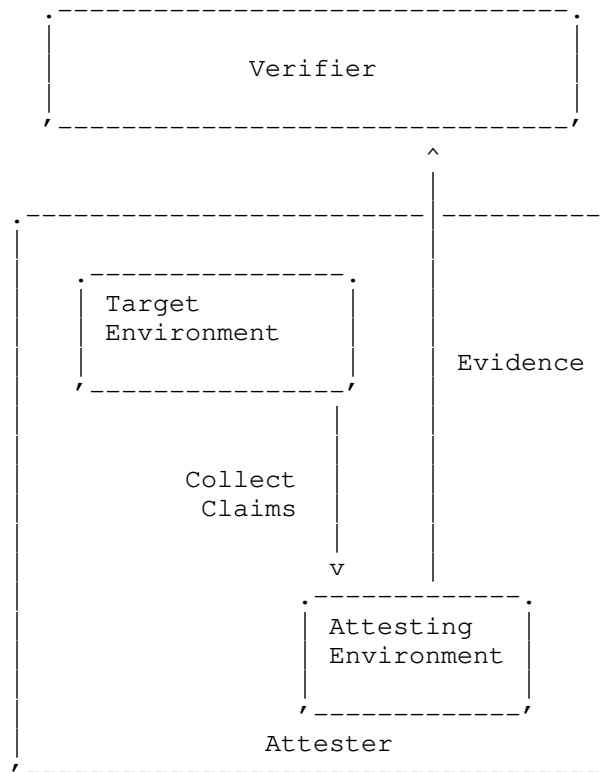


Figure 2: Two Types of Environments

Claims are collected from Target Environments. That is, Attesting Environments collect the values and the information to be represented in Claims, by reading system registers and variables, calling into subsystems, taking measurements on code, memory, or other security related assets of the Target Environment. Attesting Environments then format the Claims appropriately, and typically use key material and cryptographic functions, such as signing or cipher algorithms, to generate Evidence. There is no limit to or requirement on the types of hardware or software environments that can be used to implement an Attesting Environment, for example: Trusted Execution Environments (TEEs), embedded Secure Elements (eSEs), Trusted Platform Modules (TPMs) [TCGarch], or BIOS firmware.

An arbitrary execution environment may not, by default, be capable of Claims collection for a given Target Environment. Execution environments that are designed specifically to be capable of Claims collection are referred to in this document as Attesting Environments. For example, a TPM doesn't actively collect Claims

itself, it instead requires another component to feed various values to the TPM. Thus, an Attesting Environment in such a case would be the combination of the TPM together with whatever component is feeding it the measurements.

3.2. Layered Attestation Environments

By definition, the Attester role generates Evidence. An Attester may consist of one or more nested environments (layers). The bottom layer of an Attester has an Attesting Environment that is typically designed to be immutable or difficult to modify by malicious code. In order to appraise Evidence generated by an Attester, the Verifier needs to trust various layers, including the bottom Attesting Environment. Trust in the Attester's layers, including the bottom layer, can be established in various ways as discussed in Section 7.4.

In layered attestation, Claims can be collected from or about each layer beginning with an initial layer. The corresponding Claims can be structured in a nested fashion that reflects the nesting of the Attester's layers. Normally, Claims are not self-asserted, rather a previous layer acts as the Attesting Environment for the next layer. Claims about an initial layer typically are asserted by an Endorser.

The example device illustrated in Figure 3 includes (A) a BIOS stored in read-only memory, (B) a bootloader, and (C) an operating system kernel.

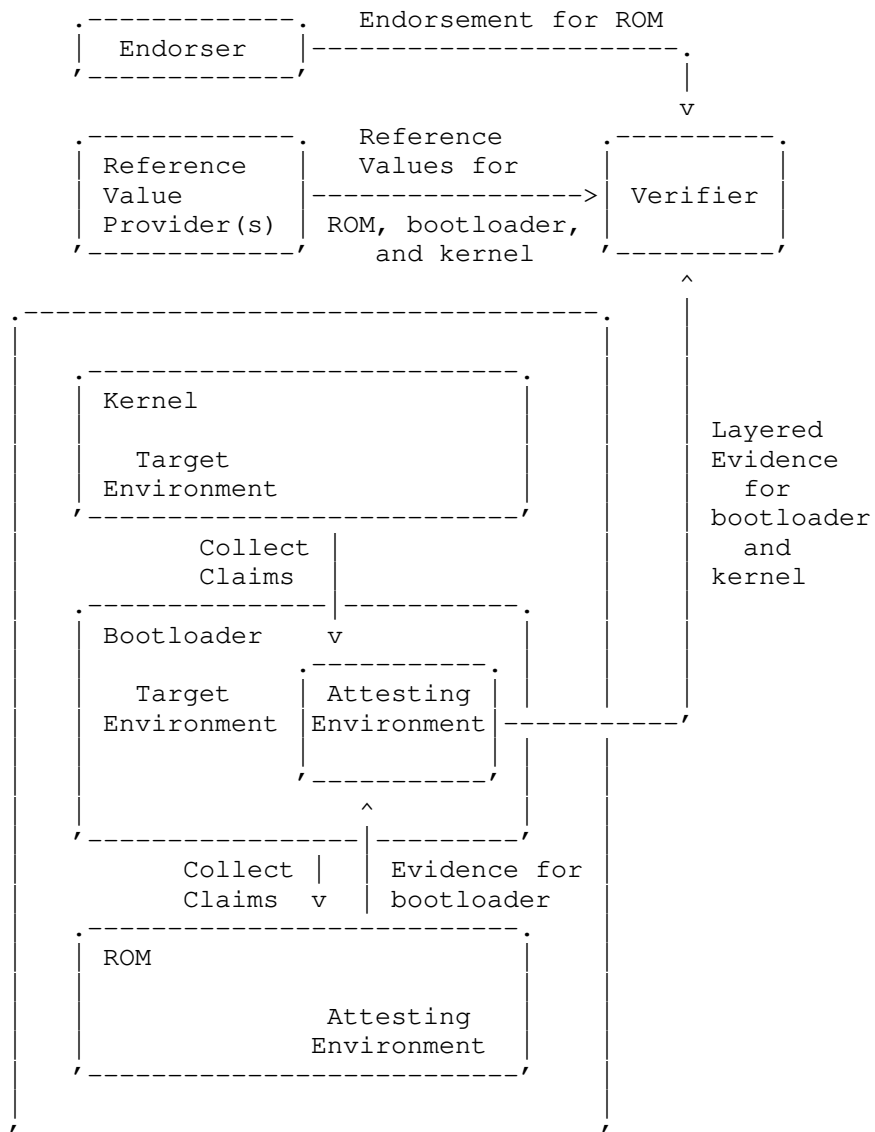


Figure 3: Layered Attester

The first Attesting Environment, the ROM in this example, has to ensure the integrity of the bootloader (the first Target Environment). There are potentially multiple kernels to boot, and the decision is up to the bootloader. Only a bootloader with intact integrity will make an appropriate decision. Therefore, the Claims relating to the integrity of the bootloader have to be measured securely. At this stage of the boot-cycle of the device, the Claims collected typically cannot be composed into Evidence.

After the boot sequence is started, the BIOS conducts the most important and defining feature of layered attestation, which is that the successfully measured bootloader now becomes (or contains) an Attesting Environment for the next layer. This procedure in layered attestation is sometimes called "staging". It is important that the bootloader not be able to alter any Claims about itself that were collected by the BIOS. This can be ensured having those Claims be either signed by the BIOS or stored in a tamper-proof manner by the BIOS.

Continuing with this example, the bootloader's Attesting Environment is now in charge of collecting Claims about the next Target Environment, which in this example is the kernel to be booted. The final Evidence thus contains two sets of Claims: one set about the bootloader as measured and signed by the BIOS, plus a set of Claims about the kernel as measured and signed by the bootloader.

This example could be extended further by making the kernel become another Attesting Environment for an application as another Target Environment. This would result in a third set of Claims in the Evidence pertaining to that application.

The essence of this example is a cascade of staged environments. Each environment has the responsibility of measuring the next environment before the next environment is started. In general, the number of layers may vary by device or implementation, and an Attesting Environment might even have multiple Target Environments that it measures, rather than only one as shown by example in Figure 3.

3.3. Composite Device

A composite device is an entity composed of multiple sub-entities such that its trustworthiness has to be determined by the appraisal of all these sub-entities.

Each sub-entity has at least one Attesting Environment collecting the Claims from at least one Target Environment, then this sub-entity generates Evidence about its trustworthiness. Therefore, each sub-

entity can be called an Attester. Among all the Attesters, there may be only some which have the ability to communicate with the Verifier while others do not.

For example, a carrier-grade router consists of a chassis and multiple slots. The trustworthiness of the router depends on all its slots' trustworthiness. Each slot has an Attesting Environment, such as a TEE, collecting the Claims of its boot process, after which it generates Evidence from the Claims.

Among these slots, only a "main" slot can communicate with the Verifier while other slots cannot. But other slots can communicate with the main slot by the links between them inside the router. So the main slot collects the Evidence of other slots, produces the final Evidence of the whole router and conveys the final Evidence to the Verifier. Therefore the router is a composite device, each slot is an Attester, and the main slot is the lead Attester.

Another example is a multi-chassis router composed of multiple single carrier-grade routers. Multi-chassis router setups create redundancy groups that provide higher throughput by interconnecting multiple routers in these groups, which can be treated as one logical router for simpler management. A multi-chassis router setup provides a management point that connects to the Verifier. Typically one router in the group is designated as the main router. Other routers in the multi-chassis setup are connected to the main router only via physical network links and are therefore managed and appraised via the main router's help. Consequently, a multi-chassis router setup is a composite device, each router is an Attester, and the main router is the lead Attester.

Figure 4 depicts the conceptual data flow for a composite device.

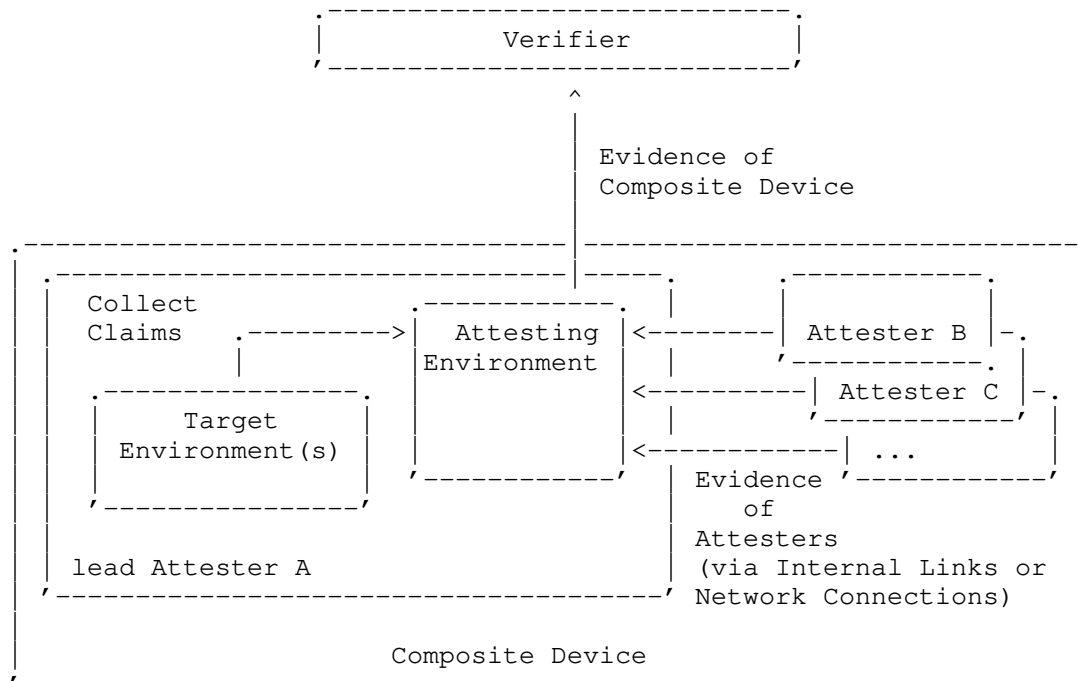


Figure 4: Composite Device

In a composite device, each Attester generates its own Evidence by its Attesting Environment(s) collecting the Claims from its Target Environment(s). The lead Attester collects Evidence from other Attesters and conveys it to a Verifier. Collection of Evidence from sub-entities may itself be a form of Claims collection that results in Evidence asserted by the lead Attester. The lead Attester generates Evidence about the layout of the whole composite device, while sub-Attesters generate Evidence about their respective (sub-)modules.

In this scenario, the trust model described in Section 7 can also be applied to an inside Verifier.

3.4. Implementation Considerations

An entity can take on multiple RATS roles (e.g., Attester, Verifier, Relying Party, etc.) at the same time. Multiple entities can cooperate to implement a single RATS role as well. In essence, the combination of roles and entities can be arbitrary. For example, in the composite device scenario, the entity inside the lead Attester can also take on the role of a Verifier, and the outer entity of

Verifier can take on the role of a Relying Party. After collecting the Evidence of other Attesters, this inside Verifier uses Endorsements and appraisal policies (obtained the same way as by any other Verifier) as part of the appraisal procedures that generate Attestation Results. The inside Verifier then conveys the Attestation Results of other Attesters to the outside Verifier, whether in the same conveyance protocol as part of the Evidence or not.

4. Terminology

This document uses the following terms.

4.1. Roles

Attester: A role performed by an entity (typically a device) whose Evidence must be appraised in order to infer the extent to which the Attester is considered trustworthy, such as when deciding whether it is authorized to perform some operation.

Produces: Evidence

Relying Party: A role performed by an entity that depends on the validity of information about an Attester, for purposes of reliably applying application specific actions. Compare /relying party/ in [RFC4949].

Consumes: Attestation Results, Appraisal Policy for Attestation Results

Verifier: A role performed by an entity that appraises the validity of Evidence about an Attester and produces Attestation Results to be used by a Relying Party.

Consumes: Evidence, Reference Values, Endorsements, Appraisal Policy for Evidence

Produces: Attestation Results

Relying Party Owner: A role performed by an entity (typically an administrator), that is authorized to configure Appraisal Policy for Attestation Results in a Relying Party.

Produces: Appraisal Policy for Attestation Results

Verifier Owner: A role performed by an entity (typically an administrator), that is authorized to configure Appraisal Policy for Evidence in a Verifier.

Produces: Appraisal Policy for Evidence

Endorser: A role performed by an entity (typically a manufacturer) whose Endorsements may help Verifiers appraise the authenticity of Evidence and infer further capabilities of the Attester.

Produces: Endorsements

Reference Value Provider: A role performed by an entity (typically a manufacturer) whose Reference Values help Verifiers appraise Evidence to determine if acceptable known Claims have been recorded by the Attester.

Produces: Reference Values

4.2. Artifacts

Claim: A piece of asserted information, often in the form of a name/value pair. Claims make up the usual structure of Evidence and other RATS artifacts. Compare /claim/ in [RFC7519].

Endorsement: A secure statement that an Endorser vouches for the integrity of an Attester's various capabilities such as Claims collection and Evidence signing.

Consumed By: Verifier

Produced By: Endorser

Evidence: A set of Claims generated by an Attester to be appraised by a Verifier. Evidence may include configuration data, measurements, telemetry, or inferences.

Consumed By: Verifier

Produced By: Attester

Attestation Result: The output generated by a Verifier, typically including information about an Attester, where the Verifier vouches for the validity of the results.

Consumed By: Relying Party

Produced By: Verifier

Appraisal Policy for Evidence: A set of rules that informs how a Verifier evaluates the validity of information about an Attester. Compare /security policy/ in [RFC4949].

Consumed By: Verifier

Produced By: Verifier Owner

Appraisal Policy for Attestation Results: A set of rules that direct how a Relying Party uses the Attestation Results regarding an Attester generated by the Verifiers. Compare /security policy/ in [RFC4949].

Consumed by: Relying Party

Produced by: Relying Party Owner

Reference Values: A set of values against which values of Claims can be compared as part of applying an Appraisal Policy for Evidence. Reference Values are sometimes referred to in other documents as known-good values, golden measurements, or nominal values, although those terms typically assume comparison for equality, whereas here Reference Values might be more general and be used in any sort of comparison.

Consumed By: Verifier

Produced By: Reference Value Provider

5. Topological Patterns

Figure 1 shows a data-flow diagram for communication between an Attester, a Verifier, and a Relying Party. The Attester conveys its Evidence to the Verifier for appraisal, and the Relying Party receives the Attestation Result from the Verifier. This section refines the data-flow diagram by describing two reference models, as well as one example composition thereof. The discussion that follows is for illustrative purposes only and does not constrain the interactions between RATS roles to the presented patterns.

5.1. Passport Model

The passport model is so named because of its resemblance to how nations issue passports to their citizens. The nature of the Evidence that an individual needs to provide to its local authority is specific to the country involved. The citizen retains control of the resulting passport document and presents it to other entities when it needs to assert a citizenship or identity Claim, such as an airport immigration desk. The passport is considered sufficient because it vouches for the citizenship and identity Claims, and it is issued by a trusted authority. Thus, in this immigration desk analogy, the citizen is the Attester, the passport issuing agency is

a Verifier, the passport application and identifying information (e.g., birth certificate) is the Evidence, the passport is an Attestation Result, and the immigration desk is a Relying Party.

In this model, an Attester conveys Evidence to a Verifier, which compares the Evidence against its appraisal policy. The Verifier then gives back an Attestation Result which the Attester treats as opaque data. The Attester does not consume the Attestation Result, but might cache it. The Attester can then present the Attestation Result (and possibly additional Claims) to a Relying Party, which then compares this information against its own appraisal policy.

Three ways in which the process may fail include:

- * First, the Verifier may not issue a positive Attestation Result due to the Evidence not passing the Appraisal Policy for Evidence.
- * The second way in which the process may fail is when the Attestation Result is examined by the Relying Party, and based upon the Appraisal Policy for Attestation Results, the result does not pass the policy.
- * The third way is when the Verifier is unreachable or unavailable.

As with any other information needed by the Relying Party to make an authorization decision, an Attestation Result can be carried in a resource access protocol between the Attester and Relying Party. In this model the details of the resource access protocol constrain the serialization format of the Attestation Result. The format of the Evidence on the other hand is only constrained by the Attester-Verifier remote attestation protocol. This implies that interoperability and standardization is more relevant for Attestation Results than it is for Evidence.

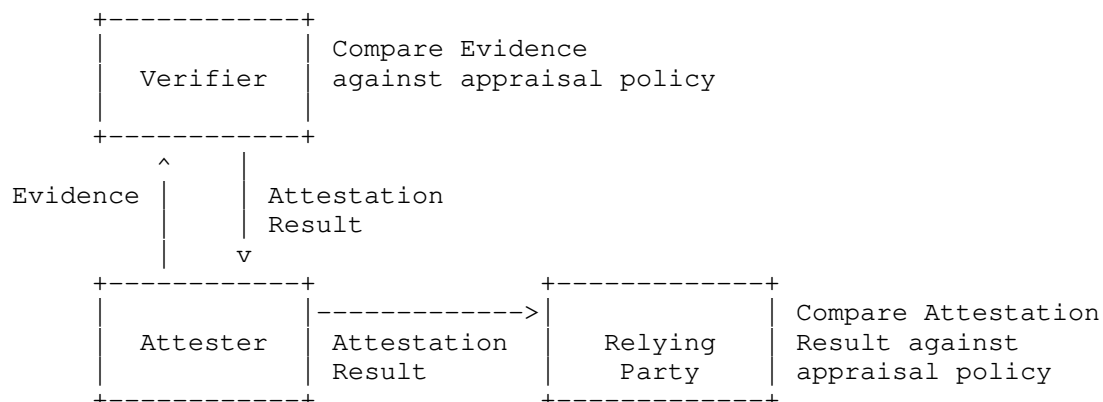


Figure 5: Passport Model

5.2. Background-Check Model

The background-check model is so named because of the resemblance of how employers and volunteer organizations perform background checks. When a prospective employee provides Claims about education or previous experience, the employer will contact the respective institutions or former employers to validate the Claim. Volunteer organizations often perform police background checks on volunteers in order to determine the volunteer's trustworthiness. Thus, in this analogy, a prospective volunteer is an Attester, the organization is the Relying Party, and the organization that issues a report is a Verifier.

In this model, an Attester conveys Evidence to a Relying Party, which treats it as opaque and simply forwards it on to a Verifier. The Verifier compares the Evidence against its appraisal policy, and returns an Attestation Result to the Relying Party. The Relying Party then compares the Attestation Result against its own appraisal policy.

The resource access protocol between the Attester and Relying Party includes Evidence rather than an Attestation Result, but that Evidence is not processed by the Relying Party. Since the Evidence is merely forwarded on to a trusted Verifier, any serialization format can be used for Evidence because the Relying Party does not need a parser for it. The only requirement is that the Evidence can be encapsulated in the format required by the resource access protocol between the Attester and Relying Party.

However, like in the Passport model, an Attestation Result is still consumed by the Relying Party. Code footprint and attack surface area can be minimized by using a serialization format for which the Relying Party already needs a parser to support the protocol between the Attester and Relying Party, which may be an existing standard or widely deployed resource access protocol. Such minimization is especially important if the Relying Party is a constrained node.

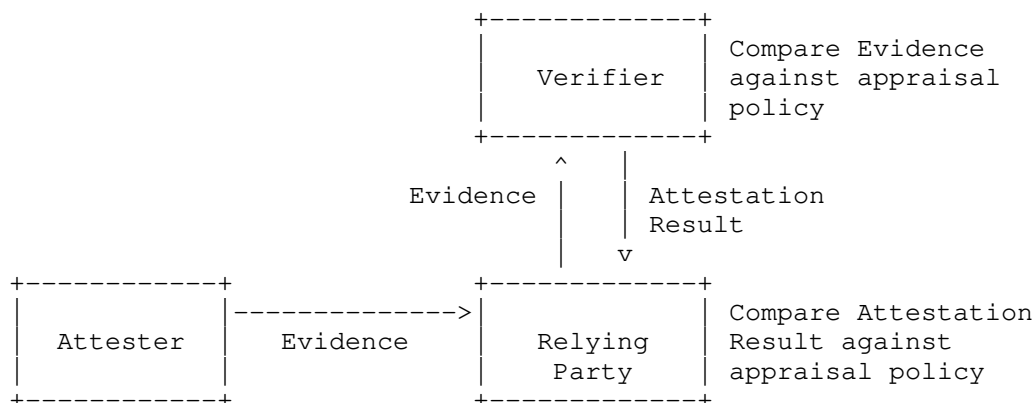


Figure 6: Background-Check Model

5.3. Combinations

One variation of the background-check model is where the Relying Party and the Verifier are on the same machine, performing both functions together. In this case, there is no need for a protocol between the two.

It is also worth pointing out that the choice of model depends on the use case, and that different Relying Parties may use different topological patterns.

The same device may need to create Evidence for different Relying Parties and/or different use cases. For instance, it would use one model to provide Evidence to a network infrastructure device to gain access to the network, and the other model to provide Evidence to a server holding confidential data to gain access to that data. As such, both models may simultaneously be in use by the same device.

Figure 7 shows another example of a combination where Relying Party 1 uses the passport-check model, whereas Relying Party 2 uses an extension of the background-check model. Specifically, in addition to the basic functionality shown in Figure 6, Relying Party 2 actually provides the Attestation Result back to the Attester, allowing the Attester to use it with other Relying Parties. This is the model that the Trusted Application Manager plans to support in the TEEP architecture [I-D.ietf-teep-architecture].

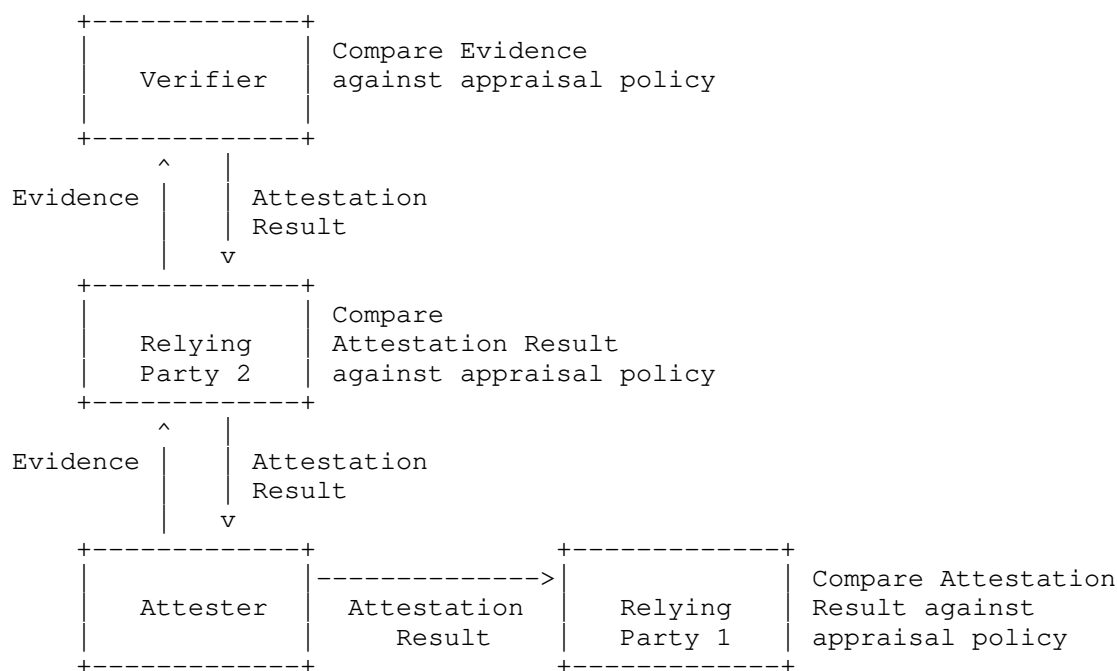


Figure 7: Example Combination

6. Roles and Entities

An entity in the RATS architecture includes at least one of the roles defined in this document.

An entity can aggregate more than one role into itself, such as being both a Verifier and a Relying Party, or being both a Reference Value Provider and an Endorser. As such, any conceptual messages (see Section 8 for more discussion) originating from such roles might also be combined. For example, Reference Values might be conveyed as part of an appraisal policy if the Verifier Owner and Reference Value Provider roles are combined. Similarly, Reference Values might be conveyed as part of an Endorsement if the Endorser and Reference Value Provider roles are combined.

Interactions between roles aggregated into the same entity do not necessarily use the Internet Protocol. Such interactions might use a loopback device or other IP-based communication between separate environments, but they do not have to. Alternative channels to convey conceptual messages include function calls, sockets, GPIO interfaces, local busses, or hypervisor calls. This type of conveyance is typically found in composite devices. Most

importantly, these conveyance methods are out-of-scope of RATS, but they are presumed to exist in order to convey conceptual messages appropriately between roles.

In essence, an entity that combines more than one role creates and consumes the corresponding conceptual messages as defined in this document.

7. Trust Model

7.1. Relying Party

This document covers scenarios for which a Relying Party trusts a Verifier that can appraise the trustworthiness of information about an Attester. Such trust is expressed by storing one or more "trust anchors" in a secure location known as a trust anchor store.

As defined in [RFC6024], "A trust anchor represents an authoritative entity via a public key and associated data. The public key is used to verify digital signatures, and the associated data is used to constrain the types of information for which the trust anchor is authoritative." The trust anchor may be a certificate or it may be a raw public key along with additional data if necessary such as its public key algorithm and parameters. In the context of this document, a trust anchor may also be a symmetric key, as in [TCG-DICE-SIBDA] or the symmetric mode described in [I-D.tschofenig-rats-psa-token].

Thus, trusting a Verifier might be expressed by having the Relying Party store the Verifier's key or certificate in its trust anchor store, or might be expressed by storing the public key or certificate of an entity (e.g., a Certificate Authority) that is in the Verifier's certificate path. For example, the Relying Party can verify that the Verifier is an expected one by out of band establishment of key material, combined with a protocol like TLS to communicate. There is an assumption that between the establishment of the trusted key material and the creation of the Evidence, that the Verifier has not been compromised.

For a stronger level of security, the Relying Party might require that the Verifier first provide information about itself that the Relying Party can use to assess the trustworthiness of the Verifier before accepting its Attestation Results. Such process would provide a stronger level of confidence in the correctness of the information provided, such as a belief that the authentic Verifier has not been compromised by malware.

For example, one explicit way for a Relying Party "A" to establish such confidence in the correctness of a Verifier "B", would be for B to first act as an Attester where A acts as a combined Verifier/Relying Party. If A then accepts B as trustworthy, it can choose to accept B as a Verifier for other Attesters.

Similarly, the Relying Party also needs to trust the Relying Party Owner for providing its Appraisal Policy for Attestation Results, and in some scenarios the Relying Party might even require that the Relying Party Owner go through a remote attestation procedure with it before the Relying Party will accept an updated policy. This can be done similarly to how a Relying Party could establish trust in a Verifier as discussed above, i.e., verifying credentials against a trust anchor store and optionally requiring Attestation Results from the Relying Party Owner.

7.2. Attester

In some scenarios, Evidence might contain sensitive information such as Personally Identifiable Information (PII) or system identifiable information. Thus, an Attester must trust entities to which it conveys Evidence, to not reveal sensitive data to unauthorized parties. The Verifier might share this information with other authorized parties, according to a governing policy that address the handling of sensitive information (potentially included in Appraisal Policies for Evidence). In the background-check model, this Evidence may also be revealed to Relying Party(s).

When Evidence contains sensitive information, an Attester typically requires that a Verifier authenticates itself (e.g., at TLS session establishment) and might even request a remote attestation before the Attester sends the sensitive Evidence. This can be done by having the Attester first act as a Verifier/Relying Party, and the Verifier act as its own Attester, as discussed above.

7.3. Relying Party Owner

The Relying Party Owner might also require that the Relying Party first act as an Attester, providing Evidence that the Owner can appraise, before the Owner would give the Relying Party an updated policy that might contain sensitive information. In such a case, authentication or attestation in both directions might be needed, in which case typically one side's Evidence must be considered safe to share with an untrusted entity, in order to bootstrap the sequence. See Section 11 for more discussion.

7.4. Verifier

The Verifier trusts (or more specifically, the Verifier's security policy is written in a way that configures the Verifier to trust) a manufacturer, or the manufacturer's hardware, so as to be able to appraise the trustworthiness of that manufacturer's devices. Such trust is expressed by storing one or more trust anchors in the Verifier's trust anchor store.

In a typical solution, a Verifier comes to trust an Attester indirectly by having an Endorser (such as a manufacturer) vouch for the Attester's ability to securely generate Evidence through Endorsements (see Section 8.2). Endorsements might describe the ways in which the Attester resists attack, protects secrets and measures Target Environments. Consequently, the Endorser's key material is stored in the Verifier's trust anchor store so that Endorsements can be authenticated and used in the Verifier's appraisal process.

In some solutions, a Verifier might be configured to directly trust an Attester by having the Verifier have the Attester's key material (rather than the Endorser's) in its trust anchor store.

Such direct trust must first be established at the time of trust anchor store configuration either by checking with an Endorser at that time, or by conducting a security analysis of the specific device. Having the Attester directly in the trust anchor store narrows the Verifier's trust to only specific devices rather than all devices the Endorser might vouch for, such as all devices manufactured by the same manufacturer in the case that the Endorser is a manufacturer.

Such narrowing is often important since physical possession of a device can also be used to conduct a number of attacks, and so a device in a physically secure environment (such as one's own premises) may be considered trusted whereas devices owned by others would not be. This often results in a desire to either have the owner run their own Endorser that would only endorse devices one owns, or to use Attesters directly in the trust anchor store. When there are many Attesters owned, the use of an Endorser enables better scalability.

That is, a Verifier might appraise the trustworthiness of an application component, operating system component, or service under the assumption that information provided about it by the lower-layer firmware or software is true. A stronger level of assurance of security comes when information can be vouched for by hardware or by ROM code, especially if such hardware is physically resistant to hardware tampering. In most cases, components that have to be vouched for via Endorsements because no Evidence is generated about them are referred to as roots of trust.

The manufacturer having arranged for an Attesting Environment to be provisioned with key material with which to sign Evidence, the Verifier is then provided with some way of verifying the signature on the Evidence. This may be in the form of an appropriate trust anchor, or the Verifier may be provided with a database of public keys (rather than certificates) or even carefully curated and secured lists of symmetric keys.

The nature of how the Verifier manages to validate the signatures produced by the Attester is critical to the secure operation of a remote attestation system, but is not the subject of standardization within this architecture.

A conveyance protocol that provides authentication and integrity protection can be used to convey Evidence that is otherwise unprotected (e.g., not signed). Appropriate conveyance of unprotected Evidence (e.g., [I-D.birkholz-rats-uccs]) relies on the following conveyance protocol's protection capabilities:

1. The key material used to authenticate and integrity protect the conveyance channel is trusted by the Verifier to speak for the Attesting Environment(s) that collected Claims about the Target Environment(s).
2. All unprotected Evidence that is conveyed is supplied exclusively by the Attesting Environment that has the key material that protects the conveyance channel
3. A trusted environment protects the conveyance channel's key material which may depend on other Attesting Environments with equivalent strength protections.

As illustrated in [I-D.birkholz-rats-uccs], an entity that receives unprotected Evidence via a trusted conveyance channel always takes on the responsibility of vouching for the Evidence's authenticity and freshness. If protected Evidence is generated, the Attester's Attesting Environments take on that responsibility. In cases where unprotected Evidence is processed by a Verifier, Relying Parties have

to trust that the Verifier is capable of handling Evidence in a manner that preserves the Evidence's authenticity and freshness. Generating and conveying unprotected Evidence always creates significant risk and the benefits of that approach have to be carefully weighed against potential drawbacks.

See Section 12 for discussion on security strength.

7.5. Endorser, Reference Value Provider, and Verifier Owner

In some scenarios, the Endorser, Reference Value Provider, and Verifier Owner may need to trust the Verifier before giving the Endorsement, Reference Values, or appraisal policy to it. This can be done similarly to how a Relying Party might establish trust in a Verifier.

As discussed in Section 7.3, authentication or attestation in both directions might be needed, in which case typically one side's identity or Evidence must be considered safe to share with an untrusted entity, in order to bootstrap the sequence. See Section 11 for more discussion.

8. Conceptual Messages

Figure 1 illustrates the flow of a conceptual messages between various roles. This section provides additional elaboration and implementation considerations. It is the responsibility of protocol specifications to define the actual data format and semantics of any relevant conceptual messages.

8.1. Evidence

Evidence is a set of Claims about the target environment that reveal operational status, health, configuration or construction that have security relevance. Evidence is appraised by a Verifier to establish its relevance, compliance, and timeliness. Claims need to be collected in a manner that is reliable such that a Target Environment cannot lie to the Attesting Environment about its trustworthiness properties. Evidence needs to be securely associated with the target environment so that the Verifier cannot be tricked into accepting Claims originating from a different environment (that may be more trustworthy). Evidence also must be protected from man-in-the-middle attackers who may observe, change or misdirect Evidence as it travels from Attester to Verifier. The timeliness of Evidence can be captured using Claims that pinpoint the time or interval when changes in operational status, health, and so forth occur.

8.2. Endorsements

An Endorsement is a secure statement that some entity (e.g., a manufacturer) vouches for the integrity of the device's various capabilities such as claims collection, signing, launching code, transitioning to other environments, storing secrets, and more. For example, if the device's signing capability is in hardware, then an Endorsement might be a manufacturer certificate that signs a public key whose corresponding private key is only known inside the device's hardware. Thus, when Evidence and such an Endorsement are used together, an appraisal procedure can be conducted based on appraisal policies that may not be specific to the device instance, but merely specific to the manufacturer providing the Endorsement. For example, an appraisal policy might simply check that devices from a given manufacturer have information matching a set of Reference Values, or an appraisal policy might have a set of more complex logic on how to appraise the validity of information.

However, while an appraisal policy that treats all devices from a given manufacturer the same may be appropriate for some use cases, it would be inappropriate to use such an appraisal policy as the sole means of authorization for use cases that wish to constrain which compliant devices are considered authorized for some purpose. For example, an enterprise using remote attestation for Network Endpoint Assessment [RFC5209] may not wish to let every healthy laptop from the same manufacturer onto the network, but instead only want to let devices that it legally owns onto the network. Thus, an Endorsement may be helpful information in authenticating information about a device, but is not necessarily sufficient to authorize access to resources which may need device-specific information such as a public key for the device or component or user on the device.

8.3. Reference Values

Reference Values used in appraisal procedures come from a Reference Value Provider and are then used by the Verifier to compare to Evidence. Reference Values with matching Evidence produces acceptable Claims. Additionally, appraisal policy may play a role in determining the acceptance of Claims.

8.4. Attestation Results

Attestation Results are the input used by the Relying Party to decide the extent to which it will trust a particular Attester, and allow it to access some data or perform some operation.

Attestation Results may carry a boolean value indicating compliance or non-compliance with a Verifier's appraisal policy, or may carry a richer set of Claims about the Attester, against which the Relying Party applies its Appraisal Policy for Attestation Results.

The quality of the Attestation Results depends upon the ability of the Verifier to evaluate the Attester. Different Attesters have a different `_Strength of Function_` [strengthoffunction], which results in the Attestation Results being qualitatively different in strength.

An Attestation Result that indicates non-compliance can be used by an Attester (in the passport model) or a Relying Party (in the background-check model) to indicate that the Attester should not be treated as authorized and may be in need of remediation. In some cases, it may even indicate that the Evidence itself cannot be authenticated as being correct.

By default, the Relying Party does not believe the Attester to be compliant. Upon receipt of an authentic Attestation Result and given the Appraisal Policy for Attestation Results is satisfied, the Attester is allowed to perform the prescribed actions or access. The simplest such appraisal policy might authorize granting the Attester full access or control over the resources guarded by the Relying Party. A more complex appraisal policy might involve using the information provided in the Attestation Result to compare against expected values, or to apply complex analysis of other information contained in the Attestation Result.

Thus, Attestation Results can contain detailed information about an Attester, which can include privacy sensitive information as discussed in section Section 11. Unlike Evidence, which is often very device- and vendor-specific, Attestation Results can be vendor-neutral, if the Verifier has a way to generate vendor-agnostic information based on the appraisal of vendor-specific information in Evidence. This allows a Relying Party's appraisal policy to be simpler, potentially based on standard ways of expressing the information, while still allowing interoperability with heterogeneous devices.

Finally, whereas Evidence is signed by the device (or indirectly by a manufacturer, if Endorsements are used), Attestation Results are signed by a Verifier, allowing a Relying Party to only need a trust relationship with one entity, rather than a larger set of entities, for purposes of its appraisal policy.

8.5. Appraisal Policies

The Verifier, when appraising Evidence, or the Relying Party, when appraising Attestation Results, checks the values of matched Claims against constraints specified in its appraisal policy. Examples of such constraints checking include:

- * comparison for equality against a Reference Value, or
- * a check for being in a range bounded by Reference Values, or
- * membership in a set of Reference Values, or
- * a check against values in other Claims.

Upon completing all appraisal policy constraints, the remaining Claims are accepted as input toward determining Attestation Results, when appraising Evidence, or as input to a Relying Party, when appraising Attestation Results.

9. Claims Encoding Formats

The following diagram illustrates a relationship to which remote attestation is desired to be added:

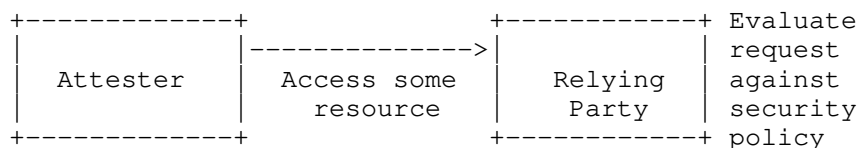


Figure 8: Typical Resource Access

In this diagram, the protocol between Attester and a Relying Party can be any new or existing protocol (e.g., HTTP(S), COAP(S), ROLIE [RFC8322], 802.1x, OPC UA [OPCUA], etc.), depending on the use case.

Typically, such protocols already have mechanisms for passing security information for authentication and authorization purposes. Common formats include JWTs [RFC7519], CWTs [RFC8392], and X.509 certificates.

Retrofitting already deployed protocols with remote attestation requires adding RATS conceptual messages to the existing data flows. This must be done in a way that does not degrade the security properties of the systems involved and should use native extension mechanisms provided by the underlying protocol. For example, if a TLS handshake is to be extended with remote attestation capabilities,

attestation Evidence may be embedded in an ad-hoc X.509 certificate extension (e.g., [TCG-DICE]), or into a new TLS Certificate Type (e.g., [I-D.tschofenig-tls-cwt]).

Especially for constrained nodes there is a desire to minimize the amount of parsing code needed in a Relying Party, in order to both minimize footprint and to minimize the attack surface. While it would be possible to embed a CWT inside a JWT, or a JWT inside an X.509 extension, etc., there is a desire to encode the information natively in a format that is already supported by the Relying Party.

This motivates having a common "information model" that describes the set of remote attestation related information in an encoding-agnostic way, and allowing multiple encoding formats (CWT, JWT, X.509, etc.) that encode the same information into the Claims format needed by the Relying Party.

The following diagram illustrates that Evidence and Attestation Results might be expressed via multiple potential encoding formats, so that they can be conveyed by various existing protocols. It also motivates why the Verifier might also be responsible for accepting Evidence that encodes Claims in one format, while issuing Attestation Results that encode Claims in a different format.

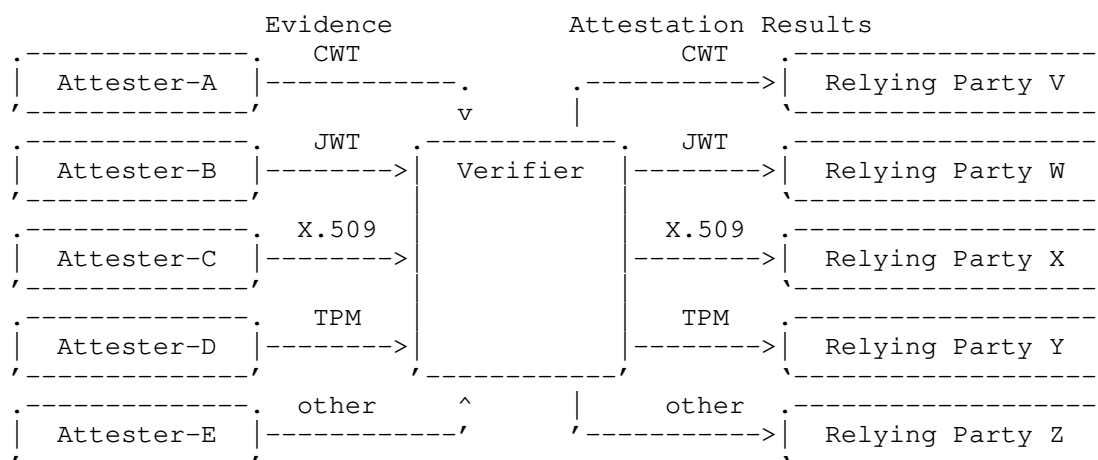


Figure 9: Multiple Attesters and Relying Parties with Different Formats

10. Freshness

A Verifier or Relying Party might need to learn the point in time (i.e., the "epoch") an Evidence or Attestation Result has been produced. This is essential in deciding whether the included Claims can be considered fresh, meaning they still reflect the latest state of the Attester, and that any Attestation Result was generated using the latest Appraisal Policy for Evidence.

This section provides a number of details. It does not however define any protocol formats, the interactions shown are abstract. This section is intended for those creating protocols and solutions to understand the options available to ensure freshness. The way in which freshness is provisioned in a protocol is an architectural decision. Provisioning of freshness has an impact on the number of needed round trips in a protocol, and therefore must be made very early in the design. Different decisions will have significant impacts on resulting interoperability, which is why this section goes into sufficient detail such that choices in freshness will be compatible across interacting protocols, such as depicted in Figure 9.

Freshness is assessed based on the Appraisal Policy for Evidence or Attestation Results that compares the estimated epoch against an "expiry" threshold defined locally to that policy. There is, however, always a race condition possible in that the state of the Attester, and the appraisal policies might change immediately after the Evidence or Attestation Result was generated. The goal is merely to narrow their recentness to something the Verifier (for Evidence) or Relying Party (for Attestation Result) is willing to accept. Some flexibility on the freshness requirement is a key component for enabling caching and reuse of both Evidence and Attestation Results, which is especially valuable in cases where their computation uses a substantial part of the resource budget (e.g., energy in constrained devices).

There are three common approaches for determining the epoch of Evidence or an Attestation Result.

10.1. Explicit Timekeeping using Synchronized Clocks

The first approach is to rely on synchronized and trustworthy clocks, and include a signed timestamp (see [I-D.birkholz-rats-tuda]) along with the Claims in the Evidence or Attestation Result. Timestamps can also be added on a per-Claim basis to distinguish the time of generation of Evidence or Attestation Result from the time that a specific Claim was generated. The clock's trustworthiness can generally be established via Endorsements and typically requires

additional Claims about the signer's time synchronization mechanism.

In some use cases, however, a trustworthy clock might not be available. For example, in many Trusted Execution Environments (TEEs) today, a clock is only available outside the TEE and so cannot be trusted by the TEE.

10.2. Implicit Timekeeping using Nonces

A second approach places the onus of timekeeping solely on the Verifier (for Evidence) or the Relying Party (for Attestation Results), and might be suitable, for example, in case the Attester does not have a trustworthy clock or time synchronization is otherwise impaired. In this approach, a non-predictable nonce is sent by the appraising entity, and the nonce is then signed and included along with the Claims in the Evidence or Attestation Result. After checking that the sent and received nonces are the same, the appraising entity knows that the Claims were signed after the nonce was generated. This allows associating a "rough" epoch to the Evidence or Attestation Result. In this case the epoch is said to be rough because:

- * The epoch applies to the entire Claim set instead of a more granular association, and
- * The time between the creation of Claims and the collection of Claims is indistinguishable.

10.3. Implicit Timekeeping using Epoch IDs

A third approach relies on having epoch identifiers (or "IDs") periodically sent to both the sender and receiver of Evidence or Attestation Results by some "Epoch ID Distributor".

Epoch IDs are different from nonces as they can be used more than once and can even be used by more than one entity at the same time. Epoch IDs are different from timestamps as they do not have to convey information about a point in time, i.e., they are not necessarily monotonically increasing integers.

Like the nonce approach, this allows associating a "rough" epoch without requiring a trustworthy clock or time synchronization in order to generate or appraise the freshness of Evidence or Attestation Results. Only the Epoch ID Distributor requires access to a clock so it can periodically send new epoch IDs.

The most recent epoch ID is included in the produced Evidence or Attestation Results, and the appraising entity can compare the epoch ID in received Evidence or Attestation Results against the latest epoch ID it received from the Epoch ID Distributor to determine if it is within the current epoch. An actual solution also needs to take into account race conditions when transitioning to a new epoch, such as by using a counter signed by the Epoch ID Distributor as the epoch ID, or by including both the current and previous epoch IDs in messages and/or checks, by requiring retries in case of mismatching epoch IDs, or by buffering incoming messages that might be associated with a epoch ID that the receiver has not yet obtained.

More generally, in order to prevent an appraising entity from generating false negatives (e.g., discarding Evidence that is deemed stale even if it is not), the appraising entity should keep an "epoch window" consisting of the most recently received epoch IDs. The depth of such epoch window is directly proportional to the maximum network propagation delay between the first to receive the epoch ID and the last to receive the epoch ID, and it is inversely proportional to the epoch duration. The appraising entity shall compare the epoch ID carried in the received Evidence or Attestation Result with the epoch IDs in its epoch window to find a suitable match.

Whereas the nonce approach typically requires the appraising entity to keep state for each nonce generated, the epoch ID approach minimizes the state kept to be independent of the number of Attesters or Verifiers from which it expects to receive Evidence or Attestation Results, as long as all use the same Epoch ID Distributor.

10.4. Discussion

Implicit and explicit timekeeping can be combined into hybrid mechanisms. For example, if clocks exist and are considered trustworthy but are not synchronized, a nonce-based exchange may be used to determine the (relative) time offset between the involved peers, followed by any number of timestamp based exchanges.

It is important to note that the actual values in Claims might have been generated long before the Claims are signed. If so, it is the signer's responsibility to ensure that the values are still correct when they are signed. For example, values generated at boot time might have been saved to secure storage until network connectivity is established to the remote Verifier and a nonce is obtained.

A more detailed discussion with examples appears in Section 16.

For a discussion on the security of epoch IDs see Section 12.3.

11. Privacy Considerations

The conveyance of Evidence and the resulting Attestation Results reveal a great deal of information about the internal state of a device as well as potentially any users of the device.

In many cases, the whole point of attestation procedures is to provide reliable information about the type of the device and the firmware/software that the device is running.

This information might be particularly interesting to many attackers. For example, knowing that a device is running a weak version of firmware provides a way to aim attacks better.

In some circumstances, if an attacker can become aware of Endorsements, Reference Values, or appraisal policies, it could potentially provide an attacker with insight into defensive mitigations. It is recommended that attention be paid to confidentiality of such information.

Additionally, many Claims in Evidence, many Claims in Attestation Results, and appraisal policies potentially contain Personally Identifying Information (PII) depending on the end-to-end use case of the remote attestation procedure. Remote attestation that includes containers and applications, e.g., a blood pressure monitor, may further reveal details about specific systems or users.

In some cases, an attacker may be able to make inferences about the contents of Evidence from the resulting effects or timing of the processing. For example, an attacker might be able to infer the value of specific Claims if it knew that only certain values were accepted by the Relying Party.

Conceptual messages (see Section 8) carrying sensitive or confidential information are expected to be integrity protected (i.e., either via signing or a secure channel) and optionally might be confidentiality protected via encryption. If there isn't confidentiality protection of conceptual messages themselves, the underlying conveyance protocol should provide these protections.

As Evidence might contain sensitive or confidential information, Attesters are responsible for only sending such Evidence to trusted Verifiers. Some Attesters might want a stronger level of assurance of the trustworthiness of a Verifier before sending Evidence to it. In such cases, an Attester can first act as a Relying Party and ask for the Verifier's own Attestation Result, and appraising it just as a Relying Party would appraise an Attestation Result for any other purpose.

Another approach to deal with Evidence is to remove PII from the Evidence while still being able to verify that the Attester is one of a large set. This approach is often called "Direct Anonymous Attestation". See [CCC-DeepDive] section 6.2 and [I-D.ietf-rats-daa] for more discussion.

12. Security Considerations

This document provides an architecture for doing remote attestation. No specific wire protocol is documented here. Without a specific proposal to compare against, it is impossible to know if the security threats listed below have been mitigated well.

The security considerations below should be read as being essentially requirements against realizations of the RATS Architecture. Some threats apply to protocols, some are against implementations (code), and some threats are against physical infrastructure (such as factories).

The fundamental purpose of the RATS architecture is to allow a Relying Party to establish a basis for trusting the Attester.

12.1. Attester and Attestation Key Protection

Implementers need to pay close attention to the protection of the Attester and the manufacturing processes for provisioning attestation key material. If either of these are compromised, intended levels of assurance for RATS are compromised because attackers can forge Evidence or manipulate the Attesting Environment. For example, a Target Environment should not be able to tamper with the Attesting Environment that measures it, by isolating the two environments from each other in some way.

Remote attestation applies to use cases with a range of security requirements, so the protections discussed here range from low to high security where low security may be limited to application or process isolation by the device's operating system, and high security may involve specialized hardware to defend against physical attacks on a chip.

12.1.1. On-Device Attester and Key Protection

It is assumed that an Attesting Environment is sufficiently isolated from the Target Environment it collects Claims about and that it signs the resulting Claims set with an attestation key, so that the Target Environment cannot forge Evidence about itself. Such an isolated environment might be provided by a process, a dedicated chip, a TEE, a virtual machine, or another secure mode of operation.

The Attesting Environment must be protected from unauthorized modification to ensure it behaves correctly. Confidentiality protection of the Attesting Environment's signing key is vital so it cannot be misused to forge Evidence.

In many cases the user or owner of a device that includes the role of Attester must not be able to modify or extract keys from the Attesting Environments, to prevent creating forged Evidence. Some common examples include the user of a mobile phone or FIDO authenticator.

Measures for a minimally protected system might include process or application isolation provided by a high-level operating system, and restricted access to root or system privileges. In contrast, For really simple single-use devices that don't use a protected mode operating system, like a Bluetooth speaker, the only factual isolation might be the sturdy housing of the device.

Measures for a moderately protected system could include a special restricted operating environment, such as a TEE. In this case, only security-oriented software has access to the Attester and key material.

Measures for a highly protected system could include specialized hardware that is used to provide protection against chip decapping attacks, power supply and clock glitching, faulting injection and RF and power side channel attacks.

12.1.2. Attestation Key Provisioning Processes

Attestation key provisioning is the process that occurs in the factory or elsewhere to establish signing key material on the device and the validation key material off the device. Sometimes this procedure is referred to as personalization or customization.

The keys generated in the factory, whether generated in the device or off-device by the factory SHOULD be generated by a Cryptographically Strong Sequence ([RFC4086], Section 6.2).

12.1.2.1. Off-Device Key Generation

One way to provision key material is to first generate it external to the device and then copy the key onto the device. In this case, confidentiality protection of the generator, as well as for the path over which the key is provisioned, is necessary. The manufacturer needs to take care to protect corresponding key material with measures appropriate for its value.

The degree of protection afforded to this key material can vary by the intended function of the device and the specific practices of the device manufacturer or integrator. The confidentiality protection is fundamentally based upon some amount of physical protection: while encryption is often used to provide confidentiality when a key is conveyed across a factory, where the attestation key is created or applied, it must be available in an unencrypted form. The physical protection can therefore vary from situations where the key is unencrypted only within carefully controlled secure enclaves within silicon, to situations where an entire facility is considered secure, by the simple means of locked doors and limited access.

The cryptography that is used to enable confidentiality protection of the attestation key comes with its own requirements to be secured. This results in recursive problems, as the key material used to provision attestation keys must again somehow have been provisioned securely beforehand (requiring an additional level of protection, and so on).

So, this is why, in general, a combination of some physical security measures and some cryptographic measures is used to establish confidentiality protection.

12.1.2.2. On-Device Key Generation

When key material is generated within a device and the secret part of it never leaves the device, then the problem may lessen. For public-key cryptography, it is, by definition, not necessary to maintain confidentiality of the public key: however integrity of the chain of custody of the public key is necessary in order to avoid attacks where an attacker is able get a key they control endorsed.

To summarize: attestation key provisioning must ensure that only valid attestation key material is established in Attesters.

12.2. Integrity Protection

Any solution that conveys information in any conceptual message (see Section 8) must support end-to-end integrity protection and replay attack prevention, and often also needs to support additional security properties, including:

- * end-to-end encryption,
- * denial of service protection,
- * authentication,

- * auditing,
- * fine grained access controls, and
- * logging.

Section 10 discusses ways in which freshness can be used in this architecture to protect against replay attacks.

To assess the security provided by a particular appraisal policy, it is important to understand the strength of the root of trust, e.g., whether it is mutable software, or firmware that is read-only after boot, or immutable hardware/ROM.

It is also important that the appraisal policy was itself obtained securely. If an attacker can configure or modify appraisal policies, Endorsements or Reference Values for a Relying Party or for a Verifier, then integrity of the process is compromised.

Security protections in RATS may be applied at different layers, whether by a conveyance protocol, or an information encoding format. This architecture expects conceptual messages to be end-to-end protected based on the role interaction context. For example, if an Attester produces Evidence that is relayed through some other entity that doesn't implement the Attester or the intended Verifier roles, then the relaying entity should not expect to have access to the Evidence.

12.3. Epoch ID-based Attestation

Epoch IDs, described in Section 10.3, can be tampered with, replayed, dropped, delayed, and reordered by an attacker.

An attacker could be either external or belong to the distribution group, for example, if one of the Attester entities have been compromised.

An attacker who is able to tamper with epoch IDs can potentially lock all the participants in a certain epoch of choice for ever, effectively freezing time. This is problematic since it destroys the ability to ascertain freshness of Evidence and Attestation Results.

To mitigate this threat, the transport should be at least integrity protected and provide origin authentication.

Selective dropping of epoch IDs is equivalent to pinning the victim node to a past epoch. An attacker could drop epoch IDs to only some entities and not others, which will typically result in a denial of service due to the permanent staleness of the Attestation Result or Evidence.

Delaying or reordering epoch IDs is equivalent to manipulating the victim's timeline at will. This ability could be used by a malicious actor (e.g., a compromised router) to mount a confusion attack where, for example, a Verifier is tricked into accepting Evidence coming from a past epoch as fresh, while in the meantime the Attester has been compromised.

Reordering and dropping attacks are mitigated if the transport provides the ability to detect reordering and drop. However, the delay attack described above can't be thwarted in this manner.

12.4. Trust Anchor Protection

As noted in Section 7, Verifiers and Relying Parties have trust anchor stores that must be secured. [RFC6024] contains more discussion of trust anchor store requirements for protecting public keys. Section 6 of [NIST-800-57-p1] contains a comprehensive treatment of the topic, including the protection of symmetric key material. Specifically, a trust anchor store must resist modification against unauthorized insertion, deletion, and modification. Additionally, if the trust anchor is a symmetric key, the trust anchor store must not allow unauthorized read.

If certificates are used as trust anchors, Verifiers and Relying Parties are also responsible for validating the entire certificate path up to the trust anchor, which includes checking for certificate revocation. See Section 6 of [RFC5280] for details.

13. IANA Considerations

This document does not require any actions by IANA.

14. Acknowledgments

Special thanks go to Joerg Borchert, Nancy Cam-Winget, Jessica Fitzgerald-McKay, Diego Lopez, Laurence Lundblade, Paul Rowe, Hannes Tschofenig, Frank Xia, and David Wooten.

15. Notable Contributions

Thomas Hardjono created initial versions of the terminology section in collaboration with Ned Smith. Eric Voit provided the conceptual separation between Attestation Provision Flows and Attestation Evidence Flows. Monty Wisemen created the content structure of the first three architecture drafts. Carsten Bormann provided many of the motivational building blocks with respect to the Internet Threat Model.

16. Appendix A: Time Considerations

Section 10 discussed various issues and requirements around freshness of evidence, and summarized three approaches that might be used by different solutions to address them. This appendix provides more details with examples to help illustrate potential approaches, to inform those creating specific solutions.

The table below defines a number of relevant events, with an ID that is used in subsequent diagrams. The times of said events might be defined in terms of an absolute clock time, such as the Coordinated Universal Time timescale, or might be defined relative to some other timestamp or timeticks counter, such as a clock resetting its epoch each time it is powered on.

ID	Event	Explanation of event
VG	Value generated	A value to appear in a Claim was created. In some cases, a value may have technically existed before an Attester became aware of it but the Attester might have no idea how long it has had that value. In such a case, the Value created time is the time at which the Claim containing the copy of the value was created.
NS	Nonce sent	A nonce not predictable to an Attester (recentness & uniqueness) is sent to an Attester.
NR	Nonce relayed	A nonce is relayed to an Attester by another entity.
IR	Epoch ID received	An epoch ID is successfully received and processed by an entity.
EG	Evidence	An Attester creates Evidence from collected

	generation	Claims.
ER	Evidence relayed	A Relying Party relays Evidence to a Verifier.
RG	Result generation	A Verifier appraises Evidence and generates an Attestation Result.
RR	Result relayed	A Relying Party relays an Attestation Result to a Relying Party.
RA	Result appraised	The Relying Party appraises Attestation Results.
OP	Operation performed	The Relying Party performs some operation requested by the Attester via a resource access protocol as depicted in Figure 8, e.g., across a session created earlier at time(RA).
RX	Result expiry	An Attestation Result should no longer be accepted, according to the Verifier that generated it.

Table 1

Using the table above, a number of hypothetical examples of how a solution might be built are illustrated below. This list is not intended to be complete, but is just representative enough to highlight various timing considerations.

All times are relative to the local clocks, indicated by an "_a" (Attester), "_v" (Verifier), or "_r" (Relying Party) suffix.

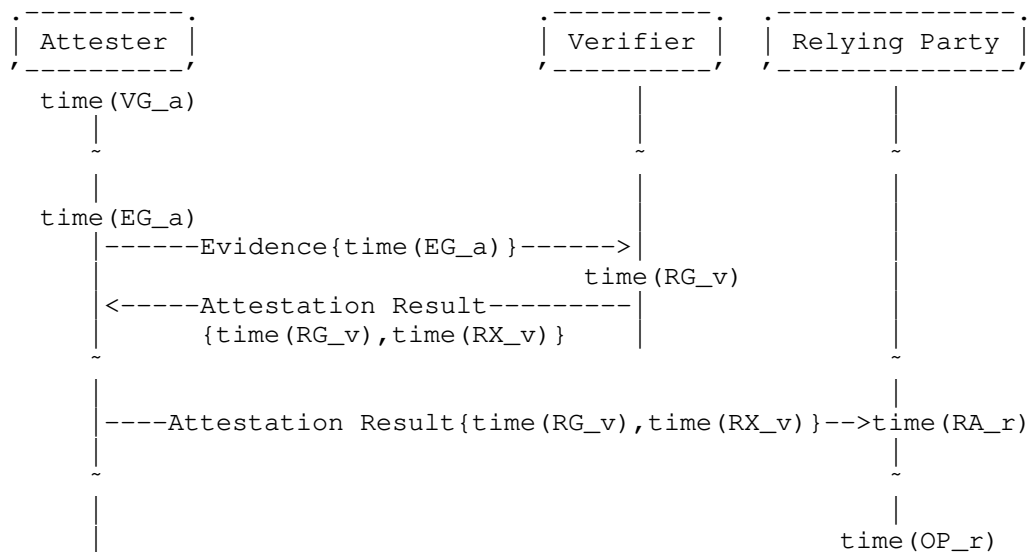
Times with an appended Prime (') indicate a second instance of the same event.

How and if clocks are synchronized depends upon the model.

In the figures below, curly braces indicate containment. For example, the notation Evidence{foo} indicates that 'foo' is contained in the Evidence and is thus covered by its signature.

16.1. Example 1: Timestamp-based Passport Model Example

The following example illustrates a hypothetical Passport Model solution that uses timestamps and requires roughly synchronized clocks between the Attester, Verifier, and Relying Party, which depends on using a secure clock synchronization mechanism. As a result, the receiver of a conceptual message containing a timestamp can directly compare it to its own clock and timestamps.



The Verifier can check whether the Evidence is fresh when appraising it at `time(RG_v)` by checking $\text{time(RG_v)} - \text{time(EG_a)} < \text{Threshold}$, where the Verifier's threshold is large enough to account for the maximum permitted clock skew between the Verifier and the Attester.

If `time(VG_a)` is also included in the Evidence along with the Claim value generated at that time, and the Verifier decides that it can trust the `time(VG_a)` value, the Verifier can also determine whether the Claim value is recent by checking $\text{time(RG_v)} - \text{time(VG_a)} < \text{Threshold}$. The threshold is decided by the Appraisal Policy for Evidence, and again needs to take into account the maximum permitted clock skew between the Verifier and the Attester.

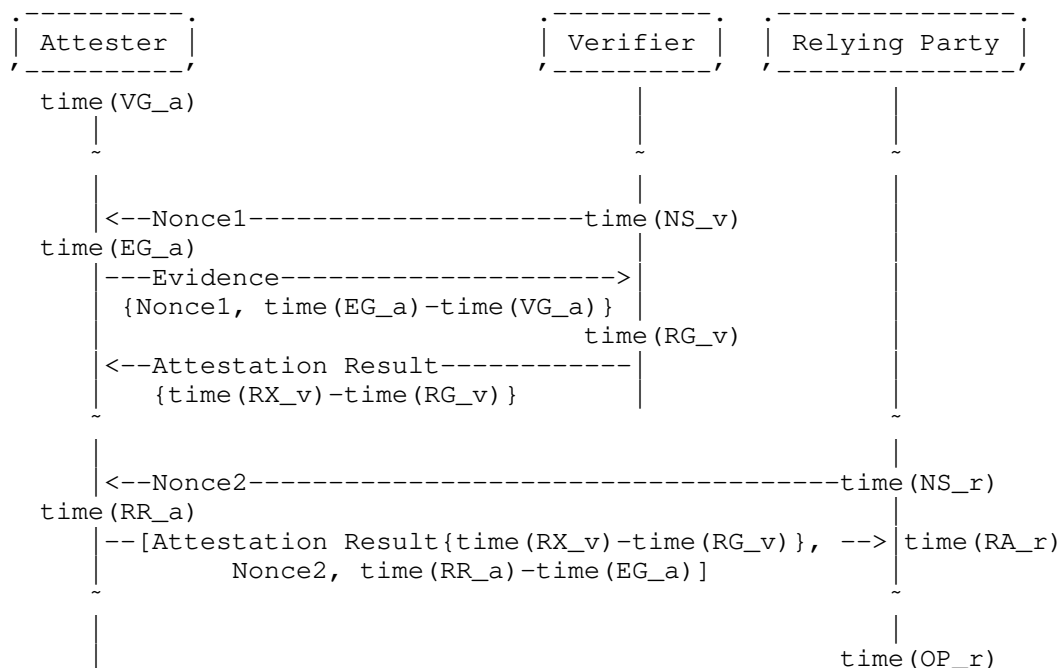
The Relying Party can check whether the Attestation Result is fresh when appraising it at `time(RA_r)` by checking $\text{time(RA_r)} - \text{time(RG_v)} < \text{Threshold}$, where the Relying Party's threshold is large enough to account for the maximum permitted clock skew between the Relying Party and the Verifier. The result might then be used for some time (e.g., throughout the lifetime of a connection established at

$\text{time}(\text{RA}_r)$). The Relying Party must be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain fresh enough. Thus, it might allow use (at $\text{time}(\text{OP}_r)$) as long as $\text{time}(\text{OP}_r) - \text{time}(\text{RG}_v) < \text{Threshold}$. However, if the Attestation Result contains an expiry time $\text{time}(\text{RX}_v)$ then it could explicitly check $\text{time}(\text{OP}_r) < \text{time}(\text{RX}_v)$.

16.2. Example 2: Nonce-based Passport Model Example

The following example illustrates a hypothetical Passport Model solution that uses nonces instead of timestamps. Compared to the timestamp-based example, it requires an extra round trip to retrieve a nonce, and requires that the Verifier and Relying Party track state to remember the nonce for some period of time.

The advantage is that it does not require that any clocks are synchronized. As a result, the receiver of a conceptual message containing a timestamp cannot directly compare it to its own clock or timestamps. Thus we use a suffix ("a" for Attester, "v" for Verifier, and "r" for Relying Party) on the IDs below indicating which clock generated them, since times from different clocks cannot be compared. Only the delta between two events from the sender can be used by the receiver.



In this example solution, the Verifier can check whether the Evidence is fresh at time(RG_v) by verifying that $\text{time(RG_v)} - \text{time(NS_v)} < \text{Threshold}$.

The Verifier cannot, however, simply rely on a Nonce to determine whether the value of a Claim is recent, since the Claim value might have been generated long before the nonce was sent by the Verifier. However, if the Verifier decides that the Attester can be trusted to correctly provide the delta time $\text{time(EG_a)} - \text{time(VG_a)}$, then it can determine recency by checking $\text{time(RG_v)} - \text{time(NS_v)} + \text{time(EG_a)} - \text{time(VG_a)} < \text{Threshold}$.

Similarly if, based on an Attestation Result from a Verifier it trusts, the Relying Party decides that the Attester can be trusted to correctly provide time deltas, then it can determine whether the Attestation Result is fresh by checking $\text{time(OP_r)} - \text{time(NS_r)} + \text{time(RR_a)} - \text{time(EG_a)} < \text{Threshold}$. Although the Nonce2 and $\text{time(RR_a)} - \text{time(EG_a)}$ values cannot be inside the Attestation Result, they might be signed by the Attester such that the Attestation Result vouches for the Attester's signing capability.

The Relying Party must still be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain valid. Thus, if the Attestation Result sends a validity lifetime in terms of $\text{time(RX_v)} - \text{time(RG_v)}$, then the Relying Party can check $\text{time(OP_r)} - \text{time(NS_r)} < \text{time(RX_v)} - \text{time(RG_v)}$.

16.3. Example 3: Epoch ID-based Passport Model Example

The example in Figure 10 illustrates a hypothetical Passport Model solution that uses epoch IDs instead of nonces or timestamps.

The Epoch ID Distributor broadcasts epoch ID I which starts a new epoch E for a protocol participant upon reception at time(IR).

The Attester generates Evidence incorporating epoch ID I and conveys it to the Verifier.

The Verifier appraises that the received epoch ID I is "fresh" according to the definition provided in Section 10.3 whereby retries are required in the case of mismatching epoch IDs, and generates an Attestation Result. The Attestation Result is conveyed to the Attester.

After the transmission of epoch ID I' a new epoch E' is established when I' is received by each protocol participant. The Attester relays the Attestation Result obtained during epoch E (associated with epoch ID I) to the Relying Party using the epoch ID for the

current epoch I' . If the Relying Party had not yet received I' , then the Attestation Result would be rejected, but in this example, it is received.

In the illustrated scenario, the epoch ID for relaying an Attestation Result to the Relying Party is current, while a previous epoch ID was used to generate Verifier evaluated evidence. This indicates that at least one epoch transition has occurred, and the Attestation Results may only be as fresh as the previous epoch. If the Relying Party remembers the previous epoch ID I during an epoch window as discussed in Section 10.3, and the message is received during that window, the Attestation Result is accepted as fresh, and otherwise it is rejected as stale.

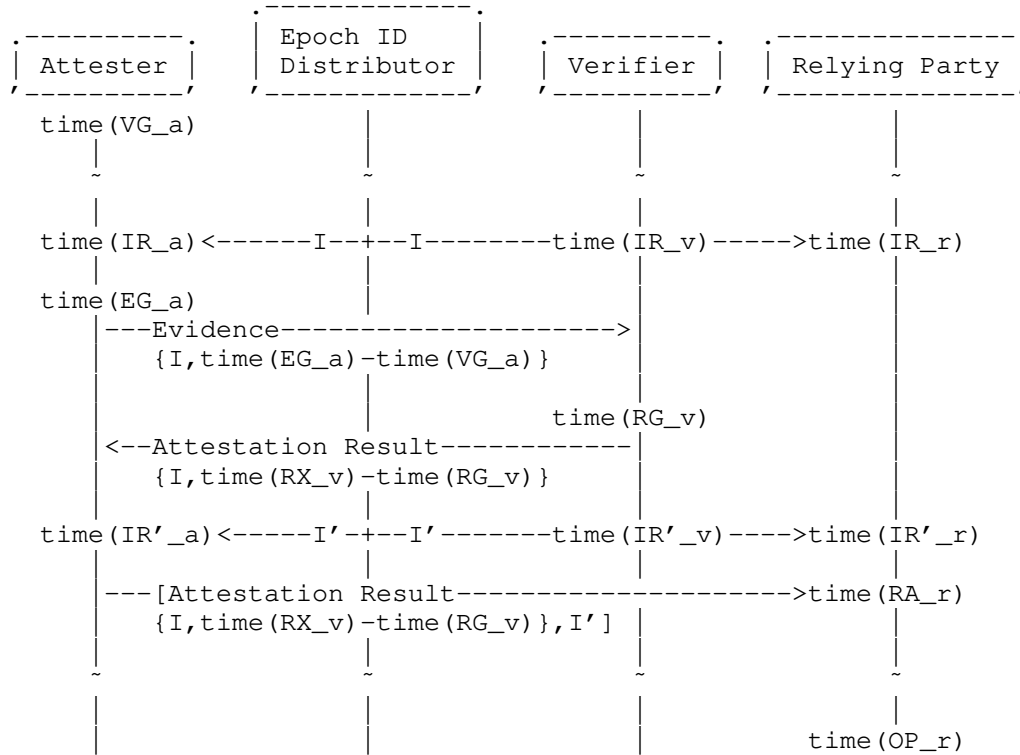
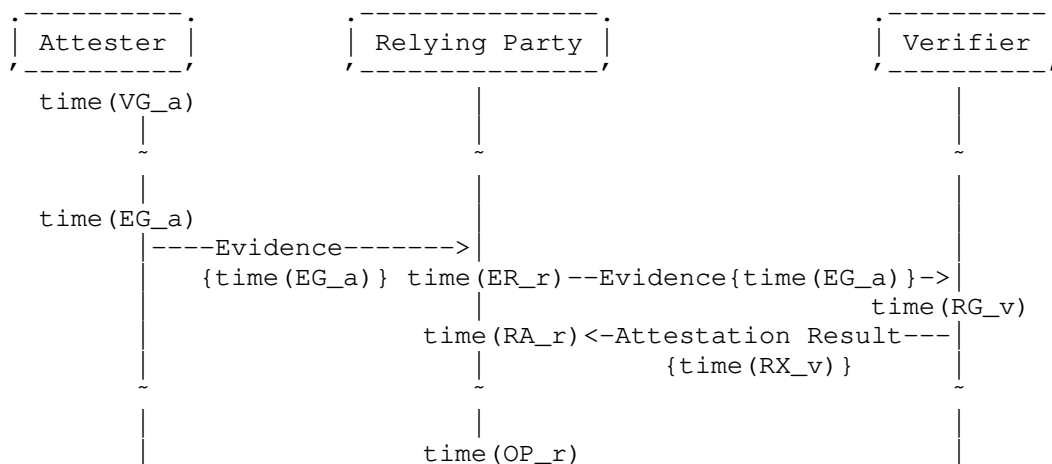


Figure 10: Epoch ID-based Passport Model

16.4. Example 4: Timestamp-based Background-Check Model Example

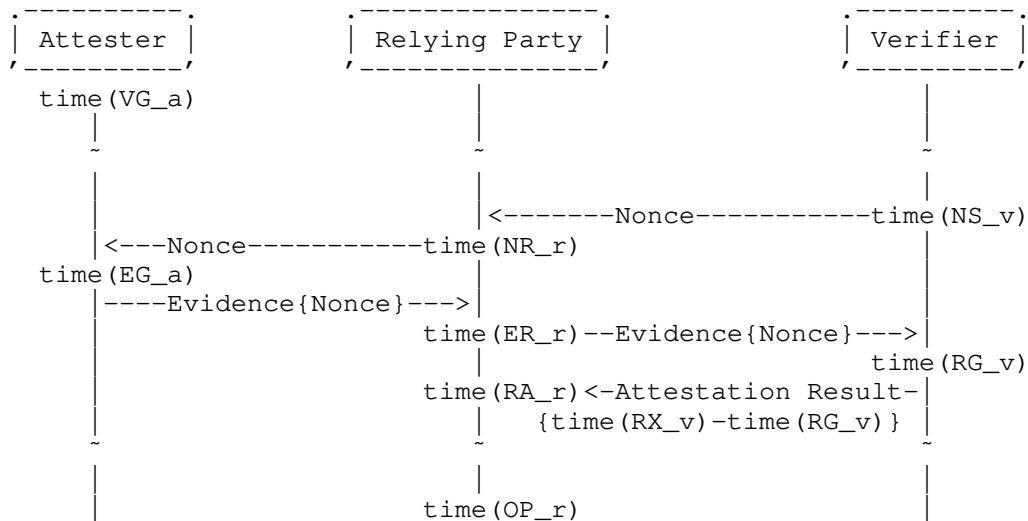
The following example illustrates a hypothetical Background-Check Model solution that uses timestamps and requires roughly synchronized clocks between the Attester, Verifier, and Relying Party.



The time considerations in this example are equivalent to those discussed under Example 1 above.

16.5. Example 5: Nonce-based Background-Check Model Example

The following example illustrates a hypothetical Background-Check Model solution that uses nonces and thus does not require that any clocks are synchronized. In this example solution, a nonce is generated by a Verifier at the request of a Relying Party, when the Relying Party needs to send one to an Attester.



The Verifier can check whether the Evidence is fresh, and whether a Claim value is recent, the same as in Example 2 above.

However, unlike in Example 2, the Relying Party can use the Nonce to determine whether the Attestation Result is fresh, by verifying that $\text{time}(\text{OP_r}) - \text{time}(\text{NR_r}) < \text{Threshold}$.

The Relying Party must still be careful, however, to not allow continued use beyond the period for which it deems the Attestation Result to remain valid. Thus, if the Attestation Result sends a validity lifetime in terms of $\text{time}(\text{RX_v}) - \text{time}(\text{RG_v})$, then the Relying Party can check $\text{time}(\text{OP_r}) - \text{time}(\text{ER_r}) < \text{time}(\text{RX_v}) - \text{time}(\text{RG_v})$.

17. References

17.1. Normative References

- [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://www.rfc-editor.org/info/rfc5280>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/info/rfc8392>>.

17.2. Informative References

- [CCC-DeepDive] Confidential Computing Consortium, "Confidential Computing Deep Dive", n.d., <<https://confidentialcomputing.io/whitepaper-02-latest>>.
- [CTAP] FIDO Alliance, "Client to Authenticator Protocol", n.d., <<https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-client-to-authenticator-protocol-v2.0-id-20180227.html>>.

[I-D.birkholz-rats-tuda]

Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-06, 12 January 2022, <<https://www.ietf.org/archive/id/draft-birkholz-rats-tuda-06.txt>>.

[I-D.birkholz-rats-uccs]

Birkholz, H., O'Donoghue, J., Cam-Winget, N., and C. Bormann, "A CBOR Tag for Unprotected CWT Claims Sets", Work in Progress, Internet-Draft, draft-birkholz-rats-uccs-03, 8 March 2021, <<https://www.ietf.org/archive/id/draft-birkholz-rats-uccs-03.txt>>.

[I-D.ietf-rats-daa]

Birkholz, H., Newton, C., Chen, L., and D. Thaler, "Direct Anonymous Attestation for the Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-daa-00, 2 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-daa-00.txt>>.

[I-D.ietf-teep-architecture]

Pei, M., Tschofenig, H., Thaler, D., and D. Wheeler, "Trusted Execution Environment Provisioning (TEEP) Architecture", Work in Progress, Internet-Draft, draft-ietf-teep-architecture-15, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-teep-architecture-15.txt>>.

[I-D.tschofenig-rats-psa-token]

Tschofenig, H., Frost, S., Brossard, M., Shaw, A., and T. Fossati, "Arm's Platform Security Architecture (PSA) Attestation Token", Work in Progress, Internet-Draft, draft-tschofenig-rats-psa-token-08, 24 March 2021, <<https://www.ietf.org/archive/id/draft-tschofenig-rats-psa-token-08.txt>>.

[I-D.tschofenig-tls-cwt]

Tschofenig, H. and M. Brossard, "Using CBOR Web Tokens (CWTs) in Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", Work in Progress, Internet-Draft, draft-tschofenig-tls-cwt-02, 13 July 2020, <<https://www.ietf.org/archive/id/draft-tschofenig-tls-cwt-02.txt>>.

- [NIST-800-57-p1] Barker, E., "Recommendation for Key Management: Part 1 - General", May 2020, <<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r5.pdf>>.
- [OPCUA] OPC Foundation, "OPC Unified Architecture Specification, Part 2: Security Model, Release 1.03", OPC 10000-2 , 25 November 2015, <<https://opcfoundation.org/developer-tools/specifications-unified-architecture/part-2-security-model/>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC5209] Sangster, P., Khosravi, H., Mani, M., Narayan, K., and J. Tardo, "Network Endpoint Assessment (NEA): Overview and Requirements", RFC 5209, DOI 10.17487/RFC5209, June 2008, <<https://www.rfc-editor.org/info/rfc5209>>.
- [RFC6024] Reddy, R. and C. Wallace, "Trust Anchor Management Requirements", RFC 6024, DOI 10.17487/RFC6024, October 2010, <<https://www.rfc-editor.org/info/rfc6024>>.
- [RFC8322] Field, J., Banghart, S., and D. Waltermire, "Resource-Oriented Lightweight Information Exchange (ROLIE)", RFC 8322, DOI 10.17487/RFC8322, February 2018, <<https://www.rfc-editor.org/info/rfc8322>>.
- [strengthoffunction] NISC, "Strength of Function", n.d., <https://csrc.nist.gov/glossary/term/strength_of_function>.
- [TCG-DICE] Trusted Computing Group, "DICE Certificate Profiles", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/DICE-Certificate-Profiles-r01_3june2020-1.pdf>.

[TCG-DICE-SIBDA]

Trusted Computing Group, "Symmetric Identity Based Device Attestation for DICE", 24 July 2019, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_DICE_SymIDAttest_v1_r0p94_pubrev.pdf>.

[TCGarch] Trusted Computing Group, "Trusted Platform Module Library - Part 1: Architecture", 8 November 2019, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_TPM2_r1p59_Part1_Architecture_pub.pdf>.

[WebAuthN] W3C, "Web Authentication: An API for accessing Public Key Credentials", n.d., <<https://www.w3.org/TR/webauthn-1/>>.

Contributors

Monty Wiseman

Email: montywiseman32@gmail.com

Liang Xia

Email: frank.xialiang@huawei.com

Laurence Lundblade

Email: lg1@island-resort.com

Eliot Lear

Email: ellear@cisco.com

Jessica Fitzgerald-McKay

Sarah C. Helbe

Andrew Guinn

Peter Loscocco

Email: pete.loscocco@gmail.com

Eric Voit

Thomas Fossati

Email: thomas.fossati@arm.com

Paul Rowe

Carsten Bormann

Email: cabo@tzi.org

Giri Mandyam

Email: mandyam@qti.qualcomm.com

Kathleen Moriarty

Email: kathleen.moriarty.ietf@gmail.com

Guy Fedorkow

Email: gfedorkow@juniper.net

Simon Frost

Email: Simon.Frost@arm.com

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Dave Thaler
Microsoft
United States of America

Email: dthaler@microsoft.com

Michael Richardson
Sandelman Software Works
Canada

Email: mcr+ietf@sandelman.ca

Ned Smith
Intel Corporation
United States of America

Email: ned.smith@intel.com

Wei Pan
Huawei Technologies

Email: william.panwei@huawei.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 9, 2021

G. Mandyam
Qualcomm Technologies Inc.
L. Lundblade
Security Theory LLC
M. Ballesteros
J. O'Donoghue
Qualcomm Technologies Inc.
June 07, 2021

The Entity Attestation Token (EAT)
draft-ietf-rats-eat-10

Abstract

An Entity Attestation Token (EAT) provides a signed (attested) set of claims that describe state and characteristics of an entity, typically a device like a phone or an IoT device. These claims are used by a relying party to determine how much it wishes to trust the entity.

An EAT is either a CWT or JWT with some attestation-oriented claims. To a large degree, all this document does is extend CWT and JWT.

Contributing

TBD

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 December 9, 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
1.1. CWT, JWT and UCCS	6
1.2. CDDL	6
1.3. Entity Overview	6
1.4. Use as Evidence and Attestation Results	7
1.5. EAT Operating Models	7
1.6. What is Not Standardized	9
1.6.1. Transmission Protocol	9
1.6.2. Signing Scheme	9
2. Terminology	10
3. The Claims	10
3.1. Token ID Claim (cti and jti)	11
3.2. Timestamp claim (iat)	11
3.3. Nonce Claim (nonce)	12
3.3.1. nonce CDDL	12
3.4. Universal Entity ID Claim (ueid)	12
3.4.1. ueid CDDL	15
3.5. Semi-permanent UEIDs (SUEIDs)	15
3.6. OEM Identification by IEEE (oemid)	16
3.6.1. oemid CDDL	16
3.7. Hardware Version Claims (hardware-version-claims)	16
3.8. Software Description and Version	17
3.9. The Security Level Claim (security-level)	17
3.9.1. security-level CDDL	18
3.10. Secure Boot Claim (secure-boot)	19
3.10.1. secure-boot CDDL	19
3.11. Debug Status Claim (debug-status)	19
3.11.1. Enabled	20
3.11.2. Disabled	21
3.11.3. Disabled Since Boot	21
3.11.4. Disabled Permanently	21

3.11.5.	Disabled Fully and Permanently	21
3.11.6.	debug-status CDDL	21
3.12.	Including Keys	22
3.13.	The Location Claim (location)	22
3.13.1.	location CDDL	23
3.14.	The Uptime Claim (uptime)	24
3.14.1.	uptime CDDL	24
3.15.	The Boot Seed Claim (boot-seed)	24
3.16.	The Intended Use Claim (intended-use)	24
3.16.1.	intended-use CDDL	25
3.17.	The Profile Claim (profile)	25
3.18.	The Software Manifests Claim (manifests)	26
3.19.	The Software Evidence Claim {swevidence}	27
3.20.	The Submodules Part of a Token (submods)	28
3.20.1.	Two Types of Submodules	28
3.20.1.1.	Non-token Submodules	29
3.20.1.2.	Nested EATs	29
3.20.1.3.	Unsecured JWTs and UCCS Tokens as Submodules	30
3.20.2.	No Inheritance	30
3.20.3.	Security Levels	31
3.20.4.	Submodule Names	31
3.20.5.	submods CDDL	31
4.	Endorsements and Verification Keys	32
4.1.	Identification Methods	33
4.1.1.	COSE/JWS Key ID	33
4.1.2.	JWS and COSE X.509 Header Parameters	34
4.1.3.	CBOR Certificate COSE Header Parameters	34
4.1.4.	Claim-Based Key Identification	34
4.2.	Other Considerations	34
5.	Profiles	35
5.1.	Format of a Profile Document	35
5.2.	List of Profile Issues	35
5.2.1.	Use of JSON, CBOR or both	35
5.2.2.	CBOR Map and Array Encoding	35
5.2.3.	CBOR String Encoding	36
5.2.4.	CBOR Preferred Serialization	36
5.2.5.	COSE/JOSE Protection	36
5.2.6.	COSE/JOSE Algorithms	36
5.2.7.	Verification Key Identification	37
5.2.8.	Endorsement Identification	37
5.2.9.	Freshness	37
5.2.10.	Required Claims	37
5.2.11.	Prohibited Claims	37
5.2.12.	Additional Claims	37
5.2.13.	Refined Claim Definition	37
5.2.14.	CBOR Tags	38
5.2.15.	Manifests and Software Evidence Claims	38
6.	Encoding	38

6.1.	Common CDDL Types	38
6.2.	CDDL for CWT-defined Claims	38
6.3.	JSON	39
6.3.1.	JSON Labels	39
6.3.2.	JSON Interoperability	40
6.4.	CBOR	41
6.4.1.	CBOR Interoperability	41
6.4.1.1.	EAT Constrained Device Serialization	41
6.5.	Collected CDDL	42
7.	IANA Considerations	47
7.1.	Reuse of CBOR Web Token (CWT) Claims Registry	47
7.2.	Claim Characteristics	48
7.2.1.	Interoperability and Relying Party Orientation	48
7.2.2.	Operating System and Technology Neutral	48
7.2.3.	Security Level Neutral	49
7.2.4.	Reuse of Extant Data Formats	49
7.2.5.	Proprietary Claims	49
7.3.	Claims Registered by This Document	49
7.3.1.	Claims for Early Assignment	50
7.3.2.	To be Assigned Claims	53
7.3.3.	Version Schemes Registered by this Document	53
8.	Privacy Considerations	53
8.1.	UEID and SUEID Privacy Considerations	53
8.2.	Location Privacy Considerations	54
9.	Security Considerations	54
9.1.	Key Provisioning	55
9.1.1.	Transmission of Key Material	55
9.2.	Transport Security	55
9.3.	Multiple EAT Consumers	56
10.	References	56
10.1.	Normative References	56
10.2.	Informative References	59
Appendix A.	Examples	61
A.1.	Very Simple EAT	61
A.2.	Example with Submodules, Nesting and Security Levels	61
Appendix B.	UEID Design Rationale	61
B.1.	Collision Probability	62
B.2.	No Use of UUID	64
Appendix C.	EAT Relation to IEEE.802.1AR Secure Device Identity (DevID)	65
C.1.	DevID Used With EAT	65
C.2.	How EAT Provides an Equivalent Secure Device Identity	66
C.3.	An X.509 Format EAT	66
C.4.	Device Identifier Permanence	67
Appendix D.	Changes from Previous Drafts	67
D.1.	From draft-rats-eat-01	67
D.2.	From draft-mandyam-rats-eat-00	67
D.3.	From draft-ietf-rats-eat-01	67

D.4.	From draft-ietf-rats-eat-02	68
D.5.	From draft-ietf-rats-eat-03	68
D.6.	From draft-ietf-rats-eat-04	68
D.7.	From draft-ietf-rats-05	69
D.8.	From draft-ietf-rats-06	69
D.9.	From draft-ietf-rats-07	69
D.10.	From draft-ietf-rats-08	69
D.11.	From draft-ietf-rats-09	69
Authors' Addresses	70

1. Introduction

Remote device attestation is a fundamental service that allows a remote device such as a mobile phone, an Internet-of-Things (IoT) device, or other endpoint to prove itself to a relying party, a server or a service. This allows the relying party to know some characteristics about the device and decide whether it trusts the device.

Remote attestation is a fundamental service that can underlie other protocols and services that need to know about the trustworthiness of the device before proceeding. One good example is biometric authentication where the biometric matching is done on the device. The relying party needs to know that the device is one that is known to do biometric matching correctly. Another example is content protection where the relying party wants to know the device will protect the data. This generalizes on to corporate enterprises that might want to know that a device is trustworthy before allowing corporate data to be accessed by it.

The notion of attestation here is large and may include, but is not limited to the following:

- o Proof of the make and model of the device hardware (HW)
- o Proof of the make and model of the device processor, particularly for security-oriented chips
- o Measurement of the software (SW) running on the device
- o Configuration and state of the device
- o Environmental characteristics of the device such as its GPS location

1.1. CWT, JWT and UCCS

For flexibility and ease of implementation in a wide variety of environments, EATs can be either CBOR [RFC8949] or JSON [ECMAScript] format. This specification simultaneously describes both formats.

An EAT is either a CWT as defined in [RFC8392], a UCCS as defined in [UCCS.Draft], or a JWT as defined in [RFC7519]. This specification extends those specifications with additional claims for attestation.

The identification of a protocol element as an EAT, whether CBOR or JSON format, follows the general conventions used by CWT, JWT and UCCS. Largely this depends on the protocol carrying the EAT. In some cases it may be by content type (e.g., MIME type). In other cases it may be through use of CBOR tags. There is no fixed mechanism across all use cases.

1.2. CDDL

This specification uses CDDL, [RFC8610], as the primary formalism to define each claim. The implementor then interprets the CDDL to come to either the CBOR [RFC8949] or JSON [ECMAScript] representation. In the case of JSON, Appendix E of [RFC8610] is followed. Additional rules are given in Section 6.3.2 of this document where Appendix E is insufficient. (Note that this is not to define a general means to translate between CBOR and JSON, but only to define enough such that the claims defined in this document can be rendered unambiguously in JSON).

The CWT specification was authored before CDDL was available and did not use it. This specification includes a CDDL definition of most of what is described in [RFC8392].

1.3. Entity Overview

An "entity" can be any device or device subassembly ("submodule") that can generate its own attestation in the form of an EAT. The attestation should be cryptographically verifiable by the EAT consumer. An EAT at the device-level can be composed of several submodule EAT's.

Modern devices such as a mobile phone have many different execution environments operating with different security levels. For example, it is common for a mobile phone to have an "apps" environment that runs an operating system (OS) that hosts a plethora of downloadable apps. It may also have a TEE (Trusted Execution Environment) that is distinct, isolated, and hosts security-oriented functionality like biometric authentication. Additionally, it may have an eSE (embedded

Secure Element) - a high security chip with defenses against HW attacks that is used to produce attestations. This device attestation format allows the attested data to be tagged at a security level from which it originates. In general, any discrete execution environment that has an identifiable security level can be considered an entity.

1.4. Use as Evidence and Attestation Results

Here, normative reference is made to [RATS-Architecture], particularly the definition of Evidence, the Verifier, Attestation Results and the Relying Party. Per Figure 1 in [RATS-Architecture], Evidence is a protocol message that goes from the Attester to the Verifier and Attestation Results a message that goes from the Verifier to the Relying Party. EAT is defined such that it can be used to represent either Evidence, Attestation Results or both. No claims defined here are considered exclusive to or are prohibited in either use. It is useful to create EAT profiles as described in Section 5 for either use.

It is useful to characterize the relationship of claims in Evidence to those in Attestation Results.

Many claims in Evidence simply will pass through the Verifier to the Relying Party without modification. They will be verified as authentic from the device by the Verifier just through normal verification of the Attester's signature. They will be protected from modification when they are conveyed to the Relying Party by whatever means is used to protect Attestation Results. (The details of that protection are out of scope of this document.)

Some claims in Evidence will be verified by the Verifier by comparison to Reference Values. In this case the claims in Evidence will not likely be conveyed to the Relying Party. Instead, some claim indicating they were checked may be added to the Attestation Results or it may be tacitly known that the Verifier always does this check.

In some cases the Verifier may provide privacy-preserving functionality by stripping or modifying claims that do not possess sufficient privacy-preserving characteristics.

1.5. EAT Operating Models

TODO: Rewrite (or eliminate) this section in light of the RATS architecture draft.

At least the following three participants exist in all EAT operating models. Some operating models have additional participants.

The Entity. This is the phone, the IoT device, the sensor, the sub-assembly or such that the attestation provides information about.

The Manufacturer. The company that made the entity. This may be a chip vendor, a circuit board module vendor or a vendor of finished consumer products.

The Relying Party. The server, service or company that makes use of the information in the EAT about the entity.

In all operating models, the manufacturer provisions some secret attestation key material (AKM) into the entity during manufacturing. This might be during the manufacturer of a chip at a fabrication facility (fab) or during final assembly of a consumer product or any time in between. This attestation key material is used for signing EATs.

In all operating models, hardware and/or software on the entity create an EAT of the format described in this document. The EAT is always signed by the attestation key material provisioned by the manufacturer.

In all operating models, the relying party must end up knowing that the signature on the EAT is valid and consistent with data from claims in the EAT. This can happen in many different ways. Here are some examples.

- o The EAT is transmitted to the relying party. The relying party gets corresponding key material (e.g. a root certificate) from the manufacturer. The relying party performs the verification.
- o The EAT is transmitted to the relying party. The relying party transmits the EAT to a verification service offered by the manufacturer. The server returns the validated claims.
- o The EAT is transmitted directly to a verification service, perhaps operated by the manufacturer or perhaps by another party. It verifies the EAT and makes the validated claims available to the relying party. It may even modify the claims in some way and re-sign the EAT (with a different signing key).

All these operating models are supported and there is no preference of one over the other. It is important to support this variety of operating models to generally facilitate deployment and to allow for some special scenarios. One special scenario has a validation

service that is monetized, most likely by the manufacturer. In another, a privacy proxy service processes the EAT before it is transmitted to the relying party. In yet another, symmetric key material is used for signing. In this case the manufacturer should perform the verification, because any release of the key material would enable a participant other than the entity to create valid signed EATs.

1.6. What is Not Standardized

The following is not standardized for EAT, just the same they are not standardized for CWT or JWT.

1.6.1. Transmission Protocol

EATs may be transmitted by any protocol the same as CWTs and JWTs. For example, they might be added in extension fields of other protocols, bundled into an HTTP header, or just transmitted as files. This flexibility is intentional to allow broader adoption. This flexibility is possible because EAT's are self-secured with signing (and possibly additionally with encryption and anti-replay). The transmission protocol is not required to fulfill any additional security requirements.

For certain devices, a direct connection may not exist between the EAT-producing device and the Relying Party. In such cases, the EAT should be protected against malicious access. The use of COSE and JOSE allows for signing and encryption of the EAT. Therefore, even if the EAT is conveyed through intermediaries between the device and Relying Party, such intermediaries cannot easily modify the EAT payload or alter the signature.

1.6.2. Signing Scheme

The term "signing scheme" is used to refer to the system that includes end-end process of establishing signing attestation key material in the entity, signing the EAT, and verifying it. This might involve key IDs and X.509 certificate chains or something similar but different. The term "signing algorithm" refers just to the algorithm ID in the COSE signing structure. No particular signing algorithm or signing scheme is required by this standard.

There are three main implementation issues driving this. First, secure non-volatile storage space in the entity for the attestation key material may be highly limited, perhaps to only a few hundred bits, on some small IoT chips. Second, the factory cost of provisioning key material in each chip or device may be high, with even millisecond delays adding to the cost of a chip. Third,

privacy-preserving signing schemes like ECDAA (Elliptic Curve Direct Anonymous Attestation) are complex and not suitable for all use cases.

Over time to facilitate interoperability, some signing schemes may be defined in EAT profiles or other documents either in the IETF or outside.

2. Terminology

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.

This document reuses terminology from JWT [RFC7519], COSE [RFC8152], and CWT [RFC8392].

Claim Name. The human-readable name used to identify a claim.

Claim Key. The CBOR map key or JSON name used to identify a claim.

Claim Value. The value portion of the claim. A claim value can be any CBOR data item or JSON value.

CWT Claims Set. The CBOR map or JSON object that contains the claims conveyed by the CWT or JWT.

Attestation Key Material (AKM). The key material used to sign the EAT token. If it is done symmetrically with HMAC, then this is a simple symmetric key. If it is done with ECC, such as an IEEE DevID [IEEE.802.1AR], then this is the private part of the EC key pair. If ECDAA is used, (e.g., as used by Enhanced Privacy ID, i.e. EPID) then it is the key material needed for ECDAA.

3. The Claims

This section describes new claims defined for attestation. It also mentions several claims defined by CWT and JWT that are particularly important for EAT.

Note also: * Any claim defined for CWT or JWT may be used in an EAT including those in the CWT [IANA.CWT.Claims] and JWT IANA [IANA.JWT.Claims] claims registries.

- o All claims are optional

- o No claims are mandatory
- o All claims that are not understood by implementations MUST be ignored

There are no default values or meanings assigned to absent claims other than they are not reported. The reason for a claim's absence may be the implementation not supporting the claim, an inability to determine its value, or a preference to report in a different way such as a proprietary claim.

CDDL along with text descriptions is used to define each claim independent of encoding. Each claim is defined as a CDDL group (the group is a general aggregation and type definition feature of CDDL). In the encoding section Section 6, the CDDL groups turn into CBOR map entries and JSON name/value pairs.

Map labels are assigned both an integer and string value. CBOR encoded tokens MUST use only integer labels. JSON encoded tokens MUST use only string labels.

TODO: add paragraph here about use for Attestation Evidence and for Results.

3.1. Token ID Claim (cti and jti)

CWT defines the "cti" claim. JWT defines the "jti" claim. These are equivalent to each other in EAT and carry a unique token identifier as they do in JWT and CWT. They may be used to defend against re use of the token but are distinct from the nonce that is used by the relying party to guarantee freshness and defend against replay.

3.2. Timestamp claim (iat)

The "iat" claim defined in CWT and JWT is used to indicate the date-of-creation of the token, the time at which the claims are collected and the token is composed and signed.

The data for some claims may be held or cached for some period of time before the token is created. This period may be long, even days. Examples are measurements taken at boot or a geographic position fix taken the last time a satellite signal was received. There are individual timestamps associated with these claims to indicate their age is older than the "iat" timestamp.

CWT allows the use floating-point for this claim. EAT disallows the use of floating-point. No token may contain an iat claim in float-point format. Any recipient of a token with a floating-point format

iat claim may consider it an error. A 64-bit integer representation of epoch time can represent a range of +/- 500 billion years, so the only point of a floating-point timestamp is to have precession greater than one second. This is not needed for EAT.

3.3. Nonce Claim (nonce)

All EATs should have a nonce to prevent replay attacks. The nonce is generated by the relying party, the end consumer of the token. It is conveyed to the entity over whatever transport is in use before the token is generated and then included in the token as the nonce claim.

This documents the nonce claim for registration in the IANA CWT claims registry. This is equivalent to the JWT nonce claim that is already registered.

The nonce must be at least 8 bytes (64 bits) as fewer are unlikely to be secure. A maximum of 64 bytes is set to limit the memory a constrained implementation uses. This size range is not set for the already-registered JWT nonce, but it should follow this size recommendation when used in an EAT.

Multiple nonces are allowed to accommodate multistage verification and consumption.

3.3.1. nonce CDDL

```
nonce-type = bstr .size (8..64)

nonce-claim = (
    nonce => nonce-type / [ 2* nonce-type ]
)
```

3.4. Universal Entity ID Claim (ueid)

UEID's identify individual manufactured entities / devices such as a mobile phone, a water meter, a Bluetooth speaker or a networked security camera. It may identify the entire device or a submodule or subsystem. It does not identify types, models or classes of devices. It is akin to a serial number, though it does not have to be sequential.

UEID's must be universally and globally unique across manufacturers and countries. UEIDs must also be unique across protocols and systems, as tokens are intended to be embedded in many different protocols and systems. No two products anywhere, even in completely different industries made by two different manufacturers in two different countries should have the same UEID (if they are not global

and universal in this way, then relying parties receiving them will have to track other characteristics of the device to keep devices distinct between manufacturers).

There are privacy considerations for UEID's. See Section 8.1.

The UEID is permanent. It never change for a given device / entity.

UEIDs are variable length. All implementations MUST be able to receive UEIDs that are 33 bytes long (1 type byte and 256 bits). The recommended maximum sent is also 33 bytes.

When the entity constructs the UEID, the first byte is a type and the following bytes the ID for that type. Several types are allowed to accommodate different industries and different manufacturing processes and to give options to avoid paying fees for certain types of manufacturer registrations.

Creation of new types requires a Standards Action [RFC8126].

Type Byte	Type Name	Specification
0x01	RAND	This is a 128, 192 or 256 bit random number generated once and stored in the device. This may be constructed by concatenating enough identifiers to make up an equivalent number of random bits and then feeding the concatenation through a cryptographic hash function. It may also be a cryptographic quality random number generated once at the beginning of the life of the device and stored. It may not be smaller than 128 bits.
0x02	IEEE EUI	This makes use of the IEEE company identification registry. An EUI is either an EUI-48, EUI-60 or EUI-64 and made up of an OUI, OUI-36 or a CID, different registered company identifiers, and some unique per-device identifier. EUIs are often the same as or similar to MAC addresses. This type includes MAC-48, an obsolete name for EUI-48. (Note that while devices with multiple network interfaces may have multiple MAC addresses, there is only one UEID for a device) [IEEE.802-2001], [OUI.Guide]
0x03	IMEI	This is a 14-digit identifier consisting of an 8-digit Type Allocation Code and a 6-digit serial number allocated by the manufacturer, which SHALL be encoded as byte string of length 14 with each byte as the digit's value (not the ASCII encoding of the digit; the digit 3 encodes as 0x03, not 0x33). The IMEI value encoded SHALL NOT include Luhn checksum or SVN information. [ThreeGPP.IMEI]

Table 1: UEID Composition Types

UEID's are not designed for direct use by humans (e.g., printing on the case of a device), so no textual representation is defined.

The consumer (the relying party) of a UEID MUST treat a UEID as a completely opaque string of bytes and not make any use of its internal structure. For example, they should not use the OUI part of a type 0x02 UEID to identify the manufacturer of the device. Instead they should use the oemid claim that is defined elsewhere. The reasons for this are:

- o UEIDs types may vary freely from one manufacturer to the next.

- o New types of UEIDs may be created. For example, a type 0x07 UEID may be created based on some other manufacturer registration scheme.
- o Device manufacturers are allowed to change from one type of UEID to another anytime they want. For example, they may find they can optimize their manufacturing by switching from type 0x01 to type 0x02 or vice versa. The main requirement on the manufacturer is that UEIDs be universally unique.

3.4.1. ueid CDDL

```
ueid-type = bstr .size (7..33)

ueid-claim = (
    ueid => ueid-type
)
```

3.5. Semi-permanent UEIDs (SUEIDs)

An SEUID is of the same format as a UEID, but it may change to a different value on device life-cycle events. Examples of these events are change of ownership, factory reset and on-boarding into an IoT device management system. A device may have both a UEID and SUEIDs, neither, one or the other.

There may be multiple SUEIDs. Each one has a text string label the purpose of which is to distinguish it from others in the token. The label may name the purpose, application or type of the SUEID. Typically, there will be few SUEIDs so there is no need for a formal labeling mechanism like a registry. The EAT profile may describe how SUEIDs should be labeled. If there is only one SUEID, the claim remains a map and there still must be a label. For example, the label for the SUEID used by FIDO Onboarding Protocol could simply be "FDO".

There are privacy considerations for SUEID's. See Section 8.1.

```
sueids-type = {
    + tstr => ueid-type
}

sueids-claim = (
    sueids => sueids-type
)
```

3.6. OEM Identification by IEEE (oemid)

The IEEE operates a global registry for MAC addresses and company IDs. This claim uses that database to identify OEMs. The contents of the claim may be either an IEEE MA-L, MA-M, MA-S or an IEEE CID [IEEE.RA]. An MA-L, formerly known as an OUI, is a 24-bit value used as the first half of a MAC address. MA-M similarly is a 28-bit value used as the first part of a MAC address, and MA-S, formerly known as OUI-36, a 36-bit value. Many companies already have purchased one of these. A CID is also a 24-bit value from the same space as an MA-L, but not for use as a MAC address. IEEE has published Guidelines for Use of EUI, OUI, and CID [OUI.Guide] and provides a lookup services [OUI.Lookup]

Companies that have more than one of these IDs or MAC address blocks should pick one and prefer that for all their devices.

Commonly, these are expressed in Hexadecimal Representation [IEEE.802-2001] also called the Canonical format. When this claim is encoded the order of bytes in the bstr are the same as the order in the Hexadecimal Representation. For example, an MA-L like "AC-DE-48" would be encoded in 3 bytes with values 0xAC, 0xDE, 0x48. For JSON encoded tokens, this is further base64url encoded.

3.6.1. oemid CDDL

```
oemid-claim = (  
    oemid => bstr  
)
```

3.7. Hardware Version Claims (hardware-version-claims)

The hardware version can be claimed at three different levels, the chip, the circuit board and the final device assembly. An EAT can include any combination these claims.

The hardware version is a simple text string the format of which is set by each manufacturer. The structure and sorting order of this text string can be specified using the version-scheme item from CoSWID [CoSWID].

The hardware version can also be given by a 13-digit [EAN-13]. A new CoSWID version scheme is registered with IANA by this document in Section 7.3.3. An EAN-13 is also known as an International Article Number or most commonly as a bar code.

```
chip-version-claim = (  
    chip-version => tstr  
)  
  
chip-version-scheme-claim = (  
    chip-version-scheme => $version-scheme  
)  
  
board-version-claim = (  
    board-version => tstr  
)  
  
board-version-scheme-claim = (  
    board-version-scheme => $version-scheme  
)  
  
device-version-claim = (  
    device-version => tstr  
)  
  
device-version-scheme-claim = (  
    device-version-scheme => $version-scheme  
)  
  
hardware-version-claims = (  
    ? chip-version-claim,  
    ? board-version-claim,  
    ? device-version-claim,  
    ? chip-version-scheme-claim,  
    ? board-version-scheme-claim,  
    ? device-version-scheme-claim,  
)
```

3.8. Software Description and Version

TODO: Add claims that reference CoSWID.

3.9. The Security Level Claim (security-level)

This claim characterizes the device/entity ability to defend against attacks aimed at capturing the signing key, forging claims and at forging EATs. This is done by defining four security levels as described below. This is similar to the key protection types defined by the Fast Identity Online (FIDO) Alliance [FIDO.Registry].

These claims describe security environment and countermeasures available on the end-entity / client device where the attestation key reside and the claims originate.

- 1 - Unrestricted There is some expectation that implementor will protect the attestation signing keys at this level. Otherwise the EAT provides no meaningful security assurances.
- 2- Restricted Entities at this level should not be general-purpose operating environments that host features such as app download systems, web browsers and complex productivity applications. It is akin to the Secure Restricted level (see below) without the security orientation. Examples include a Wi-Fi subsystem, an IoT camera, or sensor device.
- 3 - Secure Restricted Entities at this level must meet the criteria defined by FIDO Allowed Restricted Operating Environments [FIDO.AROE]. Examples include TEE's and schemes using virtualization-based security. Like the FIDO security goal, security at this level is aimed at defending well against large-scale network / remote attacks against the device.
- 4 - Hardware Entities at this level must include substantial defense against physical or electrical attacks against the device itself. It is assumed any potential attacker has captured the device and can disassemble it. Example include TPMs and Secure Elements.

The entity should claim the highest security level it achieves and no higher. This set is not extensible so as to provide a common interoperable description of security level to the relying party. If a particular implementation considers this claim to be inadequate, it can define its own proprietary claim. It may consider including both this claim as a coarse indication of security and its own proprietary claim as a refined indication.

This claim is not intended as a replacement for a proper end-device security certification schemes such as those based on FIPS 140 [FIPS-140] or those based on Common Criteria [Common.Criteria]. The claim made here is solely a self-claim made by the Entity Originator.

3.9.1. security-level CDDL


```
security-level-cbor-type = &(
    unrestricted: 1,
    restricted: 2,
    secure-restricted: 3,
    hardware: 4
)

security-level-json-type =
    "unrestricted" /
    "restricted" /
    "secure-restricted" /
    "hardware"

security-level-claim = (
    security-level => security-level-cbor-type / security-level-json-type
)
```

3.10. Secure Boot Claim (secure-boot)

The value of true indicates secure boot is enabled. Secure boot is considered enabled when base software, the firmware and operating system, are under control of the entity manufacturer identified in the oemid claimd described in Section 3.6. This may be because the software is in ROM or because it is cryptographically authenticated or some combination of the two or other.

3.10.1. secure-boot CDDL

```
secure-boot-claim = (
    secure-boot => bool
)
```

3.11. Debug Status Claim (debug-status)

This applies to system-wide or submodule-wide debug facilities of the target device / submodule like JTAG and diagnostic hardware built into chips. It applies to any software debug facilities related to root, operating system or privileged software that allow system-wide memory inspection, tracing or modification of non-system software like user mode applications.

This characterization assumes that debug facilities can be enabled and disabled in a dynamic way or be disabled in some permanent way such that no enabling is possible. An example of dynamic enabling is one where some authentication is required to enable debugging. An example of permanent disabling is blowing a hardware fuse in a chip. The specific type of the mechanism is not taken into account. For

example, it does not matter if authentication is by a global password or by per-device public keys.

As with all claims, the absence of the debug level claim means it is not reported. A conservative interpretation might assume the Not Disabled state. It could however be that it is reported in a proprietary claim.

This claim is not extensible so as to provide a common interoperable description of debug status to the relying party. If a particular implementation considers this claim to be inadequate, it can define its own proprietary claim. It may consider including both this claim as a coarse indication of debug status and its own proprietary claim as a refined indication.

The higher levels of debug disabling requires that all debug disabling of the levels below it be in effect. Since the lowest level requires that all of the target's debug be currently disabled, all other levels require that too.

There is no inheritance of claims from a submodule to a superior module or vice versa. There is no assumption, requirement or guarantee that the target of a superior module encompasses the targets of submodules. Thus, every submodule must explicitly describe its own debug state. The verifier or relying party receiving an EAT cannot assume that debug is turned off in a submodule because there is a claim indicating it is turned off in a superior module.

An individual target device / submodule may have multiple debug facilities. The use of plural in the description of the states refers to that, not to any aggregation or inheritance.

The architecture of some chips or devices may be such that a debug facility operates for the whole chip or device. If the EAT for such a chip includes submodules, then each submodule should independently report the status of the whole-chip or whole-device debug facility. This is the only way the relying party can know the debug status of the submodules since there is no inheritance.

3.11.1. Enabled

If any debug facility, even manufacturer hardware diagnostics, is currently enabled, then this level must be indicated.

3.11.2. Disabled

This level indicates all debug facilities are currently disabled. It may be possible to enable them in the future, and it may also be possible that they were enabled in the past after the target device/sub-system booted/started, but they are currently disabled.

3.11.3. Disabled Since Boot

This level indicates all debug facilities are currently disabled and have been so since the target device/sub-system booted/started.

3.11.4. Disabled Permanently

This level indicates all non-manufacturer facilities are permanently disabled such that no end user or developer cannot enable them. Only the manufacturer indicated in the OEMID claim can enable them. This also indicates that all debug facilities are currently disabled and have been so since boot/start.

3.11.5. Disabled Fully and Permanently

This level indicates that all debug capabilities for the target device/sub-module are permanently disabled.

3.11.6. debug-status CDDL

```
debug-status-cbor-type = &(
    enabled: 0,
    disabled: 1,
    disabled-since-boot: 2,
    disabled-permanently: 3,
    disabled-fully-and-permanently: 4
)

debug-status-json-type =
    "enabled" /
    "disabled" /
    "disabled-since-boot" /
    "disabled-permanently" /
    "disabled-fully-and-permanently"

debug-status-claim = (
    debug-status => debug-status-cbor-type / debug-status-json-type
)
```

3.12. Including Keys

An EAT may include a cryptographic key such as a public key. The signing of the EAT binds the key to all the other claims in the token.

The purpose for inclusion of the key may vary by use case. For example, the key may be included as part of an IoT device onboarding protocol. When the FIDO protocol includes a public key in its attestation message, the key represents the binding of a user, device and relying party. This document describes how claims containing keys should be defined for the various use cases. It does not define specific claims for specific use cases.

Keys in CBOR format tokens SHOULD be the COSE_Key format [RFC8152] and keys in JSON format tokens SHOULD be the JSON Web Key format [RFC7517]. These two formats support many common key types. Their use avoids the need to decode other serialization formats. These two formats can be extended to support further key types through their IANA registries.

The general confirmation claim format [RFC8747], [RFC7800] may also be used. It provides key encryption. It also allows for inclusion by reference through a key ID. The confirmation claim format may be employed in the definition of some new claim for a particular use case.

When the actual confirmation claim is included in an EAT, this document associates no use case semantics other than proof of possession. Different EAT use cases may choose to associate further semantics. The key in the confirmation claim MUST be protected the same as the key used to sign the EAT. That is, the same, equivalent or better hardware defenses, access controls, key generation and such must be used.

3.13. The Location Claim (location)

The location claim gives the location of the device entity from which the attestation originates. It is derived from the W3C Geolocation API [W3C.GeoLoc]. The latitude, longitude, altitude and accuracy must conform to [WGS84]. The altitude is in meters above the [WGS84] ellipsoid. The two accuracy values are positive numbers in meters. The heading is in degrees relative to true north. If the device is stationary, the heading is NaN (floating-point not-a-number). The speed is the horizontal component of the device velocity in meters per second.

When encoding floating-point numbers half-precision should not be used. It usually does not provide enough precision for a geographic location. It is not a requirement that the receiver of an EAT implement half-precision, so the receiver may not be able to decode the location.

The location may have been cached for a period of time before token creation. For example, it might have been minutes or hours or more since the last contact with a GPS satellite. Either the timestamp or age data item can be used to quantify the cached period. The timestamp data item is preferred as it a non-relative time.

The age data item can be used when the entity doesn't know what time it is either because it doesn't have a clock or it isn't set. The entity must still have a "ticker" that can measure a time interval. The age is the interval between acquisition of the location data and token creation.

See location-related privacy considerations in Section 8.2 below.

3.13.1. location CDDL

```
location-type = {  
    latitude => number,  
    longitude => number,  
    ? altitude => number,  
    ? accuracy => number,  
    ? altitude-accuracy => number,  
    ? heading => number,  
    ? speed => number,  
    ? timestamp => ~time-int,  
    ? age => uint  
}  
  
latitude = 1 / "latitude"  
longitude = 2 / "longitude"  
altitude = 3 / "altitude"  
accuracy = 4 / "accuracy"  
altitude-accuracy = 5 / "altitude-accuracy"  
heading = 6 / "heading"  
speed = 7 / "speed"  
timestamp = 8 / "timestamp"  
age = 9 / "age"  
  
location-claim = (  
    location-label => location-type  
)
```

3.14. The Uptime Claim (uptime)

The "uptime" claim contains a value that represents the number of seconds that have elapsed since the entity or submod was last booted.

3.14.1. uptime CDDL

```
uptime-claim = (  
    uptime => uint  
)
```

3.15. The Boot Seed Claim (boot-seed)

The Boot Seed claim is a random value created at system boot time that will allow differentiation of reports from different boot sessions. This value is usually public and not protected. It is not the same as a seed for a random number generator which must be kept secret.

```
boot-seed-claim = (  
    boot-seed => bytes  
)
```

3.16. The Intended Use Claim (intended-use)

EAT's may be used in the context of several different applications. The intended-use claim provides an indication to an EAT consumer about the intended usage of the token. This claim can be used as a way for an application using EAT to internally distinguish between different ways it uses EAT.

- 1 - Generic Generic attestation describes an application where the EAT consumer requires the most up-to-date security assessment of the attesting entity. It is expected that this is the most commonly-used application of EAT.
- 2- Registration Entities that are registering for a new service may be expected to provide an attestation as part of the registration process. This intended-use setting indicates that the attestation is not intended for any use but registration.
- 3 - Provisioning Entities may be provisioned with different values or settings by an EAT consumer. Examples include key material or device management trees. The consumer may require an EAT to assess device security state of the entity prior to provisioning.
- 4 - Certificate Issuance (Certificate Signing Request) Certifying authorities (CA's) may require attestations prior to the issuance

of certificates related to keypairs hosted at the entity. An EAT may be used as part of the certificate signing request (CSR).

- 5 - Proof-of-Possession An EAT consumer may require an attestation as part of an accompanying proof-of-possession (PoP) application. More precisely, a PoP transaction is intended to provide to the recipient cryptographically-verifiable proof that the sender has possession of a key. This kind of attestation may be necessary to verify the security state of the entity storing the private key used in a PoP application.

3.16.1. intended-use CDDL

```
intended-use-cbor-type = &(
    generic: 1,
    registration: 2,
    provisioning: 3,
    csr: 4,
    pop: 5
)

intended-use-json-type =
    "generic" /
    "registration" /
    "provisioning" /
    "csr" /
    "pop"

intended-use-claim = (
    intended-use => intended-use-cbor-type / intended-use-json-type
)
```

3.17. The Profile Claim (profile)

See Section 5 for the detailed description of a profile.

A profile is identified by either a URL or an OID. Typically, the URI will reference a document describing the profile. An OID is just a unique identifier for the profile. It may exist anywhere in the OID tree. There is no requirement that the named document be publicly accessible. The primary purpose of the profile claim is to uniquely identify the profile even if it is a private profile.

The OID is encoded in CBOR according to [CBOR-OID] and the URI according to [RFC8949]. Both are unwrapped and thus not tags. The OID is always absolute and never relative. If the claims CBOR type is a text string it is a URI and if a byte string it is an OID.

Note that this named "eat_profile" for JWT and is distinct from the already registered "profile" claim in the JWT claims registry.

```
oid = #6.4000(bstr) ; TODO: fill this in with correct CDDL from OID RFC

profile-claim = (
    profile => ~uri / ~oid
)
```

3.18. The Software Manifests Claim (manifests)

This claim contains descriptions of software that is present on the device. These manifests are installed on the device when the software is installed or are created as part of the installation process. Installation is anything that adds software to the device, possibly factory installation, the user installing elective applications and so on. The defining characteristic is that they are created by the software manufacturer. The purpose of these claims in an EAT is to relay them without modification to the Verifier and/or the Relying Party.

In some cases these will be signed by the software manufacturer independent of any signing for the purpose of EAT attestation. Manifest claims should include the manufacturer's signature (which will be signed over by the attestation signature). In other cases the attestation signature will be the only one.

This claim allows multiple formats for the manifest. For example the manifest may be a CBOR-format CoSWID, an XML-format SWID or other. Identification of the type of manifest is always by a CBOR tag. In many cases, for examples CoSWID, a tag will already be registered with IANA. If not, a tag MUST be registered. It can be in the first-come-first-served space which has minimal requirements for registration.

The claim is an array of one or more manifests. To facilitate hand off of the manifest to a decoding library, each manifest is contained in a byte string. This occurs for CBOR-format manifests as well as non-CBOR format manifests.

If a particular manifest type uses CBOR encoding, then the item in the array for it MUST be a byte string that contains a CBOR tag. The EAT decoder must decode the byte string and then the CBOR within it to find the tag number to identify the type of manifest. The contents of the byte string is then handed to the particular manifest processor for that type of manifest. CoSWID and SUIT manifest are examples of this.

If a particular manifest type does not use CBOR encoding, then the item in the array for it must be a CBOR tag that contains a byte string. The EAT decoder uses the tag to identify the processor for that type of manifest. The contents of the tag, the byte string, are handed to the manifest processor. Note that a byte string is used to contain the manifest whether it is a text based format or not. An example of this is an XML format ISO/IEC 19770 SWID.

It is not possible to describe the above requirements in CDDL so the type for an individual manifest is any in the CDDL below. The above text sets the encoding requirement.

This claim allows for multiple manifests in one token since multiple software packages are likely to be present. The multiple manifests may be of multiple formats. In some cases EAT submodules may be used instead of the array structure in this claim for multiple manifests.

When the [CoSWID] format is used, it MUST be a payload CoSWID, not an evidence CoSWID.

```
manifests-claim = (  
    manifests => manifests-type  
)  
  
manifests-type = [+ $manifest-formats]  
  
; Must be a CoSWID payload type  
$manifest-formats /= bytes .cbor concise-swid-tag  
  
$manifest-formats /= bytes .cbor SUIT_Envelope_Tagged
```

3.19. The Software Evidence Claim {swevidence}

This claim contains descriptions, lists, evidence or measurements of the software that exists on the device. The defining characteristic of this claim is that its contents are created by processes on the device that inventory, measure or otherwise characterize the software on the device. The contents of this claim do not originate from the software manufacturer.

In most cases the contents of this claim are signed as part of attestation signing, but independent signing in addition to the attestation signing is not ruled out when a particular evidence format supports it.

This claim uses the same mechanism for identification of the type of the swevidence as is used for the type of the manifest in the

manifests claim. It also uses the same byte string based mechanism for containing the claim and easing the hand off to a processing library. See the discussion above in the manifests claim.

When the [CoSWID] format is used, it MUST be evidence CoSWIDS, not payload CoSWIDS.

```
swevidence-claim = (  
    swevidence => swevidence-type  
)  
  
swevidence-type = [+ $swevidence-formats]  
  
; Must be a CoSWID evidence type  
$swevidence-formats /= bytes .cbor concise-swid-tag
```

3.20. The Submodules Part of a Token (submods)

Some devices are complex, having many subsystems or submodules. A mobile phone is a good example. It may have several connectivity submodules for communications (e.g., Wi-Fi and cellular). It may have subsystems for low-power audio and video playback. It may have one or more security-oriented subsystems like a TEE or a Secure Element.

The claims for each these can be grouped together in a submodule.

The submods part of a token are in a single map/object with many entries, one per submodule. There is only one submods map in a token. It is identified by its specific label. It is a peer to other claims, but it is not called a claim because it is a container for a claim set rather than an individual claim. This submods part of a token allows what might be called recursion. It allows claim sets inside of claim sets inside of claims sets...

3.20.1. Two Types of Submodules

Each entry in the submod map is one of two types:

- o A non-token submodule that is a map or object directly containing claims for the submodule.
- o A nested EAT that is a fully formed, independently signed EAT token

3.20.1.1. Non-token Submodules

This is simply a map or object containing claims about the submodule.

It may contain claims that are the same as its surrounding token or superior submodules. For example, the top-level of the token may have a UEID, a submod may have a different UEID and a further subordinate submodule may also have a UEID.

It is signed/encrypted along with the rest of the token and thus the claims are secured by the same Attester with the same signing key as the rest of the token.

If a token is in CBOR format (a CWT or a UCCS), all non-token submodules must be CBOR format. If a token is in JSON format (a JWT), all non-token submodules must be in JSON format.

When decoding, this type of submodule is recognized from the other type by being a data item of type map for CBOR or type object for JSON.

3.20.1.2. Nested EATs

This type of submodule is a fully formed secured EAT as defined in this document except that it MUST NOT be a UCCS or an unsecured JWT. A nested token that is one that is always secured using COSE or JOSE, usually by an independent Attester. When the surrounding EAT is a CWT or secured JWT, the nested token becomes securely bound with the other claims in the surrounding token.

It is allowed to have a CWT as a submodule in a JWT and vice versa, but this SHOULD be avoided unless necessary.

3.20.1.2.1. Surrounding EAT is CBOR format

The type of an EAT nested in a CWT is determined by whether the CBOR type is a text string or a byte string. If a text string, then it is a JWT. If a byte string, then it is a CWT.

A CWT nested in a CBOR-format token is always wrapped by a byte string for easier handling with standard CBOR decoders and token processing APIs that will typically take a byte buffer as input.

Nested CWTs may be either a CWT CBOR tag or a CWT Protocol Message. COSE layers in nested CWT EATs MUST be a COSE_Tagged_Message, never a COSE_Untagged_Message. If a nested EAT has more than one level of COSE, for example one that is both encrypted and signed, a COSE_Tagged_message must be used at every level.

3.20.1.2.2. Surrounding EAT is JSON format

When a CWT is nested in a JWT, it must be as a 55799 tag in order to distinguish it from a nested JWT.

When a nested EAT in a JWT is decoded, first remove the base64url encoding. Next, check to see if it starts with the bytes 0xd9d9f7. If so, then it is a CWT as a JWT will never start with these four bytes. If not if it is a JWT.

Other than the 55799 tag requirement, tag usage for CWT's nested in a JSON format token follow the same rules as for CWTs nested in CBOR-format tokens. It may be a CWT CBOR tag or a CWT Protocol Message and COSE_Tagged_Message MUST be used at all COSE layers.

3.20.1.3. Unsecured JWTs and UCCS Tokens as Submodules

To incorporate a UCCS token as a submodule, it MUST be as a non-token submodule. This can be accomplished inserting the content of the UCCS Tag into the submodule map. The content of a UCCS tag is exactly a map of claims as required for a non-token submodule. If the UCCS is not a UCCS tag, then it can just be inserted into the submodule map directly.

The definition of a nested EAT type of submodule is that it is one that is secured (signed) by an Attester. Since UCCS tokens are unsecured, they do not fulfill this definition and must be non-token submodules.

To incorporate an Unsecured JWT as a submodule, the null-security JOSE wrapping should be removed. The resulting claims set should be inserted as a non-token submodule.

To incorporate a UCCS token in a surrounding JSON token, the UCCS token claims should be translated from CBOR to JSON. To incorporate an Unsecured JWT into a surrounding CBOR-format token, the null-security JOSE should be removed and the claims translated from JSON to CBOR.

3.20.2. No Inheritance

The subordinate modules do not inherit anything from the containing token. The subordinate modules must explicitly include all of their claims. This is the case even for claims like the nonce and age.

This rule is in place for simplicity. It avoids complex inheritance rules that might vary from one type of claim to another.

3.20.3. Security Levels

The security level of the non-token subordinate modules should always be less than or equal to that of the containing modules in the case of non-token submodules. It makes no sense for a module of lesser security to be signing claims of a module of higher security. An example of this is a TEE signing claims made by the non-TEE parts (e.g. the high-level OS) of the device.

The opposite may be true for the nested tokens. They usually have their own more secure key material. An example of this is an embedded secure element.

3.20.4. Submodule Names

The label or name for each submodule in the submods map is a text string naming the submodule. No submodules may have the same name.

3.20.5. submods CDDL

```
; The part of a token that contains all the submodules. It is a peer
; with the claims in the token, but not a claim, only a map/object to
; hold all the submodules.
```

```
submods-part = (
    submods => submods-type
)
```

```
submods-type = { + submod-type }
```

```
; The type of a submodule which can either be a nested claim set or a
; nested separately signed token. Nested tokens are wrapped in a bstr
; or a tstr.
```

```
submod-type = (
    submod-name => eat-claim-set / nested-token
)
```

```
; When this is a bstr, the contents are an eat-token in CWT or UCCS
; format. When this is a tstr, the contents are an eat-token in JWT
; format.
```

```
nested-token = bstr / tstr;
```

```
; Each submodule has a unique text string name.
```

```
submod-name = tstr
```

4. Endorsements and Verification Keys

The Verifier must possess the correct key when it performs the cryptographic part of an EAT verification (e.g., verifying the COSE signature). This section describes several ways to identify the verification key. There is not one standard method.

The verification key itself may be a public key, a symmetric key or something complicated in the case of a scheme like Direct Anonymous Attestation (DAA).

RATS Architecture [RATS.Architecture] describes what is called an Endorsement. This is an input to the Verifier that is usually the basis of the trust placed in an EAT and the Attester that generated it. It may contain the public key for verification of the signature

on the EAT. It may contain Reference Values to which EAT claims are compared as part of the verification process. It may contain implied claims, those that are passed on to the Relying Party in Attestation Results.

There is not yet any standard format(s) for an Endorsement. One format that may be used for an Endorsement is an X.509 certificate. Endorsement data like Reference Values and implied claims can be carried in X.509 v3 extensions. In this use, the public key in the X.509 certificate becomes the verification key, so identification of the Endorsement is also identification of the verification key.

The verification key identification and establishment of trust in the EAT and the attester may also be by some other means than an Endorsement.

For the components (Attester, Verifier, Relying Party,...) of a particular end-end attestation system to reliably interoperate, its definition should specify how the verification key is identified. Usually, this will be in the profile document for a particular attestation system.

4.1. Identification Methods

Following is a list of possible methods of key identification. A specific attestation system may employ any one of these or one not listed here.

The following assumes Endorsements are X.509 certificates or equivalent and thus does not mention or define any identifier for Endorsements in other formats. If such an Endorsement format is created, new identifiers for them will also need to be created.

4.1.1. COSE/JWS Key ID

The COSE standard header parameter for Key ID (kid) may be used. See [RFC8152] and [RFC7515]

COSE leaves the semantics of the key ID open-ended. It could be a record locator in a database, a hash of a public key, an input to a KDF, an authority key identifier (AKI) for an X.509 certificate or other. The profile document should specify what the key ID's semantics are.

4.1.2. JWS and COSE X.509 Header Parameters

COSE X.509 [COSE.X509.Draft] and JSON Web Signature [RFC7515] define several header parameters (x5t, x5u,...) for referencing or carrying X.509 certificates any of which may be used.

The X.509 certificate may be an Endorsement and thus carrying additional input to the Verifier. It may be just an X.509 certificate, not an Endorsement. The same header parameters are used in both cases. It is up to the attestation system design and the Verifier to determine which.

4.1.3. CBOR Certificate COSE Header Parameters

Compressed X.509 and CBOR Native certificates are defined by CBOR Certificates [CBOR.Cert.Draft]. These are semantically compatible with X.509 and therefore can be used as an equivalent to X.509 as described above.

These are identified by their own header parameters (c5t, c5u,...).

4.1.4. Claim-Based Key Identification

For some attestation systems, a claim may be re-used as a key identifier. For example, the UEID uniquely identifies the device and therefore can work well as a key identifier or Endorsement identifier.

This has the advantage that key identification requires no additional bytes in the EAT and makes the EAT smaller.

This has the disadvantage that the unverified EAT must be substantially decoded to obtain the identifier since the identifier is in the COSE/JOSE payload, not in the headers.

4.2. Other Considerations

In all cases there must be some way that the verification key is itself verified or determined to be trustworthy. The key identification itself is never enough. This will always be by some out-of-band mechanism that is not described here. For example, the Verifier may be configured with a root certificate or a master key by the Verifier system administrator.

Often an X.509 certificate or an Endorsement carries more than just the verification key. For example, an X.509 certificate might have key usage constraints and an Endorsement might have Reference Values. When this is the case, the key identifier must be either a protected

header or in the payload such that it is cryptographically bound to the EAT. This is in line with the requirements in section 6 on Key Identification in JSON Web Signature [RFC7515].

5. Profiles

This EAT specification does not guarantee that implementations of it will interoperate. The variability in this specification is necessary to accommodate the widely varying use cases. An EAT profile narrows the specification for a specific use case. An ideal EAT profile will guarantee interoperability.

The profile can be named in the token using the profile claim described in Section 3.17.

5.1. Format of a Profile Document

A profile document doesn't have to be in any particular format. It may be simple text, something more formal or a combination.

In some cases CDDL may be created that replaces CDDL in this or other document to express some profile requirements. For example, to require the altitude data item in the location claim, CDDL can be written that replicates the location claim with the altitude no longer optional.

5.2. List of Profile Issues

The following is a list of EAT, CWT, UCCS, JWS, COSE, JOSE and CBOR options that a profile should address.

5.2.1. Use of JSON, CBOR or both

The profile should indicate whether the token format should be CBOR, JSON, both or even some other encoding. If some other encoding, a specification for how the CDDL described here is serialized in that encoding is necessary.

This should be addressed for the top-level token and for any nested tokens. For example, a profile might require all nested tokens to be of the same encoding of the top level token.

5.2.2. CBOR Map and Array Encoding

The profile should indicate whether definite-length arrays/maps, indefinite-length arrays/maps or both are allowed. A good default is to allow only definite-length arrays/maps.

An alternate is to allow both definite and indefinite-length arrays/maps. The decoder should accept either. Encoders that need to fit on very small hardware or be actually implement in hardware can use indefinite-length encoding.

This applies to individual EAT claims, CWT and COSE parts of the implementation.

5.2.3. CBOR String Encoding

The profile should indicate whether definite-length strings, indefinite-length strings or both are allowed. A good default is to allow only definite-length strings. As with map and array encoding, allowing indefinite-length strings can be beneficial for some smaller implementations.

5.2.4. CBOR Preferred Serialization

The profile should indicate whether encoders must use preferred serialization. The profile should indicate whether decoders must accept non-preferred serialization.

5.2.5. COSE/JOSE Protection

COSE and JOSE have several options for signed, MACed and encrypted messages. EAT/CWT has the option to have no protection using UCCS and JOSE has a NULL protection option. It is possible to implement no protection, sign only, MAC only, sign then encrypt and so on. All combinations allowed by COSE, JOSE, JWT, CWT and UCCS are allowed by EAT.

The profile should list the protections that must be supported by all decoders implementing the profile. The encoders then must implement a subset of what is listed for the decoders, perhaps only one.

Implementations may choose to sign or MAC before encryption so that the implementation layer doing the signing or MACing can be the smallest. It is often easier to make smaller implementations more secure, perhaps even implementing in solely in hardware. The key material for a signature or MAC is a private key, while for encryption it is likely to be a public key. The key for encryption requires less protection.

5.2.6. COSE/JOSE Algorithms

The profile document should list the COSE algorithms that a Verifier must implement. The Attester will select one of them. Since there

is no negotiation, the Verifier should implement all algorithms listed in the profile.

5.2.7. Verification Key Identification

Section 4 describes a number of methods for identifying a verification key. The profile document should specify one of these or one that is not described. The ones described in this document are only roughly described. The profile document should go into the full detail.

5.2.8. Endorsement Identification

Similar to, or perhaps the same as Verification Key Identification, the profile may wish to specify how Endorsements are to be identified. However note that Endorsement Identification is optional, where as key identification is not.

5.2.9. Freshness

Just about every use case will require some means of knowing the EAT is recent enough and not a replay of an old token. The profile should describe how freshness is achieved. The section on Freshness in [RATS-Architecture] describes some of the possible solutions to achieve this.

5.2.10. Required Claims

The profile can list claims whose absence results in Verification failure.

5.2.11. Prohibited Claims

The profile can list claims whose presence results in Verification failure.

5.2.12. Additional Claims

The profile may describe entirely new claims. These claims can be required or optional.

5.2.13. Refined Claim Definition

The profile may lock down optional aspects of individual claims. For example, it may require altitude in the location claim, or it may require that HW Versions always be described using EAN-13.

5.2.14. CBOR Tags

The profile should specify whether the token should be a CWT Tag or not. Similarly, the profile should specify whether the token should be a UCCS tag or not.

When COSE protection is used, the profile should specify whether COSE tags are used or not. Note that RFC 8392 requires COSE tags be used in a CWT tag.

Often a tag is unnecessary because the surrounding or carrying protocol identifies the object as an EAT.

5.2.15. Manifests and Software Evidence Claims

The profile should specify which formats are allowed for the manifests and software evidence claims. The profile may also go on to say which parts and options of these formats are used, allowed and prohibited.

6. Encoding

This makes use of the types defined in CDDL Appendix D, Standard Prelude.

Some of the CDDL included here is for claims that are defined in CWT [RFC8392] or JWT [RFC7519] or are in the IANA CWT or JWT registries. CDDL was not in use when these claims were defined.

6.1. Common CDDL Types

time-int is identical to the epoch-based time, but disallows floating-point representation.

Note that unless explicitly indicated, URIs are not the URI tag defined in [RFC8949]. They are just text strings that contain a URI.

string-or-uri = tstr

time-int = #6.1(int)

6.2. CDDL for CWT-defined Claims

This section provides CDDL for the claims defined in CWT. It is non-normative as [RFC8392] is the authoritative definition of these claims.

Note that the subject, issue and audience claims may be a text string containing a URI per [RFC8392] and [RFC7519]. These are never the URI tag defined in [RFC8949].

```
$seat-extension //= (  
  ? issuer => text,  
  ? subject => text,  
  ? audience => text,  
  ? expiration => time,  
  ? not-before => time,  
  ? issued-at => time,  
  ? cwt-id => bytes,  
)
```

```
issuer = 1  
subject = 2  
audience = 3  
expiration = 4  
not-before = 5  
issued-at = 6  
cwt-id = 7
```

6.3. JSON

6.3.1. JSON Labels

; The following are Claim Keys (labels) assigned for JSON-encoded tokens.

```
ueid /= "ueid"
sueids /= "sueids"
nonce /= "nonce"
oemid /= "oemid"
security-level /= "secllevel"
secure-boot /= "secboot"
debug-status /= "dbgstat"
location /= "location"
uptime /= "uptime"
profile /= "eat-profile"
intended-use /= "intuse"
boot-seed /= "bootseed"
submods /= "submods"
timestamp /= "timestamp"
manifests /= "manifests"
swevidence /= "swevidence"

latitude /= "lat"
longitude /= "long"
altitude /= "alt"
accuracy /= "accry"
altitude-accuracy /= "alt-accry"
heading /= "heading"
speed /= "speed"
```

6.3.2. JSON Interoperability

JSON should be encoded per RFC 8610 Appendix E. In addition, the following CDDL types are encoded in JSON as follows:

- o bstr - must be base64url encoded
- o time - must be encoded as NumericDate as described section 2 of [RFC7519].
- o string-or-uri - must be encoded as StringOrURI as described section 2 of [RFC7519].
- o uri - must be a URI [RFC3986].
- o oid - encoded as a string using the well established dotted-decimal notation (e.g., the text "1.2.250.1").

6.4. CBOR

6.4.1. CBOR Interoperability

CBOR allows data items to be serialized in more than one form. If the sender uses a form that the receiver can't decode, there will not be interoperability.

This specification gives no blanket requirements to narrow CBOR serialization for all uses of EAT. This allows individual uses to tailor serialization to the environment. It also may result in EAT implementations that don't interoperate.

One way to guarantee interoperability is to clearly specify CBOR serialization in a profile document. See Section 5 for a list of serialization issues that should be addressed.

EAT will be commonly used where the device generating the attestation is constrained and the receiver/verifier of the attestation is a capacious server. Following is a set of serialization requirements that work well for that use case and are guaranteed to interoperate. Use of this serialization is recommended where possible, but not required. An EAT profile may just reference the following section rather than spell out serialization details.

6.4.1.1. EAT Constrained Device Serialization

- o Preferred serialization described in section 4.1 of [RFC8949] is not required. The EAT decoder must accept all forms of number serialization. The EAT encoder may use any form it wishes.
- o The EAT decoder must accept indefinite length arrays and maps as described in section 3.2.2 of [RFC8949]. The EAT encoder may use indefinite length arrays and maps if it wishes.
- o The EAT decoder must accept indefinite length strings as described in section 3.2.3 of [RFC8949]. The EAT encoder may use indefinite length strings if it wishes.
- o Sorting of maps by key is not required. The EAT decoder must not rely on sorting.
- o Deterministic encoding described in Section 4.2 of [RFC8949] is not required.
- o Basic validity described in section 5.3.1 of [RFC8949] must be followed. The EAT encoder must not send duplicate map keys/labels or invalid UTF-8 strings.

6.5. Collected CDDL

```
; This is the top-level definition of the claims in EAT tokens. To
; form an actual EAT Token, this claim set is enclosed in a COSE, JOSE
; or UCCS message.
```

```
eat-claim-set = {
    ? ueid-claim,
    ? sueids-claim,
    ? nonce-claim,
    ? oemid-claim,
    ? hardware-version-claims,
    ? security-level-claim,
    ? secure-boot-claim,
    ? debug-status-claim,
    ? location-claim,
    ? intended-use-claim,
    ? profile-claim,
    ? uptime-claim,
    ? manifests-claim,
    ? swevidence-claim,
    ? submods-part,
    * $$eat-extension,
}
```

```
; This is the top-level definition of an EAT Token. It is a CWT, JWT
; or UCSS where the payload is an eat-claim-set. A JWT_Message is what
; is defined by JWT in RFC 7519. (RFC 7519 doesn't use CDDL so a there
; is no actual CDDL definition of JWT_Message).
```

```
eat-token = EAT_Tagged_Message / EAT_Untagged_Message / JWT_Message
```

```
; This is CBOR-format EAT token in the CWT or UCCS format that is a
; tag. COSE_Tagged_message is defined in RFC 8152. Tag 601 is
; proposed by the UCCS draft, but not yet assigned.
```

```
EAT_Tagged_Message = #6.61(COSE_Tagged_Message) / #6.601(eat-claim-set)
```

```
; This is a CBOR-format EAT token that is a CWT or UCSS that is not a
; tag COSE_Tagged_message and COSE_Untagged_Message are defined in RFC
; 8152.
```

```
EAT_Untagged_Message = COSE_Tagged_Message / COSE_Untagged_Message / UCCS_Untagge
d_Message
```



```
; This is an "unwrapped" UCCS tag. Unwrapping a tag means to use the
; definition of its content without the preceding type 6 tag
; integer. Since a UCCS is nothing but a tag for an unsecured CWT
; claim set, unwrapping reduces to a bare eat-claim-set.
```

```
UCCS_Untagged_Message = eat-claim-set
```

```
string-or-uri = tstr
```

```
time-int = #6.1(int)
$$eat-extension //= (
    ? issuer => text,
    ? subject => text,
    ? audience => text,
    ? expiration => time,
    ? not-before => time,
    ? issued-at => time,
    ? cwt-id => bytes,
)
```

```
issuer = 1
subject = 2
audience = 3
expiration = 4
not-before = 5
issued-at = 6
cwt-id = 7
```

```
debug-status-cbor-type = &(
    enabled: 0,
    disabled: 1,
    disabled-since-boot: 2,
    disabled-permanently: 3,
    disabled-fully-and-permanently: 4
)
```

```
debug-status-json-type =
    "enabled" /
    "disabled" /
    "disabled-since-boot" /
    "disabled-permanently" /
    "disabled-fully-and-permanently"
```

```
debug-status-claim = (
    debug-status => debug-status-cbor-type / debug-status-json-type
)
location-type = {
    latitude => number,
```

```
    longitude => number,
    ? altitude => number,
    ? accuracy => number,
    ? altitude-accuracy => number,
    ? heading => number,
    ? speed => number,
    ? timestamp => ~time-int,
    ? age => uint
}

latitude = 1 / "latitude"
longitude = 2 / "longitude"
altitude = 3 / "altitude"
accuracy = 4 / "accuracy"
altitude-accuracy = 5 / "altitude-accuracy"
heading = 6 / "heading"
speed = 7 / "speed"
timestamp = 8 / "timestamp"
age = 9 / "age"

location-claim = (
    location-label => location-type
)
nonce-type = bstr .size (8..64)

nonce-claim = (
    nonce => nonce-type / [ 2* nonce-type ]
)
oemid-claim = (
    oemid => bstr
)
chip-version-claim = (
    chip-version => tstr
)

chip-version-scheme-claim = (
    chip-version-scheme => $version-scheme
)

board-version-claim = (
    board-version => tstr
)

board-version-scheme-claim = (
    board-version-scheme => $version-scheme
)

device-version-claim = (
```

```

    device-version => tstr
)

device-version-scheme-claim = (
    device-version-scheme => $version-scheme
)

hardware-version-claims = (
    ? chip-version-claim,
    ? board-version-claim,
    ? device-version-claim,
    ? chip-version-scheme-claim,
    ? board-version-scheme-claim,
    ? device-version-scheme-claim,
)

secure-boot-claim = (
    secure-boot => bool
)

security-level-cbor-type = &(
    unrestricted: 1,
    restricted: 2,
    secure-restricted: 3,
    hardware: 4
)

security-level-json-type =
    "unrestricted" /
    "restricted" /
    "secure-restricted" /
    "hardware"

security-level-claim = (
    security-level => security-level-cbor-type / security-level-json-type
)
; The part of a token that contains all the submodules. It is a peer
; with the claims in the token, but not a claim, only a map/object to
; hold all the submodules.

submods-part = (
    submods => submods-type
)

submods-type = { + submod-type }

; The type of a submodule which can either be a nested claim set or a

```

```
; nested separately signed token. Nested tokens are wrapped in a bstr  
; or a tstr.
```

```
submod-type = (  
    submod-name => eat-claim-set / nested-token  
)
```

```
; When this is a bstr, the contents are an eat-token in CWT or UCCS  
; format. When this is a tstr, the contents are an eat-token in JWT  
; format.
```

```
nested-token = bstr / tstr;
```

```
; Each submodule has a unique text string name.
```

```
submod-name = tstr
```

```
ueid-type = bstr .size (7..33)
```

```
ueid-claim = (  
    ueid => ueid-type  
)
```

```
sueids-type = {  
    + tstr => ueid-type  
}
```

```
sueids-claim = (  
    sueids => sueids-type  
)
```

```
intended-use-chor-type = &(  
    generic: 1,  
    registration: 2,  
    provisioning: 3,  
    csr: 4,  
    pop: 5  
)
```

```
intended-use-json-type =  
    "generic" /  
    "registration" /  
    "provisioning" /  
    "csr" /  
    "pop"
```

```
intended-use-claim = (  

```

```
    intended-use => intended-use-cbor-type / intended-use-json-type
)
oid = #6.4000(bstr) ; TODO: fill this in with correct CDDL from OID RFC

uptime-claim = (
    uptime => uint
)

manifests-claim = (
    manifests => manifests-type
)

manifests-type = [+ $manifest-formats]

; Must be a CoSWID payload type
$manifest-formats /= bytes .cbor concise-swid-tag

$manifest-formats /= bytes .cbor SUIT_Envelope_Tagged

swevidence-claim = (
    swevidence => swevidence-type
)

swevidence-type = [+ $swevidence-formats]

; Must be a CoSWID evidence type
$swevidence-formats /= bytes .cbor concise-swid-tag

oid = #6.4000(bstr) ; TODO: fill this in with correct CDDL from OID RFC

profile-claim = (
    profile => ~uri / ~oid
)

boot-seed-claim = (
    boot-seed => bytes
)
```

7. IANA Considerations

7.1. Reuse of CBOR Web Token (CWT) Claims Registry

Claims defined for EAT are compatible with those of CWT so the CWT Claims Registry is re used. No new IANA registry is created. All EAT claims should be registered in the CWT and JWT Claims Registries.

7.2. Claim Characteristics

The following is design guidance for creating new EAT claims, particularly those to be registered with IANA.

Much of this guidance is generic and could also be considered when designing new CWT or JWT claims.

7.2.1. Interoperability and Relying Party Orientation

It is a broad goal that EATs can be processed by relying parties in a general way regardless of the type, manufacturer or technology of the device from which they originate. It is a goal that there be general-purpose verification implementations that can verify tokens for large numbers of use cases with special cases and configurations for different device types. This is a goal of interoperability of the semantics of claims themselves, not just of the signing, encoding and serialization formats.

This is a lofty goal and difficult to achieve broadly requiring careful definition of claims in a technology neutral way. Sometimes it will be difficult to design a claim that can represent the semantics of data from very different device types. However, the goal remains even when difficult.

7.2.2. Operating System and Technology Neutral

Claims should be defined such that they are not specific to an operating system. They should be applicable to multiple large high-level operating systems from different vendors. They should also be applicable to multiple small embedded operating systems from multiple vendors and everything in between.

Claims should not be defined such that they are specific to a SW environment or programming language.

Claims should not be defined such that they are specific to a chip or particular hardware. For example, they should not just be the contents of some HW status register as it is unlikely that the same HW status register with the same bits exists on a chip of a different manufacturer.

The boot and debug state claims in this document are an example of a claim that has been defined in this neutral way.

7.2.3. Security Level Neutral

Many use cases will have EATs generated by some of the most secure hardware and software that exists. Secure Elements and smart cards are examples of this. However, EAT is intended for use in low-security use cases the same as high-security use case. For example, an app on a mobile device may generate EATs on its own.

Claims should be defined and registered on the basis of whether they are useful and interoperable, not based on security level. In particular, there should be no exclusion of claims because they are just used only in low-security environments.

7.2.4. Reuse of Extant Data Formats

Where possible, claims should use already standardized data items, identifiers and formats. This takes advantage of the expertise put into creating those formats and improves interoperability.

Often extant claims will not be defined in an encoding or serialization format used by EAT. It is preferred to define a CBOR and JSON format for them so that EAT implementations do not require a plethora of encoders and decoders for serialization formats.

In some cases, it may be better to use the encoding and serialization as is. For example, signed X.509 certificates and CRLs can be carried as-is in a byte string. This retains interoperability with the extensive infrastructure for creating and processing X.509 certificates and CRLs.

7.2.5. Proprietary Claims

EAT allows the definition and use of proprietary claims.

For example, a device manufacturer may generate a token with proprietary claims intended only for verification by a service offered by that device manufacturer. This is a supported use case.

In many cases proprietary claims will be the easiest and most obvious way to proceed, however for better interoperability, use of general standardized claims is preferred.

7.3. Claims Registered by This Document

This specification adds the following values to the "JSON Web Token Claims" registry established by [RFC7519] and the "CBOR Web Token Claims Registry" established by [RFC8392]. Each entry below is an

addition to both registries (except for the nonce claim which is already registered for JWT, but not registered for CWT).

The "Claim Description", "Change Controller" and "Specification Documents" are common and equivalent for the JWT and CWT registries. The "Claim Key" and "Claim Value Types(s)" are for the CWT registry only. The "Claim Name" is as defined for the CWT registry, not the JWT registry. The "JWT Claim Name" is equivalent to the "Claim Name" in the JWT registry.

7.3.1. Claims for Early Assignment

RFC Editor: in the final publication this section should be combined with the following section as it will no longer be necessary to distinguish claims with early assignment. Also, the following paragraph should be removed.

The claims in this section have been (requested for / given) early assignment according to [RFC7120]. They have been assigned values and registered before final publication of this document. While their semantics is not expected to change in final publication, it is possible that they will. The JWT Claim Names and CWT Claim Keys are not expected to change.

- o Claim Name: Nonce
- o Claim Description: Nonce
- o JWT Claim Name: "nonce" (already registered for JWT)
- o Claim Key: 10
- o Claim Value Type(s): byte string
- o Change Controller: IESG
- o Specification Document(s): [OpenIDConnectCore], *this document*
- o Claim Name: UEID
- o Claim Description: The Universal Entity ID
- o JWT Claim Name: "ueid"
- o CWT Claim Key: 11
- o Claim Value Type(s): byte string

- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: OEMID
- o Claim Description: IEEE-based OEM ID
- o JWT Claim Name: "oemid"
- o Claim Key: 13
- o Claim Value Type(s): byte string
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Security Level
- o Claim Description: Characterization of the security of an Attester or submodule
- o JWT Claim Name: "secllevel"
- o Claim Key: 14
- o Claim Value Type(s): integer
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Secure Boot
- o Claim Description: Indicate whether the boot was secure
- o JWT Claim Name: "secboot"
- o Claim Key: 15
- o Claim Value Type(s): Boolean
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Debug Status

- o Claim Description: Indicate status of debug facilities
- o JWT Claim Name: "dbgstat"
- o Claim Key: 16
- o Claim Value Type(s): integer
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Location
- o Claim Description: The geographic location
- o JWT Claim Name: "location"
- o Claim Key: 17
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Profile
- o Claim Description: Indicates the EAT profile followed
- o JWT Claim Name: "eat_profile"
- o Claim Key: 18
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Submodules Section
- o Claim Description: The section containing submodules (not actually a claim)
- o JWT Claim Name: "submods"
- o Claim Key: 20

- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*

7.3.2. To be Assigned Claims

TODO: add the rest of the claims in here

7.3.3. Version Schemes Registered by this Document

IANA is requested to register a new value in the "Software Tag Version Scheme Values" established by [CoSWID].

The new value is a version scheme a 13-digit European Article Number [EAN-13]. An EAN-13 is also known as an International Article Number or most commonly as a bar code. This version scheme is the ASCII text representation of EAN-13 digits, the same ones often printed with a bar code. This version scheme must comply with the EAN allocation and assignment rules. For example, this requires the manufacturer to obtain a manufacture code from GS1.

Index	Version Scheme Name	Specification
5	ean-13	This document

8. Privacy Considerations

Certain EAT claims can be used to track the owner of an entity and therefore, implementations should consider providing privacy-preserving options dependent on the intended usage of the EAT. Examples would include suppression of location claims for EAT's provided to unauthenticated consumers.

8.1. UEID and SUEID Privacy Considerations

A UEID is usually not privacy-preserving. Any set of relying parties that receives tokens that happen to be from a single device will be able to know the tokens are all from the same device and be able to track the device. Thus, in many usage situations UEID violates governmental privacy regulation. In other usage situations a UEID will not be allowed for certain products like browsers that give privacy for the end user. It will often be the case that tokens will not have a UEID for these reasons.

An SUEID is also usually not privacy-preserving. In some cases it may have fewer privacy issues than a UEID depending on when and how and when it is generated.

There are several strategies that can be used to still be able to put UEIDs and SUEIDs in tokens:

- o The device obtains explicit permission from the user of the device to use the UEID/SUEID. This may be through a prompt. It may also be through a license agreement. For example, agreements for some online banking and brokerage services might already cover use of a UEID/SUEID.
- o The UEID/SUEID is used only in a particular context or particular use case. It is used only by one relying party.
- o The device authenticates the relying party and generates a derived UEID/SUEID just for that particular relying party. For example, the relying party could prove their identity cryptographically to the device, then the device generates a UEID just for that relying party by hashing a proofed relying party ID with the main device UEID/SUEID.

Note that some of these privacy preservation strategies result in multiple UEIDs and SUEIDs per device. Each UEID/SUEID is used in a different context, use case or system on the device. However, from the view of the relying party, there is just one UEID and it is still globally universal across manufacturers.

8.2. Location Privacy Considerations

Geographic location is most always considered personally identifiable information. Implementers should consider laws and regulations governing the transmission of location data from end user devices to servers and services. Implementers should consider using location management facilities offered by the operating system on the device generating the attestation. For example, many mobile phones prompt the user for permission when before sending location data.

9. Security Considerations

The security considerations provided in Section 8 of [RFC8392] and Section 11 of [RFC7519] apply to EAT in its CWT and JWT form, respectively. In addition, implementors should consider the following.

9.1. Key Provisioning

Private key material can be used to sign and/or encrypt the EAT, or can be used to derive the keys used for signing and/or encryption. In some instances, the manufacturer of the entity may create the key material separately and provision the key material in the entity itself. The manufacturer of any entity that is capable of producing an EAT should take care to ensure that any private key material be suitably protected prior to provisioning the key material in the entity itself. This can require creation of key material in an enclave (see [RFC4949] for definition of "enclave"), secure transmission of the key material from the enclave to the entity using an appropriate protocol, and persistence of the private key material in some form of secure storage to which (preferably) only the entity has access.

9.1.1. Transmission of Key Material

Regarding transmission of key material from the enclave to the entity, the key material may pass through one or more intermediaries. Therefore some form of protection ("key wrapping") may be necessary. The transmission itself may be performed electronically, but can also be done by human courier. In the latter case, there should be minimal to no exposure of the key material to the human (e.g. encrypted portable memory). Moreover, the human should transport the key material directly from the secure enclave where it was created to a destination secure enclave where it can be provisioned.

9.2. Transport Security

As stated in Section 8 of [RFC8392], "The security of the CWT relies upon on the protections offered by COSE". Similar considerations apply to EAT when sent as a CWT. However, EAT introduces the concept of a nonce to protect against replay. Since an EAT may be created by an entity that may not support the same type of transport security as the consumer of the EAT, intermediaries may be required to bridge communications between the entity and consumer. As a result, it is RECOMMENDED that both the consumer create a nonce, and the entity leverage the nonce along with COSE mechanisms for encryption and/or signing to create the EAT.

Similar considerations apply to the use of EAT as a JWT. Although the security of a JWT leverages the JSON Web Encryption (JWE) and JSON Web Signature (JWS) specifications, it is still recommended to make use of the EAT nonce.

9.3. Multiple EAT Consumers

In many cases, more than one EAT consumer may be required to fully verify the entity attestation. Examples include individual consumers for nested EATs, or consumers for individual claims with an EAT. When multiple consumers are required for verification of an EAT, it is important to minimize information exposure to each consumer. In addition, the communication between multiple consumers should be secure.

For instance, consider the example of an encrypted and signed EAT with multiple claims. A consumer may receive the EAT (denoted as the "receiving consumer"), decrypt its payload, verify its signature, but then pass specific subsets of claims to other consumers for evaluation ("downstream consumers"). Since any COSE encryption will be removed by the receiving consumer, the communication of claim subsets to any downstream consumer should leverage a secure protocol (e.g. one that uses transport-layer security, i.e. TLS),

However, assume the EAT of the previous example is hierarchical and each claim subset for a downstream consumer is created in the form of a nested EAT. Then transport security between the receiving and downstream consumers is not strictly required. Nevertheless, downstream consumers of a nested EAT should provide a nonce unique to the EAT they are consuming.

10. References

10.1. Normative References

[CBOR-OID]

Bormann, C., "Concise Binary Object Representation (CBOR) Tags for Object Identifiers", draft-ietf-cbor-tags-oid-06 (work in progress), March 2021.

[CBOR.Cert.Draft]

Raza, S., "CBOR Encoding of X.509 Certificates (CBOR Certificates)", 2020, <<https://tools.ietf.org/html/draft-mattsson-cose-cbor-cert-compress-05>>.

[COSE.X509.Draft]

Schaad, J., "CBOR Object Signing and Encryption (COSE): Header parameters for carrying and referencing X.509 certificates", 2020, <<https://tools.ietf.org/html/draft-ietf-cose-x509-08>>.

[CoSWID]

"Concise Software Identification Tags", November 2020, <<https://tools.ietf.org/html/draft-ietf-sacm-coswid-16>>.

- [EAN-13] GS1, "International Article Number - EAN/UPC barcodes", 2019, <<https://www.gs1.org/standards/barcodes/ean-upc>>.
- [FIDO.AROE]
The FIDO Alliance, "FIDO Authenticator Allowed Restricted Operating Environments List", November 2019, <<https://fidoalliance.org/specs/fido-uaf-v1.0-fd-20191115/fido-allowed-AROE-v1.0-fd-20191115.html>>.
- [IANA.CWT.Claims]
IANA, "CBOR Web Token (CWT) Claims", <<http://www.iana.org/assignments/cwt>>.
- [IANA.JWT.Claims]
IANA, "JSON Web Token (JWT) Claims", <<https://www.iana.org/assignments/jwt>>.
- [OpenIDConnectCore]
Sakimura, N., Bradley, J., Jones, M., Medeiros, B. D., and C. Mortimore, "OpenID Connect Core 1.0 incorporating errata set 1", November 2014, <https://openid.net/specs/openid-connect-core-1_0.html>.
- [RATS-Architecture]
Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", draft-ietf-rats-architecture-12 (work in progress), April 2021.
- [RATS.Architecture]
Birkholz, H., "Remote Attestation Procedures Architecture", 2020, <<https://tools.ietf.org/html/draft-ietf-rats-architecture-08>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.

- [RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC7800] Jones, M., Bradley, J., and H. Tschofenig, "Proof-of-Possession Key Semantics for JSON Web Tokens (JWTs)", RFC 7800, DOI 10.17487/RFC7800, April 2016, <<https://www.rfc-editor.org/info/rfc7800>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/info/rfc8392>>.
- [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://www.rfc-editor.org/info/rfc8610>>.
- [RFC8747] Jones, M., Seitz, L., Selander, G., Erdtman, S., and H. Tschofenig, "Proof-of-Possession Key Semantics for CBOR Web Tokens (CWTs)", RFC 8747, DOI 10.17487/RFC8747, March 2020, <<https://www.rfc-editor.org/info/rfc8747>>.
- [RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/RFC8949, December 2020, <<https://www.rfc-editor.org/info/rfc8949>>.

- [ThreeGPP.IMEI]
3GPP, "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification", 2019,
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=729>>.
- [UCCS.Draft]
Birkholz, H., "A CBOR Tag for Unprotected CWT Claims Sets", 2020,
<<https://tools.ietf.org/html/draft-birkholz-rats-uccs-01>>.
- [WGS84] National Imagery and Mapping Agency, "National Imagery and Mapping Agency Technical Report 8350.2, Third Edition", 2000, <<http://earth-info.nga.mil/GandG/publications/tr8350.2/wgs84fin.pdf>>.

10.2. Informative References

- [BirthdayAttack]
"Birthday attack",
<https://en.wikipedia.org/wiki/Birthday_attack>.
- [Common.Criteria]
"Common Criteria for Information Technology Security Evaluation", April 2017,
<<https://www.commoncriteriaportal.org/cc/>>.
- [ECMAScript]
"Ecma International, "ECMAScript Language Specification, 5.1 Edition", ECMA Standard 262", June 2011,
<<http://www.ecma-international.org/ecma-262/5.1/ECMA-262.pdf>>.
- [FIDO.Registry]
The FIDO Alliance, "FIDO Registry of Predefined Values", December 2019, <<https://fidoalliance.org/specs/common-specs/fido-registry-v2.1-ps-20191217.html>>.
- [FIPS-140]
National Institute of Standards, "Security Requirements for Cryptographic Modules", May 2001,
<<https://csrc.nist.gov/publications/detail/fips/140/2/final>>.

- [IEEE.802-2001]
"IEEE Standard For Local And Metropolitan Area Networks Overview And Architecture", 2007,
<<https://webstore.ansi.org/standards/ieee/ieee8022001r2007>>.
- [IEEE.802.1AR]
"IEEE Standard, "IEEE 802.1AR Secure Device Identifier"", December 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.
- [IEEE.RA] "IEEE Registration Authority",
<<https://standards.ieee.org/products-services/regauth/index.html>>.
- [OUI.Guide]
"Guidelines for Use of Extended Unique Identifier (EUI), Organizationally Unique Identifier (OUI), and Company ID (CID)", August 2017,
<<https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/tutorials/eui.pdf>>.
- [OUI.Lookup]
"IEEE Registration Authority Assignments",
<<https://regauth.standards.ieee.org/standards-ra-web/pub/view.html#registries>>.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", RFC 4122, DOI 10.17487/RFC4122, July 2005,
<<https://www.rfc-editor.org/info/rfc4122>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007,
<<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 7120, DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.
- [W3C.GeoLoc]
Worldwide Web Consortium, "Geolocation API Specification 2nd Edition", January 2018, <https://www.w3.org/TR/geolocation-API/#coordinates_interface>.

Appendix A. Examples

A.1. Very Simple EAT

This is shown in CBOR diagnostic form. Only the payload signed by COSE is shown.

```
{
  / issuer /          1: "joe",
  / nonce /           10: h'948f8860d13a463e8e',
  / UEID /            11: h'0198f50a4ff6c05861c8860d13a638ea',
  / secure-boot /     15: true,
  / debug-disable /   16: 3, / permanent-disable /
  / timestamp (iat) / 6: 1(1526542894)
}
```

A.2. Example with Submodules, Nesting and Security Levels

```
{
  / nonce /           10: h'948f8860d13a463e8e',
  / UEID /            11: h'0198f50a4ff6c05861c8860d13a638ea',
  / secure-boot /     15: true,
  / debug-disable /   16: 3, / permanent-disable /
  / timestamp (iat) / 6: 1(1526542894),
  / security-level /  14: 3, / secure restricted OS /
  / submods / 20: {
    / first submod, an Android Application /
    "Android App Foo" : {
      / security-level / 14: 1 / unrestricted /
    },

    / 2nd submod, A nested EAT from a secure element /
    "Secure Element Eat" :
      / an embedded EAT, bytes of which are not shown /
      h'420123',

    / 3rd submod, information about Linux Android /
    "Linux Android": {
      / security-level / 14: 1 / unrestricted /
    }
  }
}
```

Appendix B. UEID Design Rationale

B.1. Collision Probability

This calculation is to determine the probability of a collision of UEIDs given the total possible entity population and the number of entities in a particular entity management database.

Three different sized databases are considered. The number of devices per person roughly models non-personal devices such as traffic lights, devices in stores they shop in, facilities they work in and so on, even considering individual light bulbs. A device may have individually attested subsystems, for example parts of a car or a mobile phone. It is assumed that the largest database will have at most 10% of the world's population of devices. Note that databases that handle more than a trillion records exist today.

The trillion-record database size models an easy-to-imagine reality over the next decades. The quadrillion-record database is roughly at the limit of what is imaginable and should probably be accommodated. The 100 quadrillion database is highly speculative perhaps involving nanorobots for every person, livestock animal and domesticated bird. It is included to round out the analysis.

Note that the items counted here certainly do not have IP address and are not individually connected to the network. They may be connected to internal buses, via serial links, Bluetooth and so on. This is not the same problem as sizing IP addresses.

People	Devices / Person	Subsystems / Device	Database Portion	Database Size
10 billion	100	10	10%	trillion (10 ¹²)
10 billion	100,000	10	10%	quadrillion (10 ¹⁵)
100 billion	1,000,000	10	10%	100 quadrillion (10 ¹⁷)

This is conceptually similar to the Birthday Problem where m is the number of possible birthdays, always 365, and k is the number of people. It is also conceptually similar to the Birthday Attack where collisions of the output of hash functions are considered.

The proper formula for the collision calculation is

$$p = 1 - e^{\{-k^2/(2n)\}}$$

p Collision Probability
 n Total possible population
 k Actual population

However, for the very large values involved here, this formula requires floating point precision higher than commonly available in calculators and SW so this simple approximation is used. See [BirthdayAttack].

$$p = k^2 / 2n$$

For this calculation:

p Collision Probability
 n Total population based on number of bits in UEID
 k Population in a database

Database Size	128-bit UEID	192-bit UEID	256-bit UEID
trillion (10 ¹²)	2 * 10 ⁻¹⁵	8 * 10 ⁻³⁵	5 * 10 ⁻⁵⁵
quadrillion (10 ¹⁵)	2 * 10 ⁻⁰⁹	8 * 10 ⁻²⁹	5 * 10 ⁻⁴⁹
100 quadrillion (10 ¹⁷)	2 * 10 ⁻⁰⁵	8 * 10 ⁻²⁵	5 * 10 ⁻⁴⁵

Next, to calculate the probability of a collision occurring in one year's operation of a database, it is assumed that the database size is in a steady state and that 10% of the database changes per year. For example, a trillion record database would have 100 billion states per year. Each of those states has the above calculated probability of a collision.

This assumption is a worst-case since it assumes that each state of the database is completely independent from the previous state. In reality this is unlikely as state changes will be the addition or deletion of a few records.

The following tables gives the time interval until there is a probability of a collision based on there being one tenth the number of states per year as the number of records in the database.

$$t = 1 / ((k / 10) * p)$$

t Time until a collision
 p Collision probability for UEID size
 k Database size

Database Size	128-bit UEID	192-bit UEID	256-bit UEID
trillion (10 ¹²)	60,000 years	10 ²⁴ years	10 ⁴⁴ years
quadrillion (10 ¹⁵)	8 seconds	10 ¹⁴ years	10 ³⁴ years
100 quadrillion (10 ¹⁷)	8 microseconds	10 ¹¹ years	10 ³¹ years

Clearly, 128 bits is enough for the near future thus the requirement that UEIDs be a minimum of 128 bits.

There is no requirement for 256 bits today as quadrillion-record databases are not expected in the near future and because this time-to-collision calculation is a very worst case. A future update of the standard may increase the requirement to 256 bits, so there is a requirement that implementations be able to receive 256-bit UEIDs.

B.2. No Use of UUID

A UEID is not a UUID [RFC4122] by conscious choice for the following reasons.

UUIDs are limited to 128 bits which may not be enough for some future use cases.

Today, cryptographic-quality random numbers are available from common CPUs and hardware. This hardware was introduced between 2010 and 2015. Operating systems and cryptographic libraries give access to this hardware. Consequently, there is little need for implementations to construct such random values from multiple sources on their own.

Version 4 UUIDs do allow for use of such cryptographic-quality random numbers, but do so by mapping into the overall UUID structure of time and clock values. This structure is of no value here yet adds complexity. It also slightly reduces the number of actual bits with entropy.

UUIDs seem to have been designed for scenarios where the implementor does not have full control over the environment and uniqueness has to be constructed from identifiers at hand. UEID takes the view that

hardware, software and/or manufacturing process directly implement UEID in a simple and direct way. It takes the view that cryptographic quality random number generators are readily available as they are implemented in commonly used CPU hardware.

Appendix C. EAT Relation to IEEE.802.1AR Secure Device Identity (DevID)

This section describes several distinct ways in which an IEEE IDevID [IEEE.802.1AR] relates to EAT, particularly to UEID and SUEID.

[IEEE.802.1AR] orients around the definition of an implementation called a "DevID Module." It describes how IDevIDs and LDevIDs are stored, protected and accessed using a DevID Module. A particular level of defense against attack that should be achieved to be a DevID is defined. The intent is that IDevIDs and LDevIDs are used with an open set of network protocols for authentication and such. In these protocols the DevID secret is used to sign a nonce or similar to proof the association of the DevID certificates with the device.

By contrast, EAT defines network protocol for proving trustworthiness to a relying party, the very thing that is not defined in [IEEE.802.1AR]. Nor does not give details on how keys, data and such are stored protected and accessed. EAT is intended to work with a variety of different on-device implementations ranging from minimal protection of assets to the highest levels of asset protection. It does not define any particular level of defense against attack, instead providing a set of security considerations.

EAT and DevID can be viewed as complimentary when used together or as competing to provide a device identity service.

C.1. DevID Used With EAT

As just described, EAT defines a network protocol and [IEEE.802.1AR] doesn't. Vice versa, EAT doesn't define a device implementation and DevID does.

Hence, EAT can be the network protocol that a DevID is used with. The DevID secret becomes the attestation key used to sign EATs. The DevID and its certificate chain become the Endorsement sent to the Verifier.

In this case the EAT and the DevID are likely to both provide a device identifier (e.g. a serial number). In the EAT it is the UEID (or SUEID). In the DevID (used as an endorsement), it is a device serial number included in the subject field of the DevID certificate. It is probably a good idea in this use for them to be the same serial number or for the UEID to be a hash of the DevID serial number.

C.2. How EAT Provides an Equivalent Secure Device Identity

The UEID, SUEID and other claims like OEM ID are equivalent to the secure device identity put into the subject field of a DevID certificate. These EAT claims can represent all the same fields and values that can be put in a DevID certificate subject. EAT explicitly and carefully defines a variety of useful claims.

EAT secures the conveyance of these claims by having them signed on the device by the attestation key when the EAT is generated. EAT also signs the nonce that gives freshness at this time. Since these claims are signed for every EAT generated, they can include things that vary over time like GPS location.

DevID secures the device identity fields by having them signed by the manufacturer of the device sign them into a certificate. That certificate is created once during the manufacturing of the device and never changes so the fields cannot change.

So in one case the signing of the identity happens on the device and the other in a manufacturing facility, but in both cases the signing of the nonce that proves the binding to the actual device happens on the device.

While EAT does not specify how the signing keys, signature process and storage of the identity values should be secured against attack, an EAT implementation may have equal defenses against attack. One reason EAT uses CBOR is because it is simple enough that a basic EAT implementation can be constructed entirely in hardware. This allows EAT to be implemented with the strongest defenses possible.

C.3. An X.509 Format EAT

It is possible to define a way to encode EAT claims in an X.509 certificate. For example, the EAT claims might be mapped to X.509 v3 extensions. It is even possible to stuff a whole CBOR-encoded unsigned EAT token into a X.509 certificate.

If that X.509 certificate is an IDevID or LDevID, this becomes another way to use EAT and DevID together.

Note that the DevID must still be used with an authentication protocol that has a nonce or equivalent. The EAT here is not being used as the protocol to interact with the rely party.

C.4. Device Identifier Permanence

In terms of permanence, an IDevID is similar to a UEID in that they do not change over the life of the device. They cease to exist only when the device is destroyed.

An SUEID is similar to an LDevID. They change on device life-cycle events.

[IEEE.802.1AR] describes much of this permanence as resistant to attacks that seek to change the ID. IDevID permanence can be described this way because [IEEE.802.1AR] is oriented around the definition of an implementation with a particular level of defense against attack.

EAT is not defined around a particular implementation and must work on a range of devices that have a range of defenses against attack. EAT thus can't be defined permanence in terms of defense against attack. EAT's definition of permanence is in terms of operations and device lifecycle.

Appendix D. Changes from Previous Drafts

The following is a list of known changes from the previous drafts. This list is non-authoritative. It is meant to help reviewers see the significant differences.

D.1. From draft-rats-eat-01

- o Added UEID design rationale appendix

D.2. From draft-mandyam-rats-eat-00

This is a fairly large change in the orientation of the document, but no new claims have been added.

- o Separate information and data model using CDDL.
- o Say an EAT is a CWT or JWT
- o Use a map to structure the boot_state and location claims

D.3. From draft-ietf-rats-eat-01

- o Clarifications and corrections for OEMID claim
- o Minor spelling and other fixes

- o Add the nonce claim, clarify jti claim
- D.4. From draft-ietf-rats-eat-02
- o Roll all EUIs back into one UEID type
 - o UEIDs can be one of three lengths, 128, 192 and 256.
 - o Added appendix justifying UEID design and size.
 - o Submods part now includes nested eat tokens so they can be named and there can be more than one of them
 - o Lots of fixes to the CDDL
 - o Added security considerations
- D.5. From draft-ietf-rats-eat-03
- o Split boot_state into secure-boot and debug-disable claims
 - o Debug disable is an enumerated type rather than Booleans
- D.6. From draft-ietf-rats-eat-04
- o Change IMEI-based UEIDs to be encoded as a 14-byte string
 - o CDDL cleaned up some more
 - o CDDL allows for JWTs and UCCSs
 - o CWT format submodules are byte string wrapped
 - o Allows for JWT nested in CWT and vice versa
 - o Allows UCCS (unsigned CWTs) and JWT unsecured tokens
 - o Clarify tag usage when nesting tokens
 - o Add section on key inclusion
 - o Add hardware version claims
 - o Collected CDDL is now filled in. Other CDDL corrections.
 - o Rename debug-disable to debug-status; clarify that it is not extensible

- o Security level claim is not extensible
 - o Improve specification of location claim and added a location privacy section
 - o Add intended use claim
- D.7. From draft-ietf-rats-05
- o CDDL format issues resolved
 - o Corrected reference to Location Privacy section
- D.8. From draft-ietf-rats-06
- o Added boot-seed claim
 - o Rework CBOR interoperability section
 - o Added profiles claim and section
- D.9. From draft-ietf-rats-07
- o Filled in IANA and other sections for possible preassignment of claim keys for well understood claims
- D.10. From draft-ietf-rats-08
- o Change profile claim to be either a URL or an OID rather than a test string
- D.11. From draft-ietf-rats-09
- o Add SUEIDs
 - o Add appendix comparing IDevID to EAT
 - o Added section on use for Evidence and Attestation Results
 - o Fill in the key ID and endorsements identificaiton section
 - o Remove origination claim as it is replaced by key IDs and endorsements
 - o Added manifests and software evidence claims
 - o Add string labels non-claim labels for use with JSON (e.g. labels for members of location claim)

- o EAN-13 HW versions are no longer a separate claim. Now they are folded in as a CoSWID version scheme.

Authors' Addresses

Giridhar Mandyam
Qualcomm Technologies Inc.
5775 Morehouse Drive
San Diego, California
USA

Phone: +1 858 651 7200
EMail: mandyam@qti.qualcomm.com

Laurence Lundblade
Security Theory LLC

EMail: lgl@island-resort.com

Miguel Ballesteros
Qualcomm Technologies Inc.
5775 Morehouse Drive
San Diego, California
USA

Phone: +1 858 651 4299
EMail: mballest@qti.qualcomm.com

Jeremy O'Donoghue
Qualcomm Technologies Inc.
279 Farnborough Road
Farnborough GU14 7LS
United Kingdom

Phone: +44 1252 363189
EMail: jodonogh@qti.qualcomm.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 27, 2022

L. Lundblade
Security Theory LLC
G. Mandyam
J. O'Donoghue
Qualcomm Technologies Inc.
February 23, 2022

The Entity Attestation Token (EAT)
draft-ietf-rats-eat-12

Abstract

An Entity Attestation Token (EAT) provides an attested claims set that describes state and characteristics of an entity, a device like a phone, IoT device, network equipment or such. This claims set is used by a relying party, server or service to determine how much it wishes to trust the entity.

An EAT is either a CBOR Web Token (CWT) or JSON Web Token (JWT) with attestation-oriented claims. To a large degree, all this document does is extend CWT and JWT.

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 August 27, 2022.

Copyright Notice

Copyright (c) 2022 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
1.1. Entity Overview	6
1.2. CWT, JWT, UCCS, UJCS and DEB	7
1.3. CDDL, CBOR and JSON	8
1.4. Operating Model and RATS Architecture	8
1.4.1. Relationship between Attestation Evidence and Attestation Results	9
2. Terminology	10
3. The Claims	11
3.1. Token ID Claim (cti and jti)	11
3.2. Timestamp claim (iat)	11
3.3. Nonce Claim (nonce)	12
3.4. Universal Entity ID Claim (ueid)	12
3.5. Semi-permanent UEIDs (SUEIDs)	15
3.6. Hardware OEM Identification (oemid)	16
3.6.1. Random Number Based OEMID	16
3.6.2. IEEE Based OEMID	16
3.6.3. IANA Private Enterprise Number Based OEMID	17
3.7. Hardware Model Claim (hardware-model)	17
3.8. Hardware Version Claims (hardware-version-claims)	18
3.9. Software Name Claim	19
3.10. Software Version Claim	19
3.11. The Security Level Claim (security-level)	19
3.12. Secure Boot Claim (secure-boot)	21
3.13. Debug Status Claim (debug-status)	21
3.13.1. Enabled	22
3.13.2. Disabled	22
3.13.3. Disabled Since Boot	22
3.13.4. Disabled Permanently	22
3.13.5. Disabled Fully and Permanently	22
3.14. Including Keys	23
3.15. The Location Claim (location)	24
3.16. The Uptime Claim (uptime)	25
3.17. The Boot Odometer Claim (odometer)	25
3.18. The Boot Seed Claim (boot-seed)	25
3.19. The Intended Use Claim (intended-use)	26
3.20. The Profile Claim (profile)	27
3.21. The DLOA (Digital Letter or Approval) Claim (dloas)	27
3.22. The Software Manifests Claim (manifests)	28

3.23. The Software Evidence Claim (swevidence)	30
3.24. The SW Measurement Results Claim (swresults)	30
3.24.1. Scheme	31
3.24.2. Objective	31
3.24.3. Results	31
3.24.4. Objective Name	32
3.25. Submodules (submods)	34
3.25.1. Submodule Types	34
3.25.1.1. Submodule Claims-Set	34
3.25.1.2. Nested Token	35
3.25.1.3. Detached Submodule Digest	37
3.25.2. No Inheritance	38
3.25.3. Security Levels	38
3.25.4. Submodule Names	39
3.25.5. CDDL for submods	39
4. Unprotected JWT Claims-Sets	39
5. Detached EAT Bundles	39
6. Endorsements and Verification Keys	40
6.1. Identification Methods	41
6.1.1. COSE/JWS Key ID	41
6.1.2. JWS and COSE X.509 Header Parameters	42
6.1.3. CBOR Certificate COSE Header Parameters	42
6.1.4. Claim-Based Key Identification	42
6.2. Other Considerations	42
7. Profiles	43
7.1. Format of a Profile Document	43
7.2. List of Profile Issues	43
7.2.1. Use of JSON, CBOR or both	43
7.2.2. CBOR Map and Array Encoding	44
7.2.3. CBOR String Encoding	44
7.2.4. CBOR Preferred Serialization	44
7.2.5. COSE/JOSE Protection	44
7.2.6. COSE/JOSE Algorithms	45
7.2.7. DEB Support	45
7.2.8. Verification Key Identification	45
7.2.9. Endorsement Identification	45
7.2.10. Freshness	45
7.2.11. Required Claims	45
7.2.12. Prohibited Claims	45
7.2.13. Additional Claims	46
7.2.14. Refined Claim Definition	46
7.2.15. CBOR Tags	46
7.2.16. Manifests and Software Evidence Claims	46
8. Encoding and Collected CDDL	46
8.1. Claims-Set and CDDL for CWT and JWT	46
8.2. Encoding Data Types	47
8.2.1. Common Data Types	47
8.2.2. JSON Interoperability	47

8.2.3. Labels	48
8.3. CBOR Interoperability	48
8.3.1. EAT Constrained Device Serialization	48
8.4. Collected Common CDDL	49
8.5. Collected CDDL for CBOR	54
8.6. Collected CDDL for JSON	55
9. IANA Considerations	56
9.1. Reuse of CBOR and JSON Web Token (CWT and JWT) Claims Registries	56
9.2. Claim Characteristics	57
9.2.1. Interoperability and Relying Party Orientation . . .	57
9.2.2. Operating System and Technology Neutral	57
9.2.3. Security Level Neutral	58
9.2.4. Reuse of Extant Data Formats	58
9.2.5. Proprietary Claims	58
9.3. Claims Registered by This Document	58
9.3.1. Claims for Early Assignment	59
9.3.2. To be Assigned Claims	62
9.3.3. Version Schemes Registered by this Document	65
9.3.4. UEID URN Registered by this Document	66
9.3.5. Tag for Detached EAT Bundle	66
10. Privacy Considerations	66
10.1. UEID and SUEID Privacy Considerations	67
10.2. Location Privacy Considerations	67
10.3. Replay Protection and Privacy	68
11. Security Considerations	68
11.1. Key Provisioning	68
11.1.1. Transmission of Key Material	69
11.2. Transport Security	69
11.3. Multiple EAT Consumers	69
12. References	70
12.1. Normative References	70
12.2. Informative References	73
Appendix A. Examples	76
A.1. Simple TEE Attestation	76
A.2. Submodules for Board and Device	77
A.3. EAT Produced by Attestation Hardware Block	79
A.4. Detached EAT Bundle	79
A.5. Key / Key Store Attestation	81
A.6. SW Measurements of an IoT Device	83
A.7. Attestation Results in JSON format	86
Appendix B. UEID Design Rationale	87
B.1. Collision Probability	87
B.2. No Use of UUID	89
Appendix C. EAT Relation to IEEE.802.1AR Secure Device Identity (DevID)	90
C.1. DevID Used With EAT	90
C.2. How EAT Provides an Equivalent Secure Device Identity . .	91

C.3. An X.509 Format EAT	91
C.4. Device Identifier Permanence	92
Appendix D. Changes from Previous Drafts	92
D.1. From draft-rats-eat-01	92
D.2. From draft-mandyam-rats-eat-00	92
D.3. From draft-ietf-rats-eat-01	92
D.4. From draft-ietf-rats-eat-02	93
D.5. From draft-ietf-rats-eat-03	93
D.6. From draft-ietf-rats-eat-04	93
D.7. From draft-ietf-rats-eat-05	94
D.8. From draft-ietf-rats-eat-06	94
D.9. From draft-ietf-rats-eat-07	94
D.10. From draft-ietf-rats-eat-08	94
D.11. From draft-ietf-rats-eat-09	94
D.12. From draft-ietf-rats-eat-10	95
D.13. From draft-ietf-rats-eat-11	96
Authors' Addresses	96

1. Introduction

EAT provides the definition of a base set of claims that can be made about an entity, a device, some software and/or some hardware. This claims set is received by a relying party who uses it to decide if and how it will interact with the remote entity. It may choose to not trust the entity and not interact with it. It may choose to trust it. It may partially trust it, for example allowing monetary transactions only up to a limit.

EAT defines the encoding of the claims set in CBOR [RFC8949] and JSON [RFC7159]. EAT is an extension to CBOR Web Token (CWT) [RFC8392] and JSON Web Token (JWT) [RFC7519].

The claims set is secured in transit with the same mechanisms used by CWT and JWT, in particular CBOR Object Signing and Encryption (COSE) [RFC8152] and JSON Object Signing and Encryption (JOSE) [RFC7515] [RFC7516]. Authenticity and integrity protection must always be provided. Privacy (encryption) may additionally be provided. The key material used to sign and encrypt is specifically created and provisioned for the purpose of attestation. It is the use of this key material that make the claims set "attested" rather than just some parameters sent to the relying party by the device.

EAT is focused on authenticating, identifying and characterizing implementations where implementations are devices, chips, hardware, software and such. This is distinct from protocols like TLS [RFC8446] that authenticate and identify servers and services. It is equally distinct from protocols like SASL [RFC4422] that authenticate and identify persons.

The notion of attestation is large, ranging over a broad variety of use cases and security levels. Here are a few examples of claims:

- o Make and model of manufactured consumer device
- o Make and model of a chip or processor, particularly for a security-oriented chip
- o Identification and measurement of the software running on a device
- o Configuration and state of a device
- o Environmental characteristics of a device like its GPS location
- o Formal certifications received

EAT also supports nesting of sets of claims and EAT tokens for use with complex composite devices.

This document uses the terminology and main operational model defined in [RATS.Architecture]. In particular, it can be used for RATS Attestation Evidence and Attestation Results.

1.1. Entity Overview

The document uses the term "entity" to refer to the target of the attestation token. The claims defined in this document are claims about an entity.

An entity is an implementation in hardware, software or both.

An entity is the same as the Attester Target Environment defined in RATS Architecture.

An entity also corresponds to a "system component" as defined in the Internet Security Glossary [RFC4949]. That glossary also defines "entity" and "system entity" as something that may be a person or organization as well as a system component. Here "entity" never refers to a person or organization.

An entity is never a server or a service.

An entity may be the whole device or it may be a subsystem, a subsystem of a subsystem and so on. EAT allows claims to be organized into submodules, nested EATs and so on. See Section 3.25. The entity to which a claim applies is the submodule in which it appears, or to the top-level entity if it doesn't appear in a submodule.

Some examples of entities:

- o A Secure Element
- o A TEE
- o A card in a network router
- o A network router, perhaps with each card in the router a submodule
- o An IoT device
- o An individual process
- o An app on a smartphone
- o A smartphone with many submodules for its many subsystems
- o A subsystem in a smartphone like the modem or the camera

An entity may have strong security like defenses against hardware invasive attacks. It may also have low security, having no special security defenses. There is no minimum security requirement to be an entity.

1.2. CWT, JWT, UCCS, UJCS and DEB

An EAT is a claims set about an entity based on one of the following:

- o CBOR Web Token (CWT) [RFC8392]
- o Unprotected CWT Claims Sets (UCCS) [UCCS.Draft]
- o JSON Web Token (JWT) [RFC7519]

All definitions, requirements, creation and validation procedures, security considerations, IANA registrations and so on from these carry over to EAT.

This specification extends those specifications by defining additional claims for attestation. This specification also describes the notion of a "profile" that can narrow the definition of an EAT, ensure interoperability and fill in details for specific usage scenarios. This specification also adds some considerations for registration of future EAT-related claims.

The identification of a protocol element as an EAT, whether CBOR or JSON encoded, follows the general conventions used by CWT, JWT and

UCCS. Largely this depends on the protocol carrying the EAT. In some cases it may be by content type (e.g., MIME type). In other cases it may be through use of CBOR tags. There is no fixed mechanism across all use cases.

This specification adds two more top-level messages:

- o Unprotected JWT Claims Set (UJCS) Section 4
- o Detached EAT Bundle (DEB), Section 5

A DEB is structure to hold a collection of detached claims sets and the EAT that separately provides integrity and authenticity protection for them. It can be either CBOR or JSON encoded.

1.3. CDDL, CBOR and JSON

This document defines Concise Binary Object Representation (CBOR) [RFC8949] and Javascript Object Notation (JSON) [RFC7159] encoding for an EAT. All claims in an EAT MUST use the same encoding except where explicitly allowed. It is explicitly allowed for a nested token to be of a different encoding. Some claims explicitly contain objects and messages that may use a different encoding than the enclosing EAT.

This specification uses Concise Data Definition Language (CDDL) [RFC8610] for all definitions. The implementor interprets the CDDL to come to either the CBOR or JSON encoding. In the case of JSON, Appendix E of [RFC8610] is followed. Additional rules are given in Section 8.2.2 where Appendix E is insufficient.

The CWT and JWT specifications were authored before CDDL was available and did not use CDDL. This specification includes a CDDL definition of most of what is defined in [RFC8392]. Similarly, this specification includes CDDL for most of what is defined in [RFC7519].

The UCCS specification does not include CDDL. This specification provides CDDL for it.

1.4. Operating Model and RATS Architecture

While it is not required that EAT be used with the RATS operational model described in Figure 1 in [RATS.Architecture], or even that it be used for attestation, this document is oriented around that model.

To summarize, an Attester generates Attestation Evidence. Attestation Evidence is a claims set describing various characteristics of an entity. Attestation Evidence also is usually

signed by a key that proves the entity and the evidence it produces are authentic. The claims set includes a nonce or some other means to provide freshness. EAT is designed to carry Attestation Evidence. The Attestation Evidence goes to a Verifier where the signature is verified. Some of the claims may also be checked against Reference Values. The Verifier then produces Attestation Results which is also usually a claims set. EAT is also designed to carry Attestation Results. The Attestation Results go to the Relying Party which is the ultimate consumer of the Remote Attestation Procedure. The Relying Party uses the Attestation Results as needed for the use case, perhaps allowing an entity on the network, allowing a financial transaction or such.

Note that sometimes the Verifier and Relying Party are not separate and thus there is no need for a protocol to carry Attestation Results.

1.4.1. Relationship between Attestation Evidence and Attestation Results

Any claim defined in this document or in the IANA CWT or JWT registry may be used in Attestation Evidence or Attestation Results.

Many claims in Attestation Evidence simply will pass through the Verifier to the Relying Party without modification. They will be verified as authentic from the entity by the Verifier just through normal verification of the Attester's signature. The UEID, Section 3.4, and Location, Section 3.15, are examples of claims that may be passed through.

Some claims in Attestation Evidence will be verified by the Verifier by comparison to Reference Values. These claims will not likely be conveyed to the Relying Party. Instead, some claim indicating they were checked may be added to the Attestation Results or it may be tacitly known that the Verifier always does this check. For example, the Verifier receives the Software Evidence claim, Section 3.23, compares it to Reference Values and conveys the results to the Relying Party in a Software Measurement Results Claim, Section 3.24.

In some cases the Verifier may provide privacy-preserving functionality by stripping or modifying claims that do not possess sufficient privacy-preserving characteristics. For example, the data in the Location claim, Section 3.15, may be modified to have a precision of a few kilometers rather than a few meters.

2. Terminology

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.

This document reuses terminology from JWT [RFC7519] and CWT [RFC8392].

Claim: A piece of information asserted about a subject. A claim is represented as pair with a value and either a name or key to identify it.

Claim Name: A unique text string that identifies the claim. It is used as the claim name for JSON encoding.

Claim Key: The CBOR map key used to identify a claim.

Claim Value: The value portion of the claim. A claim value can be any CBOR data item or JSON value.

CWT/JWT Claims Set: The CBOR map or JSON object that contains the claims conveyed by the CWT or JWT.

This document reuses terminology from RATS Architecture [RATS.Architecture]

Attester: A role performed by an entity (typically a device) whose Evidence must be appraised in order to infer the extent to which the Attester is considered trustworthy, such as when deciding whether it is authorized to perform some operation.

Verifier: A role that appraises the validity of Attestation Evidence about an Attester and produces Attestation Results to be used by a Relying Party.

Relying Party: A role that depends on the validity of information about an Attester, for purposes of reliably applying application specific actions. Compare /relying party/ in [RFC4949].

Attestation Evidence: A Claims Set generated by an Attester to be appraised by a Verifier. Attestation Evidence may include configuration data, measurements, telemetry, or inferences.

Attestation Results: The output generated by a Verifier, typically including information about an Attester, where the Verifier vouches for the validity of the results

Reference Values: A set of values against which values of Claims can be compared as part of applying an Appraisal Policy for Attestation Evidence. Reference Values are sometimes referred to in other documents as known-good values, golden measurements, or nominal values, although those terms typically assume comparison for equality, whereas here Reference Values might be more general and be used in any sort of comparison.

3. The Claims

This section describes new claims defined for attestation that are to be added to the CWT [IANA.CWT.Claims] and JWT [IANA.JWT.Claims] IANA registries.

This section also describes how several extant CWT and JWT claims apply in EAT.

CDDL, along with a text description, is used to define each claim independent of encoding. Each claim is defined as a CDDL group. In Section 8 on encoding, the CDDL groups turn into CBOR map entries and JSON name/value pairs.

Each claim described has a unique text string and integer that identifies it. CBOR encoded tokens MUST use only the integer for Claim Keys. JSON encoded tokens MUST use only the text string for Claim Names.

3.1. Token ID Claim (cti and jti)

CWT defines the "cti" claim. JWT defines the "jti" claim. These are equivalent to each other in EAT and carry a unique token identifier as they do in JWT and CWT. They may be used to defend against re use of the token but are distinct from the nonce that is used by the Relying Party to guarantee freshness and defend against replay.

3.2. Timestamp claim (iat)

The "iat" claim defined in CWT and JWT is used to indicate the date-of-creation of the token, the time at which the claims are collected and the token is composed and signed.

The data for some claims may be held or cached for some period of time before the token is created. This period may be long, even days. Examples are measurements taken at boot or a geographic

position fix taken the last time a satellite signal was received. There are individual timestamps associated with these claims to indicate their age is older than the "iat" timestamp.

CWT allows the use floating-point for this claim. EAT disallows the use of floating-point. An EAT token MUST NOT contain an iat claim in float-point format. Any recipient of a token with a floating-point format iat claim MUST consider it an error. A 64-bit integer representation of epoch time can represent a range of +/- 500 billion years, so the only point of a floating-point timestamp is to have precession greater than one second. This is not needed for EAT.

3.3. Nonce Claim (nonce)

All EATs should have a nonce to prevent replay attacks. The nonce is generated by the Relying Party, the end consumer of the token. It is conveyed to the entity over whatever transport is in use before the token is generated and then included in the token as the nonce claim.

This documents the nonce claim for registration in the IANA CWT claims registry. This is equivalent to the JWT nonce claim that is already registered.

The nonce must be at least 8 bytes (64 bits) long as fewer bytes are unlikely to be secure. A maximum of 64 bytes is set to limit the memory a constrained implementation uses. This size range is not set for the already-registered JWT nonce, but it should follow this size recommendation when used in an EAT.

Multiple nonces are allowed to accommodate multistage verification and consumption.

```
$$claims-set-claims //=  
    (nonce-label => nonce-type / [ 2* nonce-type ])  
  
nonce-type = bstr .size (8..64)
```

3.4. Universal Entity ID Claim (ueid)

A UEID identifies an individual manufactured entity like a mobile phone, a water meter, a Bluetooth speaker or a networked security camera. It may identify the entire entity or a submodule. It does not identify types, models or classes of entities. It is akin to a serial number, though it does not have to be sequential.

UEIDs MUST be universally and globally unique across manufacturers and countries. UEIDs MUST also be unique across protocols and systems, as tokens are intended to be embedded in many different

protocols and systems. No two products anywhere, even in completely different industries made by two different manufacturers in two different countries should have the same UEID (if they are not global and universal in this way, then Relying Parties receiving them will have to track other characteristics of the entity to keep entities distinct between manufacturers).

There are privacy considerations for UEIDs. See Section 10.1.

The UEID is permanent. It MUST never change for a given entity.

A UEID is constructed of a single type byte followed by the bytes that are the identifier. Several types are allowed to accommodate different industries, different manufacturing processes and to have an alternative that doesn't require paying a registration fee.

Creation of new types requires a Standards Action [RFC8126].

UEIDs are variable length. All implementations MUST be able to receive UEIDs that are 33 bytes long (1 type byte and 256 bits). No UEID longer than 33 bytes SHOULD be sent.

Type Byte	Type Name	Specification
0x01	RAND	This is a 128, 192 or 256-bit random number generated once and stored in the entity. This may be constructed by concatenating enough identifiers to make up an equivalent number of random bits and then feeding the concatenation through a cryptographic hash function. It may also be a cryptographic quality random number generated once at the beginning of the life of the entity and stored. It MUST NOT be smaller than 128 bits. See the length analysis in Appendix B.
0x02	IEEE EUI	This uses the IEEE company identification registry. An EUI is either an EUI-48, EUI-60 or EUI-64 and made up of an OUI, OUI-36 or a CID, different registered company identifiers, and some unique per-entity identifier. EUIs are often the same as or similar to MAC addresses. This type includes MAC-48, an obsolete name for EUI-48. (Note that while entities with multiple network interfaces may have multiple MAC addresses, there is only one UEID for an entity) [IEEE.802-2001], [OUI.Guide].
0x03	IMEI	This is a 14-digit identifier consisting of an 8-digit Type Allocation Code and a 6-digit serial number allocated by the manufacturer, which SHALL be encoded as byte string of length 14 with each byte as the digit's value (not the ASCII encoding of the digit; the digit 3 encodes as 0x03, not 0x33). The IMEI value encoded SHALL NOT include Luhn checksum or SVN information. See [ThreeGPP.IMEI].

Table 1: UEID Composition Types

UEIDs are not designed for direct use by humans (e.g., printing on the case of a device), so no textual representation is defined.

The consumer (the Relying Party) of a UEID MUST treat a UEID as a completely opaque string of bytes and not make any use of its internal structure. For example, they should not use the OUI part of a type 0x02 UEID to identify the manufacturer of the entity. Instead, they should use the OEMID claim. See Section 3.6. The reasons for this are:

- o UEIDs types may vary freely from one manufacturer to the next.

- o New types of UEIDs may be created. For example, a type 0x07 UEID may be created based on some other manufacturer registration scheme.
- o Entity manufacturers are allowed to change from one type of UEID to another anytime they want. For example, they may find they can optimize their manufacturing by switching from type 0x01 to type 0x02 or vice versa. The essential requirement on the manufacturer is that UEIDs be universally unique.

A Device Identifier URN is registered for UEIDs. See Section 9.3.4.

```
$$claims-set-claims // = (ueid-label => ueid-type)
```

```
ueid-type = bstr .size (7..33)
```

3.5. Semi-permanent UEIDs (SUEIDs)

An SEUID is of the same format as a UEID, but it MAY change to a different value on device life-cycle events. Examples of these events are change of ownership, factory reset and on-boarding into an IoT device management system. An entity MAY have both a UEID and SUEIDs, neither, one or the other.

There MAY be multiple SUEIDs. Each one has a text string label the purpose of which is to distinguish it from others in the token. The label MAY name the purpose, application or type of the SUEID. Typically, there will be few SUEIDs so there is no need for a formal labeling mechanism like a registry. The EAT profile MAY describe how SUEIDs should be labeled. If there is only one SUEID, the claim remains a map and there still must be a label. For example, the label for the SUEID used by FIDO Onboarding Protocol could simply be "FIDO".

There are privacy considerations for SUEIDs. See Section 10.1.

A Device Identifier URN is registered for SUEIDs. See Section 9.3.4.

```
$$claims-set-claims // = (sueids-label => sueids-type)
```

```
sueids-type = {
  + tstr => ueid-type
}
```

3.6. Hardware OEM Identification (oemid)

This claim identifies the Original Equipment Manufacturer (OEM) of the hardware. Any of the three forms described below MAY be used at the convenience of the claim sender. The receiver of this claim MUST be able to handle all three forms.

3.6.1. Random Number Based OEMID

The random number based OEMID MUST always 16 bytes (128 bits).

The OEM MAY create their own ID by using a cryptographic-quality random number generator. They would perform this only once in the life of the company to generate the single ID for said company. They would use that same ID in every entity they make. This uniquely identifies the OEM on a statistical basis and is large enough should there be ten billion companies.

The OEM MAY also use a hash function like SHA-256 and truncate the output to 128 bits. The input to the hash should be somethings that have at least 96 bits of entropy, but preferably 128 bits of entropy. The input to the hash MAY be something whose uniqueness is managed by a central registry like a domain name.

In JSON format tokens this MUST be base64url encoded.

3.6.2. IEEE Based OEMID

The IEEE operates a global registry for MAC addresses and company IDs. This claim uses that database to identify OEMs. The contents of the claim may be either an IEEE MA-L, MA-M, MA-S or an IEEE CID [IEEE.RA]. An MA-L, formerly known as an OUI, is a 24-bit value used as the first half of a MAC address. MA-M similarly is a 28-bit value uses as the first part of a MAC address, and MA-S, formerly known as OUI-36, a 36-bit value. Many companies already have purchased one of these. A CID is also a 24-bit value from the same space as an MA-L, but not for use as a MAC address. IEEE has published Guidelines for Use of EUI, OUI, and CID [OUI.Guide] and provides a lookup service [OUI.Lookup].

Companies that have more than one of these IDs or MAC address blocks SHOULD select one and prefer that for all their entities.

Commonly, these are expressed in Hexadecimal Representation as described in [IEEE.802-2001]. It is also called the Canonical format. When this claim is encoded the order of bytes in the bstr are the same as the order in the Hexadecimal Representation. For

example, an MA-L like "AC-DE-48" would be encoded in 3 bytes with values 0xAC, 0xDE, 0x48.

This format is always 3 bytes in size in CBOR.

In JSON format tokens, this MUST be base64url encoded and always 4 bytes.

3.6.3. IANA Private Enterprise Number Based OEMID

IANA maintains a integer-based company registry called the Private Enterprise Number (PEN) [PEN].

PENs are often used to create an OID. That is not the case here. They are used only as an integer.

In CBOR this value MUST be encoded as a major type 0 integer and is typically 3 bytes. In JSON, this value MUST be encoded as a number.

```
oemid-pen = int
```

```
oemid-ieee = bstr .size 3
```

```
oemid-random = bstr .size 16
```

```
$$claims-set-claims //= (  
    oemid-label =>  
        oemid-random / oemid-ieee / oemid-pen  
)
```

3.7. Hardware Model Claim (hardware-model)

This claim differentiates hardware models, products and variants manufactured by a particular OEM, the one identified by OEM ID in Section 3.6.

This claim must be unique so as to differentiate the models and products for the OEM ID. This claim does not have to be globally unique, but it can be. A receiver of this claim MUST not assume it is globally unique. To globally identify a particular product, the receiver should concatenate the OEM ID and this claim.

The granularity of the model identification is for each OEM to decide. It may be very granular, perhaps including some version information. It may be very general, perhaps only indicating top-level products.

The purpose of this claim is to identify models within protocols, not for human-readable descriptions. The format and encoding of this claim should not be human-readable to discourage use other than in protocols. If this claim is to be derived from an already-in-use human-readable identifier, it can be run through a hash function.

There is no minimum length so that an OEM with a very small number of models can use a one-byte encoding. The maximum length is 32 bytes. All receivers of this claim MUST be able to receive this maximum size.

The receiver of this claim MUST treat it as a completely opaque string of bytes, even if there is some apparent naming or structure. The OEM is free to alter the internal structure of these bytes as long as the claim continues to uniquely identify its models.

```
hardware-model-type = bytes .size (1..32)
```

```
$$claims-set-claims // = (
    hardware-model-label => hardware-model-type
)
```

3.8. Hardware Version Claims (hardware-version-claims)

The hardware version is a text string the format of which is set by each manufacturer. The structure and sorting order of this text string can be specified using the version-scheme item from CoSWID [CoSWID]. It is useful to know how to sort versions so the newer can be distinguished from the older.

The hardware version can also be given by a 13-digit [EAN-13]. A new CoSWID version scheme is registered with IANA by this document in Section 9.3.3. An EAN-13 is also known as an International Article Number or most commonly as a bar code.

```
$$claims-set-claims // = (
    hardware-version-label => hardware-version-type
)
```

```
hardware-version-type = [
    version:  tstr,
    scheme:   $version-scheme
]
```

3.9. Software Name Claim

This is a free-form text claim for the name of the software for the entity or submodule. A CoSWID manifest or other type of manifest can be used instead if this claim is too limited to correctly characterize the SW for the entity or submodule.

```
$$claims-set-claims //= ( sw-name-label => tstr )
```

3.10. Software Version Claim

This makes use of the CoSWID version scheme data type to give a simple version for the software. A full CoSWID manifest or other type of manifest can be used instead if this is too simple.

```
$$claims-set-claims //= (sw-version-label => sw-version-type)
```

```
sw-version-type = [  
    version:  tstr,  
    scheme:   $version-scheme ; As defined by CoSWID  
]
```

3.11. The Security Level Claim (security-level)

This claim characterizes the entity's ability to defend against attacks aimed at capturing the signing key, forging claims and at forging EATs. This is by defining four security levels.

This claim describes the security environment and countermeasures available on the entity where the attestation key resides and the claims originate.

- 1 - Unrestricted: There is some expectation that implementor will protect the attestation signing keys at this level. Otherwise, the EAT provides no meaningful security assurances.
- 2 - Restricted: Entities at this level are not general-purpose operating environments that host features, such as app download systems, web browsers and complex applications. It is akin to the secure-restricted level (see below) without the security orientation. Examples include a Wi-Fi subsystem, an IoT camera, or sensor device. Often these can be considered more secure than unrestricted just because they are much simpler and a smaller attack surface, but this won't always be the case. Some unrestricted devices may be implemented in a way that provides poor protection of signing keys.

- 3 - Secure-Restricted: Entities at this level must meet the criteria defined in Section 4 of FIDO Allowed Restricted Operating Environments [FIDO.AROE]. Examples include TEE's and schemes using virtualization-based security. Security at this level is aimed at defending against large-scale network/remote attacks against the entity.
- 4 - Hardware: Entities at this level must include substantial defense against physical or electrical attacks against the entity itself. It is assumed the potential attacker has captured the entity and can disassemble it. Examples include TPMs and Secure Elements.

The entity should claim the highest security level it achieves and no higher. This set is not extensible so as to provide a common interoperable description of security level to the Relying Party. If a particular use case considers this claim to be inadequate, it can define its own proprietary claim. It may consider including both this claim as a coarse indication of security and its own proprietary claim as a refined indication.

This claim is not intended as a replacement for a formal security certification scheme, such as those based on FIPS 140 [FIPS-140] or those based on Common Criteria [Common.Criteria]. See Section 3.21.

```
$$claims-set-claims // = (
    security-level-label =>
        security-level-chor-type /
        security-level-json-type
)
```

```
security-level-chor-type = &(
    unrestricted: 1,
    restricted: 2,
    secure-restricted: 3,
    hardware: 4
)
```

```
security-level-json-type =
    "unrestricted" /
    "restricted" /
    "secure-restricted" /
    "hardware"
```


3.12. Secure Boot Claim (secure-boot)

The value of true indicates secure boot is enabled. Secure boot is considered enabled when the firmware and operating system, are under control of the manufacturer of the entity identified in the OEMID claim described in Section 3.6. Control by the manufacturer of the firmware and the operating system may be by it being in ROM, being cryptographically authenticated, a combination of the two or similar.

```
$$claims-set-claims // = (secure-boot-label => bool)
```

3.13. Debug Status Claim (debug-status)

This applies to entity-wide or submodule-wide debug facilities of the entity like JTAG and diagnostic hardware built into chips. It applies to any software debug facilities related to root, operating system or privileged software that allow system-wide memory inspection, tracing or modification of non-system software like user mode applications.

This characterization assumes that debug facilities can be enabled and disabled in a dynamic way or be disabled in some permanent way such that no enabling is possible. An example of dynamic enabling is one where some authentication is required to enable debugging. An example of permanent disabling is blowing a hardware fuse in a chip. The specific type of the mechanism is not taken into account. For example, it does not matter if authentication is by a global password or by per-entity public keys.

As with all claims, the absence of the debug level claim means it is not reported. A conservative interpretation might assume the enabled state.

This claim is not extensible so as to provide a common interoperable description of debug status. If a particular implementation considers this claim to be inadequate, it can define its own proprietary claim. It may consider including both this claim as a coarse indication of debug status and its own proprietary claim as a refined indication.

The higher levels of debug disabling requires that all debug disabling of the levels below it be in effect. Since the lowest level requires that all of the target's debug be currently disabled, all other levels require that too.

There is no inheritance of claims from a submodule to a superior module or vice versa. There is no assumption, requirement or guarantee that the target of a superior module encompasses the

targets of submodules. Thus, every submodule must explicitly describe its own debug state. The receiver of an EAT MUST not assume that debug is turned off in a submodule because there is a claim indicating it is turned off in a superior module.

An entity may have multiple debug facilities. The use of plural in the description of the states refers to that, not to any aggregation or inheritance.

The architecture of some chips or devices may be such that a debug facility operates for the whole chip or device. If the EAT for such a chip includes submodules, then each submodule should independently report the status of the whole-chip or whole-device debug facility. This is the only way the receiver can know the debug status of the submodules since there is no inheritance.

3.13.1. Enabled

If any debug facility, even manufacturer hardware diagnostics, is currently enabled, then this level must be indicated.

3.13.2. Disabled

This level indicates all debug facilities are currently disabled. It may be possible to enable them in the future. It may also be that they were enabled in the past, but they are currently disabled.

3.13.3. Disabled Since Boot

This level indicates all debug facilities are currently disabled and have been so since the entity booted/started.

3.13.4. Disabled Permanently

This level indicates all non-manufacturer facilities are permanently disabled such that no end user or developer can enable them. Only the manufacturer indicated in the OEMID claim can enable them. This also indicates that all debug facilities are currently disabled and have been so since boot/start.

3.13.5. Disabled Fully and Permanently

This level indicates that all debug facilities for the entity are permanently disabled.

```
$$claims-set-claims // = (
    debug-status-label =>
        debug-status-cbor-type / debug-status-json-type
)

debug-status-cbor-type = &(
    enabled: 0,
    disabled: 1,
    disabled-since-boot: 2,
    disabled-permanently: 3,
    disabled-fully-and-permanently: 4
)

debug-status-json-type =
    "enabled" /
    "disabled" /
    "disabled-since-boot" /
    "disabled-permanently" /
    "disabled-fully-and-permanently"
```

3.14. Including Keys

An EAT may include a cryptographic key such as a public key. The signing of the EAT binds the key to all the other claims in the token.

The purpose for inclusion of the key may vary by use case. For example, the key may be included as part of an IoT device onboarding protocol. When the FIDO protocol includes a public key in its attestation message, the key represents the binding of a user, device and Relying Party. This document describes how claims containing keys should be defined for the various use cases. It does not define specific claims for specific use cases.

Keys in CBOR format tokens SHOULD be the COSE_Key format [RFC8152] and keys in JSON format tokens SHOULD be the JSON Web Key format [RFC7517]. These two formats support many common key types. Their use avoids the need to decode other serialization formats. These two formats can be extended to support further key types through their IANA registries.

The general confirmation claim format [RFC8747], [RFC7800] may also be used. It provides key encryption. It also allows for inclusion by reference through a key ID. The confirmation claim format may be employed in the definition of some new claim for a particular use case.

When the actual confirmation claim is included in an EAT, this document associates no use case semantics other than proof of possession. Different EAT use cases may choose to associate further semantics. The key in the confirmation claim **MUST** be protected in the same way as the key used to sign the EAT. That is, the same, equivalent or better hardware defenses, access controls, key generation and such must be used.

3.15. The Location Claim (location)

The location claim gives the location of the entity from which the attestation originates. It is derived from the W3C Geolocation API [W3C.GeoLoc]. The latitude, longitude, altitude and accuracy must conform to [WGS84]. The altitude is in meters above the [WGS84] ellipsoid. The two accuracy values are positive numbers in meters. The heading is in degrees relative to true north. If the entity is stationary, the heading is NaN (floating-point not-a-number). The speed is the horizontal component of the entity velocity in meters per second.

When encoding floating-point numbers half-precision **SHOULD NOT** be used. They usually do not provide enough precision for a geographic location.

The location may have been cached for a period of time before token creation. For example, it might have been minutes or hours or more since the last contact with a GPS satellite. Either the timestamp or age data item can be used to quantify the cached period. The timestamp data item is preferred as it a non-relative time.

The age data item can be used when the entity doesn't know what time it is either because it doesn't have a clock or it isn't set. The entity **MUST** still have a "ticker" that can measure a time interval. The age is the interval between acquisition of the location data and token creation.

See location-related privacy considerations in Section 10.2.

```
$$claims-set-claims // = (location-label => location-type)
```

```
location-type = {  
    latitude => number,  
    longitude => number,  
    ? altitude => number,  
    ? accuracy => number,  
    ? altitude-accuracy => number,  
    ? heading => number,  
    ? speed => number,  
    ? timestamp => ~time-int,  
    ? age => uint  
}  
  
latitude = 1 / "latitude"  
longitude = 2 / "longitude"  
altitude = 3 / "altitude"  
accuracy = 4 / "accuracy"  
altitude-accuracy = 5 / "altitude-accuracy"  
heading = 6 / "heading"  
speed = 7 / "speed"  
timestamp = 8 / "timestamp"  
age = 9 / "age"
```

3.16. The Uptime Claim (uptime)

The "uptime" claim MUST contain a value that represents the number of seconds that have elapsed since the entity or submod was last booted.

```
$$claims-set-claims // = (uptime-label => uint)
```

3.17. The Boot Odometer Claim (odometer)

The "odometer" claim contains a value that represents the number of times the entity or submod has been booted. Support for this claim requires a persistent storage on the device.

```
$$claims-set-claims // = (odometer-label => uint)
```

3.18. The Boot Seed Claim (boot-seed)

The Boot Seed claim MUST contain a random value created at system boot time that will allow differentiation of reports from different boot sessions.

This value is usually public. It is not a secret and MUST NOT be used for any purpose that a secret seed is needed, such as seeding a random number generator.

`$$claims-set-claims // = (boot-seed-label => bytes)`

3.19. The Intended Use Claim (intended-use)

EAT's may be used in the context of several different applications. The intended-use claim provides an indication to an EAT consumer about the intended usage of the token. This claim can be used as a way for an application using EAT to internally distinguish between different ways it uses EAT.

- 1 - Generic: Generic attestation describes an application where the EAT consumer requires the most up-to-date security assessment of the attesting entity. It is expected that this is the most commonly-used application of EAT.
- 2- Registration: Entities that are registering for a new service may be expected to provide an attestation as part of the registration process. This intended-use setting indicates that the attestation is not intended for any use but registration.
- 3 - Provisioning: Entities may be provisioned with different values or settings by an EAT consumer. Examples include key material or device management trees. The consumer may require an EAT to assess entity security state of the entity prior to provisioning.
- 4 - Certificate Issuance Certification Authorities (CA's) may require attestations prior to the issuance of certificates related to keypairs hosted at the entity. An EAT may be used as part of the certificate signing request (CSR).
- 5 - Proof-of-Possession: An EAT consumer may require an attestation as part of an accompanying proof-of-possession (PoP) application. More precisely, a PoP transaction is intended to provide to the recipient cryptographically-verifiable proof that the sender has possession of a key. This kind of attestation may be necessary to verify the security state of the entity storing the private key used in a PoP application.

```
$$claims-set-claims //= (
    intended-use-label =>
        intended-use-cbor-type / intended-use-json-type
)

intended-use-cbor-type = &(
    generic: 1,
    registration: 2,
    provisioning: 3,
    csr: 4,
    pop: 5
)

intended-use-json-type =
    "generic" /
    "registration" /
    "provisioning" /
    "csr" /
    "pop"
```

3.20. The Profile Claim (profile)

See Section 7 for the detailed description of a profile.

A profile is identified by either a URL or an OID. Typically, the URI will reference a document describing the profile. An OID is just a unique identifier for the profile. It may exist anywhere in the OID tree. There is no requirement that the named document be publicly accessible. The primary purpose of the profile claim is to uniquely identify the profile even if it is a private profile.

The OID is always absolute and never relative. In CBOR tokens, the OID MUST be encoded according to [RFC9090] and the URI according to [RFC8949]. Both are unwrapped and thus not CBOR tags. In JSON tokens, the OID is a string of the form "X.X.X", and a URI is a normal URI string.

Note that this is named "eat_profile" for JWT and is distinct from the already registered "profile" claim in the JWT claims registry.

```
$$claims-set-claims //= (profile-label => ~uri / ~oid)
```

3.21. The DLOA (Digital Letter or Approval) Claim (dloas)

A DLOA (Digital Letter of Approval) [DLOA] is an XML document that describes a certification that an entity has received. Examples of certifications represented by a DLOA include those issued by Global Platform and those based on Common Criteria. The DLOA is unspecific

to any particular certification type or those issued by any particular organization.

This claim is typically issued by a Verifier, not an Attester. When this claim is issued by a Verifier, it MUST be because the entity has received the certification in the DLOA.

This claim MAY contain more than one DLOA. If multiple DLOAs are present, it MUST be because the entity received all of the certifications.

DLOA XML documents are always fetched from a registrar that stores them. This claim contains several data items used to construct a URL for fetching the DLOA from the particular registrar.

This claim MUST be encoded as an array with either two or three elements. The first element MUST be the URI for the registrar. The second element MUST be a platform label indicating which platform was certified. If the DLOA applies to an application, then the third element is added which MUST be an application label. The method of constructing the registrar URI, platform label and possibly application label is specified in [DLOA].

```
$$claims-set-claims // = (  
    dloas-label => [ + dloa-type ]  
)
```

```
dloa-type = [  
    dloa_registrar: ~uri  
    dloa_platform_label: text  
    ? dloa_application_label: text  
]
```

3.22. The Software Manifests Claim (manifests)

This claim contains descriptions of software present on the entity. These manifests are installed on the entity when the software is installed or are created as part of the installation process. Installation is anything that adds software to the entity, possibly factory installation, the user installing elective applications and so on. The defining characteristic is they are created by the software manufacturer. The purpose of these claims in an EAT is to relay them without modification to the Verifier and possibly to the Relying Party.

Some manifests may be signed by their software manufacturer before they are put into this EAT claim. When such manifests are put into this claim, the manufacturer's signature SHOULD be included. For

example, the manifest might be a CoSWID signed by the software manufacturer, in which case the full signed CoSWID should be put in this claim.

This claim allows multiple formats for the manifest. For example, the manifest may be a CBOR-format CoSWID, an XML-format SWID or other. Identification of the type of manifest is always by a CBOR tag. In many cases, for examples CoSWID, a tag will already be registered with IANA. If not, a tag MUST be registered. It can be in the first-come-first-served space which has minimal requirements for registration.

The claim is an array of one or more manifests. To facilitate hand off of the manifest to a decoding library, each manifest is contained in a byte string. This occurs for CBOR-format manifests as well as non-CBOR format manifests.

If a particular manifest type uses CBOR encoding, then the item in the array for it MUST be a byte string that contains a CBOR tag. The EAT decoder must decode the byte string and then the CBOR within it to find the tag number to identify the type of manifest. The contents of the byte string is then handed to the particular manifest processor for that type of manifest. CoSWID and SUIT manifest are examples of this.

If a particular manifest type does not use CBOR encoding, then the item in the array for it MUST be a CBOR tag that contains a byte string. The EAT decoder uses the tag to identify the processor for that type of manifest. The contents of the tag, the byte string, are handed to the manifest processor. Note that a byte string is used to contain the manifest whether it is a text based format or not. An example of this is an XML format ISO/IEC 19770 SWID.

It is not possible to describe the above requirements in CDDL, so the type for an individual manifest is any in the CDDL below. The above text sets the encoding requirement.

This claim allows for multiple manifests in one token since multiple software packages are likely to be present. The multiple manifests MAY be of multiple formats. In some cases EAT submodules may be used instead of the array structure in this claim for multiple manifests.

When the [CoSWID] format is used, it MUST be a payload CoSWID, not an evidence CoSWID.

```
$$claims-set-claims //= (
    manifests-label => manifests-type
)

manifests-type = [+ $$manifest-formats]

coswid-that-is-a-cbor-tag-xx = tagged-coswid<concise-swid-tag>

$$manifest-formats /= bytes .cbor coswid-that-is-a-cbor-tag-xx
```

3.23. The Software Evidence Claim (swevidence)

This claim contains descriptions, lists, evidence or measurements of the software that exists on the entity. The defining characteristic of this claim is that its contents are created by processes on the entity that inventory, measure or otherwise characterize the software on the entity. The contents of this claim do not originate from the software manufacturer.

This claim uses the same mechanism for identification of the type of the swevidence as is used for the type of the manifest in the manifests claim. It also uses the same byte string based mechanism for containing the claim and easing the hand off to a processing library. See the discussion above in the manifests claim.

When the [CoSWID] format is used, it MUST be evidence CoSWIDs, not payload CoSWIDS.

```
$$claims-set-claims //= (
    swevidence-label => swevidence-type
)

swevidence-type = [+ $$swevidence-formats]

coswid-that-is-a-cbor-tag = tagged-coswid<concise-swid-tag>
$$swevidence-formats /= bytes .cbor coswid-that-is-a-cbor-tag
```

3.24. The SW Measurement Results Claim (swresults)

This claims reports the outcome of the comparison of a measurement on some software to the expected Reference Values. It may report a successful comparison, failed comparison or other.

This claim MAY be generated by the Verifier and sent to the Relying Party. For example, it could be the results of the Verifier comparing the contents of the swevidence claim to Reference Values.

This claim MAY also be generated on the entity if the entity has the ability for one subsystem to measure another subsystem. For example, a TEE might have the ability to measure the software of the rich OS and may have the Reference Values for the rich OS.

Within an attestation target or submodule, multiple results can be reported. For example, it may be desirable to report the results for the kernel and each individual application separately.

For each software objective, the following can be reported. TODO: defined objective

3.24.1. Scheme

This is the free-form text name of the verification system or scheme that performed the verification. There is no official registry of schemes or systems. It may be the name of a commercial product or such.

3.24.2. Objective

This roughly characterizes the coverage of the software measurement software. This corresponds to the attestation target or the submodule. If all of the indicated target is not covered, the measurement must indicate partial.

- 1 - all: Indicates all the software has been verified, for example, all the software in the attestation target or the submodule
- 2 - firmware: Indicates all of and only the firmware
- 3 - kernel: Refers to all of the most-privileged software, for example the Linux kernel
- 4 - privileged: Refers to all of the software used by the root, system or administrative account
- 5 - system-libs: Refers to all of the system libraries that are broadly shared and used by applications and such
- 6 - partial: Some other partial set of the software

3.24.3. Results

This describes the result of the measurement and also the comparison to Reference Values.

- 1 - verification-not-run: Indicates that no attempt was made to run the verification
- 2 - verification-indeterminate: The verification was attempted, but it did not produce a result; perhaps it ran out of memory, the battery died or such
- 3 - verification-failed: The verification ran to completion, the comparison was completed and did not compare correctly to the Reference Values
- 4 - fully-verified: The verification ran to completion and all measurements compared correctly to Reference Values
- 5 - partially-verified: The verification ran to completion and some, but not all, measurements compared correctly to Reference Values

3.24.4. Objective Name

This is a free-form text string that describes the objective. For example, "Linux kernel" or "Facebook App"

```
$$claims-set-claims // = (swresults-label => [ + swresult-type ])  
  
verification-result-cbor-type = &(  
    verification-not-run: 1,  
    verification-indeterminate: 2,  
    verification-failed: 3,  
    fully-verified: 4,  
    partially-verified: 5,  
)  
  
verification-result-json-type =  
    "verification-not-run" /  
    "verification-indeterminate" /  
    "verification-failed" /  
    "fully-verified" /  
    "partially-verified"  
  
verification-objective-cbor-type = &(  
    all: 1,  
    firmware: 2,  
    kernel: 3,  
    privileged: 4,  
    system-libs: 5,  
    partial: 6,  
)  
  
verification-objective-json-type =  
    "all" /  
    "firmware" /  
    "kernel" /  
    "privileged" /  
    "system-libs" /  
    "partial"  
  
swresult-type = [  
    verification-system: tstr,  
    objective: verification-objective-cbor-type /  
        verification-objective-json-type,  
    result: verification-result-cbor-type /  
        verification-result-json-type,  
    ? objective-name: tstr  
]
```

3.25. Submodules (submods)

Some devices are complex, having many subsystems. A mobile phone is a good example. It may have several connectivity subsystems for communications (e.g., Wi-Fi and cellular). It may have subsystems for low-power audio and video playback. It may have multiple security-oriented subsystems like a TEE and a Secure Element.

The claims for a subsystem can be grouped together in a submodule or submod.

The submods are in a single map/object, one entry per submodule. There is only one submods map/object in a token. It is identified by its specific label. It is a peer to other claims, but it is not called a claim because it is a container for a claims set rather than an individual claim. This submods part of a token allows what might be called recursion. It allows claims sets inside of claims sets inside of claims sets...

3.25.1. Submodule Types

The following sections define the three types of submodules:

- o A submodule Claims-Set
- o A nested token, which can be any valid EAT token, CBOR or JSON
- o The digest of a detached Claims-Set

3.25.1.1. Submodule Claims-Set

This is a subordinate Claims-Set containing claims about the submodule.

The submodule claims-set is produced by the same Attester as the surrounding token. It is secured using the same mechanism as the enclosing token (e.g., it is signed by the same attestation key). It roughly corresponds to an Attester Target Environment, as described in the RATS architecture.

It may contain claims that are the same as its surrounding token or superior submodules. For example, the top-level of the token may have a UEID, a submod may have a different UEID and a further subordinate submodule may also have a UEID.

The encoding of a submodule Claims-Set MUST be the same as the encoding as the token it is part of.

This data type for this type of submodule is a map/object. It is identified when decoding by it's type being a map/object.

3.25.1.2. Nested Token

This type of submodule is a fully formed complete token. It is typically produced by a separate Attester. It is typically used by a Composite Device as described in RATS Architecture [RATS.Architecture] In being a submodule of the surrounding token, it is cryptographically bound to the surrounding token. If it was conveyed in parallel with the surrounding token, there would be no such binding and attackers could substitute a good attestation from another device for the attestation of an errant subsystem.

A nested token does not need to use the same encoding as the enclosing token. This is to allow Composite Devices to be built without regards to the encoding supported by their Attesters. Thus a CBOR-encoded token like a CWT or UCCS can have a JWT as a nested token submodule and a JSON-encoded token can have a CWT or UCCS as a nested token submodule.

The following two sections describe how to encode and decode a nested token.

3.25.1.2.1. Surrounding EAT is CBOR-Encoded

This describes the encoding and decoding of CBOR or JSON-encoded tokens nested inside a CBOR-encoded token.

If the nested token is CBOR-encoded, then it MUST be a CBOR tag and MUST be wrapped in a byte string. The tag identifies whether the nested token is a CWT, a UCCS, a CBOR-encoded DEB, or some other CBOR-format token defined in the future. A nested CBOR-encoded token that is not a CBOR tag is NOT allowed.

If the nested token is JSON-encoded, then the data item MUST be a text string. The text string MUST contain a JSON-encoded array of two items. The first item is a string identifying the type of the token. The second item is the JSON-encoded token.

The string identifying the JSON-encoded token MUST be one of the following:

"JWT": The second item MUST be a JWT formatted according to [RFC7519]

"UJCS": The second item MUST be a UJCS-Message as defined in this document.

"DEB": The second item MUST be a JSON-encoded Detached EAT Bundle as defined in this document.

The definition of additional types requires a standards action.

When decoding, if a byte string is encountered, it is known to be a nested CBOR-encoded token. The byte string wrapping is removed. The type of the token is determined by the CBOR tag.

When decoding, if a text string is encountered, it is known to be a JSON-encoded token. The two-item array is decoded and tells the type of the JSON-encoded token.

```
Nested-Token =  
    tstr / ; A JSON-encoded Nested-Token (see json-nested-token.cddl)  
    bstr .cbor Tagged-CBOR-Token
```

3.25.1.2.2. Surrounding EAT is JSON-Encoded

This describes the encoding and decoding of CBOR or JSON-encoded tokens nested inside a JSON-encoded token.

The nested token MUST be an array of two in the same format as described in the section above.

A CBOR-encoded token nested inside a JSON-encoded MUST use the same array of two, but with the type as follows:

"CBOR": Some base64url-encoded CBOR that is a tag, typically a CWT, UCCS or CBOR-encoded DEB

When decoding, the array of two is decoded. The first item indicates the type and encoding of the nested token. If the type string is not "CBOR", then the token is JSON-encoded and of the type indicated by the string.

If the type string is "CBOR", then the token is CBOR-encoded. The base64url encoding is removed. The CBOR-encoded data is then decoded. The type of nested token is determined by the CBOR-tag. It is an error if the CBOR is not a tag.


```
Nested-Token = [  
  type : "JWT" / "CBOR" / "UJCS" / "DEB",  
  nested-token : JWT-Message /  
                  B64URL-Tagged-CBOR-Token /  
                  DEB-JSON-Message /  
                  UJCS-Message  
]
```

```
B64URL-Tagged-CBOR-Token = tstr .regexp "[A-Za-z0-9_=-]+"
```

3.25.1.3. Detached Submodule Digest

This is type of submodule equivalent to a Claims-Set submodule, except the Claims-Set is conveyed separately outside of the token.

This type of submodule consists of a digest made using a cryptographic hash of a Claims-Set. The Claims-Set is not included in the token. It is conveyed to the Verifier outside of the token. The submodule containing the digest is called a detached digest. The separately conveyed Claims-Set is called a detached claims set.

The input to the digest is exactly the byte-string wrapped encoded form of the Claims-Set for the submodule. That Claims-Set can include other submodules including nested tokens and detached digests.

The primary use for this is to facilitate the implementation of a small and secure attester, perhaps purely in hardware. This small, secure attester implements COSE signing and only a few claims, perhaps just UEID and hardware identification. It has inputs for digests of submodules, perhaps 32-byte hardware registers. Software running on the device constructs larger claim sets, perhaps very large, encodes them and digests them. The digests are written into the small secure attesters registers. The EAT produced by the small secure attester only contains the UEID, hardware identification and digests and is thus simple enough to be implemented in hardware. Probably, every data item in it is of fixed length.

The integrity protection for the larger Claims Sets will not be as secure as those originating in hardware block, but the key material and hardware-based claims will be. It is possible for the hardware to enforce hardware access control (memory protection) on the digest registers so that some of the larger claims can be more secure. For example, one register may be writable only by the TEE, so the detached claims from the TEE will have TEE-level security.

The data type for this type of submodule MUST be an array. It contains two data items, an algorithm identifier and a byte string containing the digest.

When decoding a CBOR format token the detached digest type is distinguished from the other types by it being an array. In CBOR the none of other submodule types are arrays.

When decoding a JSON format token, a little more work is required because both the nested token and detached digest types are an array. To distinguish the nested token from the detached digest, the first element in the array is examined. If it is "JWT", "UJCS" or "DEB", the the submodule is a nested token. Otherwise it will contain an algorithm identifier and is a detached digest.

A DEB, described in Section 5, may be used to convey detached claims sets and the token with their detached digests. EAT, however, doesn't require use of a DEB. Any other protocols may be used to convey detached claims sets and the token with their detached digests. Note that since detached Claims-Sets are usually signed, protocols conveying them must make sure they are not modified in transit.

3.25.2. No Inheritance

The subordinate modules do not inherit anything from the containing token. The subordinate modules must explicitly include all of their claims. This is the case even for claims like the nonce.

This rule is in place for simplicity. It avoids complex inheritance rules that might vary from one type of claim to another.

3.25.3. Security Levels

The security level of the non-token subordinate modules should always be less than or equal to that of the containing modules in the case of non-token submodules. It makes no sense for a module of lesser security to be signing claims of a module of higher security. An example of this is a TEE signing claims made by the non-TEE parts (e.g. the high-level OS) of the device.

The opposite may be true for the nested tokens. They usually have their own more secure key material. An example of this is an embedded secure element.

3.25.4. Submodule Names

The label or name for each submodule in the submods map is a text string naming the submodule. No submodules may have the same name.

3.25.5. CDDL for submods

The submodule type is distinguished in the encoded bytes by its data type, map/object for a Claims-Set, string for nested token and array for a detached submodule. Nested tokens are byte-string wrapped when encoded in CBOR and base64 encoded for JSON.

```
$$claims-set-claims // = (submods-label => { + text => Submodule })
```

```
Submodule = Claims-Set / Nested-Token / Detached-Submodule-Digest
```

```
Detached-Submodule-Digest = [  
    algorithm : int / text,  
    digest : bstr  
]
```

4. Unprotected JWT Claims-Sets

This is simply the JSON equivalent of an Unprotected CWT Claims-Set [UCCS.Draft].

It has no protection of its own so protections must be provided by the protocol carrying it. These are extensively discussed in [UCCS.Draft]. All the security discussion and security considerations in [UCCS.Draft] apply to UJCS.

(Note: The EAT author is open to this definition being moved into the UCCS draft, perhaps along with the related CDDL. It is place here for now so that the current UCCS draft plus this document are complete. UJCS is needed for the same use cases that a UCCS is needed. Further, JSON will commonly be used to convey Attestation Results since JSON is common for server to server communications. Server to server communications will often have established security (e.g., TLS) therefore the signing and encryption from JWS and JWE are unnecessary and burdensome).

5. Detached EAT Bundles

A detached EAT bundle is a structure to convey a fully-formed and signed token plus detached claims set that relate to that token. It is a top-level EAT message like a CWT, JWT, UCCS and UJCS. It can be used any place that CWT, JWT, UCCS or UJCS messages are used. It may also be sent as a submodule.

A DEB has two main parts.

The first part is a full top-level token. This top-level token must have at least one submodule that is a detached digest. This top-level token may be either CBOR or JSON-encoded. It may be a CWT, JWT, UCCS or UJCS, but not a DEB. The same mechanism for distinguishing the type for nested token submodules is used here.

The second part is a map/object containing the detached Claims-Sets corresponding to the detached digests in the full token. When the DEB is CBOR-encoded, each Claims-Set is wrapped in a byte string. When the DEB is JSON-encoded, each Claims-Set is base64url encoded. All the detached Claims-Sets MUST be encoded in the same format as the DEB. No mixing of encoding formats is allowed for the Claims-Sets in a DEB.

For CBOR-encoded DEBs, tag TBD602 can be used to identify it. The normal rules apply for use or non-use of a tag. When it is sent as a submodule, it is always sent as a tag to distinguish it from the other types of nested tokens.

The digests of the detached claims sets are associated with detached claims-sets by label/name. It is up to the constructor of the detached EAT bundle to ensure the names uniquely identify the detached claims sets. Since the names are used only in the detached EAT bundle, they can be very short, perhaps one byte.

```
Detached-EAT-Bundle = [
  main-token : Nested-Token,
  detached-claims-sets: {
    + tstr => cbor-wrapped-claims-set / json-wrapped-claims-set
  }
]
```

```
json-wrapped-claims-set = tstr .regexp "[A-Za-z0-9_=-]+"
```

```
cbor-wrapped-claims-set = bstr .cbor Claims-Set
```

6. Endorsements and Verification Keys

The Verifier must possess the correct key when it performs the cryptographic part of an EAT verification (e.g., verifying the COSE/JOSE signature). This section describes several ways to identify the verification key. There is not one standard method.

The verification key itself may be a public key, a symmetric key or something complicated in the case of a scheme like Direct Anonymous Attestation (DAA).

RATS Architecture [RATS.Architecture] describes what is called an Endorsement. This is an input to the Verifier that is usually the basis of the trust placed in an EAT and the Attester that generated it. It may contain the public key for verification of the signature on the EAT. It may contain Reference Values to which EAT claims are compared as part of the verification process. It may contain implied claims, those that are passed on to the Relying Party in Attestation Results.

There is not yet any standard format(s) for an Endorsement. One format that may be used for an Endorsement is an X.509 certificate. Endorsement data like Reference Values and implied claims can be carried in X.509 v3 extensions. In this use, the public key in the X.509 certificate becomes the verification key, so identification of the Endorsement is also identification of the verification key.

The verification key identification and establishment of trust in the EAT and the attester may also be by some other means than an Endorsement.

For the components (Attester, Verifier, Relying Party,...) of a particular end-end attestation system to reliably interoperate, its definition should specify how the verification key is identified. Usually, this will be in the profile document for a particular attestation system.

6.1. Identification Methods

Following is a list of possible methods of key identification. A specific attestation system may employ any one of these or one not listed here.

The following assumes Endorsements are X.509 certificates or equivalent and thus does not mention or define any identifier for Endorsements in other formats. If such an Endorsement format is created, new identifiers for them will also need to be created.

6.1.1. COSE/JWS Key ID

The COSE standard header parameter for Key ID (kid) may be used. See [RFC8152] and [RFC7515]

COSE leaves the semantics of the key ID open-ended. It could be a record locator in a database, a hash of a public key, an input to a

KDF, an authority key identifier (AKI) for an X.509 certificate or other. The profile document should specify what the key ID's semantics are.

6.1.2. JWS and COSE X.509 Header Parameters

COSE X.509 [COSE.X509.Draft] and JSON Web Signature [RFC7515] define several header parameters (x5t, x5u,...) for referencing or carrying X.509 certificates any of which may be used.

The X.509 certificate may be an Endorsement and thus carrying additional input to the Verifier. It may be just an X.509 certificate, not an Endorsement. The same header parameters are used in both cases. It is up to the attestation system design and the Verifier to determine which.

6.1.3. CBOR Certificate COSE Header Parameters

Compressed X.509 and CBOR Native certificates are defined by CBOR Certificates [CBOR.Cert.Draft]. These are semantically compatible with X.509 and therefore can be used as an equivalent to X.509 as described above.

These are identified by their own header parameters (c5t, c5u,...).

6.1.4. Claim-Based Key Identification

For some attestation systems, a claim may be re-used as a key identifier. For example, the UEID uniquely identifies the entity and therefore can work well as a key identifier or Endorsement identifier.

This has the advantage that key identification requires no additional bytes in the EAT and makes the EAT smaller.

This has the disadvantage that the unverified EAT must be substantially decoded to obtain the identifier since the identifier is in the COSE/JOSE payload, not in the headers.

6.2. Other Considerations

In all cases there must be some way that the verification key is itself verified or determined to be trustworthy. The key identification itself is never enough. This will always be by some out-of-band mechanism that is not described here. For example, the Verifier may be configured with a root certificate or a master key by the Verifier system administrator.

Often an X.509 certificate or an Endorsement carries more than just the verification key. For example, an X.509 certificate might have key usage constraints and an Endorsement might have Reference Values. When this is the case, the key identifier must be either a protected header or in the payload such that it is cryptographically bound to the EAT. This is in line with the requirements in section 6 on Key Identification in JSON Web Signature [RFC7515].

7. Profiles

This EAT specification does not guarantee that implementations of it will interoperate. The variability in this specification is necessary to accommodate the widely varying use cases. An EAT profile narrows the specification for a specific use case. An ideal EAT profile will guarantee interoperability.

The profile can be named in the token using the profile claim described in Section 3.20.

A profile can apply to Attestation Evidence or to Attestation Results or both.

7.1. Format of a Profile Document

A profile document doesn't have to be in any particular format. It may be simple text, something more formal or a combination.

In some cases CDDL may be created that replaces CDDL in this or other document to express some profile requirements. For example, to require the altitude data item in the location claim, CDDL can be written that replicates the location claim with the altitude no longer optional.

7.2. List of Profile Issues

The following is a list of EAT, CWT, UCCS, JWS, UJCS, COSE, JOSE and CBOR options that a profile should address.

7.2.1. Use of JSON, CBOR or both

The profile should indicate whether the token format should be CBOR, JSON, both or even some other encoding. If some other encoding, a specification for how the CDDL described here is serialized in that encoding is necessary.

This should be addressed for the top-level token and for any nested tokens. For example, a profile might require all nested tokens to be of the same encoding of the top level token.

7.2.2. CBOR Map and Array Encoding

The profile should indicate whether definite-length arrays/maps, indefinite-length arrays/maps or both are allowed. A good default is to allow only definite-length arrays/maps.

An alternate is to allow both definite and indefinite-length arrays/maps. The decoder should accept either. Encoders that need to fit on very small hardware or be actually implement in hardware can use indefinite-length encoding.

This applies to individual EAT claims, CWT and COSE parts of the implementation.

7.2.3. CBOR String Encoding

The profile should indicate whether definite-length strings, indefinite-length strings or both are allowed. A good default is to allow only definite-length strings. As with map and array encoding, allowing indefinite-length strings can be beneficial for some smaller implementations.

7.2.4. CBOR Preferred Serialization

The profile should indicate whether encoders must use preferred serialization. The profile should indicate whether decoders must accept non-preferred serialization.

7.2.5. COSE/JOSE Protection

COSE and JOSE have several options for signed, MACed and encrypted messages. EAT/CWT has the option to have no protection using UCCS and JOSE has a NULL protection option. It is possible to implement no protection, sign only, MAC only, sign then encrypt and so on. All combinations allowed by COSE, JOSE, JWT, CWT, UCCS and UJCS are allowed by EAT.

The profile should list the protections that must be supported by all decoders implementing the profile. The encoders then must implement a subset of what is listed for the decoders, perhaps only one.

Implementations may choose to sign or MAC before encryption so that the implementation layer doing the signing or MACing can be the smallest. It is often easier to make smaller implementations more secure, perhaps even implementing in solely in hardware. The key material for a signature or MAC is a private key, while for encryption it is likely to be a public key. The key for encryption requires less protection.

7.2.6. COSE/JOSE Algorithms

The profile document should list the COSE algorithms that a Verifier must implement. The Attester will select one of them. Since there is no negotiation, the Verifier should implement all algorithms listed in the profile. If detached submodules are used, the COSE algorithms allowed for their digests should also be in the profile.

7.2.7. DEB Support

A Detached EAT Bundle Section 5 is a special case message that will not often be used. A profile may prohibit its use.

7.2.8. Verification Key Identification

Section 6 describes a number of methods for identifying a verification key. The profile document should specify one of these or one that is not described. The ones described in this document are only roughly described. The profile document should go into the full detail.

7.2.9. Endorsement Identification

Similar to, or perhaps the same as Verification Key Identification, the profile may wish to specify how Endorsements are to be identified. However note that Endorsement Identification is optional, where as key identification is not.

7.2.10. Freshness

Just about every use case will require some means of knowing the EAT is recent enough and not a replay of an old token. The profile should describe how freshness is achieved. The section on Freshness in [RATS.Architecture] describes some of the possible solutions to achieve this.

7.2.11. Required Claims

The profile can list claims whose absence results in Verification failure.

7.2.12. Prohibited Claims

The profile can list claims whose presence results in Verification failure.

7.2.13. Additional Claims

The profile may describe entirely new claims. These claims can be required or optional.

7.2.14. Refined Claim Definition

The profile may lock down optional aspects of individual claims. For example, it may require altitude in the location claim, or it may require that HW Versions always be described using EAN-13.

7.2.15. CBOR Tags

The profile should specify whether the token should be a CWT Tag or not. Similarly, the profile should specify whether the token should be a UCCS tag or not.

When COSE protection is used, the profile should specify whether COSE tags are used or not. Note that RFC 8392 requires COSE tags be used in a CWT tag.

Often a tag is unnecessary because the surrounding or carrying protocol identifies the object as an EAT.

7.2.16. Manifests and Software Evidence Claims

The profile should specify which formats are allowed for the manifests and software evidence claims. The profile may also go on to say which parts and options of these formats are used, allowed and prohibited.

8. Encoding and Collected CDDL

An EAT is fundamentally defined using CDDL. This document specifies how to encode the CDDL in CBOR or JSON. Since CBOR can express some things that JSON can't (e.g., tags) or that are expressed differently (e.g., labels) there is some CDDL that is specific to the encoding format.

8.1. Claims-Set and CDDL for CWT and JWT

CDDL was not used to define CWT or JWT. It was not available at the time.

This document defines CDDL for both CWT and JWT as well as UCCS. This document does not change the encoding or semantics of anything in a CWT or JWT.

A Claims-Set is the central data structure for EAT, CWT, JWT and UCCS. It holds all the claims and is the structure that is secured by signing or other means. It is not possible to define EAT, CWT, JWT or UCCS in CDDL without it. The CDDL definition of Claims-Set here is applicable to EAT, CWT, JWT and UCCS.

This document specifies how to encode a Claims-Set in CBOR or JSON.

With the exception of nested tokens and some other externally defined structures (e.g., SWIDs) an entire Claims-Set must be encoded in either CBOR or JSON, never a mixture.

CDDL for the seven claims defined by [RFC8392] and [RFC7519] is included here.

8.2. Encoding Data Types

This makes use of the types defined in [RFC8610] Appendix D, Standard Prelude.

8.2.1. Common Data Types

time-int is identical to the epoch-based time, but disallows floating-point representation.

Unless explicitly indicated, URIs are not the URI tag defined in [RFC8949]. They are just text strings that contain a URI.

string-or-uri = tstr

time-int = #6.1(int)

8.2.2. JSON Interoperability

JSON should be encoded per [RFC8610] Appendix E. In addition, the following CDDL types are encoded in JSON as follows:

- o bstr - must be base64url encoded
- o time - must be encoded as NumericDate as described section 2 of [RFC7519].
- o string-or-uri - must be encoded as StringOrURI as described section 2 of [RFC7519].
- o uri - must be a URI [RFC3986].

- o oid - encoded as a string using the well established dotted-decimal notation (e.g., the text "1.2.250.1").

8.2.3. Labels

Map labels, including Claims-Keys and Claim-Names, and enumerated-type values are always integers when encoding in CBOR and strings when encoding in JSON. There is an exception to this for naming submodules and detached claims sets in a DEB. These are strings in CBOR.

The CDDL in most cases gives both the integer label and the string label as it is not convenient to have conditional CDDL for such.

8.3. CBOR Interoperability

CBOR allows data items to be serialized in more than one form. If the sender uses a form that the receiver can't decode, there will not be interoperability.

This specification gives no blanket requirements to narrow CBOR serialization for all uses of EAT. This allows individual uses to tailor serialization to the environment. It also may result in EAT implementations that don't interoperate.

One way to guarantee interoperability is to clearly specify CBOR serialization in a profile document. See Section 7 for a list of serialization issues that should be addressed.

EAT will be commonly used where the entity generating the attestation is constrained and the receiver/Verifier of the attestation is a capacious server. Following is a set of serialization requirements that work well for that use case and are guaranteed to interoperate. Use of this serialization is recommended where possible, but not required. An EAT profile may just reference the following section rather than spell out serialization details.

8.3.1. EAT Constrained Device Serialization

- o Preferred serialization described in section 4.1 of [RFC8949] is not required. The EAT decoder must accept all forms of number serialization. The EAT encoder may use any form it wishes.
- o The EAT decoder must accept indefinite length arrays and maps as described in section 3.2.2 of [RFC8949]. The EAT encoder may use indefinite length arrays and maps if it wishes.

- o The EAT decoder must accept indefinite length strings as described in section 3.2.3 of [RFC8949]. The EAT encoder may use indefinite length strings if it wishes.
- o Sorting of maps by key is not required. The EAT decoder must not rely on sorting.
- o Deterministic encoding described in Section 4.2 of [RFC8949] is not required.
- o Basic validity described in section 5.3.1 of [RFC8949] must be followed. The EAT encoder must not send duplicate map keys/labels or invalid UTF-8 strings.

8.4. Collected Common CDDL

```
Claims-Set = {  
    * $$claims-set-claims,  
    * Claim-Label .feature "extended-label" => any  
}  
  
Claim-Label = int / text  
string-or-uri = tstr  
  
time-int = #6.1(int)  
$$claims-set-claims /= (iss-label => text)  
$$claims-set-claims /= (sub-label => text)  
$$claims-set-claims /= (aud-label => text)  
$$claims-set-claims /= (exp-label => ~time)  
$$claims-set-claims /= (nbf-label => ~time)  
$$claims-set-claims /= (iat-label => ~time)  
  
$$claims-set-claims /=  
    (nonce-label => nonce-type / [ 2* nonce-type ])  
  
nonce-type = bstr .size (8..64)  
$$claims-set-claims /= (ueid-label => ueid-type)  
  
ueid-type = bstr .size (7..33)  
$$claims-set-claims /= (sueids-label => sueids-type)  
  
sueids-type = {  
    + tstr => ueid-type  
}  
oemid-pen = int  
  
oemid-ieee = bstr .size 3
```

```
oemid-random = bstr .size 16

$$claims-set-claims //= (
    oemid-label =>
        oemid-random / oemid-ieee / oemid-pen
)
$$claims-set-claims //= (
    hardware-version-label => hardware-version-type
)

hardware-version-type = [
    version: tstr,
    scheme: $version-scheme
]
hardware-model-type = bytes .size (1..32)

$$claims-set-claims //= (
    hardware-model-label => hardware-model-type
)
$$claims-set-claims //= ( sw-name-label => tstr )
$$claims-set-claims //= (sw-version-label => sw-version-type)

sw-version-type = [
    version: tstr,
    scheme: $version-scheme ; As defined by CoSWID
]
$$claims-set-claims //= (
    security-level-label =>
        security-level-cbor-type /
        security-level-json-type
)

security-level-cbor-type = &(
    unrestricted: 1,
    restricted: 2,
    secure-restricted: 3,
    hardware: 4
)

security-level-json-type =
    "unrestricted" /
    "restricted" /
    "secure-restricted" /
    "hardware"
$$claims-set-claims //= (secure-boot-label => bool)
$$claims-set-claims //= (
    debug-status-label =>
```

```

        debug-status-cbor-type / debug-status-json-type
    )

    debug-status-cbor-type = &(
        enabled: 0,
        disabled: 1,
        disabled-since-boot: 2,
        disabled-permanently: 3,
        disabled-fully-and-permanently: 4
    )

    debug-status-json-type =
        "enabled" /
        "disabled" /
        "disabled-since-boot" /
        "disabled-permanently" /
        "disabled-fully-and-permanently"
    $$claims-set-claims // = (location-label => location-type)

    location-type = {
        latitude => number,
        longitude => number,
        ? altitude => number,
        ? accuracy => number,
        ? altitude-accuracy => number,
        ? heading => number,
        ? speed => number,
        ? timestamp => ~time-int,
        ? age => uint
    }

    latitude = 1 / "latitude"
    longitude = 2 / "longitude"
    altitude = 3 / "altitude"
    accuracy = 4 / "accuracy"
    altitude-accuracy = 5 / "altitude-accuracy"
    heading = 6 / "heading"
    speed = 7 / "speed"
    timestamp = 8 / "timestamp"
    age = 9 / "age"
    $$claims-set-claims // = (uptime-label => uint)
    $$claims-set-claims // = (boot-seed-label => bytes)
    $$claims-set-claims // = (odometer-label => uint)
    $$claims-set-claims // = (
        intended-use-label =>
            intended-use-cbor-type / intended-use-json-type
    )

```

```
intended-use-cbor-type = &(
    generic: 1,
    registration: 2,
    provisioning: 3,
    csr: 4,
    pop: 5
)

intended-use-json-type =
    "generic" /
    "registration" /
    "provisioning" /
    "csr" /
    "pop"
$$claims-set-claims // = (
    dloas-label => [ + dloa-type ]
)

dloa-type = [
    dloa_registrar: ~uri
    dloa_platform_label: text
    ? dloa_application_label: text
]
$$claims-set-claims // = (profile-label => ~uri / ~oid)
$$claims-set-claims // = (
    manifests-label => manifests-type
)

manifests-type = [+ $$manifest-formats]

coswid-that-is-a-cbor-tag-xx = tagged-coswid<concise-swid-tag>

$$manifest-formats /= bytes .cbor coswid-that-is-a-cbor-tag-xx$$claims-set-claims
// = (
    swevidence-label => swevidence-type
)

swevidence-type = [+ $$swevidence-formats]

coswid-that-is-a-cbor-tag = tagged-coswid<concise-swid-tag>
$$swevidence-formats /= bytes .cbor coswid-that-is-a-cbor-tag
$$claims-set-claims // = (swresults-label => [ + swresult-type ])

verification-result-cbor-type = &(
    verification-not-run: 1,
    verification-indeterminate: 2,
    verification-failed: 3,
    fully-verified: 4,
    partially-verified: 5,
```



```
)

verification-result-json-type =
    "verification-not-run" /
    "verification-indeterminate" /
    "verification-failed" /
    "fully-verified" /
    "partially-verified"

verification-objective-cbor-type = &(
    all: 1,
    firmware: 2,
    kernel: 3,
    privileged: 4,
    system-libs: 5,
    partial: 6,
)

verification-objective-json-type =
    "all" /
    "firmware" /
    "kernel" /
    "privileged" /
    "system-libs" /
    "partial"

swresult-type = [
    verification-system: tstr,
    objective: verification-objective-cbor-type /
        verification-objective-json-type,
    result: verification-result-cbor-type /
        verification-result-json-type,
    ? objective-name: tstr
]
$$claims-set-claims // = (submods-label => { + text => Submodule })

Submodule = Claims-Set / Nested-Token / Detached-Submodule-Digest

Detached-Submodule-Digest = [
    algorithm : int / text,
    digest : bstr
]
Detached-EAT-Bundle = [
    main-token : Nested-Token,
    detached-claims-sets: {
        + tstr => cbor-wrapped-claims-set / json-wrapped-claims-set
    }
]
```

```
    }
  ]
```

```
json-wrapped-claims-set = tstr .regexp "[A-Za-z0-9_=-]+"
```

```
cbor-wrapped-claims-set = bstr .cbor Claims-Set
```

8.5. Collected CDDL for CBOR

```
CBOR-Token = Tagged-CBOR-Token / Untagged-CBOR-Token
```

```
Tagged-CBOR-Token  = CWT-Tagged-Message
Tagged-CBOR-Token /= UCCS-Tagged-Message
Tagged-CBOR-Token /= DEB-Tagged-Message
```

```
Untagged-CBOR-Token  = CWT-Untagged-Message
Untagged-CBOR-Token /= UCCS-Untagged-Message
Untagged-CBOR-Token /= DEB-Untagged-Message
```

```
CWT-Tagged-Message = COSE_Tagged_Message
CWT-Untagged-Message = COSE_Untagged_Message
```

```
UCCS-Message = UCCS-Tagged-Message / UCCS-Untagged-Message
```

```
UCCS-Tagged-Message = #6.601(UCCS-Untagged-Message)
```

```
UCCS-Untagged-Message = Claims-Set
```

```
DEB-Tagged-Message = #6.602(DEB-Untagged-Message)
```

```
DEB-Untagged-Message = Detached-EAT-Bundle
```

```
Nested-Token =
  tstr / ; A JSON-encoded Nested-Token (see json-nested-token.cddl)
  bstr .cbor Tagged-CBOR-Token
```

```
iss-label = 1
sub-label = 2
aud-label = 3
```

```
exp-label = 4
nbf-label = 5
iat-label = 6
cti-label = 7nonce-label = 10
ueid-label = 256
sueids-label = 257
oemid-label = 258
hardware-model-label = 259
hardware-version-label = 260
secure-boot-label = 262
debug-status-label = 263
location-label = 264
profile-label = 265
submods-label = 266
security-level-label = <TBD>
uptime-label = <TBD>
boot-seed-label = <TB>
odometer-label = <TBD>
intended-use-label = <TBD>
dloas-label = <TBD>
sw-name-label = <TBD>
sw-version-label = <TBD>
manifests-label = <TBD>
swevidence-label = <TBD>
swresults-label = <TBD>
```

8.6. Collected CDDL for JSON

```
JWT-Message = text .regexp [A-Za-z0-9_=-]+\.[A-Za-z0-9_=-]+\.[A-Za-z0-9_=-]+
```

```
UJCS-Message = Claims-Set
```

```
Nested-Token = [
  type : "JWT" / "CBOR" / "UJCS" / "DEB",
  nested-token : JWT-Message /
                  B64URL-Tagged-CBOR-Token /
                  DEB-JSON-Message /
                  UJCS-Message
]
```

```
B64URL-Tagged-CBOR-Token = tstr .regexp "[A-Za-z0-9_=-]+"
iss-label = "iss"
sub-label = "sub"
aud-label = "aud"
```

```
exp-label = "exp"
nbf-label = "nbf"
iat-label = "iat"
cti-label = "cti"nonce-label /= "nonce"
```

```
ueid-label /= "ueid"
sueids-label /= "sueids"
oemid-label /= "oemid"
hardware-model-label /= "hwmodel"
hardware-version-label /= "hwversion"
security-level-label /= "seclevel"
secure-boot-label /= "secboot"
debug-status-label /= "dbgstat"
location-label /= "location"
profile-label /= "eat-profile"
uptime-label /= "uptime"
boot-seed-label /= "bootseed"
odometer-label /= "odometer"
intended-use-label /= "intuse"
dloas-label /= "dloas"
sw-name-label /= "swname"
sw-version-label /= "swversion"
manifests-label /= "manifests"
swevidence-label /= "swevidence"
swresults-label /= "swresults"
submods-label /= "submods"
```

```
latitude /= "lat"
longitude /= "long"
altitude /= "alt"
accuracy /= "accry"
altitude-accuracy /= "alt-accry"
heading /= "heading"
speed /= "speed"
```

9. IANA Considerations

9.1. Reuse of CBOR and JSON Web Token (CWT and JWT) Claims Registries

Claims defined for EAT are compatible with those of CWT and JWT so the CWT and JWT Claims Registries, [IANA.CWT.Claims] and [IANA.JWT.Claims], are re used. No new IANA registry is created.

All EAT claims defined in this document are placed in both registries. All new EAT claims defined subsequently should be placed in both registries.

9.2. Claim Characteristics

The following is design guidance for creating new EAT claims, particularly those to be registered with IANA.

Much of this guidance is generic and could also be considered when designing new CWT or JWT claims.

9.2.1. Interoperability and Relying Party Orientation

It is a broad goal that EATs can be processed by Relying Parties in a general way regardless of the type, manufacturer or technology of the device from which they originate. It is a goal that there be general-purpose verification implementations that can verify tokens for large numbers of use cases with special cases and configurations for different device types. This is a goal of interoperability of the semantics of claims themselves, not just of the signing, encoding and serialization formats.

This is a lofty goal and difficult to achieve broadly requiring careful definition of claims in a technology neutral way. Sometimes it will be difficult to design a claim that can represent the semantics of data from very different device types. However, the goal remains even when difficult.

9.2.2. Operating System and Technology Neutral

Claims should be defined such that they are not specific to an operating system. They should be applicable to multiple large high-level operating systems from different vendors. They should also be applicable to multiple small embedded operating systems from multiple vendors and everything in between.

Claims should not be defined such that they are specific to a SW environment or programming language.

Claims should not be defined such that they are specific to a chip or particular hardware. For example, they should not just be the contents of some HW status register as it is unlikely that the same HW status register with the same bits exists on a chip of a different manufacturer.

The boot and debug state claims in this document are an example of a claim that has been defined in this neutral way.

9.2.3. Security Level Neutral

Many use cases will have EATs generated by some of the most secure hardware and software that exists. Secure Elements and smart cards are examples of this. However, EAT is intended for use in low-security use cases the same as high-security use case. For example, an app on a mobile device may generate EATs on its own.

Claims should be defined and registered on the basis of whether they are useful and interoperable, not based on security level. In particular, there should be no exclusion of claims because they are just used only in low-security environments.

9.2.4. Reuse of Extant Data Formats

Where possible, claims should use already standardized data items, identifiers and formats. This takes advantage of the expertise put into creating those formats and improves interoperability.

Often extant claims will not be defined in an encoding or serialization format used by EAT. It is preferred to define a CBOR and JSON format for them so that EAT implementations do not require a plethora of encoders and decoders for serialization formats.

In some cases, it may be better to use the encoding and serialization as is. For example, signed X.509 certificates and CRLs can be carried as-is in a byte string. This retains interoperability with the extensive infrastructure for creating and processing X.509 certificates and CRLs.

9.2.5. Proprietary Claims

EAT allows the definition and use of proprietary claims.

For example, a device manufacturer may generate a token with proprietary claims intended only for verification by a service offered by that device manufacturer. This is a supported use case.

In many cases proprietary claims will be the easiest and most obvious way to proceed, however for better interoperability, use of general standardized claims is preferred.

9.3. Claims Registered by This Document

This specification adds the following values to the "JSON Web Token Claims" registry established by [RFC7519] and the "CBOR Web Token Claims Registry" established by [RFC8392]. Each entry below is an

addition to both registries (except for the nonce claim which is already registered for JWT, but not registered for CWT).

The "Claim Description", "Change Controller" and "Specification Documents" are common and equivalent for the JWT and CWT registries. The "Claim Key" and "Claim Value Types(s)" are for the CWT registry only. The "Claim Name" is as defined for the CWT registry, not the JWT registry. The "JWT Claim Name" is equivalent to the "Claim Name" in the JWT registry.

9.3.1. Claims for Early Assignment

RFC Editor: in the final publication this section should be combined with the following section as it will no longer be necessary to distinguish claims with early assignment. Also, the following paragraph should be removed.

The claims in this section have been (requested for / given) early assignment according to [RFC7120]. They have been assigned values and registered before final publication of this document. While their semantics is not expected to change in final publication, it is possible that they will. The JWT Claim Names and CWT Claim Keys are not expected to change.

In draft -06 an early allocation was described. The processing of that early allocation was never correctly completed. This early allocation assigns different numbers for the CBOR claim labels. This early allocation will presumably complete correctly

- o Claim Name: Nonce
- o Claim Description: Nonce
- o JWT Claim Name: "nonce" (already registered for JWT)
- o Claim Key: TBD (requested value 10)
- o Claim Value Type(s): byte string
- o Change Controller: IESG
- o Specification Document(s): [OpenIDConnectCore], *this document*
- o Claim Name: UEID
- o Claim Description: The Universal Entity ID
- o JWT Claim Name: "ueid"

- o CWT Claim Key: TBD (requested value 256)
- o Claim Value Type(s): byte string
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SUEIDs
- o Claim Description: Semi-permanent UEIDs
- o JWT Claim Name: "sueids"
- o CWT Claim Key: TBD (requested value 257)
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Hardware OEMID
- o Claim Description: Hardware OEM ID
- o JWT Claim Name: "oemid"
- o Claim Key: TBD (requested value 258)
- o Claim Value Type(s): byte string or integer
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Hardware Model
- o Claim Description: Model identifier for hardware
- o JWT Claim Name: "hwmodel"
- o Claim Key: TBD (requested value 259)
- o Claim Value Type(s): byte string
- o Change Controller: IESG

- o Specification Document(s): *this document*
- o Claim Name: Hardware Version
- o Claim Description: Hardware Version Identifier
- o JWT Claim Name: "hwversion"
- o Claim Key: TBD (requested value 260)
- o Claim Value Type(s): array
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Secure Boot
- o Claim Description: Indicate whether the boot was secure
- o JWT Claim Name: "secboot"
- o Claim Key: TBD (requested value 262)
- o Claim Value Type(s): Boolean
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Debug Status
- o Claim Description: Indicate status of debug facilities
- o JWT Claim Name: "dbgstat"
- o Claim Key: TBD (requested value 263)
- o Claim Value Type(s): integer or string
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Location
- o Claim Description: The geographic location

- o JWT Claim Name: "location"
- o Claim Key: TBD (requested value 264)
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Profile
- o Claim Description: Indicates the EAT profile followed
- o JWT Claim Name: "eat_profile"
- o Claim Key: TBD (requested value 265)
- o Claim Value Type(s): URI or OID
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Submodules Section
- o Claim Description: The section containing submodules
- o JWT Claim Name: "submods"
- o Claim Key: TBD (requested value 266)
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*

9.3.2. To be Assigned Claims

(Early assignment is NOT requested for these claims. Implementers should be aware they may change)

- o Claim Name: Security Level
- o Claim Description: Characterization of the security of an Attester or submodule

- o JWT Claim Name: "secclevel"
- o Claim Key: TBD
- o Claim Value Type(s): integer or string
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Uptime
- o Claim Description: Uptime
- o JWT Claim Name: "uptime"
- o Claim Key: TBD
- o Claim Value Type(s): unsigned integer
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Boot Seed
- o Claim Description: Identifies a boot cycle
- o JWT Claim Name: "bootseed"
- o Claim Key: TBD
- o Claim Value Type(s): bytes
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: Intended Use
- o Claim Description: Indicates intended use of the EAT
- o JWT Claim Name: "intuse"
- o Claim Key: TBD
- o Claim Value Type(s): integer or string

- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: DLOAs
- o Claim Description: Certifications received as Digital Letters of Approval
- o JWT Claim Name: "dloas"
- o Claim Key: TBD
- o Claim Value Type(s): array
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SW Name
- o Claim Description: The name of the SW running in the entity
- o JWT Claim Name: "swname"
- o Claim Key: TBD
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SW Version
- o Claim Description: The version of SW running in the entity
- o JWT Claim Name: "swversion"
- o Claim Key: TBD
- o Claim Value Type(s): map
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SW Manifests

- o Claim Description: Manifests describing the SW installed on the entity
- o JWT Claim Name: "manifests"
- o Claim Key: TBD
- o Claim Value Type(s): array
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SW Evidence
- o Claim Description: Measurements of the SW, memory configuration and such on the entity
- o JWT Claim Name: "swevidence"
- o Claim Key: TBD
- o Claim Value Type(s): array
- o Change Controller: IESG
- o Specification Document(s): *this document*
- o Claim Name: SW Measurment Results
- o Claim Description: The results of comparing SW measurements to reference values
- o JWT Claim Name: "swresults"
- o Claim Key: TBD
- o Claim Value Type(s): array
- o Change Controller: IESG
- o Specification Document(s): *this document*

9.3.3. Version Schemes Registered by this Document

IANA is requested to register a new value in the "Software Tag Version Scheme Values" established by [CoSWID].

The new value is a version scheme a 13-digit European Article Number [EAN-13]. An EAN-13 is also known as an International Article Number or most commonly as a bar code. This version scheme is the ASCII text representation of EAN-13 digits, the same ones often printed with a bar code. This version scheme must comply with the EAN allocation and assignment rules. For example, this requires the manufacturer to obtain a manufacture code from GS1.

Index	Version Scheme Name	Specification
5	ean-13	This document

9.3.4. UEID URN Registered by this Document

IANA is requested to register the following new subtypes in the "DEV URN Subtypes" registry under "Device Identification". See [RFC9039].

Subtype	Description	Reference
ueid	Universal Entity Identifier	This document
sueid	Semi-permanent Universal Entity Identifier	This document

9.3.5. Tag for Detached EAT Bundle

In the registry [IANA.cbor-tags], IANA is requested to allocate the following tag from the FCFS space, with the present document as the specification reference.

Tag	Data Items	Semantics
TBD602	array	Detached EAT Bundle Section 5

10. Privacy Considerations

Certain EAT claims can be used to track the owner of an entity and therefore, implementations should consider providing privacy-preserving options dependent on the intended usage of the EAT. Examples would include suppression of location claims for EAT's provided to unauthenticated consumers.

10.1. UEID and SUEID Privacy Considerations

A UEID is usually not privacy-preserving. Any set of Relying Parties that receives tokens that happen to be from a particular entity will be able to know the tokens are all from the same entity and be able to track it.

Thus, in many usage situations UEID violates governmental privacy regulation. In other usage situations a UEID will not be allowed for certain products like browsers that give privacy for the end user. It will often be the case that tokens will not have a UEID for these reasons.

An SUEID is also usually not privacy-preserving. In some cases it may have fewer privacy issues than a UEID depending on when and how and when it is generated.

There are several strategies that can be used to still be able to put UEIDs and SUEIDs in tokens:

- o The entity obtains explicit permission from the user of the entity to use the UEID/SUEID. This may be through a prompt. It may also be through a license agreement. For example, agreements for some online banking and brokerage services might already cover use of a UEID/SUEID.
- o The UEID/SUEID is used only in a particular context or particular use case. It is used only by one Relying Party.
- o The entity authenticates the Relying Party and generates a derived UEID/SUEID just for that particular Relying Party. For example, the Relying Party could prove their identity cryptographically to the entity, then the entity generates a UEID just for that Relying Party by hashing a proofed Relying Party ID with the main entity UEID/SUEID.

Note that some of these privacy preservation strategies result in multiple UEIDs and SUEIDs per entity. Each UEID/SUEID is used in a different context, use case or system on the entity. However, from the view of the Relying Party, there is just one UEID and it is still globally universal across manufacturers.

10.2. Location Privacy Considerations

Geographic location is most always considered personally identifiable information. Implementers should consider laws and regulations governing the transmission of location data from end user devices to servers and services. Implementers should consider using location

management facilities offered by the operating system on the entity generating the attestation. For example, many mobile phones prompt the user for permission when before sending location data.

10.3. Replay Protection and Privacy

EAT offers 2 primary mechanisms for token replay protection (also sometimes known as token "freshness"): the cti/jti claim and the nonce claim. The cti/jti claim in a CWT/JWT is a field that may be optionally included in the EAT and is in general derived on the same device in which the entity is instantiated. The nonce claim is based on a value that is usually derived remotely (outside of the entity). These claims can be used to extract and convey personally-identifying information either inadvertently or by intention. For instance, an implementor may choose a cti that is equivalent to a username associated with the device (e.g., account login). If the token is inspected by a 3rd-party then this information could be used to identify the source of the token or an account associated with the token (e.g., if the account name is used to derive the nonce). In order to avoid the conveyance of privacy-related information in either the cti/jti or nonce claims, these fields should be derived using a salt that originates from a true and reliable random number generator or any other source of randomness that would still meet the target system requirements for replay protection.

11. Security Considerations

The security considerations provided in Section 8 of [RFC8392] and Section 11 of [RFC7519] apply to EAT in its CWT and JWT form, respectively. In addition, implementors should consider the following.

11.1. Key Provisioning

Private key material can be used to sign and/or encrypt the EAT, or can be used to derive the keys used for signing and/or encryption. In some instances, the manufacturer of the entity may create the key material separately and provision the key material in the entity itself. The manufacturer of any entity that is capable of producing an EAT should take care to ensure that any private key material be suitably protected prior to provisioning the key material in the entity itself. This can require creation of key material in an enclave (see [RFC4949] for definition of "enclave"), secure transmission of the key material from the enclave to the entity using an appropriate protocol, and persistence of the private key material in some form of secure storage to which (preferably) only the entity has access.

11.1.1. Transmission of Key Material

Regarding transmission of key material from the enclave to the entity, the key material may pass through one or more intermediaries. Therefore some form of protection ("key wrapping") may be necessary. The transmission itself may be performed electronically, but can also be done by human courier. In the latter case, there should be minimal to no exposure of the key material to the human (e.g. encrypted portable memory). Moreover, the human should transport the key material directly from the secure enclave where it was created to a destination secure enclave where it can be provisioned.

11.2. Transport Security

As stated in Section 8 of [RFC8392], "The security of the CWT relies upon on the protections offered by COSE". Similar considerations apply to EAT when sent as a CWT. However, EAT introduces the concept of a nonce to protect against replay. Since an EAT may be created by an entity that may not support the same type of transport security as the consumer of the EAT, intermediaries may be required to bridge communications between the entity and consumer. As a result, it is RECOMMENDED that both the consumer create a nonce, and the entity leverage the nonce along with COSE mechanisms for encryption and/or signing to create the EAT.

Similar considerations apply to the use of EAT as a JWT. Although the security of a JWT leverages the JSON Web Encryption (JWE) and JSON Web Signature (JWS) specifications, it is still recommended to make use of the EAT nonce.

11.3. Multiple EAT Consumers

In many cases, more than one EAT consumer may be required to fully verify the entity attestation. Examples include individual consumers for nested EATs, or consumers for individual claims with an EAT. When multiple consumers are required for verification of an EAT, it is important to minimize information exposure to each consumer. In addition, the communication between multiple consumers should be secure.

For instance, consider the example of an encrypted and signed EAT with multiple claims. A consumer may receive the EAT (denoted as the "receiving consumer"), decrypt its payload, verify its signature, but then pass specific subsets of claims to other consumers for evaluation ("downstream consumers"). Since any COSE encryption will be removed by the receiving consumer, the communication of claim subsets to any downstream consumer should leverage a secure protocol (e.g. one that uses transport-layer security, i.e. TLS),

However, assume the EAT of the previous example is hierarchical and each claim subset for a downstream consumer is created in the form of a nested EAT. Then transport security between the receiving and downstream consumers is not strictly required. Nevertheless, downstream consumers of a nested EAT should provide a nonce unique to the EAT they are consuming.

12. References

12.1. Normative References

- [CoSWID] Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", draft-ietf-sacm-coswid-20 (work in progress), January 2022.
- [DLOA] "Digital Letter of Approval", November 2015, <https://globalplatform.org/wp-content/uploads/2015/12/GPC_DigitalLetterOfApproval_v1.0.pdf>.
- [EAN-13] GS1, "International Article Number - EAN/UPC barcodes", 2019, <<https://www.gs1.org/standards/barcodes/ean-upc>>.
- [FIDO.AROE] The FIDO Alliance, "FIDO Authenticator Allowed Restricted Operating Environments List", November 2020, <<https://fidoalliance.org/specs/fido-security-requirements/fido-authenticator-allowed-restricted-operating-environments-list-v1.2-fd-20201102.html>>.
- [IANA.cbor-tags] "IANA CBOR Tags Registry", n.d., <<https://www.iana.org/assignments/cbor-tags/cbor-tags.xhtml>>.
- [IANA.CWT.Claims] IANA, "CBOR Web Token (CWT) Claims", <<http://www.iana.org/assignments/cwt>>.
- [IANA.JWT.Claims] IANA, "JSON Web Token (JWT) Claims", <<https://www.iana.org/assignments/jwt>>.
- [OpenIDConnectCore] Sakimura, N., Bradley, J., Jones, M., Medeiros, B. D., and C. Mortimore, "OpenID Connect Core 1.0 incorporating errata set 1", November 2014, <https://openid.net/specs/openid-connect-core-1_0.html>.

- [PEN] "Private Enterprise Number (PEN) Request", n.d.,
<<https://pen.iana.org/pen/PenApplication.page>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<https://www.rfc-editor.org/info/rfc7159>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<https://www.rfc-editor.org/info/rfc7516>>.
- [RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC7800] Jones, M., Bradley, J., and H. Tschofenig, "Proof-of-Possession Key Semantics for JSON Web Tokens (JWTs)", RFC 7800, DOI 10.17487/RFC7800, April 2016, <<https://www.rfc-editor.org/info/rfc7800>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/info/rfc8392>>.
- [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://www.rfc-editor.org/info/rfc8610>>.
- [RFC8747] Jones, M., Seitz, L., Selander, G., Erdtman, S., and H. Tschofenig, "Proof-of-Possession Key Semantics for CBOR Web Tokens (CWTs)", RFC 8747, DOI 10.17487/RFC8747, March 2020, <<https://www.rfc-editor.org/info/rfc8747>>.
- [RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/RFC8949, December 2020, <<https://www.rfc-editor.org/info/rfc8949>>.
- [RFC9090] Bormann, C., "Concise Binary Object Representation (CBOR) Tags for Object Identifiers", RFC 9090, DOI 10.17487/RFC9090, July 2021, <<https://www.rfc-editor.org/info/rfc9090>>.
- [ThreeGPP.IMEI]
3GPP, "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification", 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=729>>.
- [UCCS.Draft]
Birkholz, H., O'Donoghue, J., Cam-Winget, N., and C. Bormann, "A CBOR Tag for Unprotected CWT Claims Sets", draft-ietf-rats-uccs-02 (work in progress), January 2022.
- [WGS84] National Geospatial-Intelligence Agency (NGA), "WORLD GEODETIC SYSTEM 1984, NGA.STND.0036_1.0.0_WGS84", July 2014, <<https://earth-info.nga.mil/php/download.php?file=coord-wgs84>>.

12.2. Informative References

- [BirthdayAttack]
"Birthday attack",
<https://en.wikipedia.org/wiki/Birthday_attack>.
- [CBOR.Cert.Draft]
Mattsson, J. P., Selander, G., Raza, S., Hoeglund, J., and
M. Furuheid, "CBOR Encoded X.509 Certificates (C509
Certificates)", draft-ietf-cose-chor-encoded-cert-03 (work
in progress), January 2022.
- [Common.Criteria]
"Common Criteria for Information Technology Security
Evaluation", April 2017,
<<https://www.commoncriteriaportal.org/cc/>>.
- [COSE.X509.Draft]
Schaad, J., "CBOR Object Signing and Encryption (COSE):
Header parameters for carrying and referencing X.509
certificates", draft-ietf-cose-x509-08 (work in progress),
December 2020.
- [FIPS-140]
National Institute of Standards, "Security Requirements for
Cryptographic Modules", May 2001,
<<https://csrc.nist.gov/publications/detail/fips/140/2/final>>.
- [IEEE.802-2001]
"IEEE Standard For Local And Metropolitan Area Networks
Overview And Architecture", 2007,
<<https://webstore.ansi.org/standards/ieee/ieee8022001r2007>>.
- [IEEE.802.1AR]
"IEEE Standard, "IEEE 802.1AR Secure Device Identifier",
December 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.
- [IEEE.RA]
"IEEE Registration Authority",
<<https://standards.ieee.org/products-services/regauth/index.html>>.

[OUI.Guide]

"Guidelines for Use of Extended Unique Identifier (EUI), Organizationally Unique Identifier (OUI), and Company ID (CID)", August 2017,
<<https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/tutorials/eui.pdf>>.

[OUI.Lookup]

"IEEE Registration Authority Assignments",
<<https://regauth.standards.ieee.org/standards-ra-web/pub/view.html#registries>>.

[RATS.Architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", draft-ietf-rats-architecture-15 (work in progress), February 2022.

[RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, DOI 10.17487/RFC4122, July 2005,
<<https://www.rfc-editor.org/info/rfc4122>>.

[RFC4422] Melnikov, A., Ed. and K. Zeilenga, Ed., "Simple Authentication and Security Layer (SASL)", RFC 4422, DOI 10.17487/RFC4422, June 2006,
<<https://www.rfc-editor.org/info/rfc4422>>.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007,
<<https://www.rfc-editor.org/info/rfc4949>>.

[RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 7120, DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018,
<<https://www.rfc-editor.org/info/rfc8446>>.

[RFC9039] Arkko, J., Jennings, C., and Z. Shelby, "Uniform Resource Names for Device Identifiers", RFC 9039, DOI 10.17487/RFC9039, June 2021,
<<https://www.rfc-editor.org/info/rfc9039>>.

[W3C.GeoLoc]

Worldwide Web Consortium, "Geolocation API Specification
2nd Edition", January 2018, <[https://www.w3.org/TR/
geolocation-API/#coordinates_interface](https://www.w3.org/TR/geolocation-API/#coordinates_interface)>.

Appendix A. Examples

These examples are either UCCS, shown as CBOR diagnostic, or UJCS messages. Full CWT and JWT examples with signing and encryption are not given.

All UCCS examples can be the payload of a CWT. To do so, they must be converted from the UCCS message to a Claims-Set, which is achieved by "removing" the tag.

UJCS messages can be directly used as the payload of a JWT.

WARNING: These examples use tag and label numbers not yet assigned by IANA.

A.1. Simple TEE Attestation

This is a simple attestation of a TEE that includes a manifest that is a payload CoSWID to describe the TEE's software.

/ This is a UCCS EAT that describes a simple TEE. /

```
601({
  / nonce /           10: h'948f8860d13a463e',
  / security-level / 261: 3, / secure-restricted /
  / secure-boot /    262: true,
  / debug-status /   263: 2, / disabled-since-boot /
  / manifests /      273: [
                                / This is byte-string wrapped /
                                / payload CoSWID. It gives the TEE /
                                / software name, the version and /
                                / the name of the file it is in. /
                                h' da53574944a60064336132340c01016b
                                41636d6520544545204f530d65332e31
                                2e340282a2181f6b41636d6520544545
                                204f53182101a2181f6b41636d652054
                                4545204f5318210206a111a118186e61
                                636d655f7465655f332e657865'
                                ]
  })
```


/ A payload CoSWID created by the SW vendor. All this really does /
 / is name the TEE SW, its version and lists the one file that /
 / makes up the TEE. /

```
1398229316({
  / Unique CoSWID ID /      0: "3a24",
  / tag-version /          12: 1,
  / software-name /        1: "Acme TEE OS",
  / software-version /     13: "3.1.4",
  / entity /               2: [
                                {
                                  / entity-name /      31: "Acme TEE OS",
                                  / role /              33: 1 / tag-creator /
                                },
                                {
                                  / entity-name /      31: "Acme TEE OS",
                                  / role /              33: 2 / software-creator /
                                }
                              ],
  / payload /              6: {
    / ...file /            17: {
      / ...fs-name /       24: "acme_tee_3.exe"
    }
  }
})
```

A.2. Submodules for Board and Device

```

/ This example shows use of submodules to give information /
/ about the chip, board and overall device. /
/
/ The main attestation is associated with the chip with the /
/ CPU and running the main OS. It is what has the keys and /
/ produces the token. /
/
/ The board is made by a different vendor than the chip. /
/ Perhaps it is some generic IoT board. /
/
/ The device is some specific appliance that is made by a /
/ different vendor than either the chip or the board. /
/
/ Here the board and device submodules aren't the typical /
/ target environments as described by the RATS architecture /
/ document, but they are a valid use of submodules. /

{
  / nonce /          10: h'948f8860d13a463e8e',
  / UEID /           256: h'0198f50a4ff6c05861c8860d13a638ea',
  / HW OEM ID /      258: h'894823', / IEEE OUI format OEM ID /
  / HW Model ID /    259: h'549dcecc8b987c737b44e40f7c635ce8'
                        / Hash of chip model name /,
  / HW Version /     260: ["1.3.4", 1], / Multipartnumeric version /
  / SW Name /        271: "Acme OS",
  / SW Version /     272: ["3.5.5", 1],
  / secure-boot /    262: true,
  / debug-status /   263: 3, / permanent-disable /
  / timestamp (iat) / 6: 1526542894,
  / security-level / 261: 3, / secure restricted OS /
  / submods / 266: {
    / A submodule to hold some claims about the circuit board /
    "board" : {
      / HW OEM ID /    258: h'9bef8787ebal3e2c8f6e7cb4blf4619a',
      / HW Model ID / 259: h'ee80f5a66c1fb9742999a8fdab930893'
                        / Hash of board module name /,
      / HW Version /   260: ["2.0a", 2] / multipartnumeric+suffix /
    },

    / A submodule to hold claims about the overall device /
    "device" : {
      / HW OEM ID /    258: 61234, / PEN Format OEM ID /
      / HW Version /   260: ["4012345123456", 5] / EAN-13 format (barcode) /
    }
  }
}

```

A.3. EAT Produced by Attestation Hardware Block

```

/ This is an example of a token produced by a HW block      /
/ purpose-built for attestation. Only the nonce claim changes /
/ from one attestation to the next as the rest either come   /
/ directly from the hardware or from one-time-programmable memory /
/ (e.g. a fuse). 47 bytes encoded in CBOR (8 byte nonce, 16 byte /
/ UEID). /

601({
  / nonce /          10: h'948f8860d13a463e',
  / UEID /           256: h'0198f50a4ff6c05861c8860d13a638ea',
  / OEMID /          258: 64242, / Private Enterprise Number /
  / security-level / 261: 4, / hardware level security /
  / secure-boot /    262: true,
  / debug-status /   263: 3, / disabled-permanently /
  / HW version /     260: [ "3.1", 1 ] / Type is multipartnumeric /
})

```

A.4. Detached EAT Bundle

In this DEB main token is produced by a HW attestation block. The detached Claims-Set is produced by a TEE and is largely identical to the Simple TEE examples above. The TEE digests its Claims-Set and feeds that digest to the HW block.

In a better example the attestation produced by the HW block would be a CWT and thus signed and secured by the HW block. Since the signature covers the digest from the TEE that Claims-Set is also secured.

The DEB itself can be assembled by untrusted SW.

```

/ This is a detached EAT bundle (DEB) tag.  /

602([

  / First part is a full EAT token with claims like nonce and /
  / UEID. Most importantly, it includes a submodule that is a /
  / detached digest which is the hash of the "TEE" claims set /
  / in the next section.                                     /
  /                                                         /
  / This token here is in UCCS format (unsigned). In a more /
  / realistic example, it would be a signed CWT.           /
  h'd90259a80a48948f8860d13a463e190100500198
  f50a4ff6c05861c8860d13a638ea19010219faf2
  19010504190106f5190107031901048263332e31
  0119010aa163544545822f5820e5cf95fd24fab7
  1446742dd58d43dae178e55fe2b94291a9291082
  ffc2635a0b',
  {
    / A CBOR-encoded byte-string wrapped EAT claims-set. It /
    / contains claims suitable for a TEE                       /
    "TEE" : h'a50a48948f8860d13a463e19010503190106f519
           01070219011181585dda53574944a60064336132
           340c01016b41636d6520544545204f530d65332e
           312e340282a2181f6b41636d6520544545204f53
           182101a2181f6b41636d6520544545204f531821
           0206a111a118186e61636d655f7465655f332e65
           7865'
  }
])

```

```

/ This example contains submodule that is a detached digest, /
/ which is the hash of a Claims-Set convey outside this token. /
/ Other than that is is the other example of a token from an /
/ attestation HW block /

601({
  / nonce /          10: h'948f8860d13a463e',
  / UEID /           256: h'0198f50a4ff6c05861c8860d13a638ea',
  / OEMID /          258: 64242, / Private Enterprise Number /
  / security-level / 261: 4, / hardware level security /
  / secure-boot /    262: true,
  / debug-status /   263: 3, / disabled-permanently /
  / hw version /     260: [ "3.1", 1 ], / multipartnumeric /
  / submods/         266: {
                                "TEE": [ / detached digest submod /
                                          -16, / SHA-256 /
                                          h'e5cf95fd24fab7144674
                                          2dd58d43dae178e55fe2
                                          b94291a9291082ffc2635
                                          a0b'
                                ]
  }
})

```

A.5. Key / Key Store Attestation

```

/ This is an attestation of a public key and the key store /
/ implementation that protects and manages it. The key store /
/ implementation is in a security-oriented execution /
/ environment separate from the high-level OS, for example a /
/ TEE. The key store is the Attester. /
/ /
/ There is some attestation of the high-level OS, just version /
/ and boot & debug status. It is a Claims-Set submodule because /
/ it has lower security level than the key store. The key /
/ store's implementation has access to info about the HLOS, so /
/ it is able to include it. /
/ /
/ A key and an indication of the user authentication given to /
/ allow access to the key is given. The labels for these are /
/ in the private space since this is just a hypothetical /
/ example, not part of a standard protocol. /
/ /
/ This is similar to Android Key Attestation. /

```

```
601({
```

```

/ nonce /          10: h'948f8860d13a463e',
/ security-level / 261: 3, / secure-restricted /
/ secure-boot /    262: true,
/ debug-status /   263: 2, / disabled-since-boot /
/ manifests /      273: [
                                h'da53574944a600683762623334383766
                                0c000169436172626f6e6974650d6331
                                2e320e0102a2181f75496e6475737472
                                69616c204175746f6d6174696f6e1821
                                02'
                                / Above is an encoded CoSWID      /
                                / with the following data          /
                                /   SW Name: "Carbonite"            /
                                /   SW Vers: "1.2"                  /
                                /   SW Creator:                      /
                                /     "Industrial Automation"        /
                                ],
/ expiration /      4: 1634324274, / 2021-10-15T18:57:54Z /
/ creation time /   6: 1634317080, / 2021-10-15T16:58:00Z /
-80000 : "fingerprint",
-80001 : { / The key -- A COSE_Key /
/ kty /          1: 2, / EC2, elliptic curve with x & y /
/ kid /          2: h'36675c206f96236c3f51f54637b94ced',
/ curve /        -1: 2, / curve is P-256 /
/ x-coord /      -2: h'65eda5a12577c2bae829437fe338701a
                    10aaa375e1bb5b5de108de439c08551d',
/ y-coord /      -3: h'1e52ed75701163f7f9e40ddf9f341b3d
                    c9ba860af7e0ca7ca7e9eecd0084d19c'
},

/ submods /        266 : {
                    "HLOS" : { / submod for high-level OS /
/ nonce /          10: h'948f8860d13a463e',
/ security-level / 261: 1, / unrestricted /
/ secure-boot /    262: true,
/ manifests /      273: [
                                h'da53574944a600687337
                                6537346b78380c000168
                                44726f6964204f530d65
                                52322e44320e0302a218
                                1f75496e647573747269
                                616c204175746f6d6174
                                696f6e182102'
                                / Above is an encoded CoSWID /
                                / with the following data:      /
                                /   SW Name: "Droid OS"         /
                                /   SW Vers: "R2.D2"            /
                                /   SW Creator:                  /

```

```
        /      "Industrial Automation"/  
      ]  
    }  
  }  
))
```

A.6. SW Measurements of an IoT Device

This is a simple token that might be for an IoT device. It includes CoSWID format measurements of the SW. The CoSWID is in byte-string wrapped in the token and also shown in diagnostic form.

```

/ This EAT UCCS is for an IoT device with a TEE. The attestation /
/ is produced by the TEE. There is a submodule for the IoT OS (the /
/ main OS of the IoT device that is not as secure as the TEE). The /
/ submodule contains claims for the IoT OS. The TEE also measures /
/ the IoT OS and puts the measurements in the submodule. /

```

```

601({
  / nonce / 10: h'948f8860d13a463e',
  / security-level / 261: 3, / secure-restricted /
  / secure-boot / 262: true,
  / debug-status / 263: 2, / disabled-since-boot /
  / OEMID / 258: h'8945ad', / IEEE CID based /
  / UEID / 256: h'0198f50a4ff6c05861c8860d13a638ea',
  / sumods / 266: {
    "OS" : {
      / security-level / 261: 2, / restricted /
      / secure-boot / 262: true,
      / debug-status / 263: 2, / disabled-since-boot /
      / swevidence / 274: [
        / This is a byte-string wrapped /
        / evidence CoSWID. It has /
        / hashes of the main files of /
        / the IoT OS. /
        h'da53574944a600663463613234350c
        17016d41636d6520522d496f542d4f
        530d65332e312e3402a2181f724163
        6d6520426173652041747465737465
        7218210103a11183a318187161636d
        655f725f696f745f6f732e65786514
        1a0044b349078201582005f6b327c1
        73b4192bd2c3ec248a292215eab456
        611bf7a783e25c1782479905a31818
        6d7265736f75726365732e72736314
        1a000c38b10782015820c142b9aba4
        280c4bb8c75f716a43c99526694caa
        be529571f5569bb7dc542f98a31818
        6a636f6d6d6f6e2e6c6962141a0023
        3d3b0782015820a6a9dcdfb3884da5
        f884e4e1e8e8629958c2dbc7027414
        43a913e34de9333be6'
      ]
    }
  }
})

```

```

/ An evidence CoSWID created for the "Acme R-IoT-OS" created by /
/ the "Acme Base Attester" (both fictitious names). It provides /

```



```

/ measurements of the SW (other than the attester SW) on the /
/ device. /

1398229316({
  / Unique CoSWID ID /      0: "4ca245",
  / tag-version /          12: 23, / Attester-maintained counter /
  / software-name /        1: "Acme R-IoT-OS",
  / software-version /     13: "3.1.4",
  / entity /               2: {
    / entity-name /        31: "Acme Base Attester",
    / role /               33: 1 / tag-creator /
  },
  / evidence /             3: {
    / ...file /            17: [
      {
        / ...fs-name /      24: "acme_r_iot_os.exe",
        / ...size /         20: 4502345,
        / ...hash /         7: [
          1, / SHA-256 /
          h'05f6b327c173b419
          2bd2c3ec248a2922
          15eab456611bf7a7
          83e25c1782479905'
        ]
      },
      {
        / ...fs-name /      24: "resources.rsc",
        / ...size /         20: 800945,
        / ...hash /         7: [
          1, / SHA-256 /
          h'c142b9aba4280c4b
          b8c75f716a43c995
          26694caabe529571
          f5569bb7dc542f98'
        ]
      },
      {
        / ...fs-name /      24: "common.lib",
        / ...size /         20: 2309435,
        / ...hash /         7: [
          1, / SHA-256 /
          h'a6a9dcdfb3884da5
          f884e4e1e8e86299
          58c2dbc702741443
          a913e34de9333be6'
        ]
      }
    ]
  }
}]

```

```
    }  
  })
```

A.7. Attestation Results in JSON format

This is a UJCS format token that might be the output of a Verifier that evaluated the IoT Attestation example immediately above.

This particular Verifier knows enough about the TEE Attester to be able to pass claims like security level directly through to the Relying Party. The Verifier also knows the Reference Values for the measured SW components and is able to check them. It informs the Relying Party that they were correct in the swresults claim. "Trustus Verifications" is the name of the services that verifies the SW component measurements.

This UJCS is identical to JSON-encoded Claims-Set that could be a JWT payload.

```
{  
  "nonce" : "lI+IYNE6Rj4=",  
  "secllevel" : "secure-restricted",  
  "secboot" : true,  
  "dbgstat" : "disabled-since-boot",  
  "OEMID" : "iUWt",  
  "UEID" : "AZjlCk/2wFhhyIYNE6Y4",  
  "submods" : {  
    "secllevel" : "restricted",  
    "secboot" : true,  
    "dbgstat" : "disabled-since-boot",  
    "swname" : "Acme R-IoT-OS",  
    "sw-version" : [  
      "3.1.4"  
    ],  
    "swresults" : [  
      [  
        "Trustus Verifications",  
        "all",  
        "fully-verified"  
      ]  
    ]  
  }  
}
```

Appendix B. UEID Design Rationale

B.1. Collision Probability

This calculation is to determine the probability of a collision of UEIDs given the total possible entity population and the number of entities in a particular entity management database.

Three different sized databases are considered. The number of devices per person roughly models non-personal devices such as traffic lights, devices in stores they shop in, facilities they work in and so on, even considering individual light bulbs. A device may have individually attested subsystems, for example parts of a car or a mobile phone. It is assumed that the largest database will have at most 10% of the world's population of devices. Note that databases that handle more than a trillion records exist today.

The trillion-record database size models an easy-to-imagine reality over the next decades. The quadrillion-record database is roughly at the limit of what is imaginable and should probably be accommodated. The 100 quadrillion database is highly speculative perhaps involving nanorobots for every person, livestock animal and domesticated bird. It is included to round out the analysis.

Note that the items counted here certainly do not have IP address and are not individually connected to the network. They may be connected to internal buses, via serial links, Bluetooth and so on. This is not the same problem as sizing IP addresses.

People	Devices / Person	Subsystems / Device	Database Portion	Database Size
10 billion	100	10	10%	trillion (10 ¹²)
10 billion	100,000	10	10%	quadrillion (10 ¹⁵)
100 billion	1,000,000	10	10%	100 quadrillion (10 ¹⁷)

This is conceptually similar to the Birthday Problem where m is the number of possible birthdays, always 365, and k is the number of people. It is also conceptually similar to the Birthday Attack where collisions of the output of hash functions are considered.

The proper formula for the collision calculation is

$$p = 1 - e^{\{-k^2/(2n)\}}$$

p Collision Probability
 n Total possible population
 k Actual population

However, for the very large values involved here, this formula requires floating point precision higher than commonly available in calculators and SW so this simple approximation is used. See [BirthdayAttack].

$$p = k^2 / 2n$$

For this calculation:

p Collision Probability
 n Total population based on number of bits in UEID
 k Population in a database

Database Size	128-bit UEID	192-bit UEID	256-bit UEID
trillion (10 ¹²)	2 * 10 ⁻¹⁵	8 * 10 ⁻³⁵	5 * 10 ⁻⁵⁵
quadrillion (10 ¹⁵)	2 * 10 ⁻⁰⁹	8 * 10 ⁻²⁹	5 * 10 ⁻⁴⁹
100 quadrillion (10 ¹⁷)	2 * 10 ⁻⁰⁵	8 * 10 ⁻²⁵	5 * 10 ⁻⁴⁵

Next, to calculate the probability of a collision occurring in one year's operation of a database, it is assumed that the database size is in a steady state and that 10% of the database changes per year. For example, a trillion record database would have 100 billion states per year. Each of those states has the above calculated probability of a collision.

This assumption is a worst-case since it assumes that each state of the database is completely independent from the previous state. In reality this is unlikely as state changes will be the addition or deletion of a few records.

The following tables gives the time interval until there is a probability of a collision based on there being one tenth the number of states per year as the number of records in the database.

$$t = 1 / ((k / 10) * p)$$

t Time until a collision
 p Collision probability for UEID size
 k Database size

Database Size	128-bit UEID	192-bit UEID	256-bit UEID
trillion (10 ¹²)	60,000 years	10 ²⁴ years	10 ⁴⁴ years
quadrillion (10 ¹⁵)	8 seconds	10 ¹⁴ years	10 ³⁴ years
100 quadrillion (10 ¹⁷)	8 microseconds	10 ¹¹ years	10 ³¹ years

Clearly, 128 bits is enough for the near future thus the requirement that UEIDs be a minimum of 128 bits.

There is no requirement for 256 bits today as quadrillion-record databases are not expected in the near future and because this time-to-collision calculation is a very worst case. A future update of the standard may increase the requirement to 256 bits, so there is a requirement that implementations be able to receive 256-bit UEIDs.

B.2. No Use of UUID

A UEID is not a UUID [RFC4122] by conscious choice for the following reasons.

UUIDs are limited to 128 bits which may not be enough for some future use cases.

Today, cryptographic-quality random numbers are available from common CPUs and hardware. This hardware was introduced between 2010 and 2015. Operating systems and cryptographic libraries give access to this hardware. Consequently, there is little need for implementations to construct such random values from multiple sources on their own.

Version 4 UUIDs do allow for use of such cryptographic-quality random numbers, but do so by mapping into the overall UUID structure of time and clock values. This structure is of no value here yet adds complexity. It also slightly reduces the number of actual bits with entropy.

UUIDs seem to have been designed for scenarios where the implementor does not have full control over the environment and uniqueness has to be constructed from identifiers at hand. UEID takes the view that

hardware, software and/or manufacturing process directly implement UEID in a simple and direct way. It takes the view that cryptographic quality random number generators are readily available as they are implemented in commonly used CPU hardware.

Appendix C. EAT Relation to IEEE.802.1AR Secure Device Identity (DevID)

This section describes several distinct ways in which an IEEE IDevID [IEEE.802.1AR] relates to EAT, particularly to UEID and SUEID.

[IEEE.802.1AR] orients around the definition of an implementation called a "DevID Module." It describes how IDevIDs and LDevIDs are stored, protected and accessed using a DevID Module. A particular level of defense against attack that should be achieved to be a DevID is defined. The intent is that IDevIDs and LDevIDs are used with an open set of network protocols for authentication and such. In these protocols the DevID secret is used to sign a nonce or similar to proof the association of the DevID certificates with the device.

By contrast, EAT defines network protocol for proving trustworthiness to a Relying Party, the very thing that is not defined in [IEEE.802.1AR]. Nor does not give details on how keys, data and such are stored protected and accessed. EAT is intended to work with a variety of different on-device implementations ranging from minimal protection of assets to the highest levels of asset protection. It does not define any particular level of defense against attack, instead providing a set of security considerations.

EAT and DevID can be viewed as complimentary when used together or as competing to provide a device identity service.

C.1. DevID Used With EAT

As just described, EAT defines a network protocol and [IEEE.802.1AR] doesn't. Vice versa, EAT doesn't define a device implementation and DevID does.

Hence, EAT can be the network protocol that a DevID is used with. The DevID secret becomes the attestation key used to sign EATs. The DevID and its certificate chain become the Endorsement sent to the Verifier.

In this case the EAT and the DevID are likely to both provide a device identifier (e.g. a serial number). In the EAT it is the UEID (or SUEID). In the DevID (used as an endorsement), it is a device serial number included in the subject field of the DevID certificate. It is probably a good idea in this use for them to be the same serial number or for the UEID to be a hash of the DevID serial number.

C.2. How EAT Provides an Equivalent Secure Device Identity

The UEID, SUEID and other claims like OEM ID are equivalent to the secure device identity put into the subject field of a DevID certificate. These EAT claims can represent all the same fields and values that can be put in a DevID certificate subject. EAT explicitly and carefully defines a variety of useful claims.

EAT secures the conveyance of these claims by having them signed on the device by the attestation key when the EAT is generated. EAT also signs the nonce that gives freshness at this time. Since these claims are signed for every EAT generated, they can include things that vary over time like GPS location.

DevID secures the device identity fields by having them signed by the manufacturer of the device sign them into a certificate. That certificate is created once during the manufacturing of the device and never changes so the fields cannot change.

So in one case the signing of the identity happens on the device and the other in a manufacturing facility, but in both cases the signing of the nonce that proves the binding to the actual device happens on the device.

While EAT does not specify how the signing keys, signature process and storage of the identity values should be secured against attack, an EAT implementation may have equal defenses against attack. One reason EAT uses CBOR is because it is simple enough that a basic EAT implementation can be constructed entirely in hardware. This allows EAT to be implemented with the strongest defenses possible.

C.3. An X.509 Format EAT

It is possible to define a way to encode EAT claims in an X.509 certificate. For example, the EAT claims might be mapped to X.509 v3 extensions. It is even possible to stuff a whole CBOR-encoded unsigned EAT token into a X.509 certificate.

If that X.509 certificate is an IDevID or LDevID, this becomes another way to use EAT and DevID together.

Note that the DevID must still be used with an authentication protocol that has a nonce or equivalent. The EAT here is not being used as the protocol to interact with the rely party.

C.4. Device Identifier Permanence

In terms of permanence, an IDevID is similar to a UEID in that they do not change over the life of the device. They cease to exist only when the device is destroyed.

An SUEID is similar to an LDevID. They change on device life-cycle events.

[IEEE.802.1AR] describes much of this permanence as resistant to attacks that seek to change the ID. IDevID permanence can be described this way because [IEEE.802.1AR] is oriented around the definition of an implementation with a particular level of defense against attack.

EAT is not defined around a particular implementation and must work on a range of devices that have a range of defenses against attack. EAT thus can't be defined permanence in terms of defense against attack. EAT's definition of permanence is in terms of operations and device lifecycle.

Appendix D. Changes from Previous Drafts

The following is a list of known changes from the previous drafts. This list is non-authoritative. It is meant to help reviewers see the significant differences.

D.1. From draft-rats-eat-01

- o Added UEID design rationale appendix

D.2. From draft-mandyam-rats-eat-00

This is a fairly large change in the orientation of the document, but no new claims have been added.

- o Separate information and data model using CDDL.
- o Say an EAT is a CWT or JWT
- o Use a map to structure the boot_state and location claims

D.3. From draft-ietf-rats-eat-01

- o Clarifications and corrections for OEMID claim
- o Minor spelling and other fixes

- o Add the nonce claim, clarify jti claim
- D.4. From draft-ietf-rats-eat-02
- o Roll all EUIs back into one UEID type
 - o UEIDs can be one of three lengths, 128, 192 and 256.
 - o Added appendix justifying UEID design and size.
 - o Submods part now includes nested eat tokens so they can be named and there can be more than one of them
 - o Lots of fixes to the CDDL
 - o Added security considerations
- D.5. From draft-ietf-rats-eat-03
- o Split boot_state into secure-boot and debug-disable claims
 - o Debug disable is an enumerated type rather than Booleans
- D.6. From draft-ietf-rats-eat-04
- o Change IMEI-based UEIDs to be encoded as a 14-byte string
 - o CDDL cleaned up some more
 - o CDDL allows for JWTs and UCCSs
 - o CWT format submodules are byte string wrapped
 - o Allows for JWT nested in CWT and vice versa
 - o Allows UCCS (unsigned CWTs) and JWT unsecured tokens
 - o Clarify tag usage when nesting tokens
 - o Add section on key inclusion
 - o Add hardware version claims
 - o Collected CDDL is now filled in. Other CDDL corrections.
 - o Rename debug-disable to debug-status; clarify that it is not extensible

- o Security level claim is not extensible
 - o Improve specification of location claim and added a location privacy section
 - o Add intended use claim
- D.7. From draft-ietf-rats-eat-05
- o CDDL format issues resolved
 - o Corrected reference to Location Privacy section
- D.8. From draft-ietf-rats-eat-06
- o Added boot-seed claim
 - o Rework CBOR interoperability section
 - o Added profiles claim and section
- D.9. From draft-ietf-rats-eat-07
- o Filled in IANA and other sections for possible preassignment of Claim Keys for well understood claims
- D.10. From draft-ietf-rats-eat-08
- o Change profile claim to be either a URL or an OID rather than a test string
- D.11. From draft-ietf-rats-eat-09
- o Add SUEIDs
 - o Add appendix comparing IDevID to EAT
 - o Added section on use for Evidence and Attestation Results
 - o Fill in the key ID and endorsements identificaiton section
 - o Remove origination claim as it is replaced by key IDs and endorsements
 - o Added manifests and software evidence claims
 - o Add string labels non-claim labels for use with JSON (e.g. labels for members of location claim)

- o EAN-13 HW versions are no longer a separate claim. Now they are folded in as a CoSWID version scheme.

D.12. From draft-ietf-rats-eat-10

- o Hardware version is made into an array of two rather than two claims
- o Corrections and wording improvements for security levels claim
- o Add swresults claim
- o Add dloas claim - Digital Letter of Approvals, a list of certifications
- o CDDL for each claim no longer in a separate sub section
- o Consistent use of terminology from RATS architecture document
- o Consistent use of terminology from CWT and JWT documents
- o Remove operating model and procedures; refer to CWT, JWT and RATS architecture instead
- o Some reorganization of Section 1
- o Moved a few references, including RATS Architecture, to informative.
- o Add detached submodule digests and detached eat bundles (DEBs)
- o New simpler and more universal scheme for identifying the encoding of a nested token
- o Made clear that CBOR and JSON are only mixed when nesting a token in another token
- o Clearly separate CDDL for JSON and CBOR-specific data items
- o Define UJCS (unsigned JWTs)
- o Add CDDL for a general Claims-Set used by UCCS, UJCS, CWT, JWT and EAT
- o Top level CDDL for CWT correctly refers to COSE
- o OEM ID is specifically for HW, not for SW

- o HW OEM ID can now be a PEN
- o HW OEM ID can now be a 128-bit random number
- o Expand the examples section
- o Add software and version claims as easy / JSON alternative to CoSWID

D.13. From draft-ietf-rats-eat-11

- o Add HW model claim
- o Change reference for CBOR OID draft to RFC 9090
- o Correct the iat claim in some examples
- o Make HW Version just one claim rather than 3 (device, board and chip)
- o Remove CDDL comments from CDDL blocks
- o More clearly define "entity" and use it more broadly, particularly instead of "device"
- o Re do early allocation of CBOR labels since last one didn't complete correctly
- o Lots of rewording and tightening up of section 1
- o Lots of wording improvements in section 3, particularly better use of normative language
- o Improve wording in submodules section, particularly how to distinguish types when decoding
- o Remove security-level from early allocation
- o Add boot odometer claim
- o Add privacy considerations for replay protection

Authors' Addresses

Laurence Lundblade
Security Theory LLC

EMail: lg1@securitytheory.com

Giridhar Mandyam
Qualcomm Technologies Inc.
5775 Morehouse Drive
San Diego, California
USA

Phone: +1 858 651 7200
EMail: mandyam@qti.qualcomm.com

Jeremy O'Donoghue
Qualcomm Technologies Inc.
279 Farnborough Road
Farnborough GU14 7LS
United Kingdom

Phone: +44 1252 363189
EMail: jodonogh@qti.qualcomm.com

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 14 January 2022

H. Birkholz
M. Eckel
Fraunhofer SIT
W. Pan
Huawei Technologies
E. Voit
Cisco
13 July 2021

Reference Interaction Models for Remote Attestation Procedures
draft-ietf-rats-reference-interaction-models-03

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 14 January 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. Disambiguation	4
3. Scope and Intent	4
4. Essential Requirements	5
4.1. Endorsement of Attesting Environments	5
5. Normative Prerequisites	6
6. Generic Information Elements	7
7. Interaction Models	9
7.1. Challenge/Response Remote Attestation	9
7.2. Uni-Directional Remote Attestation	11
7.3. Streaming Remote Attestation	13
8. Additional Application-Specific Requirements	15
8.1. Confidentiality	15
8.2. Mutual Authentication	15
8.3. Hardware-Enforcement/Support	15
9. Implementation Status	15
9.1. Implementer	16
9.2. Implementation Name	16
9.3. Implementation URL	16
9.4. Maturity	16
9.5. Coverage and Version Compatibility	16
9.6. License	17
9.7. Implementation Dependencies	17
9.8. Contact	17
10. Security and Privacy Considerations	17
11. Acknowledgments	17
12. References	17
12.1. Normative References	17
12.2. Informative References	18
Appendix A. CDDL Specification for a simple CoAP Challenge/ Response Interaction	19
Authors' Addresses	20

1. Introduction

Remote Attestation procedures (RATS, [I-D.ietf-rats-architecture]) 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture]: 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture] 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 [I-D.ietf-rats-architecture] 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. [I-D.ietf-rats-architecture] 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 Authentication Secret is uniquely associated with a distinguishable Attesting Environment. Consequently, an Authentication Secret ID also identifies 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture]). 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

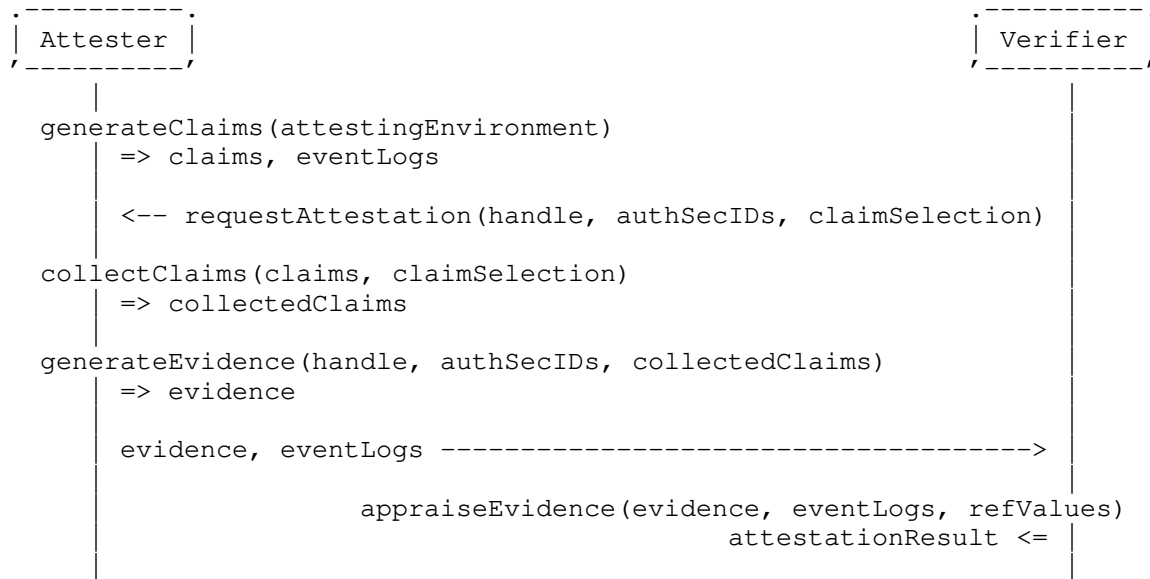
The following subsections introduce and illustrate the interaction models:

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.

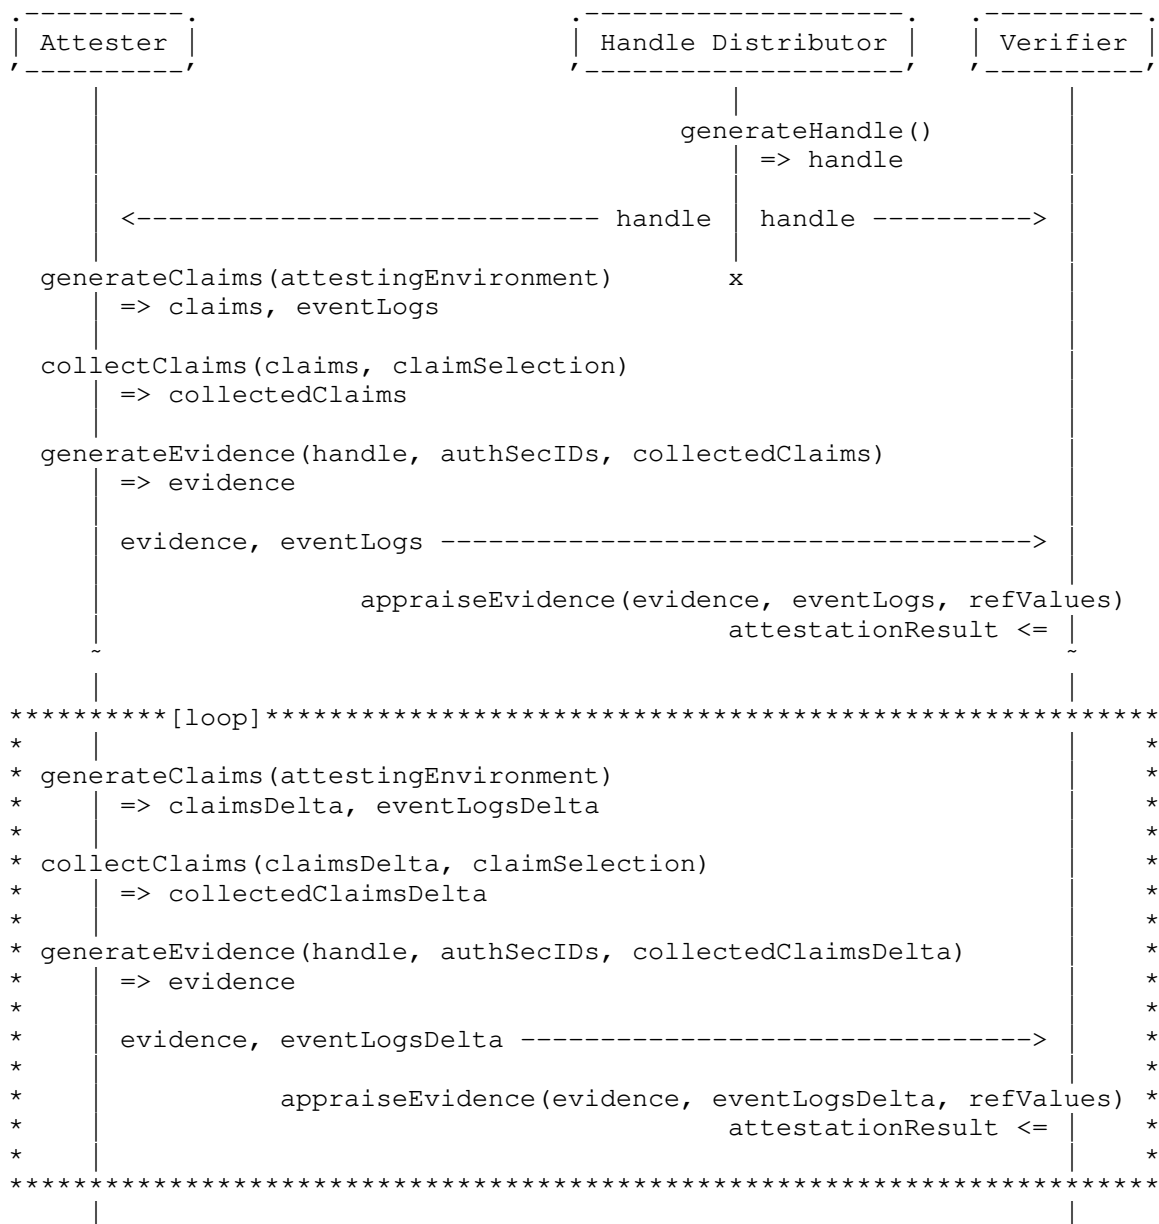
The Attester collects Claims based on the Claim Selection. With the Claim Selection the Verifier defines the set of Claims it requires. Correspondingly, collected Claims can be a subset of the produced Claims. This could be all available Claims, depending on the Claim Selection. If the Claim Selection 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.

7.2. 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.

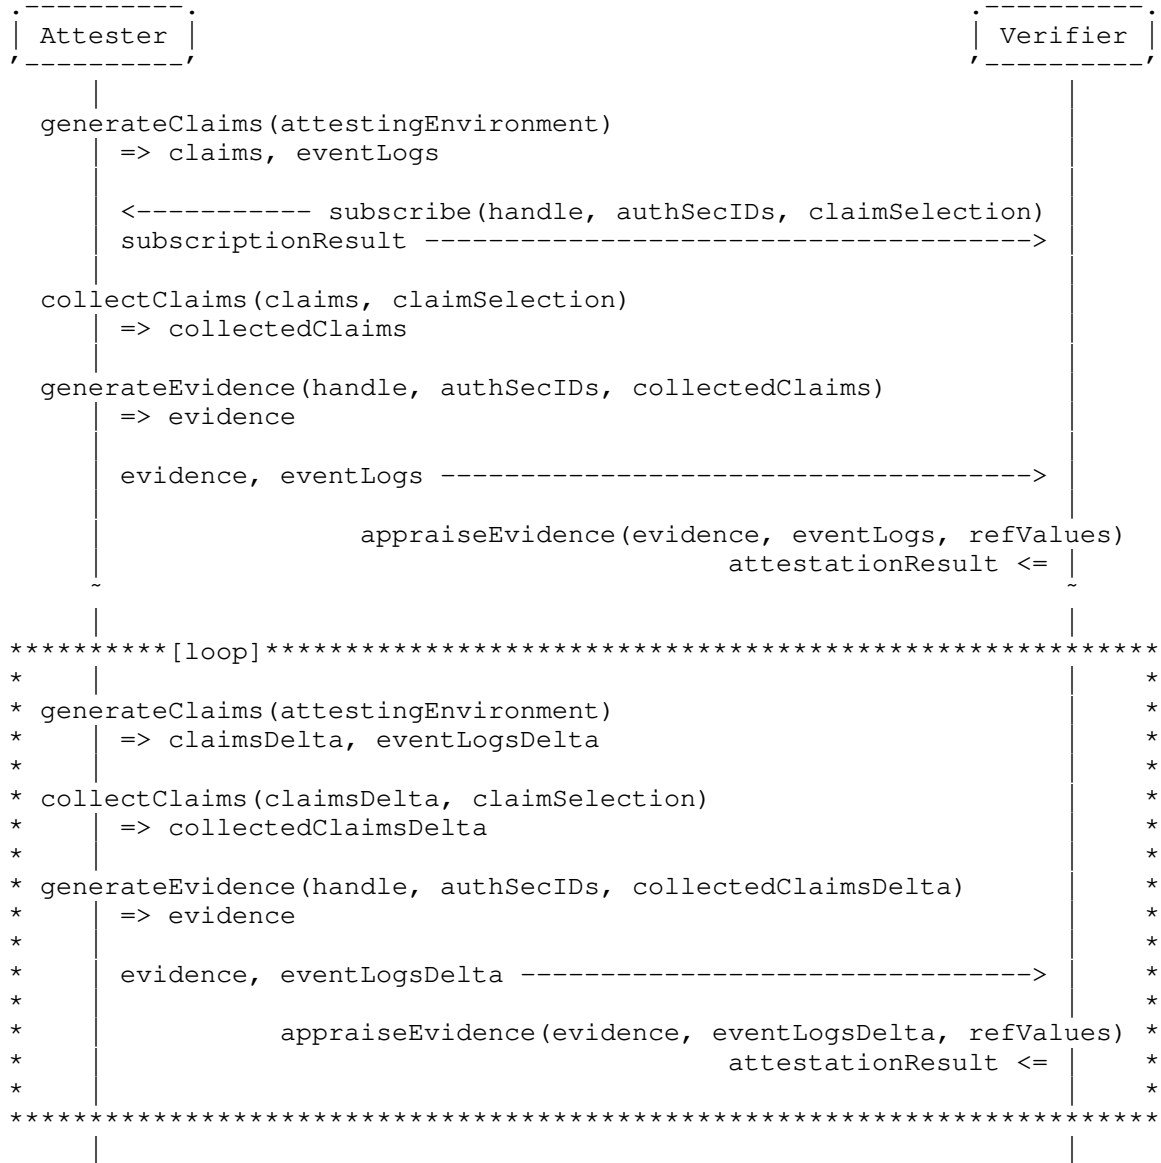
The Uni-Directional model uses the same information elements as the Challenge/Response model. In the sequence diagram above, the Attester initiates the conveyance of Evidence (comparable with a RESTful POST operation or the emission of a beacon). While a request of Evidence from the Verifier would result 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 Logs entries since the previous conveyance. This 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 procedures require the setup of 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 required Handles for producing Evidence. 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 required in this model, it is also limited

to bi-lateral subscription relationships in which each Verifier has to create and provide its individual Handle. Handles provided by a specific subscribing Verifier MUST be used in Evidence generation for that specific Verifier. The Streaming model uses the same information elements as the Challenge/Response and the Uni-Directional model. Methods to detect excessive time drift that would mandate a refreshed Handle to be conveyed via another subscribe operation are out-of-scope of this document.

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 ('1bcb469') 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 code and data is also maintained on GitHub and the project resources can be located via: <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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/rfc5280>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/rfc8610>>.

12.2. Informative References

- [DAA] Brickell, E., Camenisch, J., and L. Chen, "Direct Anonymous Attestation", page 132-145, ACM Proceedings of the 11rd ACM conference on Computer and Communications Security, 2004.
- [I-D.birkholz-rats-tuda] Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-05, 12 July 2021, <<https://www.ietf.org/archive/id/draft-birkholz-rats-tuda-05.txt>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-12.txt>>.

[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-07, 10 June 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-tpm-based-network-device-attest-07.txt>>.

[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 transfer protocol used is CoAP using the FETCH operation. The actual resource operated on can be empty. Both the Challenge 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.


```

RAIM-Bodies = CoAP-FETCH-Body / CoAP-FETCH-Response-Body

CoAP-FETCH-Body = [ hello: bool, ; if true, the AK-Cert is conveyed
                    nonce: bytes,
                    pcr-selection: [ + [ tcg-hash-alg-id: uint .size 2, ; TPM2
                                _ALG_ID
                                [ + pcr: uint .size 1 ],
                                ],
                    ],

CoAP-FETCH-Response-Body = [ attestation-evidence: TPMS_ATTEST-quote,
                             tpm-native-signature: bytes,
                             ? ak-cert: bytes, ; attestation key certificate
                             ]

TPMS_ATTEST-quote = [ qualifiediSigner: uint .size 2, ;TPM2B_NAME
                     TPMS_CLOCK_INFO,
                     firmwareVersion: uint .size 8
                     quote-responses: [ * [ pcr: uint .size 1,
                                             + [ pcr-value: bytes,
                                                ? hash-alg-id: uint .size 2,
                                                ],
                                             ],
                                             ? pcr-digest: bytes,
                                             ],
                     ]

TPMS_CLOCK_INFO = [ clock: uint .size 8,
                    resetCounter: uint .size 4,
                    restartCounter: uint .size 4,
                    save: bool,
                    ]

```

Authors' Addresses

Henk Birkholz
 Fraunhofer SIT
 Rheinstrasse 75
 64295 Darmstadt
 Germany

Email: henk.birkholz@sit.fraunhofer.de

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

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 30 July 2022

H. Birkholz
M. Eckel
Fraunhofer SIT
W. Pan
Huawei Technologies
E. Voit
Cisco
26 January 2022

Reference Interaction Models for Remote Attestation Procedures
draft-ietf-rats-reference-interaction-models-05

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 30 July 2022.

Copyright Notice

Copyright (c) 2022 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	3
2. Terminology	3
2.1. Disambiguation	4
3. Scope and Intent	4
4. Essential Requirements	5
4.1. Endorsement of Attesting Environments	5
5. Normative Prerequisites	6
6. Generic Information Elements	7
7. Interaction Models	9
7.1. Challenge/Response Remote Attestation	9
7.1.1. Models and example sequences of Challenge/Response Remote Attestation	11
7.2. Uni-Directional Remote Attestation	13
7.3. Streaming Remote Attestation	15
8. Additional Application-Specific Requirements	17
8.1. Confidentiality	17
8.2. Mutual Authentication	17
8.3. Hardware-Enforcement/Support	17
9. Implementation Status	17
9.1. Implementer	18
9.2. Implementation Name	18
9.3. Implementation URL	18
9.4. Maturity	18
9.5. Coverage and Version Compatibility	18
9.6. License	19
9.7. Implementation Dependencies	19
9.8. Contact	19
10. Security and Privacy Considerations	19
11. Acknowledgments	19
12. References	19
12.1. Normative References	19
12.2. Informative References	20
Appendix A. CDDL Specification for a simple CoAP Challenge/ Response Interaction	21
Authors' Addresses	22

1. Introduction

Remote Attestation procedures (RATS, [I-D.ietf-rats-architecture]) 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture]: 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture] 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 [I-D.ietf-rats-architecture] 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. [I-D.ietf-rats-architecture] 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 [I-D.ietf-rats-architecture]. 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 [I-D.ietf-rats-architecture]). 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

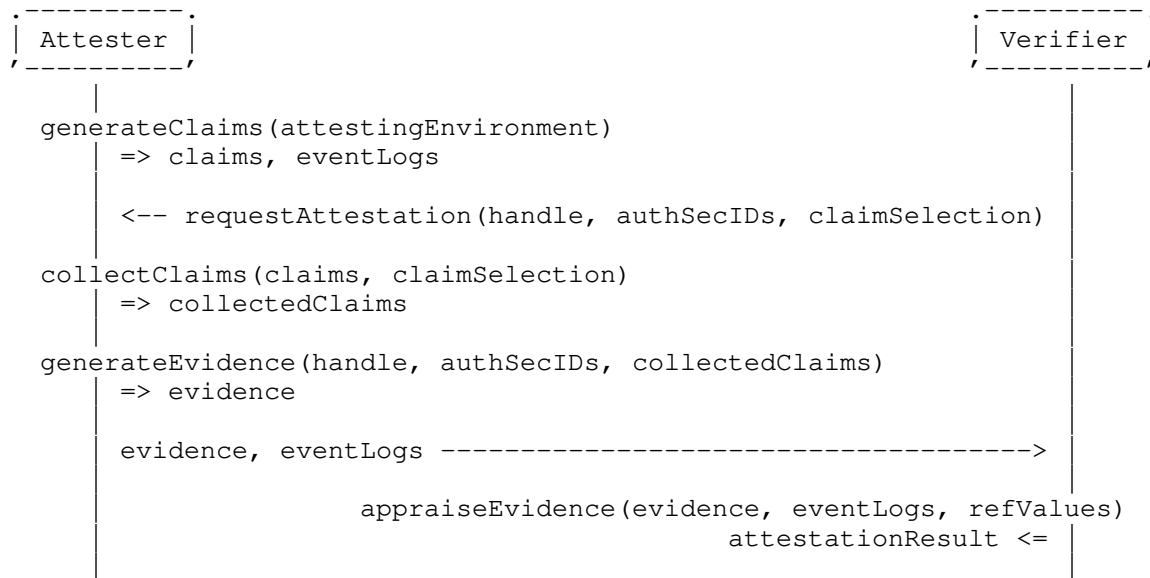
The following subsections introduce and illustrate the interaction models:

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.

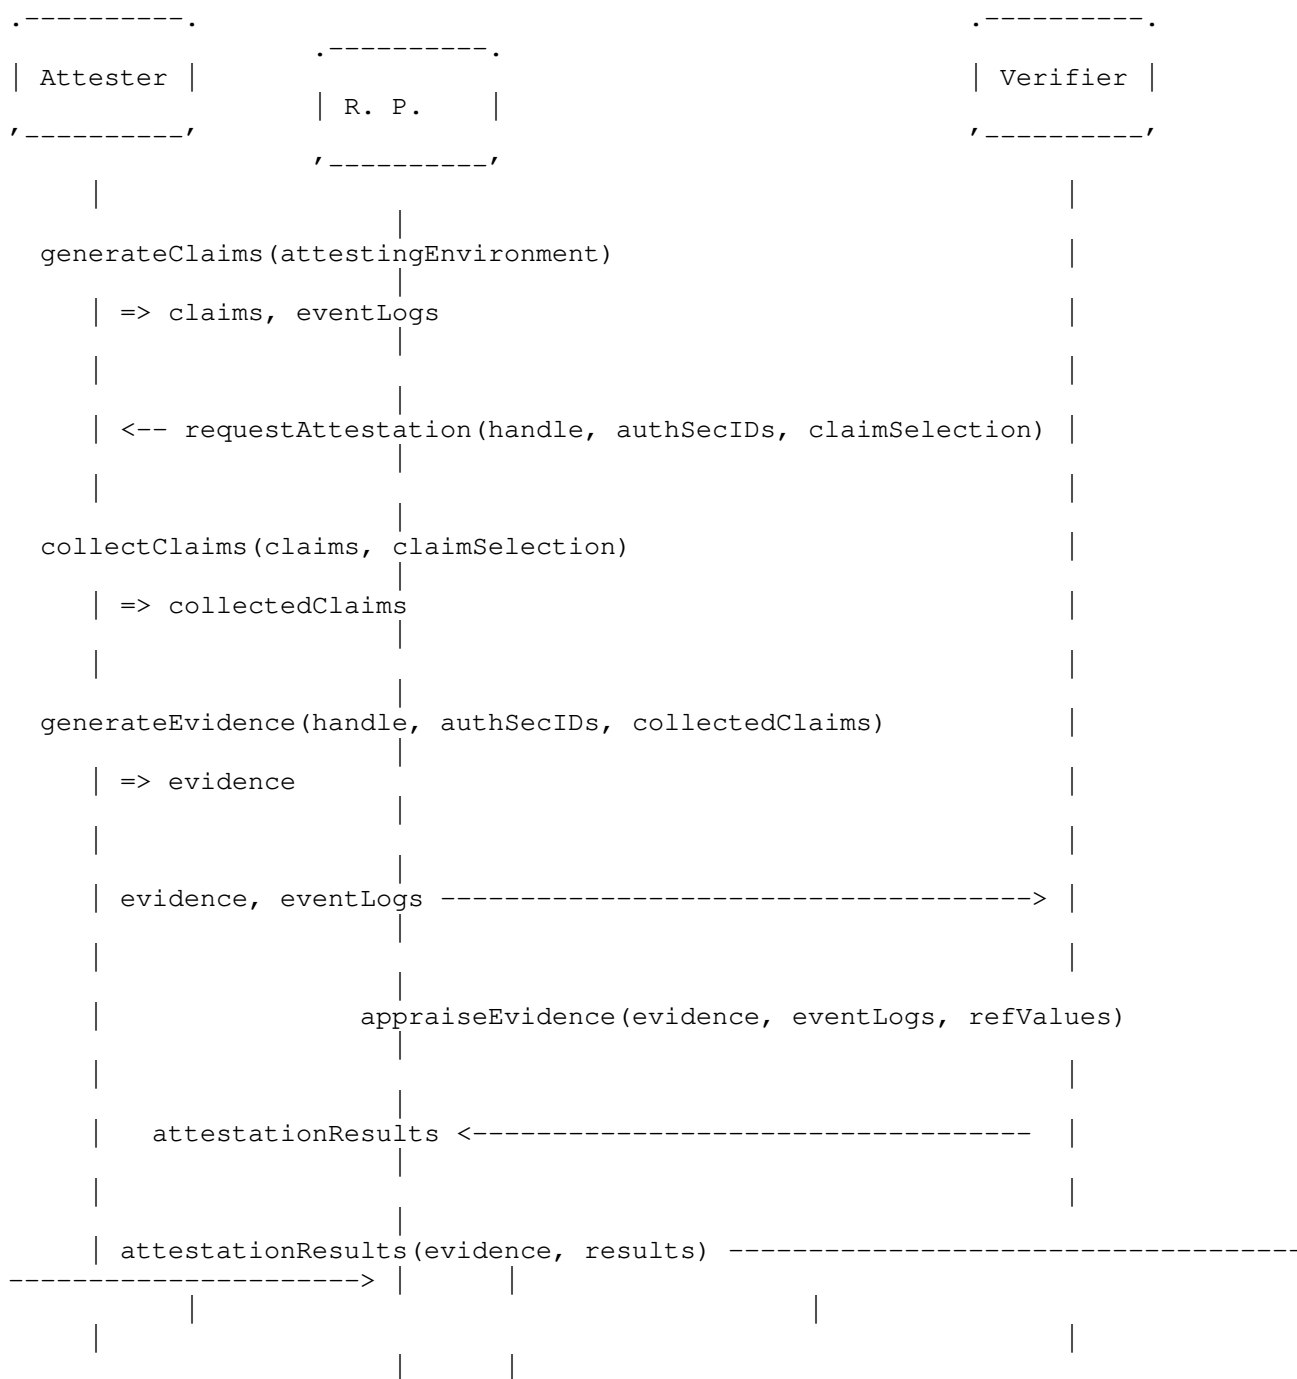
As soon as the Verifier receives the Evidence and the Event Logs, it appraises the Evidence. For this purpose, it validates the signature, the Attester Identity, and the Handle, and then appraises the Claims. Appraisal procedures are application-specific and can be conducted via comparison of the Claims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

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.

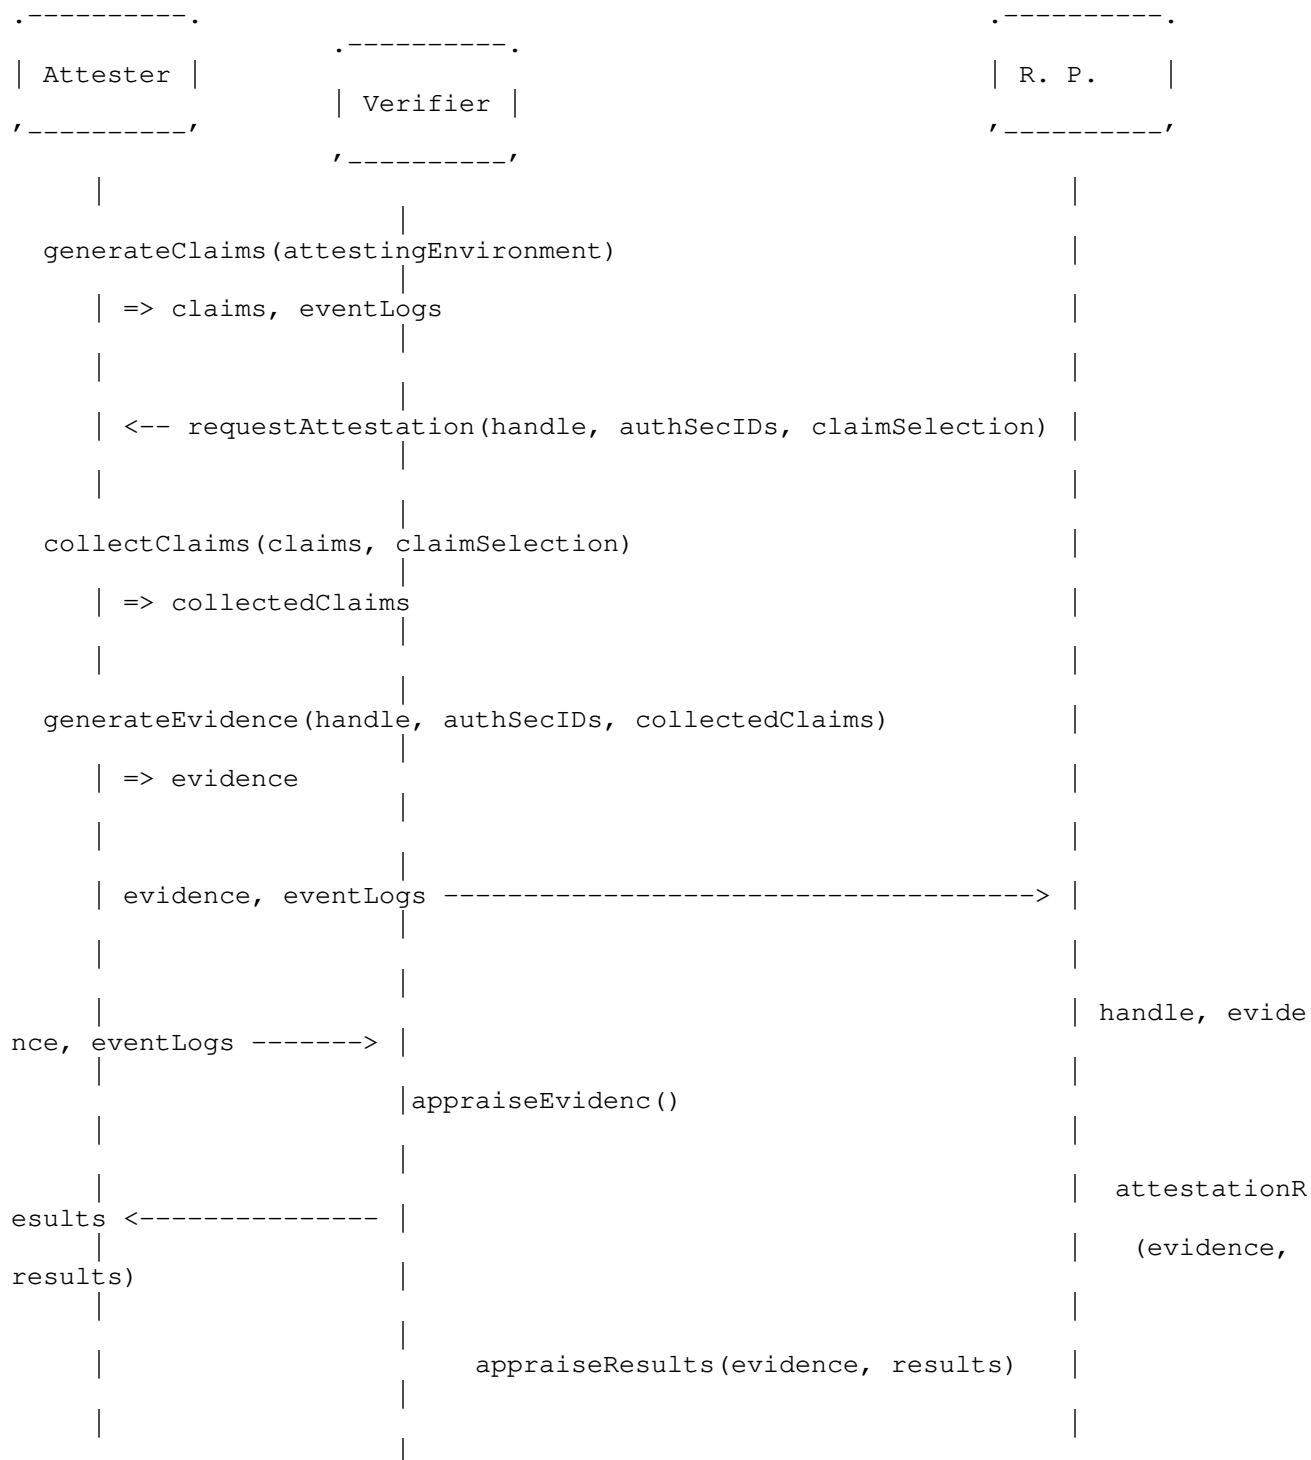


```
|
|
|                                     | appraiseResult()
|                                     |
|                                     |
```

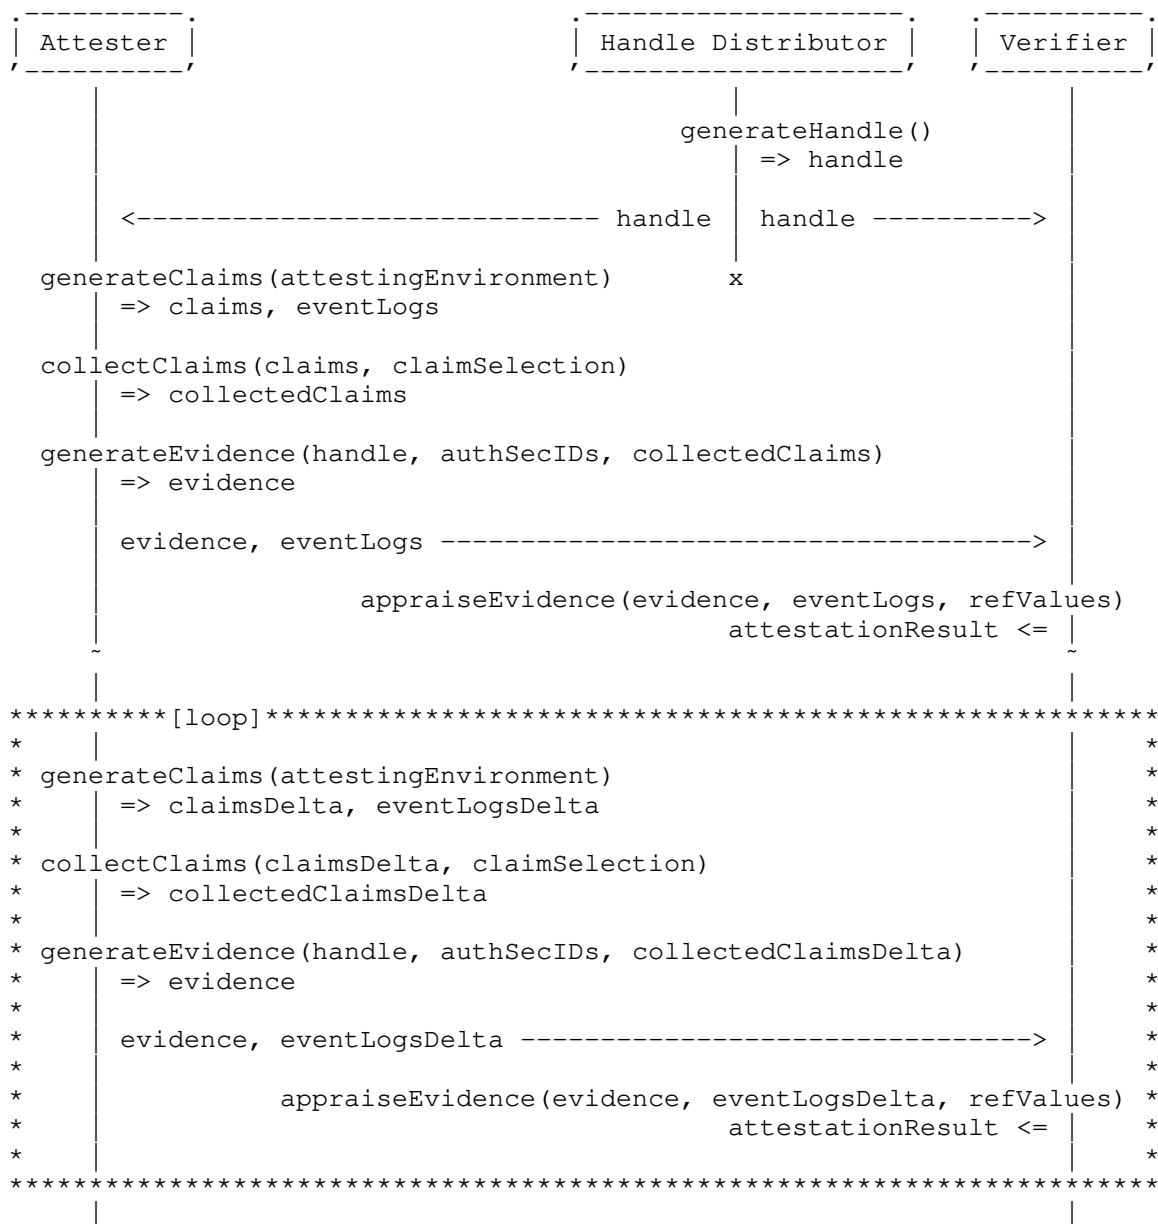
1. Background Check Model

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 realys the message and optionally check its signature against a policed trust anchor store. Upon receiving the evidence the Relying Party initiates a session with the Verifier. Once session is established, it forwards the received Evidence to the Verfier. 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.

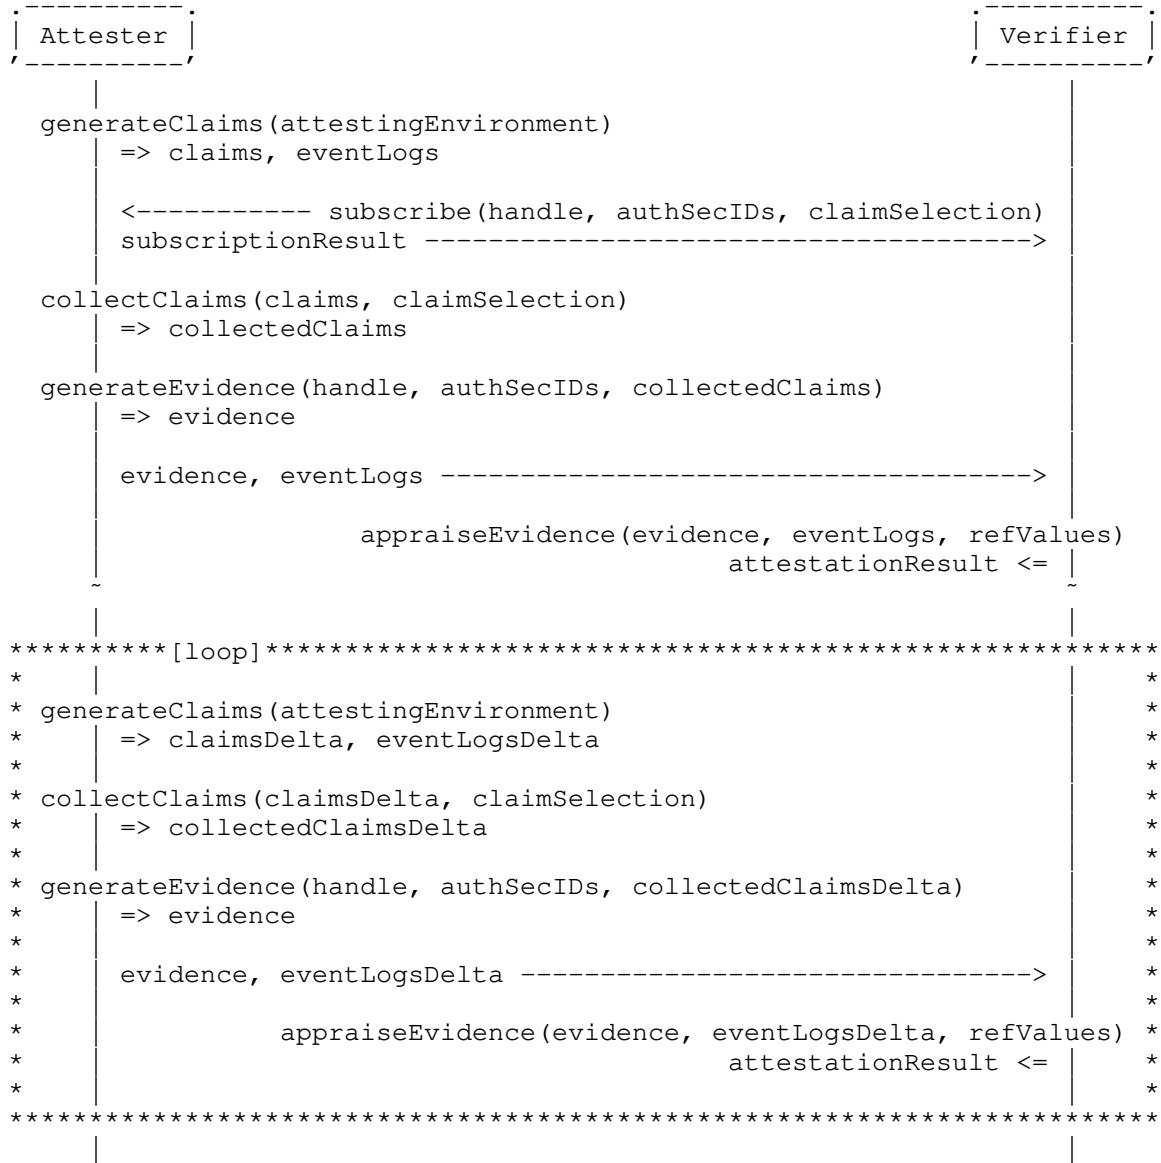
The Uni-Directional model uses the same information elements as the Challenge/Response model. In the sequence diagram above, the Attester initiates the conveyance of Evidence (comparable with a RESTful POST operation or the emission of a beacon). While a request of Evidence from the Verifier would result 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 Logs 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 procedures require the setup of 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 required Handles for producing Evidence. 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 required in this model, it is also limited

to bi-lateral subscription relationships in which each Verifier has to create and provide its individual Handle. Handles provided by a specific subscribing Verifier MUST be used in Evidence generation for that specific Verifier. The Streaming model uses the same information elements as the Challenge/Response and the Uni-Directional model. Methods to detect excessive time drift that would mandate a refreshed Handle to be conveyed via another subscribe operation are out-of-scope of this document.

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 ('1bcb469') 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 code and data is also maintained on GitHub and the project resources can be located via: <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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/rfc5280>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/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://www.rfc-editor.org/info/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://www.rfc-editor.org/info/rfc8610>>.

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.
- [I-D.birkholz-rats-tuda] Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-06, 12 January 2022, <<https://www.ietf.org/archive/id/draft-birkholz-rats-tuda-06.txt>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-14, 9 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-14.txt>>.

[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-10, 30 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-tpm-based-network-device-attest-10.txt>>.

[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 transfer protocol used is CoAP using the FETCH operation. The actual resource operated on can be empty. Both the Challenge 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.

```

RAIM-Bodies = CoAP-FETCH-Body / CoAP-FETCH-Response-Body

CoAP-FETCH-Body = [ hello: bool, ; if true, the AK-Cert is conveyed
                    nonce: bytes,
                    pcr-selection: [ + [ tcg-hash-alg-id: uint .size 2, ; TPM2
_ALG_ID
                                [ + pcr: uint .size 1 ],
                                ],
                    ],

CoAP-FETCH-Response-Body = [ attestation-evidence: TPMS_ATTEST-quote,
                             tpm-native-signature: bytes,
                             ? ak-cert: bytes, ; attestation key certificate
                             ]

TPMS_ATTEST-quote = [ qualifiediSigner: uint .size 2, ;TPM2B_NAME
                     TPMS_CLOCK_INFO,
                     firmwareVersion: uint .size 8
                     quote-responses: [ * [ pcr: uint .size 1,
                                             + [ pcr-value: bytes,
                                                ? hash-alg-id: uint .size 2,
                                                ],
                                             ],
                                             ? pcr-digest: bytes,
                                             ],
                     ]

TPMS_CLOCK_INFO = [ clock: uint .size 8,
                    resetCounter: uint .size 4,
                    restartCounter: uint .size 4,
                    save: bool,
                    ]

```

Authors' Addresses

Henk Birkholz
 Fraunhofer SIT
 Rheinstrasse 75
 64295 Darmstadt
 Germany

Email: henk.birkholz@sit.fraunhofer.de

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

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: December 12, 2021

G. Fedorkow, Ed.
Juniper Networks, Inc.
E. Voit
Cisco Systems, Inc.
J. Fitzgerald-McKay
National Security Agency
June 10, 2021

TPM-based Network Device Remote Integrity Verification
draft-ietf-rats-tpm-based-network-device-attest-07

Abstract

This document describes a workflow for remote attestation of the integrity of firmware and software installed on network devices that contain Trusted Platform Modules [TPM1.2], [TPM2.0], as defined by the Trusted Computing Group (TCG).

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 December 12, 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Document Organization	4
1.3. Goals	5
1.4. Description of Remote Integrity Verification (RIV)	5
1.5. Solution Requirements	7
1.6. Scope	8
1.6.1. Out of Scope	8
2. Solution Overview	9
2.1. RIV Software Configuration Attestation using TPM	9
2.1.1. What Does RIV Attest?	10
2.1.2. Notes on PCR Allocations	12
2.2. RIV Keying	14
2.3. RIV Information Flow	15
2.4. RIV Simplifying Assumptions	17
2.4.1. Reference Integrity Manifests (RIMs)	18
2.4.2. Attestation Logs	19
3. Standards Components	19
3.1. Prerequisites for RIV	19
3.1.1. Unique Device Identity	20
3.1.2. Keys	20
3.1.3. Appraisal Policy for Evidence	20
3.2. Reference Model for Challenge-Response	21
3.2.1. Transport and Encoding	23
3.3. Centralized vs Peer-to-Peer	24
4. Privacy Considerations	25
5. Security Considerations	26
5.1. Keys Used in RIV	26
5.2. Prevention of Spoofing and Man-in-the-Middle Attacks	28
5.3. Replay Attacks	29
5.4. Owner-Signed Keys	29
5.5. Other Factors for Trustworthy Operation	30
6. Conclusion	31
7. IANA Considerations	32
8. Acknowledgements	32
9. Appendix	32
9.1. Using a TPM for Attestation	32
9.2. Root of Trust for Measurement	33
9.3. Layering Model for Network Equipment Attester and Verifier	34
9.4. Implementation Notes	36
10. References	37
10.1. Normative References	37

10.2. Informative References	40
Authors' Addresses	43

1. Introduction

There are many aspects to consider in fielding a trusted computing device, from operating systems to applications. Mechanisms to prove that a device installed at a customer's site is authentic (i.e., not counterfeit) and has been configured with authorized software, all as part of a trusted supply chain, are just a few of the many aspects which need to be considered concurrently to have confidence that a device is truly trustworthy.

A generic architecture for remote attestation has been defined in [I-D.ietf-rats-architecture]. Additionally, the use cases for remotely attesting networking devices are discussed within Section 6 of [I-D.richardson-rats-usecases]. However, these documents do not provide sufficient guidance for network equipment vendors and operators to design, build, and deploy interoperable devices.

The intent of this document is to provide such guidance. It does this by outlining the Remote Integrity Verification (RIV) problem, and then identifies elements that are necessary to get the complete, scalable attestation procedure working with commercial networking products such as routers, switches and firewalls. An underlying assumption will be the availability within the device of a Trusted Platform Module [TPM1.2], [TPM2.0] compliant cryptoprocessor to enable the trustworthy remote assessment of the device's software and hardware.

1.1. Terminology

A number of terms are reused from [I-D.ietf-rats-architecture]. These include: Appraisal Policy for Evidence, Attestation Result, Attester, Evidence, Reference Value, Relying Party, Verifier, and Verifier Owner.

Additionally, this document defines the following term:

Attestation: the process of generating, conveying and appraising claims, backed by evidence, about device trustworthiness characteristics, including supply chain trust, identity, device provenance, software configuration, device composition, compliance to test suites, functional and assurance evaluations, etc.

The goal of attestation is simply to assure an administrator that the device configuration and software that was launched when the device was last started is authentic and untampered-with.

Within the Trusted Computing Group (TCG) context, attestation is the process by which an independent Verifier can obtain cryptographic proof as to the identity of the device in question, and evidence of the integrity of software loaded on that device when it started up, and then verify that what's there matches the intended configuration. For network equipment, a Verifier capability can be embedded in a Network Management Station (NMS), a posture collection server, or other network analytics tool (such as a software asset management solution, or a threat detection and mitigation tool, etc.). While informally referred to as attestation, this document focuses on a subset defined here as Remote Integrity Verification (RIV). RIV takes a network equipment centric perspective that includes a set of protocols and procedures for determining whether a particular device was launched with authentic software, starting from Roots of Trust. While there are many ways to accomplish attestation, RIV sets out a specific set of protocols and tools that work in environments commonly found in network equipment. RIV does not cover other device characteristics that could be attested (e.g., geographic location, connectivity; see [I-D.richardson-rats-usecases]), although it does provide evidence of a secure infrastructure to increase the level of trust in other device characteristics attested by other means (e.g., by Entity Attestation Tokens [I-D.ietf-rats-eat]).

In line with [I-D.ietf-rats-architecture] definitions, this document uses the term Endorser to refer to the role that signs identity and attestation certificates used by the Attester, while Reference Values are signed by a Reference Value Provider. Typically, the manufacturer of an embedded device would be accepted as both the Endorser and Reference Value Provider, although the choice is ultimately up to the Verifier Owner.

1.2. Document Organization

The remainder of this document is organized into several sections:

- o The remainder of this section covers goals and requirements, plus a top-level description of RIV.
- o The Solution Overview section outlines how Remote Integrity Verification works.
- o The Standards Components section links components of RIV to normative standards.
- o Privacy and Security shows how specific features of RIV contribute to the trustworthiness of the Attestation Result.
- o Supporting material is in an appendix at the end.

1.3. Goals

Network operators benefit from a trustworthy attestation mechanism that provides assurance that their network comprises authentic equipment, and has loaded software free of known vulnerabilities and unauthorized tampering. In line with the overall goal of assuring integrity, attestation can be used to assist in asset management, vulnerability and compliance assessment, plus configuration management.

The RIV attestation workflow outlined in this document is intended to meet the following high-level goals:

- o Provable Device Identity - This specification requires that an Attester (i.e., the attesting device) includes a cryptographic identifier unique to each device. Effectively this means that the TPM must be so provisioned during the manufacturing cycle.
- o Software Inventory - A key goal is to identify the software release(s) installed on the Attester, and to provide evidence that the software stored within hasn't been altered without authorization.
- o Verifiability - Verification of software and configuration of the device shows that the software that was authorized for installation by the administrator has actually been launched.

In addition, RIV is designed to operate either in a centralized environment, such as with a central authority that manages and configures a number of network devices, or 'peer-to-peer', where network devices independently verify one another to establish a trust relationship. (See Section 3.3 below, and also [I-D.voit-rats-trusted-path-routing])

1.4. Description of Remote Integrity Verification (RIV)

Attestation requires two interlocking mechanisms between the Attester network device and the Verifier:

- o Device Identity, the mechanism providing trusted identity, can reassure network managers that the specific devices they ordered from authorized manufacturers for attachment to their network are those that were installed, and that they continue to be present in their network. As part of the mechanism for Device Identity, cryptographic proof of the identity of the manufacturer is also provided.

- o Software Measurement is the mechanism that reports the state of mutable software components on the device, and can assure network managers that they have known, authentic software configured to run in their network.

Using these two interlocking mechanisms, RIV is a component in a chain of procedures that can assure a network operator that the equipment in their network can be reliably identified, and that authentic software of a known version is installed on each device. Equipment in the network includes devices that make up the network itself, such as routers, switches and firewalls.

Software used to boot a device can be described as recording a chain of measurements, anchored at the start by a Root of Trust for Measurement (see Section 9.2), each measuring the next stage, that normally ends when the system software is loaded. A measurement signifies the identity, integrity and version of each software component registered with an Attester's TPM [TPM1.2], [TPM2.0], so that a subsequent verification stage can determine if the software installed is authentic, up-to-date, and free of tampering.

RIV includes several major processes:

1. Generation of Evidence is the process whereby an Attester generates cryptographic proof (Evidence) of claims about device properties. In particular, the device identity and its software configuration are both of critical importance.
2. Device Identification refers to the mechanism assuring the Relying Party (ultimately, a network administrator) of the identity of devices that make up their network, and that their manufacturers are known.
3. Conveyance of Evidence reliably transports the collected Evidence from Attester to a Verifier to allow a management station to perform a meaningful appraisal in Step 4. The transport is typically carried out via a management network. The channel must provide integrity and authenticity, and, in some use cases, may also require confidentiality.
4. Finally, Appraisal of Evidence occurs. This is the process of verifying the Evidence received by a Verifier from the Attester, and using an Appraisal Policy to develop an Attestation Result, used to inform decision making. In practice, this means comparing the Attester's measurements reported as Evidence with the device configuration expected by the Verifier. Subsequently the Appraisal Policy for Evidence might match Evidence found

against Reference Values (aka Golden Measurements), which represent the intended configured state of the connected device.

All implementations supporting this RIV specification require the support of the following three technologies:

1. Identity: Device identity MUST be based on IEEE 802.1AR Device Identity (DevID) [IEEE-802-1AR], coupled with careful supply-chain management by the manufacturer. The Initial DevID (IDevID) certificate contains a statement by the manufacturer that establishes the identity of the device as it left the factory. Some applications with a more-complex post-manufacture supply chain (e.g., Value Added Resellers), or with different privacy concerns, may want to use alternative mechanisms for platform authentication (for example, TCG Platform Certificates [Platform-Certificates], or post-manufacture installation of Local Device ID (LDevID)).
2. Platform Attestation provides evidence of configuration of software elements present in the device. This form of attestation can be implemented with TPM Platform Configuration Registers (PCRs), Quote and Log mechanisms, which provide cryptographically authenticated evidence to report what software was started on the device through the boot cycle. Successful attestation requires an unbroken chain from a boot-time root of trust through all layers of software needed to bring the device to an operational state, in which each stage measures components of the next stage, updates the attestation log, and extends hashes into a PCR. The TPM can then report the hashes of all the measured hashes as signed evidence called a Quote (see Section 9.1 for an overview of TPM operation, or [TPM1.2] and [TPM2.0] for many more details).
3. Signed Reference Values (aka Reference Integrity Measurements) must be conveyed from the Reference Value Provider (the entity accepted as the software authority, often the manufacturer for embedded systems) to the Verifier.

1.5. Solution Requirements

Remote Integrity Verification must address the "Lying Endpoint" problem, in which malicious software on an endpoint may subvert the intended function, and also prevent the endpoint from reporting its compromised status. (See Section 5 for further Security Considerations.)

RIV attestation is designed to be simple to deploy at scale. RIV should work "out of the box" as far as possible, that is, with the

fewest possible provisioning steps or configuration databases needed at the end-user's site. Network equipment is often required to "self-configure", to reliably reach out without manual intervention to prove its identity and operating posture, then download its own configuration, a process which precludes pre-installation configuration. See [RFC8572] for an example of Secure Zero Touch Provisioning.

1.6. Scope

The need for assurance of software integrity, addressed by Remote Attestation, is a very general problem that could apply to most network-connected computing devices. However, this document includes several assumptions that limit the scope to network equipment (e.g., routers, switches and firewalls):

- o This solution is for use in non-privacy-preserving applications (for example, networking, Industrial IoT), avoiding the need for a Privacy Certificate Authority for attestation keys [AK-Enrollment] or TCG Platform Certificates [Platform-Certificates].
- o This document assumes network protocols that are common in network equipment such as YANG [RFC7950] and NETCONF [RFC6241], but not generally used in other applications.
- o The approach outlined in this document mandates the use of a compliant TPM [TPM1.2], [TPM2.0].

1.6.1. Out of Scope

- o Run-Time Attestation: The Linux Integrity Measurement Architecture attests each process launched after a device is started (and is in scope for RIV), but continuous run-time attestation of Linux or other multi-threaded operating system processes after they've started considerably expands the scope of the problem. Many researchers are working on that problem, but this document defers the problem of continuous, in-memory run-time attestation.
- o Multi-Vendor Embedded Systems: Additional coordination would be needed for devices that themselves comprise hardware and software from multiple vendors, integrated by the end user. Although out of scope for this document, these issues are accommodated in [I-D.ietf-rats-architecture].
- o Processor Sleep Modes: Network equipment typically does not "sleep", so sleep and hibernate modes are not considered. Although out of scope for RIV, Trusted Computing Group specifications do encompass sleep and hibernate states.

- o Virtualization and Containerization: In a non-virtualized system, the host OS is responsible for measuring each User Space file or process, but that's the end of the boot process. For virtualized systems, the host OS must verify the hypervisor, which then manages its own chain of trust through the virtual machine. Virtualization and containerization technologies are increasingly used in network equipment, but are not considered in this document.

2. Solution Overview

2.1. RIV Software Configuration Attestation using TPM

RIV Attestation is a process which can be used to determine the identity of software running on a specifically-identified device. Remote Attestation is broken into two phases, shown in Figure 1:

- o During system startup, each distinct software object is "measured" by the Attester. The object's identity, hash (i.e., cryptographic digest) and version information are recorded in a log. Hashes are also extended, or cryptographically folded, into the TPM, in a way that can be used to validate the log entries. The measurement process generally follows the layered chain-of-trust model used in Measured Boot, where each stage of the system measures the next one, and extends its measurement into the TPM, before launching it. See [I-D.ietf-rats-architecture], section "Layered Attestation Environments," for an architectural definition of this model.
- o Once the device is running and has operational network connectivity, a separate, trusted Verifier will interrogate the network device to retrieve the logs and a copy of the digests collected by hashing each software object, signed by an attestation private key secured by, but never released by, the TPM. The YANG model described in [I-D.ietf-rats-yang-tpm-charra] facilitates this operation.

The result is that the Verifier can verify the device's identity by checking the Subject Field and signature of certificate containing the TPM's attestation public key, and can validate the software that was launched by verifying the correctness of the logs by comparing with the signed digests from the TPM, and comparing digests in the log with Reference Values.

It should be noted that attestation and identity are inextricably linked; signed Evidence that a particular version of software was loaded is of little value without cryptographic proof of the identity of the Attester producing the Evidence.

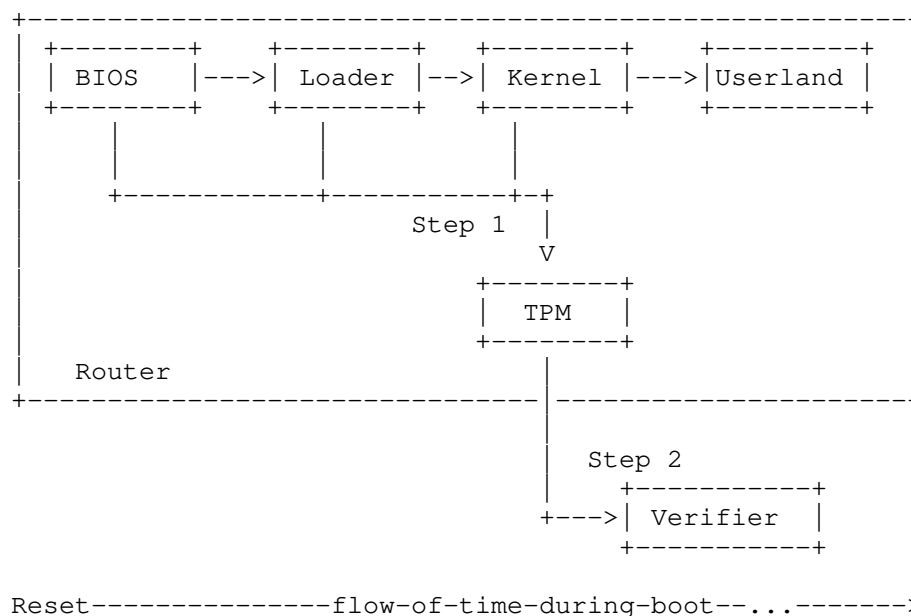


Figure 1: Layered RIV Attestation Model

In Step 1, measurements are "extended", or hashed, into the TPM as processes start, with the result that the PCR ends up containing a hash of all the measured hashes. In Step 2, signed PCR digests are retrieved from the TPM for off-box analysis after the system is operational.

2.1.1. What Does RIV Attest?

TPM attestation is strongly focused on Platform Configuration Registers (PCRs), but those registers are only vehicles for certifying accompanying Evidence, conveyed in log entries. It is the hashes in log entries that are extended into PCRs, where the final PCR values can be retrieved in the form of a structure called a Quote, signed by an Attestation key known only to the TPM. The use of multiple PCRs serves only to provide some independence between different classes of object, so that one class of objects can be updated without changing the extended hash for other classes. Although PCRs can be used for any purpose, this section outlines the objects within the scope of this document which may be extended into the TPM.

In general, assignment of measurements to PCRs is a policy choice made by the device manufacturer, selected to independently attest three classes of object:

- o Code, (i.e., instructions) to be executed by a CPU.
- o Configuration - Many devices offer numerous options controlled by non-volatile configuration variables which can impact the device's security posture. These settings may have vendor defaults, but often can be changed by administrators, who may want to verify via attestation that the operational state of the settings match their intended state.
- o Credentials - Administrators may wish to verify via attestation that public keys (and other credentials) outside the Root of Trust have not been subject to unauthorized tampering. (By definition, keys protecting the root of trust can't be verified independently.)

The TCG PC Client Platform Firmware Profile Specification [PC-Client-BIOS-TPM-2.0] gives considerable detail on what is to be measured during the boot phase of platform startup using a UEFI BIOS (www.uefi.org), but the goal is simply to measure every bit of code executed in the process of starting the device, along with any configuration information related to security posture, leaving no gap for unmeasured code to remain undetected, potentially subverting the chain.

For devices using a UEFI BIOS, [PC-Client-BIOS-TPM-2.0] gives detailed normative requirements for PCR usage. For other platform architectures, the table in Figure 2 gives non-normative guidance for PCR assignment that generalizes the specific details of [PC-Client-BIOS-TPM-2.0].

By convention, most PCRs are assigned in pairs, which the even-numbered PCR used to measure executable code, and the odd-numbered PCR used to measure whatever data and configuration are associated with that code. It is important to note that each PCR may contain results from dozens (or even thousands) of individual measurements.

Function	Assigned PCR #	
	Code	Configuration
Firmware Static Root of Trust, (i.e., initial boot firmware and drivers)	0	1
Drivers and initialization for optional or add-in devices	2	3
OS Loader code and configuration, (i.e., the code launched by firmware) to load an operating system kernel. These PCRs record each boot attempt, and an identifier for where the loader was found	4	5
Vendor Specific Measurements during boot	6	6
Secure Boot Policy. This PCR records keys and configuration used to validate the OS loader		7
Measurements made by the OS Loader (e.g GRUB2 for Linux)	8	9
Measurements made by OS (e.g., Linux IMA)	10	10

Figure 2: Attested Objects

2.1.2. Notes on PCR Allocations

It is important to recognize that PCR[0] is critical. The first measurement into PCR[0] is taken by the Root of Trust for Measurement, code which, by definition, cannot be verified by measurement. This measurement establishes the chain of trust for all subsequent measurements. If the PCR[0] measurement cannot be trusted, the validity of the entire chain is put into question.

Distinctions Between PCR[0], PCR[2], PCR[4] and PCR[8] are summarized below:

- o PCR[0] typically represents a consistent view of rarely-changed Host Platform boot components, allowing Attestation policies to be defined using the less changeable components of the transitive trust chain. This PCR typically provides a consistent view of the platform regardless of user selected options.

- o PCR[2] is intended to represent a "user configurable" environment where the user has the ability to alter the components that are measured into PCR[2]. This is typically done by adding adapter cards, etc., into user-accessible PCI or other slots. In UEFI systems these devices may be configured by Option ROMs measured into PCR[2] and executed by the BIOS.
- o PCR[4] is intended to represent the software that manages the transition between the platform's Pre-Operating System start and the state of a system with the Operating System present. This PCR, along with PCR[5], identifies the initial operating system loader (e.g., GRUB for Linux).
- o PCR[8] is used by the OS loader (e.g. GRUB) to record measurements of the various components of the operating system.

Although the TCG PC Client document specifies the use of the first eight PCRs very carefully to ensure interoperability among multiple UEFI BIOS vendors, it should be noted that embedded software vendors may have considerably more flexibility. Verifiers typically need to know which log entries are consequential and which are not (possibly controlled by local policies) but the Verifier may not need to know what each log entry means or why it was assigned to a particular PCR. Designers must recognize that some PCRs may cover log entries that a particular Verifier considers critical and other log entries that are not considered important, so differing PCR values may not on their own constitute a check for authenticity. For example, in a UEFI system, some administrators may consider booting an image from a removable drive, something recorded in a PCR, to be a security violation, while others might consider that operation an authorized recovery procedure.

Designers may allocate particular events to specific PCRs in order to achieve a particular objective with local attestation, (e.g., allowing a procedure to execute, or releasing a particular decryption key, only if a given PCR is in a given state). It may also be important to designers to consider whether streaming notification of PCR updates is required (see [I-D.birkholz-rats-network-device-subscription]). Specific log entries can only be validated if the Verifier receives every log entry affecting the relevant PCR, so (for example) a designer might want to separate rare, high-value events such as configuration changes, from high-volume, routine measurements such as IMA [IMA] logs.

2.2. RIV Keying

RIV attestation relies on two credentials:

- o An identity key pair and matching certificate is required to certify the identity of the Attester itself. RIV specifies the use of an IEEE 802.1AR Device Identity (DevID) [IEEE-802-1AR], signed by the device manufacturer, containing the device serial number.
- o An Attestation key pair and matching certificate is required to sign the Quote generated by the TPM to report evidence of software configuration.

In TPM application, both the Attestation private key and the DevID private key MUST be protected by the TPM. Depending on other TPM configuration procedures, the two keys are likely be different; some of the considerations are outlined in TCG "TPM 2.0 Keys for Device Identity and Attestation" [Platform-DevID-TPM-2.0].

The TCG TPM 2.0 Keys document [Platform-DevID-TPM-2.0] specifies further conventions for these keys:

- o When separate Identity and Attestation keys are used, the Attestation Key (AK) and its X.509 certificate should parallel the DevID, with the same device ID information as the DevID certificate (that is, the same Subject Name and Subject Alt Name, even though the key pairs are different). This allows a quote from the device, signed by an AK, to be linked directly to the device that provided it, by examining the corresponding AK certificate. If the Subject and Subject Alt Names in the AK certificate don't match the corresponding DevID certificate, or they're signed by differing authorities the Verifier may signal the detection of an Asokan-style man-in-the-middle attack (see Section 5.2).
- o Network devices that are expected to use secure zero touch provisioning as specified in [RFC8572]) MUST be shipped by the manufacturer with pre-provisioned keys (Initial DevID and Initial AK, called IDevID and IAK). IDevID and IAK certificates MUST both be signed by the Endorser (typically the device manufacturer). Inclusion of an IDevID and IAK by a vendor does not preclude a mechanism whereby an administrator can define Local Identity and Attestation Keys (LDevID and LAK) if desired.

2.3. RIV Information Flow

RIV workflow for network equipment is organized around a simple use case where a network operator wishes to verify the integrity of software installed in specific, fielded devices. This use case implies several components:

1. The Attester, the device which the network operator wants to examine.
2. A Verifier (which might be a network management station) somewhere separate from the Device that will retrieve the signed evidence and measurement logs, and analyze them to pass judgment on the security posture of the device.
3. A Relying Party, which can act on Attestation Results. Interaction between the Relying Party and the Verifier is considered out of scope for RIV.
4. Signed Reference Integrity Manifests (RIMs), containing Reference Values, can either be created by the device manufacturer and shipped along with the device as part of its software image, or alternatively, could be obtained several other ways (direct to the Verifier from the manufacturer, from a third party, from the owner's observation of what's thought to be a "known good system", etc.). Retrieving RIMs from the device itself allows attestation to be done in systems that may not have access to the public internet, or by other devices that are not management stations per se (e.g., a peer device; see Section 3.1.3). If Reference Values are obtained from multiple sources, the Verifier may need to evaluate the relative level of trust to be placed in each source in case of a discrepancy.

These components are illustrated in Figure 3.

A more-detailed taxonomy of terms is given in [I-D.ietf-rats-architecture]

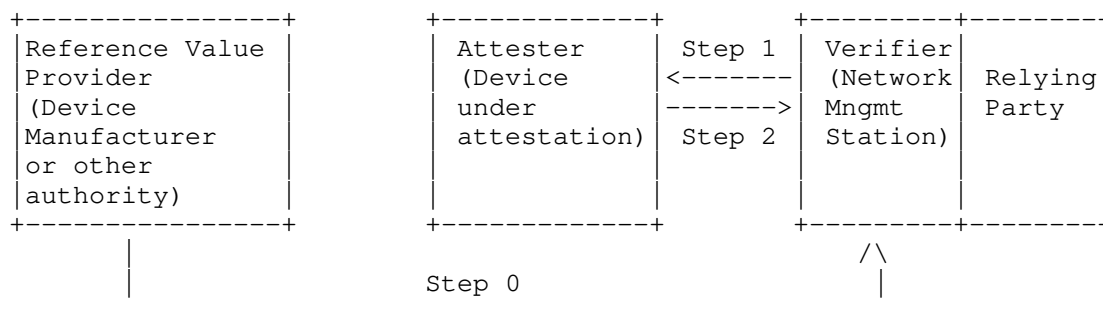


Figure 3: RIV Reference Configuration for Network Equipment

- o In Step 0, The Reference Value Provider (the device manufacturer or other authority) makes one or more Reference Integrity Manifests (RIMs), corresponding to the software image expected to be found on the device, signed by the Reference Value Provider, available to the Verifier (see Section 3.1.3 for "in-band" and "out of band" ways to make this happen).
- o In Step 1, the Verifier (Network Management Station), on behalf of a Relying Party, requests Identity, Measurement Values, and possibly RIMs, from the Attester.
- o In Step 2, the Attester responds to the request by providing a DevID, quotes (measured values, signed by the Attester), and optionally RIMs.

To achieve interoperability, the following standards components SHOULD be used:

1. TPM Keys MUST be configured according to [Platform-DevID-TPM-2.0], or [Platform-ID-TPM-1.2].
2. For devices using UEFI and Linux, measurements of firmware and bootable modules SHOULD be taken according to TCG PC Client [PC-Client-BIOS-TPM-1.2] or [PC-Client-BIOS-TPM-2.0], and Linux IMA [IMA]
3. Device Identity MUST be managed as specified in IEEE 802.1AR Device Identity certificates [IEEE-802-1AR], with keys protected by TPMs.
4. Attestation logs SHOULD be formatted according to the Canonical Event Log format [Canonical-Event-Log], although other

specialized formats may be used. UEFI-based systems MAY use the TCG UEFI BIOS event log [EFI-TPM].

5. Quotes are retrieved from the TPM according to TCG TAP Information Model [TAP]. While the TAP IM gives a protocol-independent description of the data elements involved, it's important to note that quotes from the TPM are signed inside the TPM, so MUST be retrieved in a way that does not invalidate the signature, as specified in [I-D.ietf-rats-yang-tpm-charra], to preserve the trust model. (See Section 5 Security Considerations).
6. Reference Values SHOULD be encoded as SWID or CoSWID tags, as defined in the TCG RIM document [RIM], compatible with NIST IR 8060 [NIST-IR-8060] and the IETF CoSWID draft [I-D.ietf-sacm-coswid]. See Section 2.4.1.

2.4. RIV Simplifying Assumptions

This document makes the following simplifying assumptions to reduce complexity:

- o The product to be attested MUST be shipped with an IEEE 802.1AR Device Identity and an Initial Attestation Key (IAK) with certificate. The IAK certificate MUST contain the same identity information as the DevID (specifically, the same Subject Name and Subject Alt Name, signed by the manufacturer), but it's a type of key that can be used to sign a TPM Quote, but not other objects (i.e., it's marked as a TCG "Restricted" key; this convention is described in "TPM 2.0 Keys for Device Identity and Attestation" [Platform-DevID-TPM-2.0]). For network equipment, which is generally non-privacy-sensitive, shipping a device with both an IDevID and an IAK already provisioned substantially simplifies initial startup. Privacy-sensitive applications may use the TCG Platform Certificate or other procedures to install identity credentials into the device after manufacture.
- o The product MUST be equipped with a Root of Trust for Measurement (RTM), Root of Trust for Storage and Root of Trust for Reporting (as defined in [SP800-155]) that are capable of conforming to TCG Trusted Attestation Protocol (TAP) Information Model [TAP].
- o The authorized software supplier MUST make available Reference Values in the form of signed SWID or CoSWID tags [I-D.ietf-sacm-coswid], [SWID], as described in TCG Reference Integrity Measurement Manifest Information Model [RIM].

2.4.1. Reference Integrity Manifests (RIMs)

[I-D.ietf-rats-yang-tpm-charra] focuses on collecting and transmitting evidence in the form of PCR measurements and attestation logs. But the critical part of the process is enabling the Verifier to decide whether the measurements are "the right ones" or not.

While it must be up to network administrators to decide what they want on their networks, the software supplier should supply the Reference Values, in signed Reference Integrity Manifests, that may be used by a Verifier to determine if evidence shows known good, known bad or unknown software configurations.

In general, there are two kinds of reference measurements:

1. Measurements of early system startup (e.g., BIOS, boot loader, OS kernel) are essentially single-threaded, and executed exactly once, in a known sequence, before any results could be reported. In this case, while the method for computing the hash and extending relevant PCRs may be complicated, the net result is that the software (more likely, firmware) vendor will have one known good PCR value that "should" be present in the relevant PCRs after the box has booted. In this case, the signed reference measurement could simply list the expected hashes for the given version. However, a RIM that contains the intermediate hashes can be useful in debugging cases where the expected final hash is not the one reported.
2. Measurements taken later in operation of the system, once an OS has started (for example, Linux IMA [IMA]), may be more complex, with unpredictable "final" PCR values. In this case, the Verifier must have enough information to reconstruct the expected PCR values from logs and signed reference measurements from a trusted authority.

In both cases, the expected values can be expressed as signed SWID or CoSWID tags, but the SWID structure in the second case is somewhat more complex, as reconstruction of the extended hash in a PCR may involve thousands of files and other objects.

TCG has published an information model defining elements of Reference Integrity Manifests under the title TCG Reference Integrity Manifest Information Model [RIM]. This information model outlines how SWID tags should be structured to allow attestation, and defines "bundles" of SWID tags that may be needed to describe a complete software release. The RIM contains metadata relating to the software release it belongs to, plus hashes for each individual file or other object that could be attested.

TCG has also published the PC Client Reference Integrity Measurement specification [PC-Client-RIM], which focuses on a SWID-compatible format suitable for expressing expected measurement values in the specific case of a UEFI-compatible BIOS, where the SWID focus on files and file systems is not a direct fit. While the PC Client RIM is not directly applicable to network equipment, many vendors do use a conventional UEFI BIOS to launch their network OS.

2.4.2. Attestation Logs

Quotes from a TPM can provide evidence of the state of a device up to the time the evidence was recorded, but to make sense of the quote in most cases an event log that identifies which software modules contributed which values to the quote during startup MUST also be provided. The log MUST contain enough information to demonstrate its integrity by allowing exact reconstruction of the digest conveyed in the signed quote (that is, calculating the hash of all the hashes in the log should produce the same values as contained in the PCRs; if they don't match, the log may have been tampered with. See Section 9.1).

There are multiple event log formats which may be supported as viable formats of Evidence between the Attester and Verifier:

- o IMA Event log file exports [IMA]
- o TCG UEFI BIOS event log (TCG EFI Platform Specification for TPM Family 1.1 or 1.2, Section 7) [EFI-TPM]
- o TCG Canonical Event Log [Canonical-Event-Log]

Attesters which use UEFI BIOS and Linux SHOULD use TCG Canonical Event Log [Canonical-Event-Log] and TCG UEFI BIOS event log [EFI-TPM], although the CHARRA YANG model [I-D.ietf-rats-yang-tpm-charra] has no dependence on the format of the log.

3. Standards Components

3.1. Prerequisites for RIV

The Reference Interaction Model for Challenge-Response-based Remote Attestation ([I-D.birkholz-rats-reference-interaction-model]) is based on the standard roles defined in [I-D.ietf-rats-architecture]. However additional prerequisites have been established to allow for interoperable RIV use case implementations. These prerequisites are intended to provide sufficient context information so that the

Verifier can acquire and evaluate measurements collected by the Attester.

3.1.1. Unique Device Identity

A secure Device Identity (DevID) in the form of an IEEE 802.1AR DevID certificate [IEEE-802-1AR] MUST be provisioned in the Attester's TPMs.

3.1.2. Keys

The Attestation Key (AK) and certificate MUST also be provisioned on the Attester according to [Platform-DevID-TPM-2.0], [PC-Client-BIOS-TPM-1.2], or [Platform-ID-TPM-1.2].

The Attester's TPM Keys MUST be associated with the DevID on the Verifier (see [Platform-DevID-TPM-2.0] and Section 5 Security Considerations, below).

3.1.3. Appraisal Policy for Evidence

The Verifier MUST obtain trustworthy Reference Values (encoded as SWID or CoSWID tags [I-D.ietf-sacm-coswid]). These reference measurements will eventually be compared to signed PCR Evidence ('quotes') acquired from an Attester's TPM using Attestation Policies chosen by the administrator or owner of the device.

This document does not specify the format or contents for the Appraisal Policy for Evidence, but Reference Values may be acquired in one of two ways:

1. a Verifier may obtain reference measurements directly from an Reference Value Provider chosen by the Verifier administrator (e.g., through a web site).
2. Signed reference measurements may be distributed by the Reference Value Provider to the Attester, as part of a software update. From there, the reference measurement may be acquired by the Verifier.

In either case, the Verifier Owner MUST select the source of trusted Reference Values through the Appraisal Policy for Evidence.

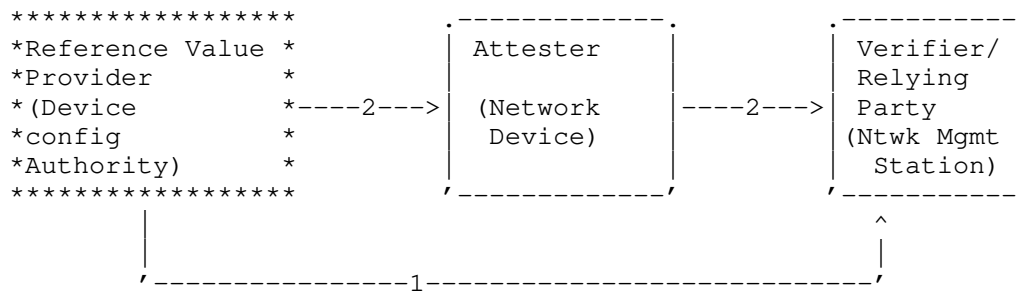


Figure 4: Appraisal Policy for Evidence Prerequisites

In either case the Reference Values must be generated, acquired and delivered in a secure way, including reference measurements of firmware and bootable modules taken according to TCG PC Client [PC-Client-BIOS-TPM-2.0] and Linux IMA [IMA]. Reference Values MUST be encoded as SWID or CoSWID tags, signed by the device manufacturer, as defined in the TCG RIM document [RIM], compatible with NIST IR 8060 [NIST-IR-8060] or the IETF CoSWID draft [I-D.ietf-sacm-coswid].

3.2. Reference Model for Challenge-Response

Once the prerequisites for RIV are met, a Verifier is able to acquire Evidence from an Attester. The following diagram illustrates a RIV information flow between a Verifier and an Attester, derived from Section 8.1 of [I-D.birkholz-rats-reference-interaction-model]. Event times shown correspond to the time types described within Appendix A of [I-D.ietf-rats-architecture]:

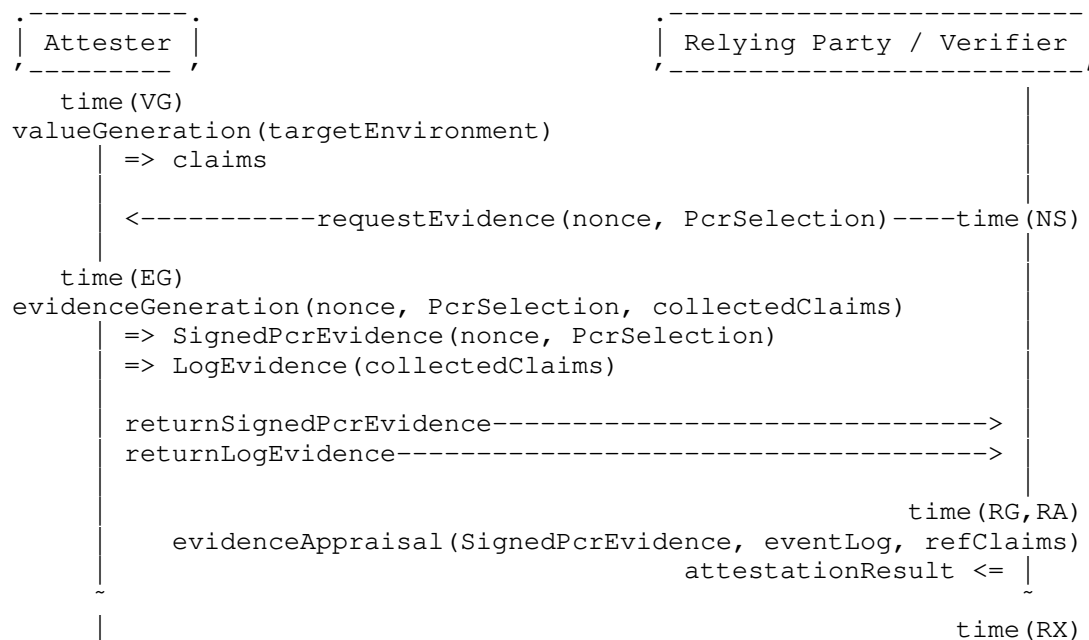


Figure 5: IETF Attestation Information Flow

- o Step 1 (time (VG)): One or more Attesting Network Device PCRs are extended with measurements. RIV provides no direct link between the time at which the event takes place and the time that it's attested, although streaming attestation as in [I-D.birkholz-rats-network-device-subscription] could.
- o Step 2 (time (NS)): The Verifier generates a unique random nonce ("number used once"), and makes a request attestation data for one or more PCRs from an Attester. For interoperability, this MUST be accomplished via a YANG [RFC7950] interface that implements the TCG TAP model (e.g., YANG Module for Basic Challenge-Response-based Remote Attestation Procedures [I-D.ietf-rats-yang-tpm-charra]).
- o Step 3 (time (EG)): On the Attester, measured values are retrieved from the Attester's TPM. This requested PCR evidence, along with the Verifier's nonce, called a Quote, is signed by the Attestation Key (AK) associated with the DevID. Quotes are retrieved according to TCG TAP Information Model [TAP]. At the same time, the Attester collects log evidence showing the values have been extended into that PCR. Section 9.1 gives more detail on how this works.

- o Step 4: Collected Evidence is passed from the Attester to the Verifier
- o Step 5 (time(RG,RA)): The Verifier reviews the Evidence and takes action as needed. As the interaction between Relying Party and Verifier is out of scope for RIV, this can be described as one step.
 - * If the signature covering TPM Evidence is not correct, the device SHOULD NOT be trusted.
 - * If the nonce in the response doesn't match the Verifier's nonce, the response may be a replay, and device SHOULD NOT be trusted.
 - * If the signed PCR values do not match the set of log entries which have extended a particular PCR, the device SHOULD NOT be trusted.
 - * If the log entries that the Verifier considers important do not match known good values, the device SHOULD NOT be trusted. We note that the process of collecting and analyzing the log can be omitted if the value in the relevant PCR is already a known-good value.
 - * If the set of log entries are not seen as acceptable by the Appraisal Policy for Evidence, the device SHOULD NOT be trusted.
 - * If time(RG)-time(NS) is greater than the Appraisal Policy for Evidence's threshold for assessing freshness, the Evidence is considered stale and SHOULD NOT be trusted.

3.2.1. Transport and Encoding

Network Management systems may retrieve signed PCR based Evidence as shown in Figure 5, and can be accomplished via NETCONF or RESTCONF, with XML, JSON, or CBOR encoded Evidence.

Implementations that use NETCONF MUST do so over a TLS or SSH secure tunnel. Implementations that use RESTCONF transport MUST do so over a TLS or SSH secure tunnel.

Retrieval of Log Evidence SHOULD be done via log interfaces specified in [I-D.ietf-rats-yang-tpm-charra]).

3.3. Centralized vs Peer-to-Peer

Figure 5 above assumes that the Verifier is trusted, while the Attester is not. In a Peer-to-Peer application such as two routers negotiating a trust relationship [I-D.void-rats-trusted-path-routing], the two peers can each ask the other to prove software integrity. In this application, the information flow is the same, but each side plays a role both as an Attester and a Verifier. Each device issues a challenge, and each device responds to the other's challenge, as shown in Figure 6. Peer-to-peer challenges, particularly if used to establish a trust relationship between routers, require devices to carry their own signed reference measurements (RIMs). Devices may also have to carry Appraisal Policy for Evidence for each possible peer device so that each device has everything needed for remote attestation, without having to resort to a central authority.

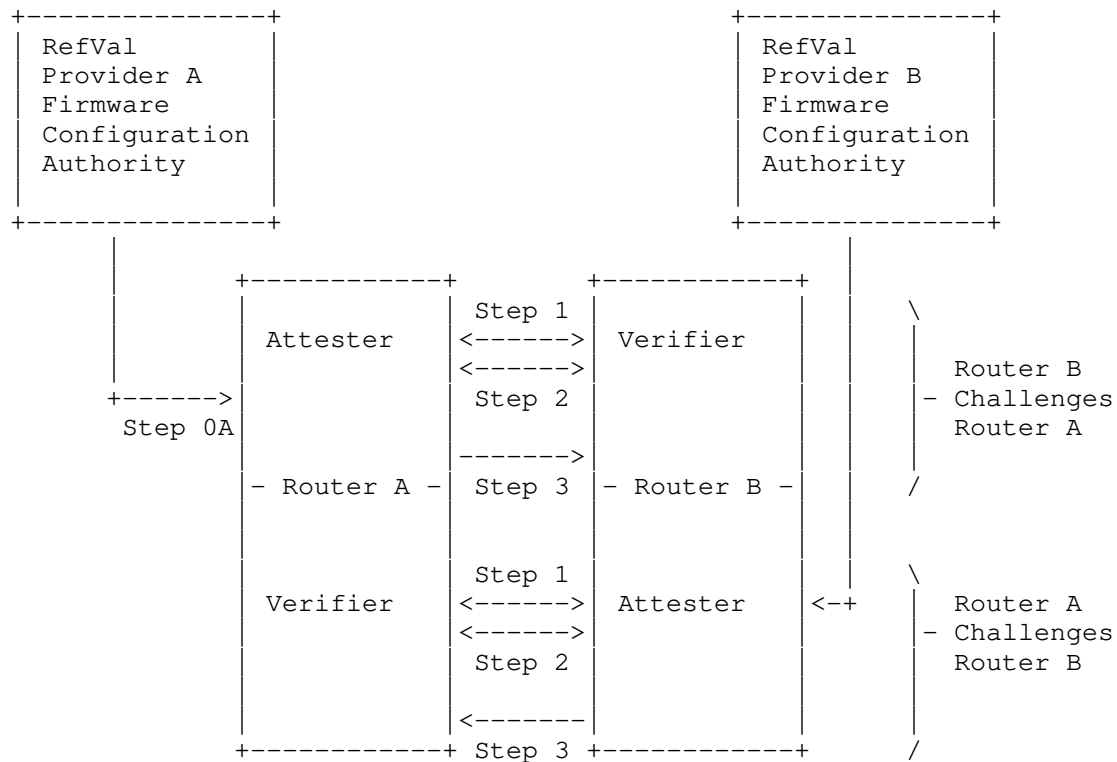


Figure 6: Peer-to-Peer Attestation Information Flow

In this application, each device may need to be equipped with signed RIMs to act as an Attester, and also an Appraisal Policy for Evidence and a selection of trusted X.509 root certificates, to allow the device to act as a Verifier. An existing link layer protocol such as 802.1X [IEEE-802.1X] or 802.1AE [IEEE-802.1AE], with Evidence being enclosed over a variant of EAP [RFC3748] or LLDP [LLDP] are suitable methods for such an exchange.

4. Privacy Considerations

Network equipment, such as routers, switches and firewalls, has a key role to play in guarding the privacy of individuals using the network. Network equipment generally adheres to several rules to protect privacy:

- o Packets passing through the device must not be sent to unauthorized destinations. For example:
 - * Routers often act as Policy Enforcement Points, where individual subscribers may be checked for authorization to access a network. Subscriber login information must not be released to unauthorized parties.
 - * Network equipment is often called upon to block access to protected resources from unauthorized users.
- o Routing information, such as the identity of a router's peers, must not be leaked to unauthorized neighbors.
- o If configured, encryption and decryption of traffic must be carried out reliably, while protecting keys and credentials.

Functions that protect privacy are implemented as part of each layer of hardware and software that makes up the networking device. In light of these requirements for protecting the privacy of users of the network, the network equipment must identify itself, and its boot configuration and measured device state (for example, PCR values), to the equipment's administrator, so there's no uncertainty as to what function each device and configuration is configured to carry out. Attestation is a component that allows the administrator to ensure that the network provides individual and peer privacy guarantees, even though the device itself may not have a right to keep its identity secret.

RIV specifically addresses the collection of information from enterprise network devices by authorized agents of the enterprise. As such, privacy is a fundamental concern for those deploying this solution, given EU GDPR, California CCPA, and many other privacy

regulations. The enterprise SHOULD implement and enforce their duty of care.

See [NetEq] for more context on privacy in networking devices.

5. Security Considerations

Attestation Evidence from the RIV procedure are subject to a number of attacks:

- o Keys may be compromised.
- o A counterfeit device may attempt to impersonate (spoof) a known authentic device.
- o Man-in-the-middle attacks may be used by a counterfeit device to attempt to deliver responses that originate in an actual authentic device.
- o Replay attacks may be attempted by a compromised device.

5.1. Keys Used in RIV

Trustworthiness of RIV attestation depends strongly on the validity of keys used for identity and attestation reports. RIV takes full advantage of TPM capabilities to ensure that evidence can be trusted.

Two sets of key-pairs are relevant to RIV attestation:

- o A DevID key-pair is used to certify the identity of the device in which the TPM is installed.
- o An Attestation Key-pair (AK) key is used to certify attestation Evidence (called 'quotes' in TCG documents), used to provide evidence for integrity of the software on the device

TPM practices usually require that these keys be different, as a way of ensuring that a general-purpose signing key cannot be used to spoof an attestation quote.

In each case, the private half of the key is known only to the TPM, and cannot be retrieved externally, even by a trusted party. To ensure that's the case, specification-compliant private/public key-pairs are generated inside the TPM, where they're never exposed, and cannot be extracted (See [Platform-DevID-TPM-2.0]).

Keeping keys safe is a critical enabler of trustworthiness, but it's just part of attestation security; knowing which keys are bound to

the device in question is just as important in an environment where private keys are never exposed.

While there are many ways to manage keys in a TPM (see [Platform-DevID-TPM-2.0]), RIV includes support for "zero touch" provisioning (also known as zero-touch onboarding) of fielded devices (e.g., Secure ZTP, [RFC8572]), where keys which have predictable trust properties are provisioned by the device vendor.

Device identity in RIV is based on IEEE 802.1AR Device Identity (DevID). This specification provides several elements:

- o A DevID requires a unique key pair for each device, accompanied by an X.509 certificate,
- o The private portion of the DevID key is to be stored in the device, in a manner that provides confidentiality (Section 6.2.5 [IEEE-802-1AR])

The X.509 certificate contains several components:

- o The public part of the unique DevID key assigned to that device allows a challenge of identity.
- o An identifying string that's unique to the manufacturer of the device. This is normally the serial number of the unit, which might also be printed on a label on the device.
- o The certificate must be signed by a key traceable to the manufacturer's root key.

With these elements, the device's manufacturer and serial number can be identified by analyzing the DevID certificate plus the chain of intermediate certificates leading back to the manufacturer's root certificate. As is conventional in TLS or SSH connections, a random nonce must be signed by the device in response to a challenge, proving possession of its DevID private key.

RIV uses the DevID to validate a TLS or SSH connection to the device as the attestation session begins. Security of this process derives from TLS or SSH security, with the DevID providing proof that the session terminates on the intended device. See [RFC8446], [RFC4253].

Evidence of software integrity is delivered in the form of a quote signed by the TPM itself. Because the contents of the quote are signed inside the TPM, any external modification (including reformatting to a different data format) after measurements have been taken will be detected as tampering. An unbroken chain of trust is

essential to ensuring that blocks of code that are taking measurements have been verified before execution (see Figure 1).

Requiring measurements of the operating software to be signed by a key known only to the TPM also removes the need to trust the device's operating software (beyond the first measurement in the RTM; see below); any changes to the quote, generated and signed by the TPM itself, made by malicious device software, or in the path back to the Verifier, will invalidate the signature on the quote.

A critical feature of the YANG model described in [I-D.ietf-rats-yang-tpm-charra] is the ability to carry TPM data structures in their native format, without requiring any changes to the structures as they were signed and delivered by the TPM. While alternate methods of conveying TPM quotes could compress out redundant information, or add an additional layer of signing using external keys, the implementation MUST preserve the TPM signing, so that tampering anywhere in the path between the TPM itself and the Verifier can be detected.

5.2. Prevention of Spoofing and Man-in-the-Middle Attacks

Prevention of spoofing attacks against attestation systems is also important. There are two cases to consider:

- o The entire device could be spoofed. If the Verifier goes to appraise a specific Attester, it might be redirected to a different Attester. Use of the 802.1AR Device Identity (DevID) in the TPM ensures that the Verifier's TLS or SSH session is in fact terminating on the right device.
- o A device with a compromised OS could return a fabricated quote providing spoofed attestation Evidence.

Protection against spoofed quotes from a device with valid identity is a bit more complex. An identity key must be available to sign any kind of nonce or hash offered by the Verifier, and consequently, could be used to sign a fabricated quote. To block a spoofed Attestation Result, the quote generated inside the TPM must be signed by a key that's different from the DevID, called an Attestation Key (AK).

Given separate Attestation and DevID keys, the binding between the AK and the same device must also be proven to prevent a man-in-the-middle attack (e.g., the 'Asokan Attack' [RFC6813]).

This is accomplished in RIV through use of an AK certificate with the same elements as the DevID (same manufacturer's serial number, signed

by the same manufacturer's key), but containing the device's unique AK public key instead of the DevID public key.

The TCG document TPM 2.0 Keys for Device Identity and Attestation [Platform-DevID-TPM-2.0] specifies OIDs for Attestation Certificates that allow the CA to mark a key as specifically known to be an Attestation key.

These two key-pairs and certificates are used together:

- o The DevID is used to validate a TLS connection terminating on the device with a known serial number.
- o The AK is used to sign attestation quotes, providing proof that the attestation evidence comes from the same device.

5.3. Replay Attacks

Replay attacks, where results of a previous attestation are submitted in response to subsequent requests, are usually prevented by inclusion of a random nonce in the request to the TPM for a quote. Each request from the Verifier includes a new random number (a nonce). The resulting quote signed by the TPM contains the same nonce, allowing the Verifier to determine freshness, (i.e., that the resulting quote was generated in response to the Verifier's specific request). Time-Based Uni-directional Attestation [I-D.birkholz-rats-tuda] provides an alternate mechanism to verify freshness without requiring a request/response cycle.

5.4. Owner-Signed Keys

Although device manufacturers MUST pre-provision devices with easily verified DevID and AK certificates if zero-touch provisioning such as described in [RFC8572] is to be supported, use of those credentials is not mandatory. IEEE 802.1AR incorporates the idea of an Initial Device ID (IDevID), provisioned by the manufacturer, and a Local Device ID (LDevID) provisioned by the owner of the device. RIV and [Platform-DevID-TPM-2.0] extends that concept by defining an Initial Attestation Key (IAK) and Local Attestation Key (LAK) with the same properties.

Device owners can use any method to provision the Local credentials.

- o TCG document [Platform-DevID-TPM-2.0] shows how the initial Attestation keys can be used to certify LDevID and LAK keys. Use of the LDevID and LAK allows the device owner to use a uniform identity structure across device types from multiple manufacturers (in the same way that an "Asset Tag" is used by many enterprises

to identify devices they own). TCG document [Provisioning-TPM-2.0] also contains guidance on provisioning Initial and Local identity keys in TPM 2.0.

- o Device owners, however, can use any other mechanism they want to assure themselves that local identity certificates are inserted into the intended device, including physical inspection and programming in a secure location, if they prefer to avoid placing trust in the manufacturer-provided keys.

Clearly, local keys can't be used for secure Zero Touch provisioning; installation of the local keys can only be done by some process that runs before the device is installed for network operation.

On the other end of the device life cycle, provision should be made to wipe local keys when a device is decommissioned, to indicate that the device is no longer owned by the enterprise. The manufacturer's Initial identity keys must be preserved, as they contain no information that's not already printed on the device's serial number plate.

5.5. Other Factors for Trustworthy Operation

In addition to trustworthy provisioning of keys, RIV depends on a number of other factors for trustworthy operation.

- o Secure identity depends on mechanisms to prevent per-device secret keys from being compromised. The TPM provides this capability as a Root of Trust for Storage.
- o Attestation depends on an unbroken chain of measurements, starting from the very first measurement. See Section 9.1 for background on TPM practices.
- o That first measurement is made by code called the Root of Trust for Measurement, typically done by trusted firmware stored in boot flash. Mechanisms for maintaining the trustworthiness of the RTM are out of scope for RIV, but could include immutable firmware, signed updates, or a vendor-specific hardware verification technique. See Section 9.2 for background on roots of trust.
- o The device owner SHOULD provide some level of physical defense for the device. If a TPM that has already been programmed with an authentic DevID is stolen and inserted into a counterfeit device, attestation of that counterfeit device may become indistinguishable from an authentic device.

RIV also depends on reliable Reference Values, as expressed by the RIM [RIM]. The definition of trust procedures for RIMs is out of scope for RIV, and the device owner is free to use any policy to validate a set of reference measurements. RIMs may be conveyed out-of-band or in-band, as part of the attestation process (see Section 3.1.3). But for embedded devices, where software is usually shipped as a self-contained package, RIMs signed by the manufacturer and delivered in-band may be more convenient for the device owner.

The validity of RIV attestation results is also influenced by procedures used to create Reference Values:

- o While the RIM itself is signed, supply-chains SHOULD be carefully scrutinized to ensure that the values are not subject to unexpected manipulation prior to signing. Insider-attacks against code bases and build chains are particularly hard to spot.
- o Designers SHOULD guard against hash collision attacks. Reference Integrity Manifests often give hashes for large objects of indeterminate size; if one of the measured objects can be replaced with an implant engineered to produce the same hash, RIV will be unable to detect the substitution. TPM1.2 uses SHA-1 hashes only, which have been shown to be susceptible to collision attack. TPM2.0 will produce quotes with SHA-256, which so far has resisted such attacks. Consequently, RIV implementations SHOULD use TPM2.0.

6. Conclusion

TCG technologies can play an important part in the implementation of Remote Integrity Verification. Standards for many of the components needed for implementation of RIV already exist:

- o Platform identity can be based on IEEE 802.1AR Device Identity, coupled with careful supply-chain management by the manufacturer.
- o Complex supply chains can be certified using TCG Platform Certificates [Platform-Certificates].
- o The TCG TAP mechanism couple with [I-D.ietf-rats-yang-tpm-charra] can be used to retrieve attestation evidence.
- o Reference Values must be conveyed from the software authority (e.g., the manufacturer) in Reference Integrity Manifests, to the system in which verification will take place. IETF and TCG SWID and CoSWID work [I-D.ietf-sacm-coswid], [RIM])) forms the basis for this function.

7. IANA Considerations

This memo includes no request to IANA.

8. Acknowledgements

The authors wish to thank numerous reviewers for generous assistance, including William Bellingrath, Mark Baushke, Ned Smith, Henk Birkholz, Tom Laffey, Dave Thaler, Wei Pan, Michael Eckel, Thomas Hardjono, Bill Sulzen, Willard (Monty) Wiseman, Kathleen Moriarty, Nancy Cam-Winget and Shwetha Bhandari

9. Appendix

9.1. Using a TPM for Attestation

The Trusted Platform Module and surrounding ecosystem provide three interlocking capabilities to enable secure collection of evidence from a remote device, Platform Configuration Registers (PCRs), a Quote mechanism, and a standardized Event Log.

Each TPM has at least eight and at most twenty-four PCRs (depending on the profile and vendor choices), each one large enough to hold one hash value (SHA-1, SHA-256, and other hash algorithms can be used, depending on TPM version). PCRs can't be accessed directly from outside the chip, but the TPM interface provides a way to "extend" a new security measurement hash into any PCR, a process by which the existing value in the PCR is hashed with the new security measurement hash, and the result placed back into the same PCR. The result is a composite fingerprint of all the security measurements extended into each PCR since the system was reset.

Every time a PCR is extended, an entry should be added to the corresponding Event Log. Logs contain the security measurement hash plus informative fields offering hints as to which event generated the security measurement. The Event Log itself is protected against accidental manipulation, but it is implicitly tamper-evident - any verification process can read the security measurement hash from the log events, compute the composite value and compare that to what ended up in the PCR. If there's no discrepancy, the logs do provide an accurate view of what was placed into the PCR.

Note that the composite hash-of-hashes recorded in PCRs is order-dependent, resulting in different PCR values for different ordering of the same set of events (e.g. Event A followed by Event B yields a different PCR value than B followed by A). For single-threaded code, where both the events and their order are fixed, a Verifier may validate a single PCR value, and use the log only to diagnose a

mismatch from Reference Values. However, operating system code is usually non-deterministic, meaning that there may never be a single "known good" PCR value. In this case, the Verifier may have to verify that the log is correct, and then analyze each item in the log to determine if it represents an authorized event.

In a conventional TPM Attestation environment, the first measurement must be made and extended into the TPM by trusted device code (called the Root of Trust for Measurement, RTM). That first measurement should cover the segment of code that is run immediately after the RTM, which then measures the next code segment before running it, and so on, forming an unbroken chain of trust. See [TCGRoT] for more on Mutable vs Immutable roots of trust.

The TPM provides another mechanism called a Quote that can read the current value of the PCRs and package them, along with the Verifier's nonce, into a TPM-specific data structure signed by an Attestation private key, known only to the TPM.

As noted above in Section 5 Security Considerations, it's important to note that the Quote data structure is signed inside the TPM. The trust model is preserved by retrieving the Quote in a way that does not invalidate the signature, as specified in [I-D.ietf-rats-yang-tpm-charra].

The Verifier uses the Quote and Log together. The Quote contains the composite hash of the complete sequence of security measurement hashes, signed by the TPM's private Attestation Key. The Log contains a record of each measurement extended into the TPM's PCRs. By computing the composite hash of all the measurements, the Verifier can verify the integrity of the Event Log, even though the Event Log itself is not signed. Each hash in the validated Event Log can then be compared to corresponding expected values in the set of Reference Values to validate overall system integrity.

A summary of information exchanged in obtaining quotes from TPM1.2 and TPM2.0 can be found in [TAP], Section 4. Detailed information about PCRs and Quote data structures can be found in [TPM1.2], [TPM2.0]. Recommended log formats include [PC-Client-BIOS-TPM-2.0] and [Canonical-Event-Log].

9.2. Root of Trust for Measurement

The measurements needed for attestation require that the device being attested is equipped with a Root of Trust for Measurement, that is, some trustworthy mechanism that can compute the first measurement in the chain of trust required to attest that each stage of system startup is verified, a Root of Trust for Storage (i.e., the TPM PCRs)

to record the results, and a Root of Trust for Reporting to report the results [TCGRoT], [SP800-155], [SP800-193].

While there are many complex aspects of a Root of Trust, two aspects that are important in the case of attestation are:

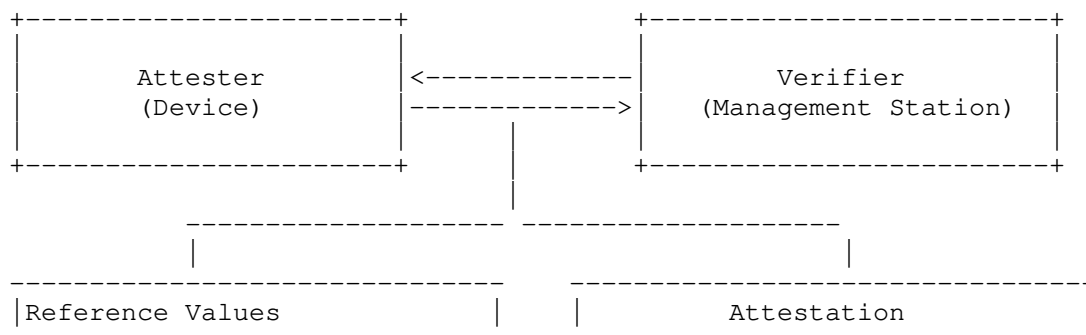
- o The first measurement computed by the Root of Trust for Measurement, and stored in the TPM's Root of Trust for Storage, must be assumed to be correct.
- o There must not be a way to reset the Root of Trust for Storage without re-entering the Root of Trust for Measurement code.

The first measurement must be computed by code that is implicitly trusted; if that first measurement can be subverted, none of the remaining measurements can be trusted. (See [SP800-155])

It's important to note that the trustworthiness of the RTM code cannot be assured by the TPM or TPM supplier - code or procedures external to the TPM must guarantee the security of the RTM.

9.3. Layering Model for Network Equipment Attester and Verifier

Retrieval of identity and attestation state uses one protocol stack, while retrieval of Reference Values uses a different set of protocols. Figure 5 shows the components involved.



 * IETF Attestation Reference Interaction Diagram *

.....
 . Reference Integrity .
 . Manifest .
 . .

.....
 . TAP (PTS2.0) Info .
 . Model and Canonical .
 . Log Format .

 * YANG SWID Module *
 * I-D.ietf-sacm-coswid *
 * *
 * *

.....
 . TCG .
 . Attestation. *
 . MIB .
 . *

 * XML, JSON, CBOR (etc) *

 * UDP *

 * RESTCONF/NETCONF *

 * RESTCONF/NETCONF *

 * TLS, SSH *

 * TLS, SSH *

Figure 7: RIV Protocol Stacks

ietf documents are captured in boxes surrounded by asterisks. TCG documents are shown in boxes surrounded by dots.

9.4. Implementation Notes

Figure 8 summarizes many of the actions needed to complete an Attestation system, with links to relevant documents. While documents are controlled by several standards organizations, the implied actions required for implementation are all the responsibility of the manufacturer of the device, unless otherwise noted. It should be noted that, while the YANG model is RECOMMENDED for attestation, this table identifies an optional SNMP MIB as well, [Attest-MIB].

Component	Controlling Specification
Make a Secure execution environment <ul style="list-style-type: none"> o Attestation depends on a secure root of trust for measurement outside the TPM, as well as roots for storage and reporting inside the TPM. o Refer to TCG Root of Trust for Measurement. o NIST SP 800-193 also provides guidelines on Roots of Trust 	TCG RoT UEFI.org
Provision the TPM as described in TCG documents.	[Platform-DevID-TPM-2.0] TCG Platform Certificate
Put a DevID or Platform Cert in the TPM <ul style="list-style-type: none"> o Install an Initial Attestation Key at the same time so that Attestation can work out of the box o Equipment suppliers and owners may want to implement Local Device ID as well as Initial Device ID 	TCG TPM DevID TCG Platform Certificate IEEE 802.1AR
Connect the TPM to the TLS stack <ul style="list-style-type: none"> o Use the DevID in the TPM to authenticate TAP connections, identifying the device 	Vendor TLS stack (This action is simply configuring TLS to use the DevID as its client certificate)
Make CoSWID tags for BIOS/LoaderLKernel objects <ul style="list-style-type: none"> o Add reference measurements into SWID tags o Manufacturer should sign the SWID tags 	IETF CoSWID ISO/IEC 19770-2 NIST IR 8060

<ul style="list-style-type: none"> o The TCG RIM-IM identifies further procedures to create signed RIM documents that provide the necessary reference information 	
Package the SWID tags with a vendor software release <ul style="list-style-type: none"> o A tag-generator plugin such as [SWID-Gen] can be used 	Retrieve tags with I-D.ietf-sacm-coswid TCG PC Client RIM
Use PC Client measurement definitions to define the use of PCRs (although Windows OS is rare on Networking Equipment, UEFI BIOS is not)	TCG PC Client BIOS
Use TAP to retrieve measurements <ul style="list-style-type: none"> o Map TAP to SNMP o Map to YANG Use Canonical Log Format	TCG SNMP MIB YANG Module for Basic Attestation TCG Canonical Log Format
Posture Collection Server (as described in IETF SACMs ECP) should request the attestation and analyze the result The Management application might be broken down to several more components: <ul style="list-style-type: none"> o A Posture Manager Server which collects reports and stores them in a database o One or more Analyzers that can look at the results and figure out what it means. 	

Figure 8: Component Status

10. References

10.1. Normative References

[Canonical-Event-Log]

Trusted Computing Group, "DRAFT Canonical Event Log Format Version: 1.0, Revision: .12", October 2018.

[I-D.ietf-rats-yang-tpm-charra]

Birkholz, H., Eckel, M., Bhandari, S., Voit, E., Sulzen, B., (Frank), L. X., Laffey, T., and G. C. Fedorkow, "A YANG Data Model for Challenge-Response-based Remote Attestation Procedures using TPMs", draft-ietf-rats-yang-tpm-charra-07 (work in progress), April 2021.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", draft-ietf-sacm-coswid-17 (work in progress), February 2021.

[IEEE-802-1AR]

Seaman, M., "802.1AR-2018 - IEEE Standard for Local and Metropolitan Area Networks - Secure Device Identity, IEEE Computer Society", August 2018.

[PC-Client-BIOS-TPM-1.2]

Trusted Computing Group, "TCG PC Client Specific Implementation Specification for Conventional BIOS, Specification Version 1.21 Errata, Revision 1.00", February 2012,
<<https://trustedcomputinggroup.org/resource/pc-client-work-group-specific-implementation-specification-for-conventional-bios/>>.

[PC-Client-BIOS-TPM-2.0]

Trusted Computing Group, "PC Client Specific Platform Firmware Profile Specification Family "2.0", Level 00 Revision 1.04", June 2019,
<<https://trustedcomputinggroup.org/pc-client-specific-platform-firmware-profile-specification>>.

[PC-Client-RIM]

Trusted Computing Group, "DRAFT: TCG PC Client Reference Integrity Manifest Specification, v.09", December 2019,
<https://trustedcomputinggroup.org/wp-content/uploads/TCG_PC_Client_RIM_r0p15_15june2020.pdf>.

[Platform-DevID-TPM-2.0]

Trusted Computing Group, "TPM 2.0 Keys for Device Identity and Attestation, Specification Version 1.0, Revision 2", September 2020,
<<https://trustedcomputinggroup.org/resource/tpm-2-0-keys-for-device-identity-and-attestation/>>.

- [Platform-ID-TPM-1.2] Trusted Computing Group, "TPM Keys for Platform Identity for TPM 1.2, Specification Version 1.0, Revision 3", August 2015, <<https://trustedcomputinggroup.org/resource/tpm-keys-for-platform-identity-for-tpm-1-2-2/>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", RFC 4253, DOI 10.17487/RFC4253, January 2006, <<https://www.rfc-editor.org/info/rfc4253>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8572] Watsen, K., Farrer, I., and M. Abrahamsson, "Secure Zero Touch Provisioning (SZTP)", RFC 8572, DOI 10.17487/RFC8572, April 2019, <<https://www.rfc-editor.org/info/rfc8572>>.
- [RIM] Trusted Computing Group, "DRAFT: TCG Reference Integrity Manifest Information Model", June 2019, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_RIM_Model_v1-r13_2feb20.pdf>.
- [SWID] The International Organization for Standardization/ International Electrotechnical Commission, "Information Technology Software Asset Management Part 2: Software Identification Tag, ISO/IEC 19770-2", October 2015, <<https://www.iso.org/standard/65666.html>>.

- [TAP] Trusted Computing Group, "TCG Trusted Attestation Protocol (TAP) Information Model for TPM Families 1.2 and 2.0 and DICE Family 1.0, Version 1.0, Revision 0.36", October 2018, <<https://trustedcomputinggroup.org/resource/tcg-tap-information-model/>>.

10.2. Informative References

- [AK-Enrollment] Trusted Computing Group, "TCG Infrastructure Working Group - A CMC Profile for AIK Certificate Enrollment Version 1.0, Revision 7", March 2011, <<https://trustedcomputinggroup.org/resource/tcg-infrastructure-working-group-a-cmc-profile-for-aik-certificate-enrollment/>>.
- [Attest-MIB] Trusted Computing Group, "SNMP MIB for TPM-Based Attestation, Version 0.8Revision 0.02", May 2018, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_SNMP_MIB_for_TPM-Based_Attestation_v0.8r2_PUBLIC_REVIEW.pdf>.
- [EFI-TPM] Trusted Computing Group, "TCG EFI Platform Specification for TPM Family 1.1 or 1.2, Specification Version 1.22, Revision 15", January 2014, <<https://trustedcomputinggroup.org/resource/tcg-efi-platform-specification/>>.
- [I-D.birkholz-rats-network-device-subscription] Birkholz, H., Voit, E., and W. Pan, "Attestation Event Stream Subscription", draft-birkholz-rats-network-device-subscription-02 (work in progress), March 2021.
- [I-D.birkholz-rats-reference-interaction-model] Birkholz, H., Eckel, M., Newton, C., and L. Chen, "Reference Interaction Models for Remote Attestation Procedures", draft-birkholz-rats-reference-interaction-model-03 (work in progress), July 2020.
- [I-D.birkholz-rats-tuda] Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", draft-birkholz-rats-tuda-04 (work in progress), January 2021.

- [I-D.ietf-rats-architecture]
Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", draft-ietf-rats-architecture-12 (work in progress), April 2021.
- [I-D.ietf-rats-eat]
Mandyam, G., Lundblade, L., Ballesteros, M., and J. O'Donoghue, "The Entity Attestation Token (EAT)", draft-ietf-rats-eat-09 (work in progress), March 2021.
- [I-D.richardson-rats-usecases]
Richardson, M., Wallace, C., and W. Pan, "Use cases for Remote Attestation common encodings", draft-richardson-rats-usecases-08 (work in progress), November 2020.
- [I-D.voit-rats-trusted-path-routing]
Voit, E., "Trusted Path Routing", draft-voit-rats-trusted-path-routing-02 (work in progress), June 2020.
- [IEEE-802.1AE]
Seaman, M., "802.1AE MAC Security (MACsec)", 2018, <<https://1.ieee802.org/security/802-1ae/>>.
- [IEEE-802.1X]
IEEE Computer Society, "802.1X-2020 - IEEE Standard for Local and Metropolitan Area Networks--Port-Based Network Access Control", February 2020, <https://standards.ieee.org/standard/802_1X-2020.html>.
- [IMA]
and , "Integrity Measurement Architecture", June 2019, <<https://sourceforge.net/p/linux-ima/wiki/Home/>>.
- [LLDP]
IEEE Computer Society, "802.1AB-2016 - IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery", March 2016, <https://standards.ieee.org/standard/802_1AB-2016.html>.
- [NetEq]
Trusted Computing Group, "TCG Guidance for Securing Network Equipment, Version 1.0, Revision 29", January 2018, <<https://trustedcomputinggroup.org/resource/tcg-guidance-securing-network-equipment/>>.

[NIST-IR-8060]

National Institute for Standards and Technology,
"Guidelines for the Creation of Interoperable Software
Identification (SWID) Tags", April 2016,
<[https://nvlpubs.nist.gov/nistpubs/ir/2016/
NIST.IR.8060.pdf](https://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8060.pdf)>.

[Platform-Certificates]

Trusted Computing Group, "TCG Platform Attribute
Credential Profile, Specification Version 1.0, Revision
16", January 2018,
<[https://trustedcomputinggroup.org/resource/tcg-platform-
attribute-credential-profile/](https://trustedcomputinggroup.org/resource/tcg-platform-attribute-credential-profile/)>.

[Provisioning-TPM-2.0]

Trusted Computing Group, "TCG TPM v2.0 Provisioning
Guidance, Version 1.0, Revision 1.0", March 2015,
<[https://trustedcomputinggroup.org/wp-content/uploads/TCG-
TPM-v2.0-Provisioning-Guidance-Published-v1r1.pdf](https://trustedcomputinggroup.org/wp-content/uploads/TCG-TPM-v2.0-Provisioning-Guidance-Published-v1r1.pdf)>.

[RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H.
Levkowetz, Ed., "Extensible Authentication Protocol
(EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004,
<<https://www.rfc-editor.org/info/rfc3748>>.

[RFC6813] Salowey, J. and S. Hanna, "The Network Endpoint Assessment
(NEA) Asokan Attack Analysis", RFC 6813,
DOI 10.17487/RFC6813, December 2012,
<<https://www.rfc-editor.org/info/rfc6813>>.

[SP800-155]

National Institute of Standards and Technology, "BIOS
Integrity Measurement Guidelines (Draft)", December 2011,
<[https://csrc.nist.gov/csrc/media/publications/sp/800-
155/draft/documents/draft-sp800-155_dec2011.pdf](https://csrc.nist.gov/csrc/media/publications/sp/800-155/draft/documents/draft-sp800-155_dec2011.pdf)>.

[SP800-193]

National Institute for Standards and Technology, "NIST
Special Publication 800-193: Platform Firmware Resiliency
Guidelines", April 2018,
<[https://nvlpubs.nist.gov/nistpubs/SpecialPublications/
NIST.SP.800-193.pdf](https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-193.pdf)>.

[SWID-Gen]

Labs64, Munich, Germany, "SoftWare IDentification (SWID)
Tags Generator (Maven Plugin)", n.d.,
<<https://github.com/Labs64/swid-maven-plugin>>.

- [TCGRoT] Trusted Computing Group, "DRAFT: TCG Roots of Trust Specification", October 2018,
<https://trustedcomputinggroup.org/wp-content/uploads/TCG_Roots_of_Trust_Specification_v0p20_PUBLIC_REVIEW.pdf>.
- [TPM1.2] Trusted Computing Group, "TPM Main Specification Level 2 Version 1.2, Revision 116", March 2011,
<<https://trustedcomputinggroup.org/resource/tpm-main-specification/>>.
- [TPM2.0] Trusted Computing Group, "Trusted Platform Module Library Specification, Family "2.0", Level 00, Revision 01.59", November 2019,
<<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.

Authors' Addresses

Guy Fedorkow (editor)
Juniper Networks, Inc.
US

Email: gfedorkow@juniper.net

Eric Voit
Cisco Systems, Inc.
US

Email: evoit@cisco.com

Jessica Fitzgerald-McKay
National Security Agency
US

Email: jmfitz2@nsa.gov

RATS Working Group
Internet-Draft
Intended status: Informational
Expires: 23 September 2022

G. C. Fedorkow, Ed.
Juniper Networks, Inc.
E. Voit
Cisco
J. Fitzgerald-McKay
National Security Agency
22 March 2022

TPM-based Network Device Remote Integrity Verification
draft-ietf-rats-tpm-based-network-device-attest-14

Abstract

This document describes a workflow for remote attestation of the integrity of firmware and software installed on network devices that contain Trusted Platform Modules [TPM1.2], [TPM2.0], as defined by the Trusted Computing Group (TCG)), or equivalent hardware implementations that include the protected capabilities, as provided by TPMs.

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 23 September 2022.

Copyright Notice

Copyright (c) 2022 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	3
1.1. Requirements notation	3
1.2. Terminology	3
1.3. Document Organization	5
1.4. Goals	5
1.5. Description of Remote Integrity Verification (RIV)	6
1.6. Solution Requirements	8
1.7. Scope	8
1.7.1. Out of Scope	9
2. Solution Overview	9
2.1. RIV Software Configuration Attestation using TPM	9
2.1.1. What Does RIV Attest?	11
2.1.2. Notes on PCR Allocations	13
2.2. RIV Keying	15
2.3. RIV Information Flow	16
2.4. RIV Simplifying Assumptions	18
2.4.1. Reference Integrity Manifests (RIMs)	18
2.4.2. Attestation Logs	20
3. Standards Components	20
3.1. Prerequisites for RIV	20
3.1.1. Unique Device Identity	20
3.1.2. Keys	21
3.1.3. Appraisal Policy for Evidence	21
3.2. Reference Model for Challenge-Response	21
3.2.1. Transport and Encoding	23
3.3. Centralized vs Peer-to-Peer	24
4. Privacy Considerations	25
5. Security Considerations	26
5.1. Keys Used in RIV	26
5.2. Prevention of Spoofing and Person-in-the-Middle Attacks	28
5.3. Replay Attacks	29
5.4. Owner-Signed Keys	30
5.5. Other Factors for Trustworthy Operation	30
6. IANA Considerations	32
7. Conclusion	32
8. Acknowledgements	32
9. Appendix	32
9.1. Using a TPM for Attestation	32
9.2. Root of Trust for Measurement	34
9.3. Layering Model for Network Equipment Attester and Verifier	35

9.4. Implementation Notes	37
10. References	38
10.1. Normative References	38
10.2. Informative References	41
Authors' Addresses	44

1. Introduction

There are many aspects to consider in fielding a trusted computing device, from operating systems to applications. Mechanisms to prove that a device installed at a customer's site is authentic (i.e., not counterfeit) and has been configured with authorized software, all as part of a trusted supply chain, are just a few of the many aspects which need to be considered concurrently to have confidence that a device is truly trustworthy.

A generic architecture for remote attestation has been defined in [I-D.ietf-rats-architecture]. Additionally, use cases for remotely attesting networking devices are discussed within Section 6 of [I-D.richardson-rats-usecases]. However, these documents do not provide sufficient guidance for network equipment vendors and operators to design, build, and deploy interoperable devices.

The intent of this document is to provide such guidance. It does this by outlining the Remote Integrity Verification (RIV) problem, and then identifies elements that are necessary to get the complete, scalable attestation procedure working with commercial networking products such as routers, switches and firewalls. An underlying assumption will be the availability within the device of a Trusted Platform Module [TPM1.2], [TPM2.0] compatible cryptoprocessor to enable the trustworthy remote assessment of the device's software and hardware.

1.1. Requirements notation

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.

1.2. Terminology

A number of terms are reused from [I-D.ietf-rats-architecture]. These include: Appraisal Policy for Evidence, Attestation Result, Attester, Evidence, Reference Value, Relying Party, Verifier, and Verifier Owner.

Additionally, this document defines the following term:

Attestation: the process of generating, conveying and appraising claims, backed by evidence, about device trustworthiness characteristics, including supply chain trust, identity, device provenance, software configuration, device composition, compliance to test suites, functional and assurance evaluations, etc.

The goal of attestation is simply to assure an administrator or auditor that the device configuration and software that was launched when the device was last started is authentic and untampered-with. The determination of software authenticity is not prescribed in this document, but it's typically taken to mean a software image generated by an authority trusted by the administrator, such as the device manufacturer.

Within the Trusted Computing Group (TCG) context, the scope of attestation is typically narrowed to describe the process by which an independent Verifier can obtain cryptographic proof as to the identity of the device in question, and evidence of the integrity of software loaded on that device when it started up, and then verify that what's there matches the intended configuration. For network equipment, a Verifier capability can be embedded in a Network Management Station (NMS), a posture collection server, or other network analytics tool (such as a software asset management solution, or a threat detection and mitigation tool, etc.). While informally referred to as attestation, this document focuses on a specific subset of attestation tasks, defined here as Remote Integrity Verification (RIV). RIV in this document takes a network-equipment-centric perspective that includes a set of protocols and procedures for determining whether a particular device was launched with authentic software, starting from Roots of Trust. While there are many ways to accomplish attestation, RIV sets out a specific set of protocols and tools that work in environments commonly found in network equipment. RIV does not cover other device characteristics that could be attested (e.g., geographic location, connectivity; see [I-D.richardson-rats-usecases]), although it does provide evidence of a secure infrastructure to increase the level of trust in other device characteristics attested by other means (e.g., by Entity Attestation Tokens [I-D.ietf-rats-eat]).

In line with [I-D.ietf-rats-architecture] definitions, this document uses the term Endorser to refer to the role that signs identity and attestation certificates used by the Attester, while Reference Values are signed by a Reference Value Provider. Typically, the manufacturer of a network device would be accepted as both the Endorser and Reference Value Provider, although the choice is ultimately up to the Verifier Owner.

1.3. Document Organization

The remainder of this document is organized into several sections:

- * The remainder of this section covers goals and requirements, plus a top-level description of RIV.
- * The Solution Overview section outlines how Remote Integrity Verification works.
- * The Standards Components section links components of RIV to normative standards.
- * Privacy and Security shows how specific features of RIV contribute to the trustworthiness of the Attestation Result.
- * Supporting material is in an appendix at the end.

1.4. Goals

Network operators benefit from a trustworthy attestation mechanism that provides assurance that their network comprises authentic equipment, and has loaded software free of known vulnerabilities and unauthorized tampering. In line with the overall goal of assuring integrity, attestation can be used to assist in asset management, vulnerability and compliance assessment, plus configuration management.

The RIV attestation workflow outlined in this document is intended to meet the following high-level goals:

- * Provable Device Identity - This specification requires that an Attester (i.e., the attesting device) includes a cryptographic identifier unique to each device. Effectively this means that the device's TPM must be so provisioned during the manufacturing cycle.
- * Software Inventory - A key goal is to identify the software release(s) installed on the Attester, and to provide evidence that the software stored within hasn't been altered without authorization.
- * Verifiability - Verification of software and configuration of the device shows that the software that the administrator authorized for use was actually launched.

In addition, RIV is designed to operate either in a centralized environment, such as with a central authority that manages and configures a number of network devices, or 'peer-to-peer', where network devices independently verify one another to establish a trust relationship. (See Section 3.3 below)

1.5. Description of Remote Integrity Verification (RIV)

Attestation requires two interlocking mechanisms between the Attester network device and the Verifier:

- * Device Identity, the mechanism providing trusted identity, can reassure network managers that the specific devices they ordered from authorized manufacturers for attachment to their network are those that were installed, and that they continue to be present in their network. As part of the mechanism for Device Identity, cryptographic proof of the identity of the manufacturer is also provided.
- * Software Measurement is the mechanism that reports the state of mutable software components on the device, and can assure administrators that they have known, authentic software configured to run in their network.

Using these two interlocking mechanisms, RIV is a component in a chain of procedures that can assure a network operator that the equipment in their network can be reliably identified, and that authentic software of a known version is installed on each device. Equipment in the network includes devices that make up the network itself, such as routers, switches and firewalls.

Software used to boot a device can be identified by a chain of measurements, anchored at the start by a Root of Trust for Measurement (see Section 9.2), each measuring the next stage and recording the result in tamper-resistant storage, normally ending when the system software is fully loaded. A measurement signifies the identity, integrity and version of each software component registered with an Attester's TPM [TPM1.2], [TPM2.0], so that a subsequent verification stage can determine if the software installed is authentic, up-to-date, and free of tampering.

RIV includes several major processes, split between the Attester and Verifier:

1. Generation of Evidence is the process whereby an Attester generates cryptographic proof (Evidence) of claims about device properties. In particular, the device identity and its software configuration are both of critical importance.

2. Device Identification refers to the mechanism assuring the Relying Party (ultimately, a network administrator) of the identity of devices that make up their network, and that their manufacturers are known.
3. Conveyance of Evidence reliably transports the collected Evidence from Attester to a Verifier to allow a management station to perform a meaningful appraisal in Step 4. The transport is typically carried out via a management network. While not required for reliable attestation, an encrypted channel may be used to provide integrity, authenticity, or confidentiality once attestation is complete. It should be noted that critical attestation evidence from the TPM is signed by a key known only to TPM, and is not dependent on encryption carried out as part of a reliable transport.
4. Finally, Appraisal of Evidence occurs. This is the process of verifying the Evidence received by a Verifier from the Attester, and using an Appraisal Policy to develop an Attestation Result, used to inform decision-making. In practice, this means comparing the Attester's measurements reported as Evidence with the device configuration expected by the Verifier. Subsequently, the Appraisal Policy for Evidence might match Evidence found against Reference Values (aka Golden Measurements), which represent the intended configured state of the connected device.

All implementations supporting this RIV specification require the support of the following three technologies:

1. Identity: Device identity in RIV is based on IEEE 802.1AR Device Identity (DevID) [IEEE-802-1AR], coupled with careful supply-chain management by the manufacturer. The Initial DevID (IDevID) certificate contains a statement by the manufacturer that establishes the identity of the device as it left the factory. Some applications with a more-complex post-manufacture supply chain (e.g., Value Added Resellers), or with different privacy concerns, may want to use alternative mechanisms for platform authentication (for example, TCG Platform Certificates [Platform-Certificates], or post-manufacture installation of Local Device ID (LDevID)).
2. Platform Attestation provides evidence of configuration of software elements present in the device. This form of attestation can be implemented with TPM Platform Configuration Registers (PCRs), Quote and Log mechanisms, which provide cryptographically authenticated evidence to report what software was started on the device through the boot cycle. Successful attestation requires an unbroken chain from a boot-time root of

trust through all layers of software needed to bring the device to an operational state, in which each stage computes the hash of components of the next stage, then updates the attestation log and the TPM. The TPM can then report the hashes of all the measured hashes as signed evidence called a Quote (see Section 9.1 for an overview of TPM operation, or [TPM1.2] and [TPM2.0] for many more details).

3. Signed Reference Values (aka Reference Integrity Measurements) must be conveyed from the Reference Value Provider (the entity accepted as the software authority, often the manufacturer of the network device) to the Verifier.

1.6. Solution Requirements

Remote Integrity Verification must address the "Lying Endpoint" problem, in which malicious software on an endpoint may subvert the intended function, and also prevent the endpoint from reporting its compromised status. (See Section 5 for further Security Considerations.)

RIV attestation is designed to be simple to deploy at scale. RIV should work "out of the box" as far as possible, that is, with the fewest possible provisioning steps or configuration databases needed at the end-user's site. Network equipment is often required to "self-configure", to reliably reach out without manual intervention to prove its identity and operating posture, then download its own configuration, a process which precludes pre-installation configuration. See [RFC8572] for an example of Secure Zero Touch Provisioning.

1.7. Scope

The need for assurance of software integrity, addressed by Remote Attestation, is a very general problem that could apply to most network-connected computing devices. However, this document includes several assumptions that limit the scope to network equipment (e.g., routers, switches and firewalls):

- * This solution is for use in non-privacy-preserving applications (for example, networking, Industrial IoT), avoiding the need for a Privacy Certificate Authority (also called an Attestation CA) for attestation keys [AK-Enrollment] or TCG Platform Certificates [Platform-Certificates].
- * This document assumes network protocols that are common in network equipment such as YANG [RFC7950] and NETCONF [RFC6241], but not generally used in other applications.

- * The approach outlined in this document mandates the use of a TPM [TPM1.2], [TPM2.0], or a compatible cryptoprocessor.

1.7.1. Out of Scope

- * **Run-Time Attestation:** The Linux Integrity Measurement Architecture [IMA] attests each process launched after a device is started (and is in scope for RIV in general), but continuous run-time attestation of Linux or other multi-threaded operating system processes after the OS has started considerably expands the scope of the problem. Many researchers are working on that problem, but this document defers the problem of continuous, in-memory run-time attestation.
- * **Multi-Vendor Embedded Systems:** Additional coordination would be needed for devices that themselves comprise hardware and software from multiple vendors, integrated by the end user. Although out of scope for this document, these issues are accommodated in [I-D.ietf-rats-architecture].
- * **Processor Sleep Modes:** Network equipment typically does not "sleep", so sleep and hibernate modes are not considered. Although out of scope for RIV in this document, Trusted Computing Group specifications do encompass sleep and hibernate states, which could be incorporated into remote attestation for network equipment in the future, given a compelling need.
- * **Virtualization and Containerization:** In a non-virtualized system, the host OS is responsible for measuring each User Space file or process throughout the operational lifetime of the system. For virtualized systems, the host OS must verify the hypervisor, but then the hypervisor must manage its own chain of trust through the virtual machine. Virtualization and containerization technologies are increasingly used in network equipment, but are not considered in this document.

2. Solution Overview

2.1. RIV Software Configuration Attestation using TPM

RIV Attestation is a process which can be used to determine the identity of software running on a specifically-identified device. The Remote Attestation steps of Section 1.5 are broken into two phases, shown in Figure 1:

- * During system startup, or boot phase, each distinct software object is "measured" by the Attester. The object's identity, hash (i.e., cryptographic digest) and version information are recorded

in a log. Hashes are also extended into the TPM (see Section 9.1 for more on 'extending hashes'), in a way that can be used to validate the log entries. The measurement process generally follows the layered chain-of-trust model used in Measured Boot, where each stage of the system measures the next one, and extends its measurement into the TPM, before launching it. See [I-D.ietf-rats-architecture], section "Layered Attestation Environments," for an architectural definition of this model.

- * Once the device is running and has operational network connectivity, verification can take place. A separate Verifier, running in its own trusted environment, will interrogate the network device to retrieve the logs and a copy of the digests collected by hashing each software object, signed by an attestation private key secured by, but never released by, the TPM. The YANG model described in [I-D.ietf-rats-yang-tpm-charra] facilitates this operation.

The result is that the Verifier can verify the device's identity by checking the subject[RFC5280] and signature of the certificate containing the TPM's attestation public key, and can validate the software that was launched by verifying the correctness of the logs by comparing with the signed digests from the TPM, and comparing digests in the log with Reference Values.

It should be noted that attestation and identity are inextricably linked; signed Evidence that a particular version of software was loaded is of little value without cryptographic proof of the identity of the Attester producing the Evidence.

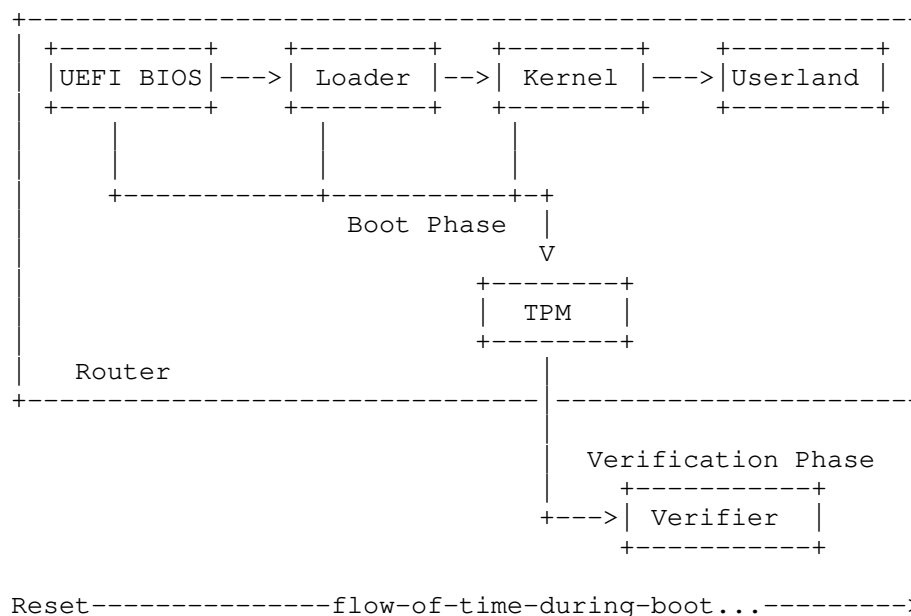


Figure 1: Layered RIV Attestation Model

In the Boot phase, measurements are "extended", or hashed, into the TPM as processes start, with the result that the TPM ends up containing hashes of all the measured hashes. Later, once the system is operational, during the Verification phase, signed digests are retrieved from the TPM for off-box analysis.

2.1.1.1. What Does RIV Attest?

TPM attestation is focused on Platform Configuration Registers (PCRs), but those registers are only vehicles for certifying accompanying Evidence, conveyed in log entries. It is the hashes in log entries that are extended into PCRs, where the final PCR values can be retrieved in the form of a structure called a Quote, signed by an Attestation key known only to the TPM. The use of multiple PCRs serves only to provide some independence between different classes of object, so that one class of objects can be updated without changing the extended hash for other classes. Although PCRs can be used for any purpose, this section outlines the objects within the scope of this document which may be extended into the TPM.

In general, assignment of measurements to PCRs is a policy choice made by the device manufacturer, selected to independently attest three classes of object:

- * Code, (i.e., instructions) to be executed by a CPU.
- * Configuration - Many devices offer numerous options controlled by non-volatile configuration variables which can impact the device's security posture. These settings may have vendor defaults, but often can be changed by administrators, who may want to verify via attestation that the operational state of the settings match their intended state.
- * Credentials - Administrators may wish to verify via attestation that public keys and credentials outside the Root of Trust have not been subject to unauthorized tampering. (By definition, keys protecting the root of trust can't be verified independently.)

The TCG PC Client Platform Firmware Profile Specification [PC-Client-BIOS-TPM-2.0] gives considerable detail on what is to be measured during the boot phase of platform startup using a UEFI BIOS (www.uefi.org), but the goal is simply to measure every bit of code executed in the process of starting the device, along with any configuration information related to security posture, leaving no gap for unmeasured code to remain undetected, potentially subverting the chain.

For devices using a UEFI BIOS, [PC-Client-BIOS-TPM-2.0] and [PC-Client-EFI-TPM-1.2] give detailed normative requirements for PCR usage. For other platform architectures, where TCG normative requirements currently do not exist, the table in Figure 2 gives non-normative guidance for PCR assignment that generalizes the specific details of [PC-Client-BIOS-TPM-2.0].

By convention, most PCRs are assigned in pairs, which the even-numbered PCR used to measure executable code, and the odd-numbered PCR used to measure whatever data and configuration are associated with that code. It is important to note that each PCR may contain results from dozens (or even thousands) of individual measurements.

Function	Assigned PCR #	
	Code	Configuration
Firmware Static Root of Trust, (i.e., initial boot firmware and drivers)	0	1
Drivers and initialization for optional or add-in devices	2	3
OS Loader code and configuration, (i.e., the code launched by firmware) to load an operating system kernel. These PCRs record each boot attempt, and an identifier for where the loader was found	4	5
Vendor Specific Measurements during boot	6	6
Secure Boot Policy. This PCR records keys and configuration used to validate the OS loader		7
Measurements made by the OS Loader (e.g. GRUB2 for Linux)	8	9
Measurements made by OS (e.g., Linux IMA)	10	10

Figure 2: Attested Objects

2.1.2. Notes on PCR Allocations

It is important to recognize that PCR[0] is critical. The first measurement into PCR[0] is taken by the Root of Trust for Measurement, code which, by definition, cannot be verified by measurement. This measurement establishes the chain of trust for all subsequent measurements. If the PCR[0] measurement cannot be trusted, the validity of the entire chain is put into question.

Distinctions Between PCR[0], PCR[2], PCR[4] and PCR[8] are summarized below:

- * PCR[0] typically represents a consistent view of rarely-changed Host Platform boot components, allowing Attestation policies to be defined using the less changeable components of the transitive trust chain. This PCR typically provides a consistent view of the platform regardless of user selected options.

- * PCR[2] is intended to represent a "user configurable" environment where the user has the ability to alter the components that are measured into PCR[2]. This is typically done by adding adapter cards, etc., into user-accessible PCI or other slots. In UEFI systems these devices may be configured by Option ROMs measured into PCR[2] and executed by the UEFI BIOS.
- * PCR[4] is intended to represent the software that manages the transition between the platform's Pre-Operating System start and the state of a system with the Operating System present. This PCR, along with PCR[5], identifies the initial operating system loader (e.g., GRUB for Linux).
- * PCR[8] is used by the OS loader (e.g. GRUB) to record measurements of the various components of the operating system.

Although the TCG PC Client document specifies the use of the first eight PCRs very carefully to ensure interoperability among multiple UEFI BIOS vendors, it should be noted that embedded software vendors may have considerably more flexibility. Verifiers typically need to know which log entries are consequential and which are not (possibly controlled by local policies) but the Verifier may not need to know what each log entry means or why it was assigned to a particular PCR. Designers must recognize that some PCRs may cover log entries that a particular Verifier considers critical and other log entries that are not considered important, so differing PCR values may not on their own constitute a check for authenticity. For example, in a UEFI system, some administrators may consider booting an image from a removable drive, something recorded in a PCR, to be a security violation, while others might consider that operation an authorized recovery procedure.

Designers may allocate particular events to specific PCRs in order to achieve a particular objective with local attestation, (e.g., allowing a procedure to execute, or releasing a particular decryption key, only if a given PCR is in a given state). It may also be important to designers to consider whether streaming notification of PCR updates is required (see [I-D.birkholz-rats-network-device-subscription]). Specific log entries can only be validated if the Verifier receives every log entry affecting the relevant PCR, so (for example) a designer might want to separate rare, high-value events such as configuration changes, from high-volume, routine measurements such as IMA [IMA] logs.

2.2. RIV Keying

RIV attestation relies on two credentials:

- * An identity key pair and matching certificate is required to certify the identity of the Attester itself. RIV specifies the use of an IEEE 802.1AR Device Identity (DevID) [IEEE-802-1AR], signed by the device manufacturer, containing the device serial number. This requirement goes slightly beyond 802.1AR; see Section 2.4 for notes.
- * An Attestation key pair and matching certificate is required to sign the Quote generated by the TPM to report evidence of software configuration.

In a TPM application, both the Attestation private key and the DevID private key MUST be protected by the TPM. Depending on other TPM configuration procedures, the two keys are likely to be different; some of the considerations are outlined in TCG "TPM 2.0 Keys for Device Identity and Attestation" [Platform-DevID-TPM-2.0].

The TCG TPM 2.0 Keys document [Platform-DevID-TPM-2.0] specifies further conventions for these keys:

- * When separate Identity and Attestation keys are used, the Attestation Key (AK) and its X.509 certificate should parallel the DevID, with the same unique device identification as the DevID certificate (that is, the same subject and subjectAltName (if present), even though the key pairs are different). This allows a quote from the device, signed by an AK, to be linked directly to the device that provided it, by examining the corresponding AK certificate. If the subject in the AK certificate doesn't match the corresponding DevID certificate, or they're signed by differing authorities the Verifier may signal the detection of an Asokan-style person-in-the-middle attack (see Section 5.2).
- * Network devices that are expected to use secure zero touch provisioning as specified in [RFC8572] MUST be shipped by the manufacturer with pre-provisioned keys (Initial DevID and Initial AK, called IDevID and IAK). IDevID and IAK certificates MUST both be signed by the Endorser (typically the device manufacturer). Inclusion of an IDevID and IAK by a vendor does not preclude a mechanism whereby an administrator can define Local Identity and Attestation Keys (LDevID and LAK) if desired.

2.3. RIV Information Flow

RIV workflow for network equipment is organized around a simple use case where a network operator wishes to verify the integrity of software installed in specific, fielded devices. A normative taxonomy of terms is given in [I-D.ietf-rats-architecture], but as a reminder, this use case implies several roles and objects:

1. The Attester, the device which the network operator wants to examine.
2. A Verifier (which might be a network management station) somewhere separate from the Device that will retrieve the signed evidence and measurement logs, and analyze them to pass judgment on the security posture of the device.
3. A Relying Party, which can act on Attestation Results. Interaction between the Relying Party and the Verifier is considered out of scope for RIV.
4. Signed Reference Integrity Manifests (RIMs), containing Reference Values, can either be created by the device manufacturer and shipped along with the device as part of its software image, or alternatively, could be obtained several other ways (direct to the Verifier from the manufacturer, from a third party, from the owner's observation of what's thought to be a "known good system", etc.). Retrieving RIMs from the device itself allows attestation to be done in systems that may not have access to the public internet, or by other devices that are not management stations per se (e.g., a peer device; see Section 3.1.3). If Reference Values are obtained from multiple sources, the Verifier may need to evaluate the relative level of trust to be placed in each source in case of a discrepancy.

These components are illustrated in Figure 3.

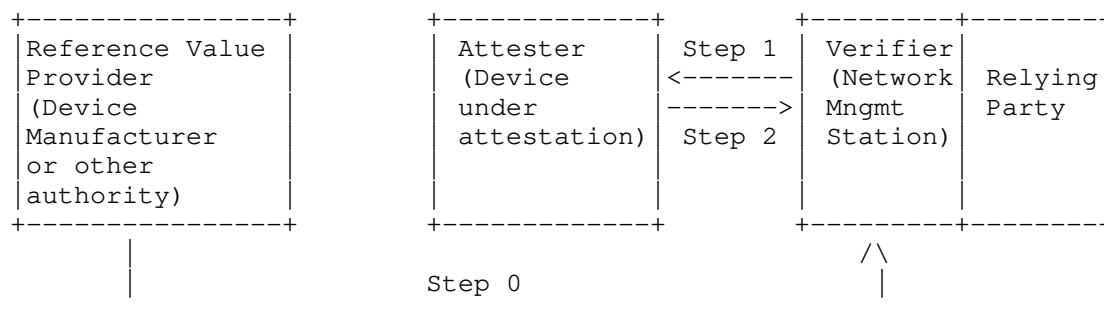


Figure 3: RIV Reference Configuration for Network Equipment

- * In Step 0, The Reference Value Provider (the device manufacturer or other authority) makes one or more Reference Integrity Manifests (RIMs), corresponding to the software image expected to be found on the device, signed by the Reference Value Provider, available to the Verifier (see Section 3.1.3 for "in-band" and "out of band" ways to make this happen).
- * In Step 1, the Verifier (Network Management Station), on behalf of a Relying Party, requests Identity, Measurement Values, and possibly RIMs, from the Attester.
- * In Step 2, the Attester responds to the request by providing a DevID, quotes (measured values, signed by the Attester), and optionally RIMs.

Use of the following standards components allows for interoperability:

1. TPM Keys MUST be configured according to [Platform-DevID-TPM-2.0], or [Platform-ID-TPM-1.2].
2. For devices using UEFI and Linux, measurements of firmware and bootable modules MUST be taken according to TCG PC Client [PC-Client-EFI-TPM-1.2] or [PC-Client-BIOS-TPM-2.0], and Linux IMA [IMA].
3. Device Identity MUST be managed as specified in IEEE 802.1AR Device Identity certificates [IEEE-802-1AR], with keys protected by TPMs.
4. Attestation logs from Linux-based systems MUST be formatted according to the Canonical Event Log format [Canonical-Event-Log]. UEFI-based systems MUST use the TCG UEFI BIOS event log [PC-Client-EFI-TPM-1.2] for TPM1.2 systems, and TCG PC Client Platform Firmware Profile [PC-Client-BIOS-TPM-2.0] for TPM2.0.
5. Quotes MUST be retrieved from the TPM according to TCG TAP Information Model [TAP] and the CHARRA YANG model [I-D.ietf-rats-yang-tpm-charra]. While the TAP IM gives a protocol-independent description of the data elements involved, it's important to note that quotes from the TPM are signed inside the TPM, and MUST be retrieved in a way that does not invalidate the signature, to preserve the trust model. The [I-D.ietf-rats-yang-tpm-charra] is used for this purpose. (See Section 5 Security Considerations).

6. Reference Values MUST be encoded as defined in the TCG RIM document [RIM], typically using SWID [SWID], [NIST-IR-8060] or CoSWID tags [I-D.ietf-sacm-coswid].

2.4. RIV Simplifying Assumptions

This document makes the following simplifying assumptions to reduce complexity:

- * The product to be attested MUST be shipped by the equipment vendor with both an IEEE 802.1AR Device Identity and an Initial Attestation Key (IAK), with certificates in place. The IAK certificate must contain the same identity information as the DevID (specifically, the same subject and subjectAltName (if used), signed by the manufacturer). The IAK is a type of key that can be used to sign a TPM Quote, but not other objects (i.e., it's marked as a TCG "Restricted" key; this convention is described in "TPM 2.0 Keys for Device Identity and Attestation" [Platform-DevID-TPM-2.0]). For network equipment, which is generally non-privacy-sensitive, shipping a device with both an IDevID and an IAK already provisioned substantially simplifies initial startup.
- * IEEE 802.1AR does not require a product serial number as part of the subject, but RIV-compliant devices MUST include their serial numbers in the DevID/IAK certificates to simplify tracking logistics for network equipment users. All other optional 802.1AR fields remain optional in RIV.

It should be noted that 802.1AR use of X.509 certificate fields is not identical to those described in [RFC6125] for representation of application service identity.

- * The product MUST be equipped with a Root of Trust for Measurement (RTM), Root of Trust for Storage and Root of Trust for Reporting (as defined in [SP800-155]) which together are capable of conforming to TCG Trusted Attestation Protocol Information Model [TAP].
- * The authorized software supplier MUST make available Reference Values in the form of signed SWID or CoSWID tags.

2.4.1. Reference Integrity Manifests (RIMs)

[I-D.ietf-rats-yang-tpm-charra] focuses on collecting and transmitting evidence in the form of PCR measurements and attestation logs. But the critical part of the process is enabling the Verifier to decide whether the measurements are "the right ones" or not.

While it must be up to network administrators to decide what they want on their networks, the software supplier should supply the Reference Values, in signed Reference Integrity Manifests, that may be used by a Verifier to determine if evidence shows known good, known bad or unknown software configurations.

In general, there are two kinds of reference measurements:

1. Measurements of early system startup (e.g., BIOS, boot loader, OS kernel) are essentially single-threaded, and executed exactly once, in a known sequence, before any results could be reported. In this case, while the method for computing the hash and extending relevant PCRs may be complicated, the net result is that the software (more likely, firmware) vendor will have one known good PCR value that "should" be present in the relevant PCRs after the box has booted. In this case, the signed reference measurement could simply list the expected hashes for the given version. However, a RIM that contains the intermediate hashes can be useful in debugging cases where the expected final hash is not the one reported.
2. Measurements taken later in operation of the system, once an OS has started (for example, Linux IMA [IMA]), may be more complex, with unpredictable "final" PCR values. In this case, the Verifier must have enough information to reconstruct the expected PCR values from logs and signed reference measurements from a trusted authority.

In both cases, the expected values can be expressed as signed SWID or CoSWID tags, but the SWID structure in the second case is somewhat more complex, as reconstruction of the extended hash in a PCR may involve thousands of files and other objects.

TCG has published an information model defining elements of Reference Integrity Manifests under the title TCG Reference Integrity Manifest Information Model [RIM]. This information model outlines how SWID tags should be structured to allow attestation, and defines "bundles" of SWID tags that may be needed to describe a complete software release. The RIM contains metadata relating to the software release it belongs to, plus hashes for each individual file or other object that could be attested.

Many network equipment vendors use a UEFI BIOS to launch their network operating system. These vendors may want to also use the TCG PC Client Reference Integrity Measurement specification [PC-Client-RIM], which focuses specifically on a SWID-compatible format suitable for expressing measurement values expected from a UEFI BIOS.

2.4.2. Attestation Logs

Quotes from a TPM can provide evidence of the state of a device up to the time the evidence was recorded, but to make sense of the quote in cases where several events are extended into one PCR an event log that identifies which software modules contributed which values to the quote during startup must also be provided. When required, the log MUST contain enough information to demonstrate its integrity by allowing exact reconstruction of the digest conveyed in the signed quote (that is, calculating the hash of all the hashes in the log should produce the same values as contained in the PCRs; if they don't match, the log may have been tampered with. See Section 9.1).

There are multiple event log formats which may be supported as viable formats of Evidence between the Attester and Verifier, but to simplify interoperability, RIV focuses on just three:

- * TCG UEFI BIOS event log for TPM 2.0 (TCG PC Client Platform Firmware Profile) [PC-Client-BIOS-TPM-2.0]
- * TCG UEFI BIOS event log for TPM 1.2 (TCG EFI Platform Specification for TPM Family 1.1 or 1.2, Section 7) [PC-Client-EFI-TPM-1.2]
- * TCG Canonical Event Log [Canonical-Event-Log]

3. Standards Components

3.1. Prerequisites for RIV

The Reference Interaction Model for Challenge-Response-based Remote Attestation ([I-D.birkholz-rats-reference-interaction-model]) is based on the standard roles defined in [I-D.ietf-rats-architecture]. However, additional prerequisites have been established to allow for interoperable RIV use case implementations. These prerequisites are intended to provide sufficient context information so that the Verifier can acquire and evaluate measurements collected by the Attester.

3.1.1. Unique Device Identity

A secure Device Identity (DevID) in the form of an IEEE 802.1AR DevID certificate [IEEE-802-1AR] must be provisioned in the Attester's TPMs.

3.1.2. Keys

The Attestation Key (AK) and certificate must also be provisioned on the Attester according to [Platform-DevID-TPM-2.0], or [Platform-ID-TPM-1.2].

It MUST be possible for the Verifier to determine that the Attester's Attestation keys are resident in the same TPM as its DevID keys (see Section 2.2 and Section 5 Security Considerations).

3.1.3. Appraisal Policy for Evidence

As noted in Section 2.3, the Verifier may obtain Reference Values from several sources. In addition, administrators may make authorized, site-specific changes (e.g. keys in key databases) that could impact attestation results. As such, there could be conflicts, omissions or ambiguities between some Reference Values and collected Evidence.

The Verifier MUST have an Appraisal Policy for Evidence to evaluate the significance of any discrepancies between different reference sources, or between reference values and evidence from logs and quotes. While there must be an Appraisal Policy, this document does not specify the format or mechanism to convey the intended policy, nor does RIV specify mechanisms by which the results of applying the policy are communicated to the Relying Party.

3.2. Reference Model for Challenge-Response

Once the prerequisites for RIV are met, a Verifier is able to acquire Evidence from an Attester. The following diagram illustrates a RIV information flow between a Verifier and an Attester, derived from Section 7.1 of [I-D.birkholz-rats-reference-interaction-model]. In this diagram, each event with its input and output parameters is shown as "Event(input-params)=>(outputs)". Event times shown correspond to the time types described within Appendix A of [I-D.ietf-rats-architecture]:

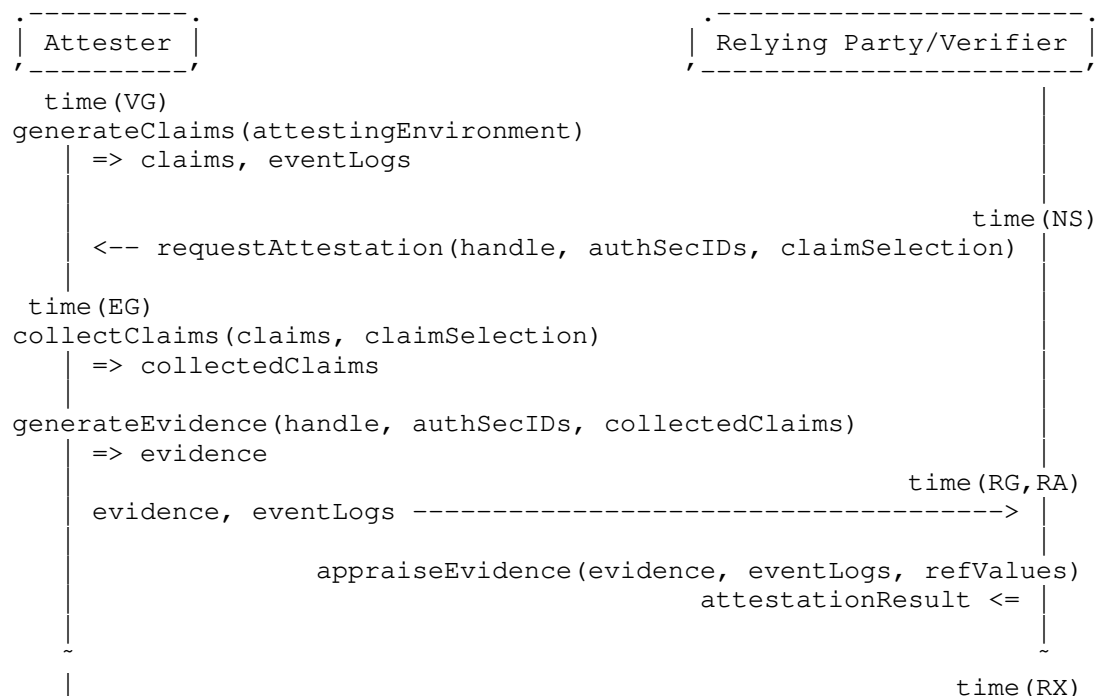


Figure 4: IETF Attestation Information Flow

- * Step 1 (time(VG)): One or more Attesting Network Device PCRs are extended with measurements. RIV provides no direct link between the time at which the event takes place and the time that it's attested, although streaming attestation as in [I-D.birkholz-rats-network-device-subscription] could.
- * Step 2 (time(NS)): The Verifier generates a unique random nonce ("number used once"), and makes a request for one or more PCRs from an Attester. For interoperability, this must be accomplished as specified in the YANG Data Model for Challenge-Response-based Remote Attestation Procedures using TPMs [I-D.ietf-rats-yang-tpm-charra]. TPM1.2 and TPM2.0 both allow nonces as large as the operative digest size (i.e., 20 or 32 bytes; see [TPM1.2] Part 2, Section 5.5 and [TPM2.0] Part 2, Section 10.4.4).
- * Step 3 (time(EG)): On the Attester, measured values are retrieved from the Attester's TPM. This requested PCR evidence, along with the Verifier's nonce, called a Quote, is signed by the Attestation Key (AK) associated with the DevID. Quotes are retrieved according to CHARRA YANG model [I-D.ietf-rats-yang-tpm-charra].

At the same time, the Attester collects log evidence showing the values have been extended into that PCR. Section 9.1 gives more detail on how this works, including references to the structure and contents of quotes in TPM documents.

- * Step 4: Collected Evidence is passed from the Attester to the Verifier
- * Step 5 (time(RG,RA)): The Verifier reviews the Evidence and takes action as needed. As the interaction between Relying Party and Verifier is out of scope for RIV, this can be described as one step.
 - If the signature covering TPM Evidence is not correct, the device SHOULD NOT be trusted.
 - If the nonce in the response doesn't match the Verifier's nonce, the response may be a replay, and device SHOULD NOT be trusted.
 - If the signed PCR values do not match the set of log entries which have extended a particular PCR, the device SHOULD NOT be trusted.
 - If the log entries that the Verifier considers important do not match known good values, the device SHOULD NOT be trusted. We note that the process of collecting and analyzing the log can be omitted if the value in the relevant PCR is already a known-good value.
 - If the set of log entries are not seen as acceptable by the Appraisal Policy for Evidence, the device SHOULD NOT be trusted.
 - If $\text{time(RG)} - \text{time(NS)}$ is greater than the Appraisal Policy for Evidence's threshold for assessing freshness, the Evidence is considered stale and SHOULD NOT be trusted.

3.2.1. Transport and Encoding

Network Management systems may retrieve signed PCR based Evidence using NETCONF or RESTCONF with [I-D.ietf-rats-yang-tpm-charra]. In either case, implementations must do so using a secure tunnel.

Log Evidence MUST be retrieved via log interfaces specified in [I-D.ietf-rats-yang-tpm-charra].

3.3. Centralized vs Peer-to-Peer

Figure 4 above assumes that the Verifier is trusted, while the Attester is not. In a Peer-to-Peer application such as two routers negotiating a trust relationship, the two peers can each ask the other to prove software integrity. In this application, the information flow is the same, but each side plays a role both as an Attester and a Verifier. Each device issues a challenge, and each device responds to the other's challenge, as shown in Figure 5. Peer-to-peer challenges, particularly if used to establish a trust relationship between routers, require devices to carry their own signed reference measurements (RIMs). Devices may also have to carry Appraisal Policy for Evidence for each possible peer device so that each device has everything needed for remote attestation, without having to resort to a central authority.

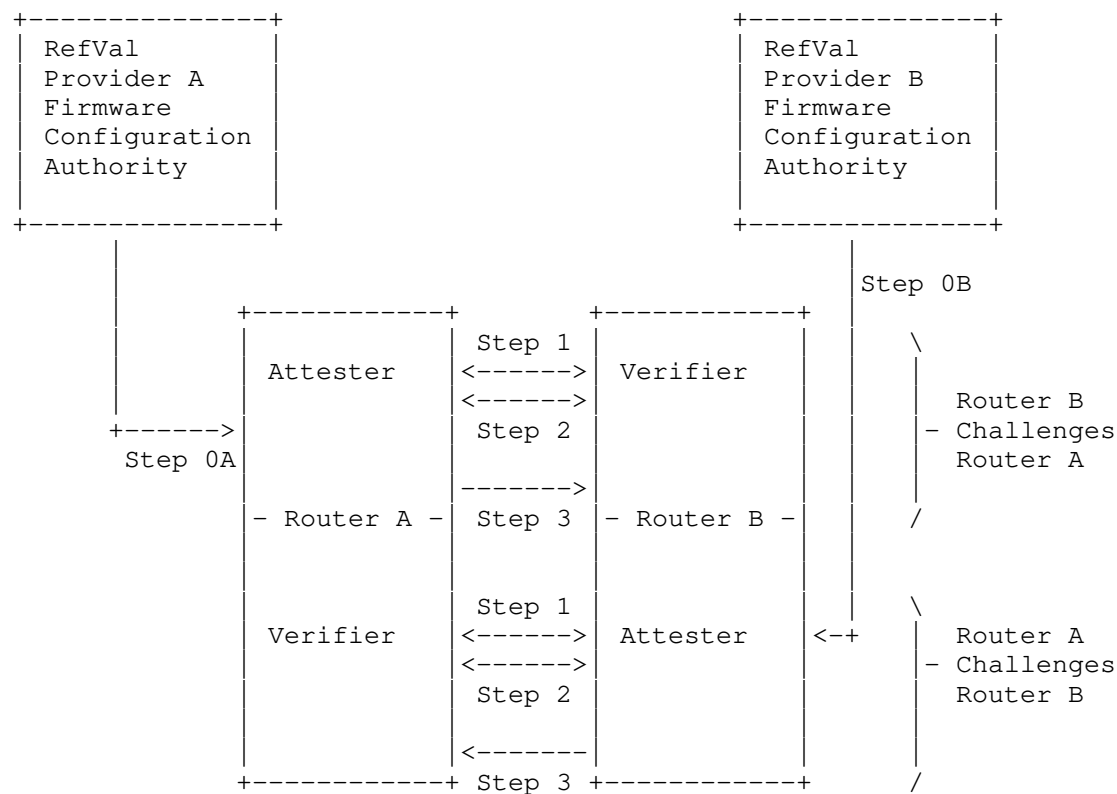


Figure 5: Peer-to-Peer Attestation Information Flow

In this application, each device may need to be equipped with signed RIMs to act as an Attester, and also an Appraisal Policy for Evidence and a selection of trusted X.509 root certificates, to allow the device to act as a Verifier. An existing link layer protocol such as 802.1X [IEEE-802.1X] or 802.1AE [IEEE-802.1AE], with Evidence being enclosed over a variant of EAP [RFC3748] or LLDP [LLDP] are suitable methods for such an exchange. Details of peer-to-peer operation are out of scope for this document.

4. Privacy Considerations

Network equipment, such as routers, switches and firewalls, has a key role to play in guarding the privacy of individuals using the network. Network equipment generally adheres to several rules to protect privacy:

- * Packets passing through the device must not be sent to unauthorized destinations. For example:
 - Routers often act as Policy Enforcement Points, where individual subscribers may be checked for authorization to access a network. Subscriber login information must not be released to unauthorized parties.
 - Network equipment is often called upon to block access to protected resources from unauthorized users.
- * Routing information, such as the identity of a router's peers, must not be leaked to unauthorized neighbors.
- * If configured, encryption and decryption of traffic must be carried out reliably, while protecting keys and credentials.

Functions that protect privacy are implemented as part of each layer of hardware and software that makes up the networking device. In light of these requirements for protecting the privacy of users of the network, the network equipment must identify itself, and its boot configuration and measured device state (for example, PCR values), to the equipment's administrator, so there's no uncertainty as to what function each device and configuration is configured to carry out. Attestation is a component that allows the administrator to ensure that the network provides individual and peer privacy guarantees, even though the device itself may not have a right to keep its identity secret.

See [NetEq] for more context on privacy in networking devices.

While attestation information from network devices is not likely to contain privacy-sensitive content regarding network users, administrators may want to keep attestation records confidential to avoid disclosing versions of software loaded on the device, information which could facilitate attacks against known vulnerabilities.

5. Security Considerations

Specifications such as [RFC8446] (TLS) and [RFC7950] (YANG) contain considerable advice on keeping network-connected systems secure. This section outlines specific risks and mitigations related to attestation.

Attestation Evidence obtained by the RIV procedure is subject to a number of attacks:

- * Keys may be compromised.
- * A counterfeit device may attempt to impersonate (spoof) a known authentic device.
- * Person-in-the-middle attacks may be used by a compromised device to attempt to deliver responses that originate in an authentic device.
- * Replay attacks may be attempted by a compromised device.

5.1. Keys Used in RIV

Trustworthiness of RIV attestation depends strongly on the validity of keys used for identity and attestation reports. RIV takes full advantage of TPM capabilities to ensure that evidence can be trusted.

Two sets of key-pairs are relevant to RIV attestation:

- * A DevID key-pair is used to certify the identity of the device in which the TPM is installed.
- * An Attestation Key-pair (AK) key is used to certify attestation Evidence (called 'quotes' in TCG documents), used to provide evidence for integrity of the software on the device

TPM practices usually require that these keys be different, as a way of ensuring that a general-purpose signing key cannot be used to spoof an attestation quote.

In each case, the private half of the key is known only to the TPM, and cannot be retrieved externally, even by a trusted party. To ensure that's the case, specification-compliant private/public key-pairs are generated inside the TPM, where they are never exposed, and cannot be extracted (See [Platform-DevID-TPM-2.0]).

Keeping keys safe is a critical enabler of trustworthiness, but it's just part of attestation security; knowing which keys are bound to the device in question is just as important in an environment where private keys are never exposed.

While there are many ways to manage keys in a TPM (see [Platform-DevID-TPM-2.0]), RIV includes support for "zero touch" provisioning (also known as zero-touch onboarding) of fielded devices (e.g., Secure ZTP, [RFC8572]), where keys which have predictable trust properties are provisioned by the device vendor.

Device identity in RIV is based on IEEE 802.1AR Device Identity (DevID). This specification provides several elements:

- * A DevID requires a unique key pair for each device, accompanied by an X.509 certificate,
- * The private portion of the DevID key is to be stored in the device, in a manner that provides confidentiality (Section 6.2.5 [IEEE-802-1AR])

The X.509 certificate contains several components:

- * The public part of the unique DevID key assigned to that device allows a challenge of identity.
- * An identifying string that's unique to the manufacturer of the device. This is normally the serial number of the unit, which might also be printed on a label on the device.
- * The certificate must be signed by a key traceable to the manufacturer's root key.

With these elements, the device's manufacturer and serial number can be identified by analyzing the DevID certificate plus the chain of intermediate certificates leading back to the manufacturer's root certificate. As is conventional in TLS or SSH connections, a random nonce must be signed by the device in response to a challenge, proving possession of its DevID private key.

RIV uses the DevID to validate a TLS or SSH connection to the device as the attestation session begins. Security of this process derives from TLS or SSH security, with the DevID, containing a device serial number, providing proof that the session terminates on the intended device. See [RFC8446], [RFC4253].

Evidence of software integrity is delivered in the form of a quote signed by the TPM itself, accompanied by an IAK certificate containing the same identity information as the DevID. Because the contents of the quote are signed inside the TPM, any external modification (including reformatting to a different data format) after measurements have been taken will be detected as tampering. An unbroken chain of trust is essential to ensuring that blocks of code that are taking measurements have been verified before execution (see Figure 1).

Requiring measurements of the operating software to be signed by a key known only to the TPM also removes the need to trust the device's operating software (beyond the first measurement in the RTM; see below); any changes to the quote, generated and signed by the TPM itself, made by malicious device software, or in the path back to the Verifier, will invalidate the signature on the quote.

A critical feature of the YANG model described in [I-D.ietf-rats-yang-tpm-charra] is the ability to carry TPM data structures in their TCG-defined format, without requiring any changes to the structures as they were signed and delivered by the TPM. While alternate methods of conveying TPM quotes could compress out redundant information, or add another layer of signing using external keys, the implementation MUST preserve the TPM signing, so that tampering anywhere in the path between the TPM itself and the Verifier can be detected.

5.2. Prevention of Spoofing and Person-in-the-Middle Attacks

Prevention of spoofing attacks against attestation systems is also important. There are several cases to consider:

- * The entire device could be spoofed. If the Verifier goes to appraise a specific Attester, it might be redirected to a different Attester.
- * A compromised device could have a valid DevID, but substitute a quote from a known-good device, instead of returning its own, as described in [RFC6813].
- * A device with a compromised OS could return a fabricated quote providing spoofed attestation Evidence.

Use of the 802.1AR Device Identity (DevID) in the TPM provides protection against the case of a spoofed device, by ensuring that the Verifier's TLS or SSH session is in fact terminating on the right device.

Protection against spoofed quotes from a device with valid identity is a bit more complex. An identity key must be available to sign any kind of nonce or hash offered by the Verifier, and consequently, could be used to sign a fabricated quote. To block a spoofed Attestation Result, the quote generated inside the TPM must be signed by a key that's different from the DevID, called an Attestation Key (AK).

Given separate Attestation and DevID keys, the binding between the AK and the same device must also be proven to prevent a person-in-the-middle attack (e.g., the 'Asokan Attack' [RFC6813]).

This is accomplished in RIV through use of an AK certificate with the same elements as the DevID (same manufacturer's serial number, signed by the same manufacturer's key), but containing the device's unique AK public key instead of the DevID public key. This binding between DevID and AK certificates is critical to reliable attestation.

The TCG document TPM 2.0 Keys for Device Identity and Attestation [Platform-DevID-TPM-2.0] specifies OIDs for Attestation Certificates that allow the CA to mark a key as specifically known to be an Attestation key.

These two key-pairs and certificates are used together:

- * The DevID is used to validate a TLS connection terminating on the device with a known serial number.
- * The AK is used to sign attestation quotes, providing proof that the attestation evidence comes from the same device.

5.3. Replay Attacks

Replay attacks, where results of a previous attestation are submitted in response to subsequent requests, are usually prevented by inclusion of a random nonce in the request to the TPM for a quote. Each request from the Verifier includes a new random number (a nonce). The resulting quote signed by the TPM contains the same nonce, allowing the Verifier to determine freshness, (i.e., that the resulting quote was generated in response to the Verifier's specific request). Time-Based Uni-directional Attestation [I-D.birkholz-rats-tuda] provides an alternate mechanism to verify freshness without requiring a request/response cycle.

5.4. Owner-Signed Keys

Although device manufacturers must pre-provision devices with easily verified DevID and AK certificates if zero-touch provisioning such as described in [RFC8572] is to be supported, use of those credentials is not mandatory. IEEE 802.1AR incorporates the idea of an Initial Device ID (IDevID), provisioned by the manufacturer, and a Local Device ID (LDevID) provisioned by the owner of the device. RIV and [Platform-DevID-TPM-2.0] extends that concept by defining an Initial Attestation Key (IAK) and Local Attestation Key (LAK) with the same properties.

Device owners can use any method to provision the Local credentials.

- * TCG document [Platform-DevID-TPM-2.0] shows how the initial Attestation keys can be used to certify LDevID and LAK keys. Use of the LDevID and LAK allows the device owner to use a uniform identity structure across device types from multiple manufacturers (in the same way that an "Asset Tag" is used by many enterprises to identify devices they own). TCG document [Provisioning-TPM-2.0] also contains guidance on provisioning Local identity keys in TPM 2.0. Owners should follow the same practice of binding Local DevID and Local AK as the manufacturer would for IDevID and IAK. See Section 2.2.
- * Device owners, however, can use any other mechanism they want to assure themselves that local identity certificates are inserted into the intended device, including physical inspection and programming in a secure location, if they prefer to avoid placing trust in the manufacturer-provided keys.

Clearly, local keys can't be used for secure Zero Touch provisioning; installation of the local keys can only be done by some process that runs before the device is installed for network operation, or using procedures such as those outlined in Bootstrapping Remote Secure Key Infrastructure (BRSKI) [RFC8995].

On the other end of the device life cycle, provision should be made to wipe local keys when a device is decommissioned, to indicate that the device is no longer owned by the enterprise. The manufacturer's Initial identity keys must be preserved, as they contain no information that's not already printed on the device's serial number plate.

5.5. Other Factors for Trustworthy Operation

In addition to trustworthy provisioning of keys, RIV depends on a number of other factors for trustworthy operation.

- * Secure identity depends on mechanisms to prevent per-device secret keys from being compromised. The TPM provides this capability as a Root of Trust for Storage.
- * Attestation depends on an unbroken chain of measurements, starting from the very first measurement. See Section 9.1 for background on TPM practices.
- * That first measurement is made by code called the Root of Trust for Measurement, typically done by trusted firmware stored in boot flash. Mechanisms for maintaining the trustworthiness of the RTM are out of scope for RIV, but could include immutable firmware, signed updates, or a vendor-specific hardware verification technique. See Section 9.2 for background on roots of trust.
- * The device owner SHOULD provide some level of physical defense for the device. If a TPM that has already been programmed with an authentic DevID is stolen and inserted into a counterfeit device, attestation of that counterfeit device may become indistinguishable from an authentic device.

RIV also depends on reliable Reference Values, as expressed by the RIM [RIM]. The definition of trust procedures for RIMs is out of scope for RIV, and the device owner is free to use any policy to validate a set of reference measurements. It should also be noted that, while RIV can provide a reliable indication that a known software package is in use by the device, and that the package has not been tampered, it is the device owner's responsibility to determine that it's the correct package for the application.

RIMs may be conveyed out-of-band or in-band, as part of the attestation process (see Section 3.1.3). But for network devices, where software is usually shipped as a self-contained package, RIMs signed by the manufacturer and delivered in-band may be more convenient for the device owner.

The validity of RIV attestation results is also influenced by procedures used to create Reference Values:

- * While the RIM itself is signed, supply-chains SHOULD be carefully scrutinized to ensure that the values are not subject to unexpected manipulation prior to signing. Insider-attacks against code bases and build chains are particularly hard to spot.
- * Designers SHOULD guard against hash collision attacks. Reference Integrity Manifests often give hashes for large objects of indeterminate size; if one of the measured objects can be replaced with an implant engineered to produce the same hash, RIV will be

unable to detect the substitution. TPM1.2 uses SHA-1 hashes only, which have been shown to be susceptible to collision attack. TPM2.0 will produce quotes with SHA-256, which so far has resisted such attacks. Consequently, RIV implementations SHOULD use TPM2.0.

6. IANA Considerations

This document has no IANA actions.

7. Conclusion

TCG technologies can play an important part in the implementation of Remote Integrity Verification. Standards for many of the components needed for implementation of RIV already exist:

- * Platform identity can be based on IEEE 802.1AR Device Identity, coupled with careful supply-chain management by the manufacturer.
- * Complex supply chains can be certified using TCG Platform Certificates [Platform-Certificates].
- * The TCG TAP mechanism coupled with [I-D.ietf-rats-yang-tpm-charra] can be used to retrieve attestation evidence.
- * Reference Values must be conveyed from the software authority (e.g., the manufacturer) in Reference Integrity Manifests, to the system in which verification will take place. IETF and TCG SWID and CoSWID work ([I-D.ietf-sacm-coswid], [RIM]) forms the basis for this function.

8. Acknowledgements

The authors wish to thank numerous reviewers for generous assistance, including William Bellingrath, Mark Baushke, Ned Smith, Henk Birkholz, Tom Laffey, Dave Thaler, Wei Pan, Michael Eckel, Thomas Hardjono, Bill Sulzen, Willard (Monty) Wiseman, Kathleen Moriarty, Nancy Cam-Winget and Shwetha Bhandari

9. Appendix

9.1. Using a TPM for Attestation

The Trusted Platform Module and surrounding ecosystem provide three interlocking capabilities to enable secure collection of evidence from a remote device, Platform Configuration Registers (PCRs), a Quote mechanism, and a standardized Event Log.

Each TPM has at least eight and at most twenty-four PCRs (depending on the profile and vendor choices), each one large enough to hold one hash value (SHA-1, SHA-256, and other hash algorithms can be used, depending on TPM version). PCRs can't be accessed directly from outside the chip, but the TPM interface provides a way to "extend" a new security measurement hash into any PCR, a process by which the existing value in the PCR is hashed with the new security measurement hash, and the result placed back into the same PCR. The result is a composite fingerprint comprising the hash of all the security measurements extended into each PCR since the system was reset.

Every time a PCR is extended, an entry should be added to the corresponding Event Log. Logs contain the security measurement hash plus informative fields offering hints as to which event generated the security measurement. The Event Log itself is protected against accidental manipulation, but it is implicitly tamper-evident - any verification process can read the security measurement hash from the log events, compute the composite value and compare that to what ended up in the PCR. If there's no discrepancy, the logs do provide an accurate view of what was placed into the PCR.

Note that the composite hash-of-hashes recorded in PCRs is order-dependent, resulting in different PCR values for different ordering of the same set of events (e.g. Event A followed by Event B yields a different PCR value than B followed by A). For single-threaded code, where both the events and their order are fixed, a Verifier may validate a single PCR value, and use the log only to diagnose a mismatch from Reference Values. However, operating system code is usually non-deterministic, meaning that there may never be a single "known good" PCR value. In this case, the Verifier may have to verify that the log is correct, and then analyze each item in the log to determine if it represents an authorized event.

In a conventional TPM Attestation environment, the first measurement must be made and extended into the TPM by trusted device code (called the Root of Trust for Measurement, RTM). That first measurement should cover the segment of code that is run immediately after the RTM, which then measures the next code segment before running it, and so on, forming an unbroken chain of trust. See [TCGRoT] for more on Mutable vs Immutable roots of trust.

The TPM provides another mechanism called a Quote that can read the current value of the PCRs and package them, along with the Verifier's nonce, into a TPM-specific data structure signed by an Attestation private key, known only to the TPM.

As noted above in Section 5 Security Considerations, it's important to note that the Quote data structure is signed inside the TPM. The trust model is preserved by retrieving the Quote in a way that does not invalidate the signature, as specified in [I-D.ietf-rats-yang-tpm-charra]. The structure of the command and response for a quote, including its signature, as generated by the TPM, can be seen in [TPM1.2] Part 3, Section 16.5, and [TPM2.0] Section 18.4.2.

The Verifier uses the Quote and Log together. The Quote contains the composite hash of the complete sequence of security measurement hashes, signed by the TPM's private Attestation Key. The Log contains a record of each measurement extended into the TPM's PCRs. By computing the composite hash of all the measurements, the Verifier can verify the integrity of the Event Log, even though the Event Log itself is not signed. Each hash in the validated Event Log can then be compared to corresponding expected values in the set of Reference Values to validate overall system integrity.

A summary of information exchanged in obtaining quotes from TPM1.2 and TPM2.0 can be found in [TAP], Section 4. Detailed information about PCRs and Quote data structures can be found in [TPM1.2], [TPM2.0]. Recommended log formats include [PC-Client-BIOS-TPM-2.0], and [Canonical-Event-Log].

9.2. Root of Trust for Measurement

The measurements needed for attestation require that the device being attested is equipped with a Root of Trust for Measurement, that is, some trustworthy mechanism that can compute the first measurement in the chain of trust required to attest that each stage of system startup is verified, a Root of Trust for Storage (i.e., the TPM PCRs) to record the results, and a Root of Trust for Reporting to report the results.

While there are many complex aspects of Roots of Trust ([TCGRoT], [SP800-155], [SP800-193]), two aspects that are important in the case of attestation are:

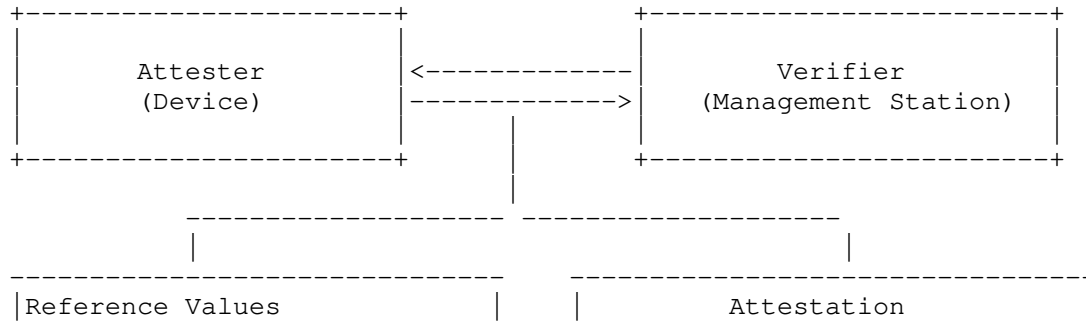
- * The first measurement computed by the Root of Trust for Measurement, and stored in the TPM's Root of Trust for Storage, must be assumed to be correct.
- * There must not be a way to reset the Root of Trust for Storage without re-entering the Root of Trust for Measurement code.

The first measurement must be computed by code that is implicitly trusted; if that first measurement can be subverted, none of the remaining measurements can be trusted. (See [SP800-155])

It's important to note that the trustworthiness of the RTM code cannot be assured by the TPM or TPM supplier - code or procedures external to the TPM must guarantee the security of the RTM.

9.3. Layering Model for Network Equipment Attester and Verifier

Retrieval of identity and attestation state uses one protocol stack, while retrieval of Reference Values uses a different set of protocols. Figure 5 shows the components involved.



```

*****
*           IETF Attestation Reference Interaction Diagram           *
*****

```

```

.....
. Reference Integrity .
.   Manifest         .
.                   .
.....

```

```

*****
* YANG SWID Module      *
* I-D.ietf-sacm-coswid *
*                   *
*                   *
*****

```

```

*****
* XML, JSON, CBOR (etc) *
*****

```

```

*****
* RESTCONF/NETCONF      *
*****

```

```

*****
*      TLS, SSH         *
*****

```

```

.....
. TAP (PTS2.0) Info    .
. Model and Canonical  .
.   Log Format          .
.....

```

```

*****
* YANG Attestation     *
* Module               *
* I-D.ietf-rats-       *
* yang-tpm-charra      *
*****

```

```

*****
* XML, JSON, CBOR (etc) *
*****

```

```

*****
* RESTCONF/NETCONF      *
*****

```

```

*****
*      TLS, SSH         *
*****

```

Figure 6: RIV Protocol Stacks

IETF documents are captured in boxes surrounded by asterisks. TCG documents are shown in boxes surrounded by dots.

9.4. Implementation Notes

Figure 7 summarizes many of the actions needed to complete an Attestation system, with links to relevant documents. While documents are controlled by several standards organizations, the implied actions required for implementation are all the responsibility of the manufacturer of the device, unless otherwise noted.

As noted, SWID tags can be generated many ways, but one possible tool is [SWID-Gen]

Component	Controlling Specification
Make a Secure execution environment <ul style="list-style-type: none"> o Attestation depends on a secure root of trust for measurement outside the TPM, as well as roots for storage and reporting inside the TPM. o Refer to TCG Root of Trust for Measurement. o NIST SP 800-193 also provides guidelines on Roots of Trust 	TCG RoT UEFI.org
Provision the TPM as described in TCG documents.	[Platform-DevID-TPM-2.0] TCG Platform Certificate
Put a DevID or Platform Cert in the TPM <ul style="list-style-type: none"> o Install an Initial Attestation Key at the same time so that Attestation can work out of the box o Equipment suppliers and owners may want to implement Local Device ID as well as Initial Device ID 	TCG TPM DevID TCG Platform Certificate ----- IEEE 802.1AR
Connect the TPM to the TLS stack <ul style="list-style-type: none"> o Use the DevID in the TPM to authenticate TAP connections, identifying the device 	Vendor TLS stack (This action is configuring TLS to use the DevID as its client certificate)
Make CoSWID tags for BIOS/Loader/Kernel objects <ul style="list-style-type: none"> o Add reference measurements into SWID tags o Manufacturer should sign the SWID tags 	IETF CoSWID ISO/IEC 19770-2 NIST IR 8060

<ul style="list-style-type: none"> o The TCG RIM-IM identifies further procedures to create signed RIM documents that provide the necessary reference information 	
Package the SWID tags with a vendor software release <ul style="list-style-type: none"> o A tag-generator plugin such as [SWID-Gen] can be used 	Retrieve tags with I-D.ietf-sacm-coswid TCG PC Client RIM
Use PC Client measurement definitions to define the use of PCRs (although Windows OS is rare on Networking Equipment, UEFI BIOS is not)	TCG PC Client BIOS
Use TAP to retrieve measurements <ul style="list-style-type: none"> o Map to YANG Use Canonical Log Format	YANG Module for Basic Attestation TCG Canonical Log Format
Posture Collection Server (as described in IETF SACMs ECP) should request the attestation and analyze the result The Management application might be broken down to several more components: <ul style="list-style-type: none"> o A Posture Manager Server which collects reports and stores them in a database o One or more Analyzers that can look at the results and figure out what it means. 	

Figure 7: Component Status

10. References

10.1. Normative References

[Canonical-Event-Log]

Trusted Computing Group, "Canonical Event Log Format Version 1.0 Revision .41, February 25, 2022", December 2020, <<https://trustedcomputinggroup.org/resource/canonical-event-log-format/>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-15, 8 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-15.txt>>.

[I-D.ietf-rats-yang-tpm-charra]

Birkholz, H., Eckel, M., Bhandari, S., Voit, E., Sulzen, B., (Frank), L. X., Laffey, T., and G. C. Fedorkow, "A YANG Data Model for Challenge-Response-based Remote Attestation Procedures using TPMs", Work in Progress, Internet-Draft, draft-ietf-rats-yang-tpm-charra-18, 20 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-yang-tpm-charra-18.txt>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-21, 7 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-sacm-coswid-21.txt>>.

[IEEE-802-1AR]

Seaman, M., "802.1AR-2018 - IEEE Standard for Local and Metropolitan Area Networks - Secure Device Identity, IEEE Computer Society", August 2018.

[IMA]

dsafford, kds_etu, mzohar, reinersailer, and serge_hallyn, "Integrity Measurement Architecture", June 2019, <<https://sourceforge.net/p/linux-ima/wiki/Home/>>.

[PC-Client-BIOS-TPM-2.0]

Trusted Computing Group, "PC Client Specific Platform Firmware Profile Specification Family "2.0", Level 00 Revision 1.05 Revision 23, May 7, 2021", May 2021, <<https://trustedcomputinggroup.org/resource/pc-client-specific-platform-firmware-profile-specification/>>.

[PC-Client-EFI-TPM-1.2]

Trusted Computing Group, "TCG EFI Platform Specification for TPM Family 1.1 or 1.2, Specification Version 1.22, Revision 15", January 2014, <<https://trustedcomputinggroup.org/resource/tcg-efi-platform-specification/>>.

[PC-Client-RIM]

Trusted Computing Group, "TCG PC Client Reference Integrity Manifest Specification, v1.04, Nov 4, 2020", December 2019, <<https://trustedcomputinggroup.org/resource/tcg-pc-client-reference-integrity-manifest-specification/>>.

[Platform-DevID-TPM-2.0]

Trusted Computing Group, "TPM 2.0 Keys for Device Identity and Attestation, Specification Version 1.0, Revision 2", September 2020, <<https://trustedcomputinggroup.org/resource/tpm-2-0-keys-for-device-identity-and-attestation/>>.

[Platform-ID-TPM-1.2]

Trusted Computing Group, "TPM Keys for Platform Identity for TPM 1.2, Specification Version 1.0, Revision 3", August 2015, <<https://trustedcomputinggroup.org/resource/tpm-keys-for-platform-identity-for-tpm-1-2-2/>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", RFC 4253, DOI 10.17487/RFC4253, January 2006, <<https://www.rfc-editor.org/info/rfc4253>>.

[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://www.rfc-editor.org/info/rfc5280>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

- [RIM] Trusted Computing Group, "TCG Reference Integrity Manifest (RIM) Information Model, v1.0, Revision 0.16, Nov 12, 2020", June 2019, <<https://trustedcomputinggroup.org/resource/tcg-reference-integrity-manifest-rim-information-model/>>.
- [SWID] The International Organization for Standardization/International Electrotechnical Commission, "Information Technology Software Asset Management Part 2: Software Identification Tag, ISO/IEC 19770-2", October 2015, <<https://www.iso.org/standard/65666.html>>.
- [TAP] Trusted Computing Group, "TCG Trusted Attestation Protocol (TAP) Information Model for TPM Families 1.2 and 2.0 and DICE Family 1.0, Version 1.0, Revision 0.36", October 2018, <<https://trustedcomputinggroup.org/resource/tcg-tap-information-model/>>.

10.2. Informative References

- [AK-Enrollment] Trusted Computing Group, "TCG Infrastructure Working Group - A CMC Profile for AIK Certificate Enrollment Version 1.0, Revision 7", March 2011, <<https://trustedcomputinggroup.org/resource/tcg-infrastructure-working-group-a-cmc-profile-for-aik-certificate-enrollment/>>.
- [I-D.birkholz-rats-network-device-subscription] Birkholz, H., Voit, E., and W. Pan, "Attestation Event Stream Subscription", Work in Progress, Internet-Draft, draft-birkholz-rats-network-device-subscription-03, 17 August 2021, <<https://www.ietf.org/archive/id/draft-birkholz-rats-network-device-subscription-03.txt>>.
- [I-D.birkholz-rats-reference-interaction-model] Birkholz, H., Eckel, M., Newton, C., and L. Chen, "Reference Interaction Models for Remote Attestation Procedures", Work in Progress, Internet-Draft, draft-birkholz-rats-reference-interaction-model-03, 7 July 2020, <<https://www.ietf.org/archive/id/draft-birkholz-rats-reference-interaction-model-03.txt>>.

[I-D.birkholz-rats-tuda]

Fuchs, A., Birkholz, H., McDonald, I. E., and C. Bormann, "Time-Based Uni-Directional Attestation", Work in Progress, Internet-Draft, draft-birkholz-rats-tuda-06, 12 January 2022, <<https://www.ietf.org/archive/id/draft-birkholz-rats-tuda-06.txt>>.

[I-D.ietf-rats-eat]

Lundblade, L., Mandyam, G., and J. O'Donoghue, "The Entity Attestation Token (EAT)", Work in Progress, Internet-Draft, draft-ietf-rats-eat-12, 24 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-eat-12.txt>>.

[I-D.richardson-rats-usecases]

Richardson, M., Wallace, C., and W. Pan, "Use cases for Remote Attestation common encodings", Work in Progress, Internet-Draft, draft-richardson-rats-usecases-08, 2 November 2020, <<https://www.ietf.org/archive/id/draft-richardson-rats-usecases-08.txt>>.

[IEEE-802.1AE]

Seaman, M., "802.1AE MAC Security (MACsec)", 2018, <<https://1.ieee802.org/security/802-1ae/>>.

[IEEE-802.1X]

IEEE Computer Society, "802.1X-2020 - IEEE Standard for Local and Metropolitan Area Networks--Port-Based Network Access Control", February 2020, <https://standards.ieee.org/standard/802_1X-2020.html>.

[LLDP]

IEEE Computer Society, "802.1AB-2016 - IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery", March 2016, <https://standards.ieee.org/standard/802_1AB-2016.html>.

[NetEq]

Trusted Computing Group, "TCG Guidance for Securing Network Equipment, Version 1.0, Revision 29", January 2018, <<https://trustedcomputinggroup.org/resource/tcg-guidance-securing-network-equipment/>>.

[NIST-IR-8060]

National Institute for Standards and Technology, "Guidelines for the Creation of Interoperable Software Identification (SWID) Tags", April 2016, <<https://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8060.pdf>>.

[Platform-Certificates]

Trusted Computing Group, "TCG Platform Attribute Credential Profile, Specification Version 1.0, Revision 16", January 2018, <<https://trustedcomputinggroup.org/resource/tcg-platform-attribute-credential-profile/>>.

[Provisioning-TPM-2.0]

Trusted Computing Group, "TCG TPM v2.0 Provisioning Guidance, Version 1.0, Revision 1.0", March 2015, <<https://trustedcomputinggroup.org/wp-content/uploads/TCG-TPM-v2.0-Provisioning-Guidance-Published-v1r1.pdf>>.

[RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<https://www.rfc-editor.org/info/rfc3748>>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.

[RFC6813] Salowey, J. and S. Hanna, "The Network Endpoint Assessment (NEA) Asokan Attack Analysis", RFC 6813, DOI 10.17487/RFC6813, December 2012, <<https://www.rfc-editor.org/info/rfc6813>>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

[RFC8572] Watsen, K., Farrer, I., and M. Abrahamsson, "Secure Zero Touch Provisioning (SZTP)", RFC 8572, DOI 10.17487/RFC8572, April 2019, <<https://www.rfc-editor.org/info/rfc8572>>.

[RFC8995] Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructure (BRSKI)", RFC 8995, DOI 10.17487/RFC8995, May 2021, <<https://www.rfc-editor.org/info/rfc8995>>.

[SP800-155]

National Institute of Standards and Technology, "BIOS Integrity Measurement Guidelines (Draft)", December 2011, <https://csrc.nist.gov/csrc/media/publications/sp/800-155/draft/documents/draft-sp800-155_dec2011.pdf>.

[SP800-193]

National Institute for Standards and Technology, "NIST Special Publication 800-193: Platform Firmware Resiliency Guidelines", April 2018, <<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-193.pdf>>.

[SWID-Gen] Labs64, Munich, Germany, "SoftWare IDentification (SWID) Tags Generator (Maven Plugin)", n.d., <<https://github.com/Labs64/swid-maven-plugin>>.

[TCGRoT] Trusted Computing Group, "DRAFT: TCG Roots of Trust Specification", October 2018, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_Roots_of_Trust_Specification_v0p20_PUBLIC_REVIEW.pdf>.

[TPM1.2] Trusted Computing Group, "TPM Main Specification Level 2 Version 1.2, Revision 116", March 2011, <<https://trustedcomputinggroup.org/resource/tpm-main-specification/>>.

[TPM2.0] Trusted Computing Group, "Trusted Platform Module Library Specification, Family "2.0", Level 00, Revision 01.59", November 2019, <<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.

Authors' Addresses

Guy Fedorkow (editor)
Juniper Networks, Inc.
10 Technology Park Drive
Westford, Massachusetts 01886
United States of America
Email: gfedorkow@juniper.net

Eric Voit
Cisco Systems
Email: evoit@cisco.com

Jessica Fitzgerald-McKay
National Security Agency
9800 Savage Road
Ft. Meade, Maryland 20755
United States of America
Email: jmfitz2@nsa.gov

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 5 December 2021

H. Birkholz
M. Eckel
Fraunhofer SIT
S. Bhandari
ThoughtSpot
E. Voit
B. Sulzen
Cisco
L. Xia
Huawei
T. Laffey
HPE
G. Fedorkow
Juniper
3 June 2021

A YANG Data Model for Challenge-Response-based Remote Attestation
Procedures using TPMs
draft-ietf-rats-yang-tpm-charra-08

Abstract

This document defines YANG RPCs and a small number of configuration nodes required to retrieve attestation evidence about integrity measurements from a device, following the operational context defined in TPM-based Network Device Remote Integrity Verification. Complementary measurement logs are also provided by the YANG RPCs, originating from one or more roots of trust for measurement (RTMs). The module defined requires at least one TPM 1.2 or TPM 2.0 as well as a corresponding TPM Software Stack (TSS), included in the device components of the composite device the YANG server is running on.

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 December 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements notation	3
2. The YANG Module for Basic Remote Attestation Procedures . . .	3
2.1. YANG Modules	3
2.1.1. 'ietf-tpm-remote-attestation'	3
2.1.2. 'ietf-tcg-algs'	32
3. IANA Considerations	47
4. Security Considerations	47
5. Acknowledgements	48
6. Change Log	48
7. References	49
7.1. Normative References	49
7.2. Informative References	51
Authors' Addresses	51

1. Introduction

This document is based on the general terminology defined in the [I-D.ietf-rats-architecture] and uses the operational context defined in [I-D.ietf-rats-tpm-based-network-device-attest] as well as the interaction model and information elements defined in [I-D.ietf-rats-reference-interaction-models]. The currently supported hardware security modules (HSMs) are the Trusted Platform Modules (TPMs) [TPM1.2] and [TPM2.0] as specified by the Trusted Computing Group (TCG). One or more TPMs embedded in the components of a Composite Device are required in order to use the YANG module defined in this document. A TPM is used as a root of trust for reporting (RTR) in order to retrieve attestation Evidence from a composite device (_TPM Quote_ primitive operation). Additionally, it is used as a root of trust for storage (RTS) in order to retain

shielded secrets and store system measurements using a folding hash function (`_TPM_PCR_Extend_` primitive operation).

Specific terms imported from [I-D.ietf-rats-architecture] and used in this document include: Attester, Composite Device, Evidence.

Specific terms imported from [TPM2.0-Key] and used in this document include: Endorsement Key (EK), Initial Attestation Key (IAK), Local Attestation Key (LAK).

1.1. Requirements notation

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. The YANG Module for Basic Remote Attestation Procedures

One or more TPMs MUST be embedded in a Composite Device that provides attestation evidence via the YANG module defined in this document. The `ietf-basic-remote-attestation` YANG module enables a composite device to take on the role of an Attester, in accordance with the Remote Attestation Procedures (RATS) architecture [I-D.ietf-rats-architecture], and the corresponding challenge-response interaction model defined in the [I-D.ietf-rats-reference-interaction-models] document. A fresh nonce with an appropriate amount of entropy MUST be supplied by the YANG client in order to enable a proof-of-freshness with respect to the attestation Evidence provided by the Attester running the YANG datastore. Further, this nonce is used to prevent replay attacks. The functions of this YANG module are restricted to 0-1 TPMs per hardware component.

2.1. YANG Modules

In this section the several YANG modules are defined.

2.1.1. 'ietf-tpm-remote-attestation'

This YANG module imports modules from [RFC6991], [RFC8348], [I-D.ietf-netconf-keystore], and "ietf-tcg-algs.yang" Section 2.1.2.3.

2.1.1.1. Features

This module supports the following features:

- * 'TPMs': Indicates that multiple TPMs on the device can support remote attestation. This feature is applicable in cases where multiple line cards are present, each with its own TPM.
- * 'bios': Indicates that the device supports the retrieval of BIOS/UEFI event logs.
- * 'ima': Indicates that the device supports the retrieval of event logs from the Linux Integrity Measurement Architecture (IMA).
- * 'netequip_boot': Indicates that the device supports the retrieval of netequip boot event logs.

2.1.1.2. Identities

This module supports the following types of attestation event logs: 'bios', 'ima', and 'netequip_boot'.

2.1.1.3. Remote Procedure Calls (RPCs)

In the following, RPCs for both TPM 1.2 and TPM 2.0 attestation procedures are defined.

2.1.1.3.1. 'tpm12-challenge-response-attestation'

This RPC allows a Verifier to request signed TPM PCRs (_TPM Quote_ operation) from a TPM 1.2 compliant cryptoprocessor. Where the feature 'TPMs' is active, and one or more 'certificate-name' is not provided, all TPM 1.2 compliant cryptoprocessors will respond. A YANG tree diagram of this RPC is as follows:

```
+---x tpm12-challenge-response-attestation {taa:TPM12}?
  +---w input
    |   +---w tpm12-attestation-challenge
    |       +---w pcr-index*          pcr
    |       +---w nonce-value         binary
    |       +---w certificate-name*    certificate-name-ref {tpm:TPMs}?
  +---ro output
    +---ro tpm12-attestation-response* []
      +---ro certificate-name    certificate-name-ref
      +---ro up-time?           uint32
      +---ro TPM_QUOTE2?        binary
```

2.1.1.3.2. 'tpm20-challenge-response-attestation'

This RPC allows a Verifier to request signed TPM PCRs (_TPM Quote_ operation) from a TPM 2.0 compliant cryptoprocessor. Where the feature 'TPMs' is active, and one or more 'certificate-name' is not provided, all TPM 2.0 compliant cryptoprocessors will respond. A YANG tree diagram of this RPC is as follows:

```
+---x tpm20-challenge-response-attestation {taa:TPM20}?
  +---w input
    +---w tpm20-attestation-challenge
      +---w nonce-value          binary
      +---w tpm20-pcr-selection* []
        +---w TPM20-hash-algo?  identityref
        +---w pcr-index*        tpm:pcr
      +---w certificate-name*    certificate-name-ref {tpm:TPMs}?
  +---ro output
    +---ro tpm20-attestation-response* []
      +---ro certificate-name      certificate-name-ref
      +---ro TPMS_QUOTE_INFO      binary
      +---ro quote-signature?     binary
      +---ro up-time?             uint32
      +---ro unsigned-pcr-values* []
        +---ro TPM20-hash-algo?  identityref
        +---ro pcr-values* [pcr-index]
          +---ro pcr-index      pcr
          +---ro pcr-value?     binary
```

An example of an RPC challenge requesting PCRs 0-7 from a SHA-256 bank could look like the following:


```

<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <tpm20-challenge-response-attestation>
    xmlns="urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation"
    <certificate-name>
      (identifier of a TPM signature key with which the Verifier is supposed
      to sign the attestation data)
    </certificate-name>
    <nonce>
      0xe041307208d9f78f5b1bbecd19e2d152ad49de2fc5a7d8dbf769f6b8ffdeab9d
    </nonce>
    <tpm20-pcr-selection>
      <tpm20-hash-algo>
        xmlns:taa="urn:ietf:params:xml:ns:yang:ietf-tcg-algs"
        taa:TPM_ALG_SHA256
      </tpm20-hash-algo>
      <pcr-index>0</pcr-index>
      <pcr-index>1</pcr-index>
      <pcr-index>2</pcr-index>
      <pcr-index>3</pcr-index>
      <pcr-index>4</pcr-index>
      <pcr-index>5</pcr-index>
      <pcr-index>6</pcr-index>
      <pcr-index>7</pcr-index>
    </tpm20-pcr-selection>
  </tpm20-challenge-response-attestation>
</rpc>

```

A successful response could be formatted as follows:

```

<rpc-reply message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <tpm20-attestation-response>
    xmlns="urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation"
    <certificate-name>
      xmlns:ks="urn:ietf:params:xml:ns:yang:ietf-keystore"
      ks:(instance of Certificate name in the Keystore)
    </certificate-name>
    <attestation-data>
      (raw attestation data, i.e. the TPM quote; this includes
      a composite digest of requested PCRs, the nonce,
      and TPM 2.0 time information.)
    </attestation-data>
    <quote-signature>
      (signature over attestation-data using the TPM key
      identified by sig-key-id)
    </quote-signature>
  </tpm20-attestation-response>
</rpc-reply>

```

2.1.1.4. 'log-retrieval'

This RPC allows a Verifier to acquire the evidence which was extended into specific TPM PCRs. A YANG tree diagram of this RPC is as follows:

```

+---x log-retrieval
+---w input
|   +---w log-selector* []
|   |   +---w name* string
|   |   +---w (index-type)?
|   |   |   +---:(last-entry)
|   |   |   |   +---w last-entry-value? binary
|   |   |   |   +---:(index)
|   |   |   |   |   +---w last-index-number? uint64
|   |   |   |   |   +---:(timestamp)
|   |   |   |   |   |   +---w timestamp? yang:date-and-time
|   |   |   |   |   |   +---w log-entry-quantity? uint16
|   |   +---w log-type identityref
+---ro output
+---ro system-event-logs
+---ro node-data* []
+---ro name? string
+---ro up-time? uint32
+---ro log-result
+---ro (attested_event_log_type)
+---:(bios) {bios}?
+---ro bios-event-logs
+---ro bios-event-entry* [event-number]
+---ro event-number uint32
+---ro event-type? uint32
+---ro pcr-index? pcr
+---ro digest-list* []
+---ro hash-algo? identityref
+---ro digest* binary
+---ro event-size? uint32
+---ro event-data* uint8
+---:(ima) {ima}?
+---ro ima-event-logs
+---ro ima-event-entry* [event-number]
+---ro event-number uint64
+---ro ima-template? string
+---ro filename-hint? string
+---ro filedata-hash? binary
+---ro filedata-hash-algorithm? string
+---ro template-hash-algorithm? string
+---ro template-hash? binary
+---ro pcr-index? pcr

```

```

|          +---ro signature?                      binary
+---:(netequip_boot) {netequip_boot}?
  +---ro boot-event-logs
    +---ro boot-event-entry* [event-number]
      +---ro event-number                        uint64
      +---ro ima-template?                      string
      +---ro filename-hint?                    string
      +---ro filedata-hash?                    binary
      +---ro filedata-hash-algorithm?          string
      +---ro template-hash-algorithm?          string
      +---ro template-hash?                    binary
      +---ro pcr-index?                        pcr
      +---ro signature?                        binary

```

2.1.1.5. Data Nodes

This section provides a high level description of the data nodes containing the configuration and operational objects with the YANG model. For more details, please see the YANG model itself in Figure 1.

Container 'rats-support-structures': This houses the set of information relating to a device's TPM(s).

Container 'tpms': Provides configuration and operational details for each supported TPM, including the tpm-firmware-version, PCRs which may be quoted, certificates which are associated with that TPM, and the current operational status. Of note are the certificates which are associated with that TPM. As a certificate is associated with a particular TPM attestation key, knowledge of the certificate allows a specific TPM to be identified.

```

+--rw tpms
  +--rw tpm* [name]
    +--rw name string
    +--ro hardware-based? boolean
    +--ro physical-index? int32 {ietfhw:entity-mib}?
    +--ro path? string
    +--ro compute-node compute-node-ref {tpm:tpms}?
    +--ro manufacturer? string
    +--rw firmware-version identityref
    +--rw tpm12-hash-algo? identityref
    +--rw tpm12-pcrs* pcr
    +--rw tpm20-pcr-bank* [tpm20-hash-algo]
      | +--rw tpm20-hash-algo identityref
      | +--rw pcr-index* tpm:pcr
    +--ro status enumeration
    +--rw certificates
      +--rw certificate* [name]
        +--rw name string
        +--rw keystore-ref? leafref
        +--rw type? enumeration

```

container 'attester-supported-algos' - Identifies which TCG hash algorithms are available for use on the Attesting platform. This allows an operator to limit algorithms available for use by RPCs to just a desired set from the universe of all allowed hash algorithms by the TCG.

```

+--rw attester-supported-algos
  +--rw tpm12-asymmetric-signing* identityref
  +--rw tpm12-hash* identityref
  +--rw tpm20-asymmetric-signing* identityref
  +--rw tpm20-hash* identityref

```

container 'compute-nodes' - When there is more than one TPM supported, this container maintains the set of information related to the compute node associated with a specific TPM. This allows each specific TPM to identify to which 'compute-node' it belongs.

```

+--rw compute-nodes {tpm:TPMs}?
  +--ro compute-node* [node-id]
    +--ro node-id string
    +--ro node-physical-index? int32 {ietfhw:entity-mib}?
    +--ro node-name? string
    +--ro node-location? string

```

2.1.1.6. YANG Module

```
<CODE BEGINS> file "ietf-tpm-remote-attestation@2021-05-11.yang"
module ietf-tpm-remote-attestation {
  namespace "urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation";
  prefix tpm;

  import ietf-yang-types {
    prefix yang;
  }
  import ietf-hardware {
    prefix ietfhw;
  }
  import ietf-keystore {
    prefix ks;
  }
  import ietf-tcg-algs {
    prefix taa;
  }

  organization
    "IETF RATS (Remote ATtestation procedureS) Working Group";
  contact
    "WG Web : <http://datatracker.ietf.org/wg/rats/>
    WG List : <mailto:rats@ietf.org>
    Author : Eric Voit <evoit@cisco.com>
    Author : Henk Birkholz <henk.birkholz@sit.fraunhofer.de>
    Author : Michael Eckel <michael.eckel@sit.fraunhofer.de>
    Author : Shwetha Bhandari <shwetha.bhandari@thoughtspot.com>
    Author : Bill Sulzen <bsulzen@cisco.com>
    Author : Liang Xia (Frank) <frank.xialiang@huawei.com>
    Author : Tom Laffey <tom.laffey@hpe.com>
    Author : Guy Fedorkow <gfredorkow@juniper.net>";
  description
    "A YANG module to enable a TPM 1.2 and TPM 2.0 based
    remote attestation procedure using a challenge-response
    interaction model and the TPM 1.2 and TPM 2.0 Quote
    primitive operations.
    Copyright (c) 2021 IETF Trust and the persons identified
    as authors of the code. All rights reserved.
    Redistribution and use in source and binary forms, with
    or without modification, is permitted pursuant to, and
    subject to the license terms contained in, the Simplified
    BSD License set forth in Section 4.c of the IETF Trust's
    Legal Provisions Relating to IETF Documents
    (https://trustee.ietf.org/license-info).
    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); see the RFC
    itself for full legal notices.
```

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 (RFC 2119) (RFC 8174) when, and only when, they appear in all capitals, as shown here.";

```
revision 2021-05-11 {
  description
    "Initial version";
  reference
    "draft-ietf-rats-yang-tpm-charra";
}

/*****/
/*   Features   */
/*****/

feature tpms {
  description
    "The device supports the remote attestation of multiple
    TPM based cryptoprocessors.";
}

feature bios {
  description
    "The device supports the bios logs.";
  reference
    "https://trustedcomputinggroup.org/wp-content/uploads/
    PC-ClientSpecific_Platform_Profile_for_TPM_2p0_Systems_v51.pdf
    Section 9.4.5.2";
}

feature ima {
  description
    "The device supports Integrity Measurement Architecture logs.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TCG_IWG_CEL_v1_r0p30_13feb2021.pdf Section 4.3";
}

feature netequip_boot {
  description
    "The device supports the netequip_boot logs.";
}

/*****/
/*   Typedefs   */
/*****/
```

```

/*****/

typedef pcr {
  type uint8 {
    range "0..31";
  }
  description
    "Valid index number for a PCR. At this point 0-31 is viable.";
}

typedef compute-node-ref {
  type leafref {
    path "/tpm:rats-support-structures/tpm:compute-nodes"
      + "/tpm:compute-node/tpm:node-name";
  }
  description
    "This type is used to reference a hardware node. It is quite
    possible this leafref will eventually point to another YANG
    module's node.";
}

typedef certificate-name-ref {
  type leafref {
    path "/tpm:rats-support-structures/tpm:tpms/tpm:tpm"
      + "/tpm:certificates/tpm:certificate/tpm:name";
  }
  description
    "A type which allows identification of a TPM based certificate.";
}

/*****/
/* Identities */
/*****/

identity attested_event_log_type {
  description
    "Base identity allowing categorization of the reasons why and
    attested measurement has been taken on an Attester.";
}

identity ima {
  base attested_event_log_type;
  description
    "An event type recorded in IMA.";
}

identity bios {
  base attested_event_log_type;
}
```

```
    description
      "An event type associated with BIOS/UEFI.";
  }

  identity netequip_boot {
    base attested_event_log_type;
    description
      "An event type associated with Network Equipment Boot.";
  }

  /*****/
  /*  Groupings  */
  /*****/

  grouping tpm20-hash-algo {
    description
      "The cryptographic algorithm used to hash the TPM2 PCRs.  This
      must be from the list of platform supported options.";
    leaf tpm20-hash-algo {
      type identityref {
        base taa:hash;
      }
      must '/tpm:rats-support-structures/tpm:attester-supported-algos'
        + '/tpm:tpm20-hash' {
        error-message "This platform does not support tpm20-hash-algo";
      }
      default "taa:TPM_ALG_SHA256";
      description
        "The hash scheme that is used to hash a TPM1.2 PCR.  This
        must be one of those supported by a platform.";
    }
  }

  grouping tpml2-hash-algo {
    description
      "The cryptographic algorithm used to hash the TPM1.2 PCRs.";
    leaf tpml2-hash-algo {
      type identityref {
        base taa:hash;
      }
      must '/tpm:rats-support-structures/tpm:attester-supported-algos'
        + '/tpm:tpml2-hash' {
        error-message "This platform does not support tpml2-hash-algo";
      }
      default "taa:TPM_ALG_SHA1";
      description
        "The hash scheme that is used to hash a TPM1.2 PCR.  This
        MUST be one of those supported by a platform.  This assumes
```



```
        that an algorithm other than SHA1 can be supported on some
        TPM1.2 cryptoprocessor variant.";
    }
}

grouping nonce {
    description
        "A random number intended to be used once to show freshness
        and to allow the detection of replay attacks.";
    leaf nonce-value {
        type binary;
        mandatory true;
        description
            "A cryptographically generated random number which should
            not be predictable prior to its issuance from a random
            number generation function. The random number MUST be
            derived from an entropy source external to the Attester.";
    }
}

grouping tpm12-pcr-selection {
    description
        "A Verifier can request one or more PCR values using its
        individually created Attestation Key Certificate (AC).
        The corresponding selection filter is represented in this
        grouping.
        Requesting a PCR value that is not in scope of the AC used,
        detailed exposure via error msg should be avoided.";
    leaf-list pcr-index {
        type pcr;
        must '/tpm:rats-support-structures/tpm:tpms'
            + '/tpm:tpm[name = current()] and '
            + '/tpm:rats-support-structures/tpm:tpms'
            + '/tpm:tpm[tpm12-pcrs = current()]' {
            error-message "Acquiring this PCR index is not supported";
        }
        description
            "The numbers/indexes of the PCRs. At the moment this is limited
            to 32.";
    }
}

grouping tpm20-pcr-selection {
    description
        "A Verifier can acquire one or more PCR values, which are hashed
        together in a TPM2B_DIGEST coming from the TPM2. The selection
        list of desired PCRs and the Hash Algorithm is represented in
        this grouping.";
```

```
list tpm20-pcr-selection {
  unique "tpm20-hash-algo";
  description
    "Specifies the list of PCRs and Hash Algorithms that can be
    returned within a TPM2B_DIGEST.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.9.7";
  uses tpm20-hash-algo;
  leaf-list pcr-index {
    type pcr;
    must '/tpm:rats-support-structures/tpm:tpms'
      + '/tpm:tpm[name = current()] and '
      + '/tpm:rats-support-structures/tpm:tpms/tpm:tpm'
      + '/tpm:tpm20-pcr-bank[pcr-index = current()]' {
      error-message "Acquiring this PCR index is not supported";
    }
    description
      "The numbers of the PCRs that which are being tracked
      with a hash based on the tpm20-hash-algo.";
  }
}

grouping certificate-name-ref {
  description
    "Identifies a certificate in a keystore.";
  leaf certificate-name {
    type certificate-name-ref;
    mandatory true;
    description
      "Identifies a certificate in a keystore.";
  }
}

grouping tpm-name {
  description
    "A unique TPM on a device.";
  leaf name {
    type string;
    description
      "Unique system generated name for a TPM on a device.";
  }
}

grouping tpm-name-selector {
  description
    "One or more TPM on a device.";
```

```
    leaf-list name {
      type string;
      config false;
      description
        "Name of one or more unique TPMs on a device.  If this object
        exists, a selection should pull only the objects related to
        these TPM(s).  If it does not exist, all qualifying TPMs that
        are 'hardware-based' equals true on the device are selected.";
    }
  }

  grouping node-uptime {
    description
      "Uptime in seconds of the node.";
    leaf up-time {
      type uint32;
      description
        "Uptime in seconds of this node reporting its data";
    }
  }

  grouping tpm12-attestation {
    description
      "Contains an instance of TPM1.2 style signed cryptoprocessor
      measurements.  It is supplemented by unsigned Attester
      information.";
    uses node-uptime;
    leaf TPM_QUOTE2 {
      type binary;
      description
        "Result of a TPM1.2 Quote2 operation.  This includes PCRs,
        signatures, locality, the provided nonce and other data which
        can be further parsed to appraise the Attester.";
      reference
        "TPM1.2 commands rev116 July 2007, Section 16.5";
    }
  }

  grouping tpm20-attestation {
    description
      "Contains an instance of TPM2 style signed cryptoprocessor
      measurements.  It is supplemented by unsigned Attester
      information.";
    leaf TPMS_QUOTE_INFO {
      type binary;
      mandatory true;
      description
        "A hash of the latest PCR values (and the hash algorithm used)";
    }
  }
}
```

```
        which have been returned from a Verifier for the selected PCRs
        and Hash Algorithms.";
    reference
        "https://www.trustedcomputinggroup.org/wp-content/uploads/
        TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.12.1";
}
leaf quote-signature {
    type binary;
    description
        "Quote signature returned by TPM Quote. The signature was
        generated using the key associated with the
        certificate 'name'.";
    reference
        "https://www.trustedcomputinggroup.org/wp-content/uploads/
        TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 11.2.1";
}
uses node-uptime;
list unsigned-pcr-values {
    description
        "PCR values in each PCR bank. This might appear redundant with
        the TPM2B_DIGEST, but that digest is calculated across multiple
        PCRs. Having to verify across multiple PCRs does not
        necessarily make it easy for a Verifier to appraise just the
        minimum set of PCR information which has changed since the last
        received TPM2B_DIGEST. Put another way, why should a Verifier
        reconstruct the proper value of all PCR Quotes when only a
        single PCR has changed?
        To help this happen, if the Attester does know specific PCR
        values, the Attester can provide these individual values via
        'unsigned-pcr-values'. By comparing this information to the
        what has previously been validated, it is possible for a
        Verifier to confirm the Attester's signature while eliminating
        significant processing.";
    uses tpm20-hash-algo;
    list pcr-values {
        key "pcr-index";
        description
            "List of one PCR bank.";
        leaf pcr-index {
            type pcr;
            description
                "PCR index number.";
        }
        leaf pcr-value {
            type binary;
            description
                "PCR value.";
            reference
```

```
        "https://www.trustedcomputinggroup.org/wp-content/uploads/
        TPM-Rev-2.0-Part-2-Structures-01.38.pdf  Section 10.9.7";
    }
}
}

grouping log-identifier {
  description
    "Identifier for type of log to be retrieved.";
  leaf log-type {
    type identityref {
      base attested_event_log_type;
    }
    mandatory true;
    description
      "The corresponding measurement log type identity.";
  }
}

grouping boot-event-log {
  description
    "Defines an event log corresponding to the event that extended
    the PCR";
  leaf event-number {
    type uint32;
    description
      "Unique event number of this event";
  }
  leaf event-type {
    type uint32;
    description
      "log event type";
  }
  leaf pcr-index {
    type pcr;
    description
      "Defines the PCR index that this event extended";
  }
  list digest-list {
    description
      "Hash of event data";
    leaf hash-algo {
      type identityref {
        base taa:hash;
      }
      description
        "The hash scheme that is used to compress the event data in
        each of the leaf-list digest items.";
    }
  }
}
```

```
    }
    leaf-list digest {
      type binary;
      description
        "The hash of the event data using the algorithm of the
        'hash-algo' against 'event data'.";
    }
  }
  leaf event-size {
    type uint32;
    description
      "Size of the event data";
  }
  leaf-list event-data {
    type uint8;
    description
      "The event data size determined by event-size";
  }
}
grouping bios-event-log {
  description
    "Measurement log created by the BIOS/UEFI.";
  list bios-event-entry {
    key event-number;
    description
      "Ordered list of TCG described event log
      that extended the PCRs in the order they
      were logged";
    uses boot-event-log;
  }
}
grouping ima-event {
  description
    "Defines an hash log extend event for IMA measurements";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TCG_IWG_CEL_v1_r0p30_13feb2021.pdf Section 4.3";
  leaf event-number {
    type uint64;
    description
      "Unique number for this event for sequencing";
  }
  leaf ima-template {
    type string;
    description
      "Name of the template used for event logs
      for e.g. ima, ima-ng, ima-sig";
  }
}
```

```
    leaf filename-hint {
      type string;
      description
        "File that was measured";
    }
    leaf filedata-hash {
      type binary;
      description
        "Hash of filedata";
    }
    leaf filedata-hash-algorithm {
      type string;
      description
        "Algorithm used for filedata-hash";
    }
    leaf template-hash-algorithm {
      type string;
      description
        "Algorithm used for template-hash";
    }
    leaf template-hash {
      type binary;
      description
        "hash(filedata-hash, filename-hint)";
    }
    leaf pcr-index {
      type pcr;
      description
        "Defines the PCR index that this event extended";
    }
    leaf signature {
      type binary;
      description
        "The file signature";
    }
  }
  grouping ima-event-log {
    description
      "Measurement log created by IMA.";
    list ima-event-entry {
      key event-number;
      description
        "Ordered list of ima event logs by event-number";
      uses ima-event;
    }
  }
  grouping network-equipment-boot-event-log {
```

```
description
  "Measurement log created by Network Equipment Boot. The Network
  Equipment Boot format is identical to the IMA format. In
  contrast to the IMA log, the Network Equipment Boot log
  includes every measurable event from an Attester, including
  the boot stages of BIOS, Bootloader, etc. In essence, the scope
  of events represented in this format combines the scope of BIOS
  events and IMA events.";
list boot-event-entry {
  key event-number;
  description
    "Ordered list of Network Equipment Boot event logs
    by event-number, using the IMA event format.";
  uses ima-event;
}
}
grouping event-logs {
  description
    "A selector for the log and its type.";
  choice attested_event_log_type {
    mandatory true;
    description
      "Event log type determines the event logs content.";
    case bios {
      if-feature "bios";
      description
        "BIOS/UEFI event logs";
      container bios-event-logs {
        description
          "BIOS/UEFI event logs";
        uses bios-event-log;
      }
    }
    case ima {
      if-feature "ima";
      description
        "IMA event logs.";
      container ima-event-logs {
        description
          "IMA event logs.";
        uses ima-event-log;
      }
    }
  }
  case netequip_boot {
    if-feature "netequip_boot";
    description
      "Network Equipment Boot event logs";
    container boot-event-logs {
```



```

        description
            "Network equipment boot event logs.";
        uses network-equipment-boot-event-log;
    }
}
}

/*****
/*   RPC operations   */
*****/

rpc tpm12-challenge-response-attestation {
    if-feature "taa:tpm12";
    description
        "This RPC accepts the input for TSS TPM 1.2 commands made to the
        attesting device.";
    input {
        container tpm12-attestation-challenge {
            description
                "This container includes every information element defined
                in the reference challenge-response interaction model for
                remote attestation. Corresponding values are based on
                TPM 1.2 structure definitions";
            uses tpm12-pcr-selection;
            uses nonce;
            leaf-list certificate-name {
                if-feature "tpm:tpms";
                type certificate-name-ref;
                must "/tpm:rats-support-structures/tpm:tpms"
                    + "/tpm:tpm[tpm:firmware-version='taa:tpm12']"
                    + "/tpm:certificates/"
                    + "/tpm:certificate[name=current()]" {
                    error-message "Not an available TPM1.2 AIK certificate.";
                }
            }
            description
                "When populated, the RPC will only get a Quote for the
                TPMs associated with these certificate(s).";
        }
    }
    output {
        list tpm12-attestation-response {
            unique "certificate-name";
            description
                "The binary output of TPM 1.2 TPM_Quote/TPM_Quote2, including
                the PCR selection and other associated attestation evidence
                metadata";
        }
    }
}

```

```

        uses certificate-name-ref {
            description
                "Certificate associated with this tpm12-attestation.";
        }
        uses tpm12-attestation;
    }
}

rpc tpm20-challenge-response-attestation {
    if-feature "taa:tpm20";
    description
        "This RPC accepts the input for TSS TPM 2.0 commands of the
        managed device. ComponentIndex from the hardware manager YANG
        module to refer to dedicated TPM in composite devices,
        e.g. smart NICs, is still a TODO.";
    input {
        container tpm20-attestation-challenge {
            description
                "This container includes every information element defined
                in the reference challenge-response interaction model for
                remote attestation. Corresponding values are based on
                TPM 2.0 structure definitions";
            uses nonce;
            uses tpm20-pcr-selection;
            leaf-list certificate-name {
                if-feature "tpm:tpms";
                type certificate-name-ref;
                must "/tpm:rats-support-structures/tpm:tpms"
                    + "/tpm:tpm[tpm:firmware-version='taa:tpm20']"
                    + "/tpm:certificates/"
                    + "/tpm:certificate[name=current()]" {
                    error-message "Not an available TPM2.0 AIK certificate.";
                }
            }
            description
                "When populated, the RPC will only get a Quote for the
                TPMs associated with the certificates.";
        }
    }
}

output {
    list tpm20-attestation-response {
        unique "certificate-name";
        description
            "The binary output of TPM2b_Quote in one TPM chip of the
            node which identified by node-id. An TPMS_ATTEST structure
            including a length, encapsulated in a signature";
        uses certificate-name-ref {

```

```
        description
            "Certificate associated with this tpm20-attestation.";
    }
    uses tpm20-attestation;
}
}
}

rpc log-retrieval {
    description
        "Logs Entries are either identified via indices or via providing
        the last line received. The number of lines returned can be
        limited. The type of log is a choice that can be augmented.";
    input {
        list log-selector {
            description
                "Selection of log entries to be reported.";
            uses tpm-name-selector;
            choice index-type {
                description
                    "Last log entry received, log index number, or timestamp.";
                case last-entry {
                    description
                        "The last entry of the log already retrieved.";
                    leaf last-entry-value {
                        type binary;
                        description
                            "Content of an log event which matches 1:1 with a
                            unique event record contained within the log. Log
                            entries subsequent to this will be passed to the
                            requester. Note: if log entry values are not unique,
                            this MUST return an error.";
                    }
                }
                case index {
                    description
                        "Numeric index of the last log entry retrieved, or
                        zero.";
                    leaf last-index-number {
                        type uint64;
                        description
                            "The last numeric index number of a log entry.
                            Zero means to start at the beginning of the log.
                            Entries subsequent to this will be passed to the
                            requester.";
                    }
                }
            }
        }
        case timestamp {
```

```

        leaf timestamp {
            type yang:date-and-time;
            description
                "Timestamp from which to start the extraction. The
                 next log entry subsequent to this timestamp is to
                 be sent.";
        }
        description
            "Timestamp from which to start the extraction.";
    }
}
leaf log-entry-quantity {
    type uint16;
    description
        "The number of log entries to be returned. If omitted, it
         means all of them.";
}
}
uses log-identifier;
}
output {
    container system-event-logs {
        description
            "The requested data of the measurement event logs";
        list node-data {
            unique "name";
            description
                "Event logs of a node in a distributed system
                 identified by the node name";
            uses tpm-name;
            uses node-uptime;
            container log-result {
                description
                    "The requested entries of the corresponding log.";
                uses event-logs;
            }
        }
    }
}
}

/*****/
/*    Config & Oper accessible nodes    */
/*****/

container rats-support-structures {
    description
        "The datastore definition enabling verifiers or relying

```

```
    parties to discover the information necessary to use the
    remote attestation RPCs appropriately.";
container compute-nodes {
  if-feature "tpm:tpms";
  description
    "Holds the set device subsystems/components in this composite
    device that support TPM operations.";
  list compute-node {
    key "node-id";
    config false;
    min-elements 2;
    description
      "A component within this composite device which
      supports TPM operations.";
    leaf node-id {
      type string;
      description
        "ID of the compute node, such as Board Serial Number.";
    }
    leaf node-physical-index {
      if-feature "ietfhw:entity-mib";
      type int32 {
        range "1..2147483647";
      }
      config false;
      description
        "The entPhysicalIndex for the compute node.";
      reference
        "RFC 6933: Entity MIB (Version 4) - entPhysicalIndex";
    }
    leaf node-name {
      type string;
      description
        "Name of the compute node.";
    }
    leaf node-location {
      type string;
      description
        "Location of the compute node, such as slot number.";
    }
  }
}
container tpms {
  description
    "Holds the set of TPMs within an Attester.";
  list tpm {
    key "name";
    unique "path";
```

```
description
  "A list of TPMs in this composite device that RATS
   can be conducted with.";
uses tpm-name;
leaf hardware-based {
  type boolean;
  config false;
  description
    "Answers the question: is this TPM is a hardware based
     TPM?";
}
leaf physical-index {
  if-feature "ietfhw:entity-mib";
  type int32 {
    range "1..2147483647";
  }
  config false;
  description
    "The entPhysicalIndex for the TPM.";
  reference
    "RFC 6933: Entity MIB (Version 4) - entPhysicalIndex";
}
leaf path {
  type string;
  config false;
  description
    "Path to a unique TPM on a device. This can change across
     reboots.";
}
leaf compute-node {
  if-feature "tpm:tpms";
  type compute-node-ref;
  config false;
  mandatory true;
  description
    "Indicates the compute node measured by this TPM.";
}
leaf manufacturer {
  type string;
  config false;
  description
    "TPM manufacturer name.";
}
leaf firmware-version {
  type identityref {
    base taa:cryptoprocessor;
  }
  mandatory true;
}
```

```
description
  "Identifies the cryptoprocessor API set supported.  This
   is automatically configured by the device and should not
   be changed.";
}
uses tpm12-hash-algo {
  when "firmware-version = 'taa:tpm12'";
  refine "tpm12-hash-algo" {
    description
      "The hash algorithm overwrites the default used for PCRs
       on this TPM1.2 compliant cryptoprocessor.";
  }
}
leaf-list tpm12-pcrs {
  when "../firmware-version = 'taa:tpm12'";
  type pcr;
  description
    "The PCRs which may be extracted from this TPM1.2
     compliant cryptoprocessor.";
}
list tpm20-pcr-bank {
  when "../firmware-version = 'taa:tpm20'";
  key "tpm20-hash-algo";
  description
    "Specifies the list of PCRs that may be extracted for
     a specific Hash Algorithm on this TPM2 compliant
     cryptoprocessor.  A bank is a set of PCRs which are
     extended using a particular hash algorithm.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
     TPM-Rev-2.0-Part-2-Structures-01.38.pdf  Section 10.9.7";
  leaf tpm20-hash-algo {
    type identityref {
      base taa:hash;
    }
    must '/tpm:rats-support-structures'
      + '/tpm:attester-supported-algos'
      + '/tpm:tpm20-hash' {
      error-message
        "This platform does not support tpm20-hash-algo";
    }
    description
      "The hash scheme actively being used to hash a
       one or more TPM2.0 PCRs.";
  }
}
leaf-list pcr-index {
  type tpm:pcr;
  description
```

```
        "Defines what TPM2 PCRs are available to be extracted.";
    }
}
leaf status {
    type enumeration {
        enum operational {
            value 0;
            description
                "The TPM currently is currently running normally and
                is ready to accept and process TPM quotes.";
            reference
                "TPM-Rev-2.0-Part-1-Architecture-01.07-2014-03-13.pdf
                Section 12";
        }
        enum non-operational {
            value 1;
            description
                "TPM is in a state such as startup or shutdown which
                precludes the processing of TPM quotes.";
        }
    }
    config false;
    mandatory true;
    description
        "TPM chip self-test status.";
}
container certificates {
    description
        "The TPM's certificates, including EK certificates
        and AK certificates.";
    list certificate {
        key "name";
        description
            "Three types of certificates can be accessed via
            this statement, including Initial Attestation
            Key Certificate, Local Attestation Key Certificate or
            Endorsement Key Certificate.";
        leaf name {
            type string;
            description
                "An arbitrary name uniquely identifying a certificate
                associated within key within a TPM.";
        }
        leaf keystore-ref {
            type leafref {
                path "/ks:keystore/ks:asymmetric-keys/ks:asymmetric-key"
                + "/ks:certificates/ks:certificate/ks:name";
            }
        }
    }
}
```



```

        description
            "A reference to a specific certificate of an
              asymmetric key in the Keystore.";
    }
    leaf type {
        type enumeration {
            enum endorsement-certificate {
                value 0;
                description
                    "Endorsement Key (EK) Certificate type.";
                reference
                    "https://trustedcomputinggroup.org/wp-content/
                      uploads/TCG_IWG_DevID_v1r2_02dec2020.pdf
                      Section 3.11";
            }
            enum initial-attestation-certificate {
                value 1;
                description
                    "Initial Attestation key (IAK) Certificate type.";
                reference
                    "https://trustedcomputinggroup.org/wp-content/
                      uploads/TCG_IWG_DevID_v1r2_02dec2020.pdf
                      Section 3.2";
            }
            enum local-attestation-certificate {
                value 2;
                description
                    "Local Attestation Key (LAK) Certificate type.";
                reference
                    "https://trustedcomputinggroup.org/wp-content/
                      uploads/TCG_IWG_DevID_v1r2_02dec2020.pdf
                      Section 3.2";
            }
        }
        description
            "Function supported by this certificate from within the
              TPM.";
    }
}
}
}
}
container attester-supported-algos {
    description
        "Identifies which TPM algorithms are available for use on an
          attesting platform.";
    leaf-list tpml2-asymmetric-signing {
        when "../..//tpm:tpms"

```

```

        + "/tpm:tpm[tpm:firmware-version='taa:tpm12']";
    type identityref {
        base taa:asymmetric;
    }
    description
        "Platform Supported TPM12 asymmetric algorithms.";
}
leaf-list tpm12-hash {
    when "../../tpm:tpms"
        + "/tpm:tpm[tpm:firmware-version='taa:tpm12']";
    type identityref {
        base taa:hash;
    }
    description
        "Platform supported TPM12 hash algorithms.";
}
leaf-list tpm20-asymmetric-signing {
    when "../../tpm:tpms"
        + "/tpm:tpm[tpm:firmware-version='taa:tpm20']";
    type identityref {
        base taa:asymmetric;
    }
    description
        "Platform Supported TPM20 asymmetric algorithms.";
}
leaf-list tpm20-hash {
    when "../../tpm:tpms"
        + "/tpm:tpm[tpm:firmware-version='taa:tpm20']";
    type identityref {
        base taa:hash;
    }
    description
        "Platform supported TPM20 hash algorithms.";
}
}
}
}
<CODE ENDS>

```

Figure 1

2.1.2. 'ietf-tcg-algs'

Cryptographic algorithm types were initially included within -v14 NETCONF's iana-crypto-types.yang. Unfortunately, all this content including the algorithms needed here failed to make the -v15 used WGLC. As a result, this document has encoded the TCG Algorithm definitions of [TCG-Algos], revision 1.32. By including this full table as a separate YANG file within this document, it is possible for other YANG models to leverage the contents of this model.

2.1.2.1. Features

There are two types of features supported: 'TPM12' and 'TPM20'. Support for either of these features indicates that a cryptoprocessor supporting the corresponding type of TCG TPM API is present on an Attester. Most commonly, only one type of cryptoprocessor will be available on an Attester.

2.1.2.2. Identities

There are three types of identities in this model:

1. *Cryptographic functions* supported by a TPM algorithm; these include: 'asymmetric', 'symmetric', 'hash', 'signing', 'anonymous_signing', 'encryption_mode', 'method', and 'object_type'. The definitions of each of these are in Table 2 of [TCG-Algos].
2. *API specifications* for TPMs: 'tpm12' and 'tpm20'
3. *Specific algorithm types*: Each algorithm type defines what cryptographic functions may be supported, and on which type of API specification. It is not required that an implementation of a specific TPM will support all algorithm types. The contents of each specific algorithm mirrors what is in Table 3 of [TCG-Algos].

2.1.2.3. YANG Module

```
<CODE BEGINS> file "ietf-tcg-algs@2021-05-11.yang"
module ietf-tcg-algs {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-tcg-algs";
  prefix taa;

  organization
    "IETF RATS Working Group";
```

contact

"WG Web: <<http://datatracker.ietf.org/wg/rats/>>
WG List: <<mailto:rats@ietf.org>>
Author: Eric Voit <<mailto:evoit@cisco.com>>";

description

"This module defines a identities for asymmetric algorithms.

Copyright (c) 2021 IETF Trust and the persons identified
as authors of the code. All rights reserved.
Redistribution and use in source and binary forms, with
or without modification, is permitted pursuant to, and
subject to the license terms contained in, the Simplified
BSD License set forth in Section 4.c of the IETF Trust's
Legal Provisions Relating to IETF Documents
(<https://trustee.ietf.org/license-info>).
This version of this YANG module is part of RFC XXXX
(<https://www.rfc-editor.org/info/rfcXXXX>); see the RFC
itself for full legal notices.
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 (RFC 2119)
(RFC 8174) when, and only when, they appear in all
capitals, as shown here.";

```
revision 2021-05-11 {  
  description  
    "Initial version";  
  reference  
    "RFC XXXX: tbd";  
}
```

```
/*  
*****  
/* Features */  
*****  
*/
```

```
feature tpm12 {  
  description  
    "This feature indicates algorithm support for the TPM 1.2 API  
    as per TPM-main-1.2-Rev94-part-2, Section 4.8.";  
}
```

```
feature tpm20 {  
  description  
    "This feature indicates algorithm support for the TPM 2.0 API  
    as per TPM-Rev-2.0-Part-1-Architecture-01.38 Section 11.4.";  
}
```

```

/*****/
/* Identities */
/*****/

/* There needs to be collapsing/verification of some of the identity
   types between the various algorithm types listed below */

identity asymmetric {
  description
    "A TCG recognized asymmetric algorithm with a public and
    private key.";
  reference
    "http://trustedcomputinggroup.org/resource/tcg-algorithm-registry/
    TCG_Algorithm_Registry_rlp32_pub Table 2";
}

identity symmetric {
  description
    "A TCG recognized symmetric algorithm with only a private key.";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 2";
}

identity hash {
  description
    "A TCG recognized hash algorithm that compresses input data to
    a digest value or indicates a method that uses a hash.";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 2";
}

identity signing {
  description
    "A TCG recognized signing algorithm";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 2";
}

identity anonymous_signing {
  description
    "A TCG recognized anonymous signing algorithm.";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 2";
}

identity encryption_mode {
  description
    "A TCG recognized encryption mode.";
}
```

```
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 2";
  }

  identity method {
    description
      "A TCG recognized method such as a mask generation function.";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 2";
  }

  identity object_type {
    description
      "A TCG recognized object type.";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 2";
  }

  identity cryptoprocessor {
    description
      "Base identity identifying a cryptoprocessor.";
  }

  identity tpm12 {
    if-feature "tpm12";
    base cryptoprocessor;
    description
      "Supportable by a TPM1.2.";
    reference
      "TPM-Main-Part-2-TPM-Structures_v1.2_rev116_01032011.pdf
       TPM_ALGORITHM_ID values, page 18";
  }

  identity tpm20 {
    if-feature "tpm12";
    base cryptoprocessor;
    description
      "Supportable by a TPM2.";
    reference
      "TPM-Rev-2.0-Part-2-Structures-01.38.pdf
       The TCG Algorithm Registry. Table 9";
  }

  identity TPM_ALG_RSA {
    if-feature "tpm12 or tpm20";
    base tpm12;
    base tpm20;
    base asymmetric;
```

```
    base object_type;
    description
      "RSA algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      RFC 8017. ALG_ID: 0x0001";
  }

  identity TPM_ALG_TDES {
    if-feature "tpm12";
    base tpm12;
    base symmetric;
    description
      "Block cipher with various key sizes (Triple Data Encryption
      Algorithm, commonly called Triple Data Encryption Standard)
      Note: was banned in TPM1.2 v94";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      ISO/IEC 18033-3. ALG_ID: 0x0003";
  }

  identity TPM_ALG_SHA1 {
    if-feature "tpm12 or tpm20";
    base hash;
    base tpm12;
    base tpm20;
    description
      "SHA1 algorithm - Deprecated due to insufficient cryptographic
      protection. However it is still useful for hash algorithms
      where protection is not required.";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      ISO/IEC 10118-3. ALG_ID: 0x0004";
  }

  identity TPM_ALG_HMAC {
    if-feature "tpm12 or tpm20";
    base tpm12;
    base tpm20;
    base hash;
    base signing;
    description
      "Hash Message Authentication Code (HMAC) algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3,
      ISO/IEC 9797-2 and RFC2014. ALG_ID: 0x0005";
  }
```

```
identity TPM_ALG_AES {
  if-feature "tpm12";
  base tpm12;
  base symmetric;
  description
    "The AES algorithm with various key sizes";
  reference
    "TCG Algorithm_Registry_r1p32_pub Table 3 and
    ISO/IEC 18033-3. ALG_ID: 0x0006";
}

identity TPM_ALG_MGF1 {
  if-feature "tpm20";
  base tpm20;
  base hash;
  base method;
  description
    "hash-based mask-generation function";
  reference
    "TCG Algorithm_Registry_r1p32_pub Table 3,
    IEEE Std 1363-2000 and IEEE Std 1363a -2004.
    ALG_ID: 0x0007";
}

identity TPM_ALG_KEYEDHASH {
  if-feature "tpm20";
  base tpm20;
  base hash;
  base object_type;
  description
    "An encryption or signing algorithm using a keyed hash. These
    may use XOR for encryption or an HMAC for signing and may
    also refer to a data object that is neither signing nor
    encrypting.";
  reference
    "TCG Algorithm_Registry_r1p32_pub Table 3 and
    TCG TPM 2.0 library specification. . ALG_ID: 0x0008";
}

identity TPM_ALG_XOR {
  if-feature "tpm12 or tpm20";
  base tpm12;
  base tpm20;
  base hash;
  base symmetric;
  description
    "The XOR encryption algorithm.";
  reference
```



```
    "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      TCG TPM 2.0 library specification. ALG_ID: 0x000A";
  }

  identity TPM_ALG_SHA256 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 256 algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        ISO/IEC 10118-3. ALG_ID: 0x000B";
  }

  identity TPM_ALG_SHA384 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 384 algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        ISO/IEC 10118-3. ALG_ID: 0x000C";
  }

  identity TPM_ALG_SHA512 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 512 algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        ISO/IEC 10118-3. ALG_ID: 0x000D";
  }

  identity TPM_ALG_NULL {
    if-feature "tpm20";
    base tpm20;
    description
      "NULL algorithm";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        TCG TPM 2.0 library specification. ALG_ID: 0x0010";
  }

  identity TPM_ALG_SM3_256 {
```

```
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
        "The SM3 hash algorithm.";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        GM/T 0004-2012 - SM3_256. ALG_ID: 0x0012";
}

identity TPM_ALG_SM4 {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    description
        "SM4 symmetric block cipher";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        GB/T 32907-2016. ALG_ID: 0x0013";
}

identity TPM_ALG_RSASSA {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
        "Signature algorithm defined in section 8.2 (RSASSAPKCS1-v1_5)";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and RFC 8017.
        ALG_ID: 0x0014";
}

identity TPM_ALG_RSAES {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base encryption_mode;
    description
        "Signature algorithm defined in section 7.2 (RSAES-PKCS1-v1_5)";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and RFC 8017
        ALG_ID: 0x0015";
}

identity TPM_ALG_RSAPSS {
    if-feature "tpm20";
    base tpm20;
```

```
    base asymmetric;
    base signing;
    description
      "Padding algorithm defined in section 8.1 (RSASSA PSS)";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and RFC 8017.
      ALG_ID: 0x0016";
  }

  identity TPM_ALG_OAEP {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base encryption_mode;
    description
      "Padding algorithm defined in section 7.1 (RSASSA OAEP)";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and RFC 8017.
      ALG_ID: 0x0017";
  }

  identity TPM_ALG_ECDSA {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
      "Signature algorithm using elliptic curve cryptography (ECC)";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      ISO/IEC 14888-3. ALG_ID: 0x0018";
  }

  identity TPM_ALG_ECDH {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base method;
    description
      "Secret sharing using ECC";
    reference
      "TCG_Algorithm_Registry_rlp32_pub Table 3 and
      NIST SP800-56A and RFC 7748. ALG_ID: 0x0019";
  }

  identity TPM_ALG_ECDAA {
    if-feature "tpm20";
    base tpm20;
```

```
    base asymmetric;
    base signing;
    base anonymous_signing;
    description
        "Elliptic-curve based anonymous signing scheme";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        TCG TPM 2.0 library specification. ALG_ID: 0x001A";
}

identity TPM_ALG_SM2 {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    base encryption_mode;
    base method;
    description
        "SM2 - depending on context, either an elliptic-curve based,
        signature algorithm, an encryption scheme, or a key exchange
        protocol";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        A GM/T 0003.1-2012, GM/T 0003.2-2012, GM/T 0003.3-2012,
        GM/T 0003.5-2012 SM2. ALG_ID: 0x001B";
}

identity TPM_ALG_ECSCHNORR {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
        "Elliptic-curve based Schnorr signature";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        TCG TPM 2.0 library specification. ALG_ID: 0x001C";
}

identity TPM_ALG_ECMQV {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base method;
    description
        "Two-phase elliptic-curve key";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
```

```
        NIST SP800-56A. ALG_ID: 0x001D";
    }

    identity TPM_ALG_KDF1_SP800_56A {
        if-feature "tpm20";
        base tpm20;
        base hash;
        base method;
        description
            "Concatenation key derivation function";
        reference
            "TCG_Algorithm_Registry_rlp32_pub Table 3 and
            NIST SP800-56A (approved alternative1) section 5.8.1.
            ALG_ID: 0x0020";
    }

    identity TPM_ALG_KDF2 {
        if-feature "tpm20";
        base tpm20;
        base hash;
        base method;
        description
            "Key derivation function";
        reference
            "TCG_Algorithm_Registry_rlp32_pub Table 3 and
            IEEE 1363a-2004 KDF2 section 13.2. ALG_ID: 0x0021";
    }

    identity TPM_ALG_KDF1_SP800_108 {
        base TPM_ALG_KDF2;
        description
            "A key derivation method";
        reference
            "TCG_Algorithm_Registry_rlp32_pub Table 3 and
            NIST SP800-108 - Section 5.1 KDF. ALG_ID: 0x0022";
    }

    identity TPM_ALG_ECC {
        if-feature "tpm20";
        base tpm20;
        base asymmetric;
        base object_type;
        description
            "Prime field ECC";
        reference
            "TCG_Algorithm_Registry_rlp32_pub Table 3 and
            ISO/IEC 15946-1. ALG_ID: 0x0023";
    }
```

```
identity TPM_ALG_SYMCIPHER {
  if-feature "tpm20";
  base tpm20;
  description
    "Object type for a symmetric block cipher";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 3 and
    TCG TPM 2.0 library specification. ALG_ID: 0x0025";
}

identity TPM_ALG_CAMELLIA {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  description
    "The Camellia algorithm";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 3 and
    ISO/IEC 18033-3. ALG_ID: 0x0026";
}

identity TPM_ALG_SHA3_256 {
  if-feature "tpm20";
  base tpm20;
  base hash;
  description
    "ISO/IEC 10118-3 - the SHA 256 algorithm";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 3 and
    NIST PUB FIPS 202. ALG_ID: 0x0027";
}

identity TPM_ALG_SHA3_384 {
  if-feature "tpm20";
  base tpm20;
  base hash;
  description
    "The SHA 384 algorithm";
  reference
    "TCG_Algorithm_Registry_rlp32_pub Table 3 and
    NIST PUB FIPS 202. ALG_ID: 0x0028";
}

identity TPM_ALG_SHA3_512 {
  if-feature "tpm20";
  base tpm20;
  base hash;
  description
```

```
    "The SHA 512 algorithm";
  reference
    "TCG_Algorithm_Registry_r1p32_pub Table 3 and
    NIST PUB FIPS 202. ALG_ID: 0x0029";
}

identity TPM_ALG_CMAC {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  base signing;
  description
    "block Cipher-based Message Authentication Code (CMAC)";
  reference
    "TCG_Algorithm_Registry_r1p32_pub Table 3 and
    ISO/IEC 9797-1:2011 Algorithm 5. ALG_ID: 0x003F";
}

identity TPM_ALG_CTR {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  base encryption_mode;
  description
    "Counter mode";
  reference
    "TCG_Algorithm_Registry_r1p32_pub Table 3 and
    ISO/IEC 10116. ALG_ID: 0x0040";
}

identity TPM_ALG_OFB {
  base tpm20;
  base symmetric;
  base encryption_mode;
  description
    "Output Feedback mode";
  reference
    "TCG_Algorithm_Registry_r1p32_pub Table 3 and
    ISO/IEC 10116. ALG_ID: 0x0041";
}

identity TPM_ALG_CBC {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  base encryption_mode;
  description
    "Cipher Block Chaining mode";
```

```
    reference
      "TCG_Algorithm_Registry_r1p32_pub Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0042";
  }

  identity TPM_ALG_CFB {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Cipher Feedback mode";
    reference
      "TCG_Algorithm_Registry_r1p32_pub Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0043";
  }

  identity TPM_ALG_ECB {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Electronic Codebook mode";
    reference
      "TCG_Algorithm_Registry_r1p32_pub Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0044";
  }

  identity TPM_ALG_CCM {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    base encryption_mode;
    description
      "Counter with Cipher Block Chaining-Message Authentication
      Code (CCM)";
    reference
      "TCG_Algorithm_Registry_r1p32_pub Table 3 and
      NIST SP800-38C. ALG_ID: 0x0050";
  }

  identity TPM_ALG_GCM {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
```



```
    base encryption_mode;
    description
        "Galois/Counter Mode (GCM)";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        NIST SP800-38D. ALG_ID: 0x0051";
}

identity TPM_ALG_KW {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    base encryption_mode;
    description
        "AES Key Wrap (KW)";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        NIST SP800-38F. ALG_ID: 0x0052";
}

identity TPM_ALG_KWP {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    base encryption_mode;
    description
        "AES Key Wrap with Padding (KWP)";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        NIST SP800-38F. ALG_ID: 0x0053";
}

identity TPM_ALG_EAX {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    base encryption_mode;
    description
        "Authenticated-Encryption Mode";
    reference
        "TCG_Algorithm_Registry_rlp32_pub Table 3 and
        NIST SP800-38F. ALG_ID: 0x0054";
}

identity TPM_ALG_EDDSA {
```

```
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
      "Edwards-curve Digital Signature Algorithm (PureEdDSA)";
    reference
      "TCG_Algorithm_Registry_r1p32_pub Table 3 and
      RFC 8032. ALG_ID: 0x0060";
  }
}
<CODE ENDS>
```

Note that not all cryptographic functions are required for use by "ietf-tpm-remote-attestation.yang". However the full definition of Table 3 of [TCG-Algos] will allow use by additional YANG specifications.

3. IANA Considerations

This document will include requests to IANA:

To be defined yet. But keeping up with changes to "ietf-tcg-algs.yang" will be necessary.

4. Security Considerations

The YANG module specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., `_config true`, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., `_edit-config`) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes as well as their sensitivity/vulnerability:

Container `'/rats-support-structures/attester-supported-algos'`: `'tpm12-asymmetric-signing'`, `'tpm12-hash'`, `'tpm20-asymmetric-signing'`, and `'tpm20-hash'`. All could be populated with algorithms that are not supported by the underlying physical TPM installed by the equipment vendor.

Container: `'/rats-support-structures/tpms'`: `'name'`: Although shown as `'rw'`, it is system generated. Therefore it should not be possible for an operator to add or remove a TPM from the configuration.

`'tpm20-pcr-bank'`: It is possible to configure PCRs for extraction which are not being extended by system software. This could unnecessarily use TPM resources.

`'certificates'`: It is possible to provision a certificate which does not correspond to an Attestation Identity Key (AIK) within the TPM 1.2, or an Attestation Key (AK) within the TPM 2.0 respectively.

RPC `'tpm12-challenge-response-attestation'`: It must be verified that the certificate is for an active AIK, i. e. the certificate provided is able to support Attestation on the targeted TPM 1.2.

RPC `'tpm20-challenge-response-attestation'`: It must be verified that the certificate is for an active AK, i. e. the certificate provided is able to support Attestation on the targeted TPM 2.0.

RPC `'log-retrieval'`: Pulling lots of logs can chew up system resources.

5. Acknowledgements

Not yet.

6. Change Log

Changes from version 05 to version 06:

- * More YANG Dr comments covered

Changes from version 04 to version 05:

- * YANG Dr comments covered

Changes from version 03 to version 04:

- * TPM1.2 Quotel eliminated
- * YANG model simplifications so redundant info isn't exposed

Changes from version 02 to version 03:

- * moved to tcg-algs

- * cleaned up model to eliminate sources of errors
- * removed key establishment RPC
- * added lots of XPATH which must all be scrubbed still
- * Descriptive text added on model contents.

Changes from version 01 to version 02:

- * Extracted Crypto-types into a separate YANG file
- * Makes the algorithms explicit, not strings
- * Hash Algo as key the selected TPM2 PCRs
- * PCR numbers are their own type
- * Eliminated nested keys for node-id plus tpm-name
- * Eliminated TPM-Name of "ALL"
- * Added TPM-Path

Changes from version 00 to version 01:

- * Addressed author's comments
- * Extended complementary details about attestation-certificates
- * Relabeled chunk-size to log-entry-quantity
- * Relabeled location with compute-node or tpm-name where appropriate
- * Added a valid entity-mib physical-index to compute-node and tpm-name to map it back to hardware inventory
- * Relabeled name to tpm_name
- * Removed event-string in last-entry

7. References

7.1. Normative References

[I-D.ietf-netconf-keystore]

Watsen, K., "A YANG Data Model for a Keystore", Work in Progress, Internet-Draft, draft-ietf-netconf-keystore-22, 18 May 2021, <<https://www.ietf.org/archive/id/draft-ietf-netconf-keystore-22.txt>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-12, 23 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-12.txt>>.

[I-D.ietf-rats-reference-interaction-models]

Birkholz, H., Eckel, M., Pan, W., and E. Voit, "Reference Interaction Models for Remote Attestation Procedures", Work in Progress, Internet-Draft, draft-ietf-rats-reference-interaction-models-02, 25 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-reference-interaction-models-02.txt>>.

[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-06, 7 December 2020, <<https://www.ietf.org/archive/id/draft-ietf-rats-tpm-based-network-device-attest-06.txt>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8348] Bierman, A., Bjorklund, M., Dong, J., and D. Romascanu, "A YANG Data Model for Hardware Management", RFC 8348, DOI 10.17487/RFC8348, March 2018, <<https://www.rfc-editor.org/info/rfc8348>>.

- [TCG-Algos] "TCG_Algorithm_Registry_r1p32_pub", n.d.,
<<https://trustedcomputinggroup.org/resource/tcg-algorithm-registry/>>.
- [TPM1.2] TCG, ., "TPM 1.2 Main Specification", 2 October 2003,
<<https://trustedcomputinggroup.org/resource/tpm-main-specification/>>.
- [TPM2.0] TCG, ., "TPM 2.0 Library Specification", 15 March 2013,
<<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.
- [TPM2.0-Key] TCG, ., "TPM 2.0 Keys for Device Identity and Attestation, Rev10", 14 April 2021, <https://trustedcomputinggroup.org/wp-content/uploads/TCG_IWG_DevID_v1r2_02dec2020.pdf>.

7.2. Informative References

- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de

Michael Eckel
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: michael.eckel@sit.fraunhofer.de

Shwetha Bhandari
ThoughtSpot

Email: shwetha.bhandari@thoughtspot.com

Eric Voit
Cisco Systems

Email: evoit@cisco.com

Bill Sulzen
Cisco Systems

Email: bsulzen@cisco.com

Liang Xia (Frank)
Huawei Technologies
101 Software Avenue, Yuhuatai District
Nanjing
Jiangsu, 210012
China

Email: Frank.Xialiang@huawei.com

Tom Laffey
Hewlett Packard Enterprise

Email: tom.laffey@hpe.com

Guy C. Fedorkow
Juniper Networks
10 Technology Park Drive
Westford

Email: gfedorkow@juniper.net

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 17 October 2022

H. Birkholz
M. Eckel
Fraunhofer SIT
S. Bhandari
ThoughtSpot
E. Voit
B. Sulzen
Cisco
L. Xia
Huawei
T. Laffey
HPE
G. Fedorkow
Juniper
15 April 2022

A YANG Data Model for Challenge-Response-based Remote Attestation
Procedures using TPMs
draft-ietf-rats-yang-tpm-charra-19

Abstract

This document defines YANG RPCs and a few configuration nodes required to retrieve attestation evidence about integrity measurements from a device, following the operational context defined in TPM-based Network Device Remote Integrity Verification. Complementary measurement logs are also provided by the YANG RPCs, originating from one or more roots of trust for measurement (RTMs). The module defined requires at least one TPM 1.2 or TPM 2.0 as well as a corresponding TPM Software Stack (TSS), or equivalent hardware implementations that include the protected capabilities as provided by TPMs as well as a corresponding software stack, included in the device components of the composite device the YANG server is running on.

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 17 October 2022.

Copyright Notice

Copyright (c) 2022 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
1.1. Requirements notation	3
2. The YANG Module for Basic Remote Attestation Procedures	3
2.1. YANG Modules	3
2.1.1. 'ietf-tpm-remote-attestation'	4
2.1.2. 'ietf-tcg-algs'	33
3. IANA Considerations	48
4. Security Considerations	49
5. References	51
5.1. Normative References	51
5.2. Informative References	56
Appendix A. Integrity Measurement Architecture (IMA)	56
Appendix B. IMA for Network Equipment Boot Logs	57
Authors' Addresses	58

1. Introduction

This document is based on the general terminology defined in the [I-D.ietf-rats-architecture] and uses the operational context defined in [I-D.ietf-rats-tpm-based-network-device-attest] as well as the interaction model and information elements defined in [I-D.ietf-rats-reference-interaction-models]. The currently supported hardware security modules (HSMs) are the Trusted Platform Modules (TPMs) [TPM1.2] and [TPM2.0] as specified by the Trusted Computing Group (TCG). One TPM, or multiple TPMs in the case of a

Composite Device, are required in order to use the YANG module defined in this document. Each TPM is used as a root of trust for storage (RTS) in order to store system security measurement Evidence. And each TPM is used as a root of trust for reporting (RTR) in order to retrieve attestation Evidence. This is done by using a YANG RPC to request a quote which exposes a rolling hash of the security measurements held internally within the TPM.

Specific terms imported from [I-D.ietf-rats-architecture] and used in this document include: Attester, Composite Device, Evidence.

Specific terms imported from [TPM2.0-Key] and used in this document include: Endorsement Key (EK), Initial Attestation Key (IAK), Attestation Identity Key (AIK), Local Attestation Key (LAK).

1.1. Requirements notation

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. The YANG Module for Basic Remote Attestation Procedures

One or more TPMs MUST be embedded in a Composite Device that provides attestation evidence via the YANG module defined in this document. The ietf-tpm-remote-attestation YANG module enables a composite device to take on the role of an Attester, in accordance with the Remote Attestation Procedures (RATS) architecture [I-D.ietf-rats-architecture], and the corresponding challenge-response interaction model defined in the [I-D.ietf-rats-reference-interaction-models] document. A fresh nonce with an appropriate amount of entropy [NIST-915121] MUST be supplied by the YANG client in order to enable a proof-of-freshness with respect to the attestation Evidence provided by the Attester running the YANG datastore. Further, this nonce is used to prevent replay attacks. The method for communicating the relationship of each individual TPM to specific measured component within the Composite Device is out of the scope of this document.

2.1. YANG Modules

In this section the several YANG modules are defined.

2.1.1.1. 'ietf-tpm-remote-attestation'

This YANG module imports modules from [RFC6991] with prefix 'yang', [RFC8348] with prefix 'hw', [I-D.ietf-netconf-keystore] with prefix 'ks', and 'ietf-tcg-algs.yang' Section 2.1.2.3 with prefix 'taa'. Additionally, references are made to [RFC8032], [RFC8017], [RFC6933], [TPM1.2-Commands], [TPM2.0-Arch], [TPM2.0-Structures], [TPM2.0-Key], [TPM1.2-Structures], [bios-log], [BIOS-Log-Event-Type], as well as Appendix A and Appendix B.

2.1.1.1.1. Features

This module supports the following features:

- * 'mtpm': Indicates that multiple TPMs on the device can support remote attestation. For example, this feature could be used in cases where multiple line cards are present, each with its own TPM.
- * 'bios': Indicates that the device supports the retrieval of BIOS/UEFI event logs. [bios-log]
- * 'ima': Indicates that the device supports the retrieval of event logs from the Linux Integrity Measurement Architecture (IMA, see Appendix A).
- * 'netequip_boot': Indicates that the device supports the retrieval of netequip boot event logs. See Appendix A and Appendix B.

2.1.1.1.2. Identities

This module supports the following types of attestation event logs: 'bios', 'ima', and 'netequip_boot'.

2.1.1.1.3. Remote Procedure Calls (RPCs)

In the following, RPCs for both TPM 1.2 and TPM 2.0 attestation procedures are defined.

2.1.1.1.3.1. 'tpm12-challenge-response-attestation'

This RPC allows a Verifier to request signed TPM PCRs (_TPM Quote_ operation) from a TPM 1.2 compliant cryptoprocessor. Where the feature 'mtpm' is active, and one or more 'certificate-name' is not provided, all TPM 1.2 compliant cryptoprocessors will respond. A YANG tree diagram of this RPC is as follows:

```

+---x tpm12-challenge-response-attestation {taa:tpm12}?
+---w input
|   +---w tpm12-attestation-challenge
|       +---w pcr-index*          pcr
|       +---w nonce-value         binary
|       +---w certificate-name*    certificate-name-ref
|                                   {tpm:mtpm}?
+---ro output
+---ro tpm12-attestation-response* []
+---ro certificate-name            certificate-name-ref
+---ro up-time?                   uint32
+---ro TPM_QUOTE2?                binary

```

2.1.1.3.2. 'tpm20-challenge-response-attestation'

This RPC allows a Verifier to request signed TPM PCRs (`_TPM Quote_` operation) from a TPM 2.0 compliant cryptoprocessor. Where the feature 'mtpm' is active, and one or more 'certificate-name' is not provided, all TPM 2.0 compliant cryptoprocessors will respond. A YANG tree diagram of this RPC is as follows:

```

+---x tpm20-challenge-response-attestation {taa:tpm20}?
+---w input
|   +---w tpm20-attestation-challenge
|       +---w nonce-value         binary
|       +---w tpm20-pcr-selection* []
|           +---w tpm20-hash-algo? identityref
|           +---w pcr-index*      pcr
|       +---w certificate-name*    certificate-name-ref
|                                   {tpm:mtpm}?
+---ro output
+---ro tpm20-attestation-response* []
+---ro certificate-name            certificate-name-ref
+---ro TPMS_QUOTE_INFO            binary
+---ro quote-signature?          binary
+---ro up-time?                  uint32
+---ro unsigned-pcr-values* []
|   +---ro tpm20-hash-algo?      identityref
|   +---ro pcr-values* [pcr-index]
|       +---ro pcr-index         pcr
|       +---ro pcr-value?       binary

```

An example of an RPC challenge requesting PCRs 0-7 from a SHA-256 bank could look like the following:

```

<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <tpm20-challenge-response-attestation>
    xmlns="urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation">
      <certificate-name>
        (identifier of a TPM signature key with which the Verifier is
        supposed to sign the attestation data)
      </certificate-name>
      <nonce>
        0xe041307208d9f78f5b1bbecd19e2d152ad49de2fc5a7d8dbf769f6b8ffdeab9
      </nonce>
      <tpm20-pcr-selection>
        <tpm20-hash-algo
          xmlns="urn:ietf:params:xml:ns:yang:ietf-tcg-algs">
            TPM_ALG_SHA256
          </tpm20-hash-algo>
        <pcr-index>0</pcr-index>
        <pcr-index>1</pcr-index>
        <pcr-index>2</pcr-index>
        <pcr-index>3</pcr-index>
        <pcr-index>4</pcr-index>
        <pcr-index>5</pcr-index>
        <pcr-index>6</pcr-index>
        <pcr-index>7</pcr-index>
      </tpm20-pcr-selection>
    </tpm20-challenge-response-attestation>
  </rpc>

```

A successful response could be formatted as follows:

```

<rpc-reply message-id="101"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <tpm20-attestation-response
    xmlns="urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation">
    <certificate-name
      xmlns="urn:ietf:params:xml:ns:yang:ietf-keystore">
      (instance of Certificate name in the Keystore)
    </certificate-name>
    <attestation-data>
      (raw attestation data, i.e. the TPM quote; this includes
      a composite digest of requested PCRs, the nonce,
      and TPM 2.0 time information.)
    </attestation-data>
    <quote-signature>
      (signature over attestation-data using the TPM key
      identified by sig-key-id)
    </quote-signature>
  </tpm20-attestation-response>
</rpc-reply>

```

2.1.1.4. 'log-retrieval'

This RPC allows a Verifier to acquire the evidence which was extended into specific TPM PCRs. A YANG tree diagram of this RPC is as follows:

```

+---x log-retrieval
  +---w input
    +---w log-type          identityref
    +---w log-selector* []
      +---w name*           string
      +---w (index-type)?
        +---:(last-entry)
          +---w last-entry-value?  binary
        +---:(index)
          +---w last-index-number? uint64
        +---:(timestamp)
          +---w timestamp?         yang:date-and-time
    +---w log-entry-quantity?  uint16
  +---ro output
    +---ro system-event-logs
      +---ro node-data* []
        +---ro name?          string
        +---ro up-time?       uint32
        +---ro log-result
          +---ro (attested_event_log_type)
            +---:(bios) {bios}?
              +---ro bios-event-logs
                +---ro bios-event-entry* [event-number]
                  +---ro event-number    uint32
                  +---ro event-type?     uint32
                  +---ro pcr-index?      pcr
                  +---ro digest-list* []
                    +---ro hash-algo?    identityref
                    +---ro digest*       binary
                  +---ro event-size?     uint32
                  +---ro event-data*     binary
            +---:(ima) {ima}?
              +---ro ima-event-logs
                +---ro ima-event-entry* [event-number]
                  +---ro event-number    uint64
                  +---ro ima-template?   string
                  +---ro filename-hint?  string
                  +---ro filedata-hash?  binary
                  +---ro filedata-hash-algorithm? string
                  +---ro template-hash-algorithm? string
                  +---ro template-hash?  binary
                  +---ro pcr-index?      pcr

```

```

|          +---ro signature?                      binary
+---:(netequip_boot) {netequip_boot}?
  +---ro boot-event-logs
    +---ro boot-event-entry* [event-number]
      +---ro event-number                        uint64
      +---ro ima-template?                      string
      +---ro filename-hint?                    string
      +---ro filedata-hash?                    binary
      +---ro filedata-hash-algorithm?          string
      +---ro template-hash-algorithm?          string
      +---ro template-hash?                    binary
      +---ro pcr-index?                        pcr
      +---ro signature?                        binary

```

2.1.1.5. Data Nodes

This section provides a high level description of the data nodes containing the configuration and operational objects with the YANG model. For more details, please see the YANG model itself in Figure 1.

Container 'rats-support-structures': This houses the set of information relating to remote attestation for a device. This includes specific device TPM(s), the compute nodes (such as line cards) on which the TPM(s) reside, and the algorithms supported across the platform.

Container 'tpms': Provides configuration and operational details for each supported TPM, including the tpm-firmware-version, PCRs which may be quoted, certificates which are associated with that TPM, and the current operational status. Of note are the certificates which are associated with that TPM. As a certificate is associated with a particular TPM attestation key, knowledge of the certificate allows a specific TPM to be identified.


```

+--rw tpms
  +--rw tpm* [name]
    +--rw name string
    +--ro hardware-based boolean
    +--ro physical-index? int32 {hw:entity-mib}?
    +--ro path? string
    +--ro compute-node compute-node-ref {tpm:mtpm}?
    +--ro manufacturer? string
    +--rw firmware-version identityref
    +--rw tpm12-hash-algo? identityref
    +--rw tpm12-pcrs* pcr
    +--rw tpm20-pcr-bank* [tpm20-hash-algo]
      | +--rw tpm20-hash-algo identityref
      | +--rw pcr-index* tpm:pcr
    +--ro status enumeration
    +--rw certificates
      +--rw certificate* [name]
        +--rw name string
        +--rw keystore-ref? leafref {ks:asymmetric-keys}?
        +--rw type? enumeration

```

container 'attester-supported-algos' - Identifies which TCG hash algorithms are available for use on the Attesting platform. An operator will use this information to limit algorithms available for use by RPCs to just a desired set from the universe of all allowed hash algorithms by the TCG.

```

+--rw attester-supported-algos
  +--rw tpm12-asymmetric-signing* identityref
  +--rw tpm12-hash* identityref
  +--rw tpm20-asymmetric-signing* identityref
  +--rw tpm20-hash* identityref

```

container 'compute-nodes' - When there is more than one TPM supported, this container maintains the set of information related to the compute node associated with a specific TPM. This allows each specific TPM to identify to which 'compute-node' it belongs.

```

+--rw compute-nodes {tpm:mtpm}?
  +--ro compute-node* [node-id]
    +--ro node-id string
    +--ro node-physical-index? int32 {hw:entity-mib}?
    +--ro node-name? string
    +--ro node-location? string

```

2.1.1.6. YANG Module

```
<CODE BEGINS> file "ietf-tpm-remote-attestation@2022-03-23.yang"
module ietf-tpm-remote-attestation {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation";
  prefix tpm;

  import ietf-yang-types {
    prefix yang;
  }
  import ietf-hardware {
    prefix hw;
  }
  import ietf-keystore {
    prefix ks;
  }
  import ietf-tcg-algs {
    prefix taa;
  }

  organization
    "IETF RATS (Remote ATtestation procedureS) Working Group";
  contact
    "WG Web : <https://datatracker.ietf.org/wg/rats/>
    WG List : <mailto:rats@ietf.org>
    Author : Eric Voit <evoit@cisco.com>
    Author : Henk Birkholz <henk.birkholz@sit.fraunhofer.de>
    Author : Michael Eckel <michael.eckel@sit.fraunhofer.de>
    Author : Shwetha Bhandari <shwetha.bhandari@thoughtspot.com>
    Author : Bill Sulzen <bsulzen@cisco.com>
    Author : Liang Xia (Frank) <frank.xialiang@huawei.com>
    Author : Tom Laffey <tom.laffey@hpe.com>
    Author : Guy Fedorkow <gfedorkow@juniper.net>";
  description
    "A YANG module to enable a TPM 1.2 and TPM 2.0 based
    remote attestation procedure using a challenge-response
    interaction model and the TPM 1.2 and TPM 2.0 Quote
    primitive operations.

    Copyright (c) 2022 IETF Trust and the persons identified
    as authors of the code. All rights reserved.
    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject to
    the license terms contained in, the Simplified BSD License set
    forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
```

(<https://www.rfc-editor.org/info/rfcXXXX>); see the RFC itself for full legal notices.

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 (RFC 2119) (RFC 8174) when, and only when, they appear in all capitals, as shown here.";

```
revision 2022-03-23 {
  description
    "Initial version";
  reference
    "RFC XXXX: A YANG Data Model for Challenge-Response-based Remote
      Attestation Procedures using TPMs";
}

/*****
/*   Features   */
*****/

feature mtpm {
  description
    "The device supports the remote attestation of multiple
      TPM based cryptoprocessors.";
}

feature bios {
  description
    "The device supports the bios logs.";
  reference
    "bios-log:
      https://trustedcomputinggroup.org/wp-content/uploads/
        PC-ClientSpecific\_Platform\_Profile\_for\_TPM\_2p0\_Systems\_v51.pdf
      Section 9.4.5.2";
}

feature ima {
  description
    "The device supports Integrity Measurement Architecture logs.
      Many variants of IMA logs exist in the deployment. Each encodes
      the log entry contents as the specific measurements which get
      hashed into a PCRs as Evidence. See the reference below for
      one example of such an encoding.";
  reference
    "ima-log:
      https://www.trustedcomputinggroup.org/wp-content/uploads/
        TCG\_IWG\_CEL\_v1\_r0p41\_pub.pdf Section 5.1.6";
}
```

```
}

feature netequip_boot {
  description
    "The device supports the netequip_boot logs.";
  reference
    "netequip-boot-log:
    RFC XXXX Appendix B";
}

/*****/
/*  Typedefs  */
/*****/

typedef pcr {
  type uint8 {
    range "0..31";
  }
  description
    "Valid index number for a PCR. A {{TPM2.0}} compliant PCR index
    extends from 0-31. At this time a typical TPM would have no
    more than 32 PCRS.";
}

typedef compute-node-ref {
  type leafref {
    path "/tpm:rats-support-structures/tpm:compute-nodes"
      + "/tpm:compute-node/tpm:node-id";
  }
  description
    "This type is used to reference a hardware node. Note that an
    implementer might include an alternative leafref pointing to a
    different YANG module node specifying hardware structures.";
}

typedef certificate-name-ref {
  type leafref {
    path "/tpm:rats-support-structures/tpm:tpms/tpm:tpm"
      + "/tpm:certificates/tpm:certificate/tpm:name";
  }
  description
    "A type which allows identification of a TPM based certificate.";
}

/*****/
/*  Identities  */
/*****/
```

```
identity attested_event_log_type {
  description
    "Base identity allowing categorization of the reasons why an
    attested measurement has been taken on an Attester.";
}

identity ima {
  base attested_event_log_type;
  description
    "An event type recorded in IMA.";
}

identity bios {
  base attested_event_log_type;
  description
    "An event type associated with BIOS/UEFI.";
}

identity netequip_boot {
  base attested_event_log_type;
  description
    "An event type associated with Network Equipment Boot.";
}

/*****/
/*  Groupings  */
/*****/

grouping tpm20-hash-algo {
  description
    "The cryptographic algorithm used to hash the TPM2 PCRs. This
    must be from the list of platform supported options.";
  leaf tpm20-hash-algo {
    type identityref {
      base taa:hash;
    }
  }
  must '. = /tpm:rats-support-structures'
    + '/tpm:attester-supported-algos/tpm:tpm20-hash' {
    error-message "This platform does not support tpm20-hash-algo";
  }
  description
    "The hash scheme that is used to hash a TPM2.0 PCR. This
    must be one of those supported by a platform.
    Where this object does not appear, the default value of
    'taa:TPM_ALG_SHA256' will apply.";
}
}
```

```
grouping tpm12-hash-algo {
  description
    "The cryptographic algorithm used to hash the TPM1.2 PCRs.";
  leaf tpm12-hash-algo {
    type identityref {
      base taa:hash;
    }
    must '. = /tpm:rats-support-structures'
      + '/tpm:attester-supported-algos/tpm:tpm12-hash' {
      error-message "This platform does not support tpm12-hash-algo";
    }
    description
      "The hash scheme that is used to hash a TPM1.2 PCR. This
      MUST be one of those supported by a platform.
      Where this object does not appear, the default value of
      'taa:TPM_ALG_SHA1' will apply.";
  }
}

grouping nonce {
  description
    "A random number intended to guarantee freshness and for use
    as part of a replay-detection mechanism.";
  leaf nonce-value {
    type binary;
    mandatory true;
    description
      "A cryptographically generated random number which should
      not be predictable prior to its issuance from a random
      number generation function. The random number MUST be
      derived from an entropy source external to the Attester.

      Note that a nonce sent into a TPM will typically be 160 or 256
      binary digits long. (This is 20 or 32 bytes.) So if fewer
      binary digits are sent, this nonce object will be padded
      with leading zeros within Quotes returned from the TPM.
      Additionally if more bytes are sent, the nonce will be trimmed
      to the most significant binary digits.";
  }
}

grouping tpm12-pcr-selection {
  description
    "A Verifier can request one or more PCR values using its
    individually created Attestation Key Certificate (AC).
    The corresponding selection filter is represented in this
    grouping.";
  leaf-list pcr-index {
```

```
type pcr;
description
  "The numbers/indexes of the PCRs. In addition, any selection
  of PCRs MUST verify that the set of PCRs requested are a
  subset the set of PCRs exposed by in the leaf-list
  /tpm:rats-support-structures
  /tpm:tpms/tpm:tpm[name=current()]/tpm:tpm12-pcrs";
}
}

grouping tpm20-pcr-selection {
  description
    "A Verifier can acquire one or more PCR values, which are hashed
    together in a TPM2B_DIGEST coming from the TPM2. The selection
    list of desired PCRs and the Hash Algorithm is represented in
    this grouping.";
  list tpm20-pcr-selection {
    unique "tpm20-hash-algo";
    description
      "Specifies the list of PCRs and Hash Algorithms that can be
      returned within a TPM2B_DIGEST.";
    reference
      "TPM2.0-Structures:
      https://www.trustedcomputinggroup.org/wp-content/uploads/
      TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.9.7";
    uses tpm20-hash-algo;
    leaf-list pcr-index {
      type pcr;
      must '/tpm:rats-support-structures/tpm:tpms'
        + '/tpm:tpm[name = current()]'
        + '/tpm:tpm20-pcr-bank[pcr-index = current()]' {
        error-message "Acquiring this PCR index is not supported";
      }
      description
        "The numbers of the PCRs that which are being tracked
        with a hash based on the tpm20-hash-algo. In addition,
        any selection of PCRs MUST verify that the set of PCRs
        requested are a subset the set of PCR indexes exposed
        within /tpm:rats-support-structures/tpm:tpms
        /tpm:tpm[name=current()]/tpm:tpm20-pcr-bank
        /tpm:pcr-index";
    }
  }
}

grouping certificate-name-ref {
  description
    "Identifies a certificate in a keystore.";
```

```
    leaf certificate-name {
      type certificate-name-ref;
      mandatory true;
      description
        "Identifies a certificate in a keystore.";
    }
  }

  grouping tpm-name {
    description
      "A unique TPM on a device.";
    leaf name {
      type string;
      description
        "Unique system generated name for a TPM on a device.";
    }
  }

  grouping node-uptime {
    description
      "Uptime in seconds of the node.";
    leaf up-time {
      type uint32;
      description
        "Uptime in seconds of this node reporting its data";
    }
  }

  grouping tpml2-attestation {
    description
      "Contains an instance of TPM1.2 style signed cryptoprocessor
      measurements. It is supplemented by unsigned Attester
      information.";
    uses node-uptime;
    leaf TPM_QUOTE2 {
      type binary;
      description
        "Result of a TPM1.2 Quote2 operation. This includes PCRs,
        signatures, locality, the provided nonce and other data which
        can be further parsed to appraise the Attester.";
      reference
        "TPM1.2-Commands:
        TPM1.2 commands rev116 July 2007, Section 16.5
        https://trustedcomputinggroup.org/wp-content/uploads/TPM-Main-Part-3-Commands\_v1.2\_rev116\_01032011.pdf";
    }
  }
}
```



```
grouping tpm20-attestation {
  description
    "Contains an instance of TPM2 style signed cryptoprocessor
    measurements. It is supplemented by unsigned Attester
    information.";
  leaf TPMS_QUOTE_INFO {
    type binary;
    mandatory true;
    description
      "A hash of the latest PCR values (and the hash algorithm used)
      which have been returned from a Verifier for the selected PCRs
      and Hash Algorithms.";
    reference
      "TPM2.0-Structures:
      https://www.trustedcomputinggroup.org/wp-content/uploads/
      TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.12.1";
  }
  leaf quote-signature {
    type binary;
    description
      "Quote signature returned by TPM Quote. The signature was
      generated using the key associated with the
      certificate 'name'.";
    reference
      "TPM2.0-Structures:
      https://www.trustedcomputinggroup.org/wp-content/uploads/
      TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 11.2.1";
  }
  uses node-uptime;
  list unsigned-pcr-values {
    description
      "PCR values in each PCR bank. This might appear redundant with
      the TPM2B_DIGEST, but that digest is calculated across multiple
      PCRs. Having to verify across multiple PCRs does not
      necessarily make it easy for a Verifier to appraise just the
      minimum set of PCR information which has changed since the last
      received TPM2B_DIGEST. Put another way, why should a Verifier
      reconstruct the proper value of all PCR Quotes when only a
      single PCR has changed?
      To help this happen, if the Attester does know specific PCR
      values, the Attester can provide these individual values via
      'unsigned-pcr-values'. By comparing this information to
      what has previously been validated, it is possible for a
      Verifier to confirm the Attester's signature while eliminating

      significant processing. Note that there should never be a
      result where an unsigned PCR value differs from what may be
      reconstructed from the within the PCR quote and the event logs.
```

```
        If there is a difference, a signed result which has been
        verified from retrieved logs is considered definitive.";
uses tpm20-hash-algo;
list pcr-values {
  key "pcr-index";
  description
    "List of one PCR bank.";
  leaf pcr-index {
    type pcr;
    description
      "PCR index number.";
  }
  leaf pcr-value {
    type binary;
    description
      "PCR value.";
    reference
      "TPM2.0-Structures:
      https://www.trustedcomputinggroup.org/wp-content/uploads/
      TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.9.7";
  }
}
}
}

grouping log-identifier {
  description
    "Identifier for type of log to be retrieved.";
  leaf log-type {
    type identityref {
      base attested_event_log_type;
    }
    mandatory true;
    description
      "The corresponding measurement log type identity.";
  }
}

grouping boot-event-log {
  description
    "Defines a specific instance of an event log entry
    and corresponding to the information used to
    extend the PCR";
  leaf event-number {
    type uint32;
    description
      "Unique event number of this event which monotonically
      increases within a given event log. The maximum event
```

```
        number should not be reached, nor is wrapping back to
        an earlier number supported.";
    }
    leaf event-type {
        type uint32;
        description
            "BIOS Log Event Type:
            https://trustedcomputinggroup.org/wp-content/uploads/
            TCG_PCCClient_PFP_r1p05_v23_pub.pdf Section 10.4.1";
    }
    leaf pcr-index {
        type pcr;
        description
            "Defines the PCR index that this event extended";
    }
    list digest-list {
        description
            "Hash of event data";
        leaf hash-algo {
            type identityref {
                base taa:hash;
            }
            description
                "The hash scheme that is used to compress the event data in
                each of the leaf-list digest items.";
        }
        leaf-list digest {
            type binary;
            description
                "The hash of the event data using the algorithm of the
                'hash-algo' against 'event data'.";
        }
    }
    leaf event-size {
        type uint32;
        description
            "Size of the event data";
    }
    leaf-list event-data {
        type binary;
        description
            "The event data. This is a binary structure
            of size 'event-size'. For more on what
            might be recorded within this object
            see [bios-log] Section 9 which details
            viable events which might be recorded.";
    }
}
```

```
grouping bios-event-log {
  description
    "Measurement log created by the BIOS/UEFI.";
  list bios-event-entry {
    key "event-number";
    description
      "Ordered list of TCG described event log
       that extended the PCRs in the order they
       were logged";
    uses boot-event-log;
  }
}

grouping ima-event {
  description
    "Defines a hash log extend event for IMA measurements";
  reference
    "ima-log:
     https://www.trustedcomputinggroup.org/wp-content/uploads/TCG\_IWG\_CEL\_v1\_r0p41\_pub.pdf Section 4.3";
  leaf event-number {
    type uint64;
    description
      "Unique event number of this event which monotonically
       increases. The maximum event number should not be
       reached, nor is wrapping back to an earlier number
       supported.";
  }
  leaf ima-template {
    type string;
    description
      "Name of the template used for event logs
       for e.g. ima, ima-ng, ima-sig";
  }
  leaf filename-hint {
    type string;
    description
      "File name (including the path) that was measured.";
  }
  leaf filedata-hash {
    type binary;
    description
      "Hash of filedata as updated based upon the
       filedata-hash-algorithm";
  }
  leaf filedata-hash-algorithm {
    type string;
    description
```

```
        "Algorithm used for filedata-hash";
    }
    leaf template-hash-algorithm {
        type string;
        description
            "Algorithm used for template-hash";
    }
    leaf template-hash {
        type binary;
        description
            "hash(filedata-hash, filename-hint)";
    }
    leaf pcr-index {
        type pcr;
        description
            "Defines the PCR index that this event extended";
    }
    leaf signature {
        type binary;
        description
            "Digital file signature which provides a
            fingerprint for the file being measured.";
    }
}

grouping ima-event-log {
    description
        "Measurement log created by IMA.";
    list ima-event-entry {
        key "event-number";
        description
            "Ordered list of ima event logs by event-number";
        uses ima-event;
    }
}

grouping network-equipment-boot-event-log {
    description
        "Measurement log created by Network Equipment Boot. The Network
        Equipment Boot format is identical to the IMA format. In
        contrast to the IMA log, the Network Equipment Boot log
        includes every measurable event from an Attester, including
        the boot stages of BIOS, Bootloader, etc. In essence, the scope
        of events represented in this format combines the scope of BIOS
        events and IMA events.";
    list boot-event-entry {
        key "event-number";
        description
```

```

        "Ordered list of Network Equipment Boot event logs
        by event-number, using the IMA event format.";
    uses ima-event;
}
}

grouping event-logs {
    description
        "A selector for the log and its type.";
    choice attested_event_log_type {
        mandatory true;
        description
            "Event log type determines the event logs content.";
        case bios {
            if-feature "bios";
            description
                "BIOS/UEFI event logs";
            container bios-event-logs {
                description
                    "BIOS/UEFI event logs";
                uses bios-event-log;
            }
        }
        case ima {
            if-feature "ima";
            description
                "IMA event logs.";
            container ima-event-logs {
                description
                    "IMA event logs.";
                uses ima-event-log;
            }
        }
        case netequip_boot {
            if-feature "netequip_boot";
            description
                "Network Equipment Boot event logs";
            container boot-event-logs {
                description
                    "Network equipment boot event logs.";
                uses network-equipment-boot-event-log;
            }
        }
    }
}

/*****/
/*   RPC operations   */

```

```

/*****/

rpc tpm12-challenge-response-attestation {
  if-feature "taa:tpm12";
  description
    "This RPC accepts the input for TSS TPM 1.2 commands made to the
    attesting device.";
  input {
    container tpm12-attestation-challenge {
      description
        "This container includes every information element defined
        in the reference challenge-response interaction model for
        remote attestation. Corresponding values are based on
        TPM 1.2 structure definitions";
      uses tpm12-pcr-selection;
      uses nonce;
      leaf-list certificate-name {
        if-feature "tpm:mtpm";
        type certificate-name-ref;
        must "/tpm:rats-support-structures/tpm:tpms"
          + "/tpm:tpm[tpm:firmware-version='taa:tpm12']"
          + "/tpm:certificates/"
          + "/tpm:certificate[name=current()]" {
          error-message "Not an available TPM1.2 AIK certificate.";
        }
        description
          "When populated, the RPC will only get a Quote for the
          TPMs associated with these certificate(s).";
      }
    }
  }
  output {
    list tpm12-attestation-response {
      unique "certificate-name";
      description
        "The binary output of TPM 1.2 TPM_Quote/TPM_Quote2, including
        the PCR selection and other associated attestation evidence
        metadata";
      uses certificate-name-ref {
        description
          "Certificate associated with this tpm12-attestation.";
      }
      uses tpm12-attestation;
    }
  }
}

rpc tpm20-challenge-response-attestation {

```

```
if-feature "taa:tpm20";
description
  "This RPC accepts the input for TSS TPM 2.0 commands of the
  managed device. ComponentIndex from the hardware manager YANG
  module is used to refer to dedicated TPM in composite devices,
  e.g. smart NICs, is not covered.";
input {
  container tpm20-attestation-challenge {
    description
      "This container includes every information element defined
      in the reference challenge-response interaction model for
      remote attestation. Corresponding values are based on
      TPM 2.0 structure definitions";
    uses nonce;
    uses tpm20-pcr-selection;
    leaf-list certificate-name {
      if-feature "tpm:mtpm";
      type certificate-name-ref;
      must "/tpm:rats-support-structures/tpm:tpms"
        + "/tpm:tpm[tpm:firmware-version='taa:tpm20']"
        + "/tpm:certificates/"
        + "/tpm:certificate[name=current()]" {
        error-message "Not an available TPM2.0 AIK certificate.";
      }
      description
        "When populated, the RPC will only get a Quote for the
        TPMs associated with the certificates.";
    }
  }
}
output {
  list tpm20-attestation-response {
    unique "certificate-name";
    description
      "The binary output of TPM2b_Quote from one TPM of the
      node which identified by node-id. An TPMS_ATTEST structure
      including a length, encapsulated in a signature";
    uses certificate-name-ref {
      description
        "Certificate associated with this tpm20-attestation.";
    }
    uses tpm20-attestation;
  }
}

rpc log-retrieval {
  description
```



```
"Logs Entries are either identified via indices or via providing
the last line received. The number of lines returned can be
limited. The type of log is a choice that can be augmented.";
input {
  uses log-identifier;
  list log-selector {
    description
      "Only log entries which meet all the selection criteria
      provided are to be returned by the RPC output.";
    leaf-list name {
      type string;
      description
        "Name of one or more unique TPMs on a device. If this
        object exists, a selection should pull only the objects
        related to these TPM(s). If it does not exist, all
        qualifying TPMs that are 'hardware-based' equals true
        on the device are selected. When this selection
        criteria is provided, it will be considered as a logical
        AND with any other selection criteria provided.";
    }
    choice index-type {
      description
        "Last log entry received, log index number, or timestamp.";
      case last-entry {
        description
          "The last entry of the log already retrieved.";
        leaf last-entry-value {
          type binary;
          description
            "Content of a log event which matches 1:1 with a
            unique event record contained within the log. Log
            entries after this will be passed to the
            requester. Note: if log entry values are not unique,
            this MUST return an error.";
        }
      }
      case index {
        description
          "Numeric index of the last log entry retrieved, or
          zero.";
        leaf last-index-number {
          type uint64;
          description
            "The last numeric index number of a log entry.
            Zero means to start at the beginning of the log.
            Entries after this will be passed to the
            requester.";
        }
      }
    }
  }
}
```

```

    }
    case timestamp {
      leaf timestamp {
        type yang:date-and-time;
        description
          "Timestamp from which to start the extraction. The
           next log entry after this timestamp is to
           be sent.";
      }
      description
        "Timestamp from which to start the extraction.";
    }
  }
  leaf log-entry-quantity {
    type uint16;
    description
      "The number of log entries to be returned. If omitted, it
       means all of them.";
  }
}
}
output {
  container system-event-logs {
    description
      "The requested data of the measurement event logs";
    list node-data {
      unique "name";
      description
        "Event logs of a node in a distributed system
         identified by the node name";
      uses tpm-name;
      uses node-uptime;
      container log-result {
        description
          "The requested entries of the corresponding log.";
        uses event-logs;
      }
    }
  }
}
}

/*****/
/*  Config & Oper accessible nodes  */
/*****/

container rats-support-structures {
  description

```

```
    "The datastore definition enabling verifiers or relying
    parties to discover the information necessary to use the
    remote attestation RPCs appropriately.";
container compute-nodes {
  if-feature "tpm:mtpm";
  description
    "Holds the set of device subsystems/components in this
    composite device that support TPM operations.";
  list compute-node {
    key "node-id";
    unique "node-name";
    config false;
    min-elements 2;
    description
      "A component within this composite device which
      supports TPM operations.";
    leaf node-id {
      type string;
      description
        "ID of the compute node, such as Board Serial Number.";
    }
    leaf node-physical-index {
      if-feature "hw:entity-mib";
      type int32 {
        range "1..2147483647";
      }
      config false;
      description
        "The entPhysicalIndex for the compute node.";
      reference
        "RFC 6933: Entity MIB (Version 4) - entPhysicalIndex";
    }
    leaf node-name {
      type string;
      description
        "Name of the compute node.";
    }
    leaf node-location {
      type string;
      description
        "Location of the compute node, such as slot number.";
    }
  }
}
container tpms {
  description
    "Holds the set of TPMs within an Attester.";
  list tpm {
```

```
key "name";
unique "path";
description
  "A list of TPMs in this composite device that RATS
   can be conducted with.";
uses tpm-name;
leaf hardware-based {
  type boolean;
  config false;
  mandatory true;
  description
    "System generated indication of whether this is a
     hardware based TPM.";
}
leaf physical-index {
  if-feature "hw:entity-mib";
  type int32 {
    range "1..2147483647";
  }
  config false;
  description
    "The entPhysicalIndex for the TPM.";
  reference
    "RFC 6933: Entity MIB (Version 4) - entPhysicalIndex";
}
leaf path {
  type string;
  config false;
  description
    "Device path to a unique TPM on a device. This can change
     across reboots.";
}
leaf compute-node {
  if-feature "tpm:mtpm";
  type compute-node-ref;
  config false;
  mandatory true;
  description
    "Indicates the compute node measured by this TPM.";
}
leaf manufacturer {
  type string;
  config false;
  description
    "TPM manufacturer name.";
}
leaf firmware-version {
  type identityref {
```

```
    base taa:cryptoprocessor;
  }
  mandatory true;
  description
    "Identifies the cryptoprocessor API set supported. This
    is automatically configured by the device and should not
    be changed.";
}
uses tpm12-hash-algo {
  when "derived-from-or-self(firmware-version, 'taa:tpm12')";
  refine "tpm12-hash-algo" {
    description
      "The hash algorithm overwrites the default used for PCRs
      on this TPM1.2 compliant cryptoprocessor.";
  }
}
leaf-list tpm12-pcrs {
  when
    "derived-from-or-self(..../firmware-version, 'taa:tpm12')";
  type pcr;
  description
    "The PCRs which may be extracted from this TPM1.2
    compliant cryptoprocessor.";
}
list tpm20-pcr-bank {
  when
    "derived-from-or-self(..../firmware-version, 'taa:tpm20')";
  key "tpm20-hash-algo";
  description
    "Specifies the list of PCRs that may be extracted for
    a specific Hash Algorithm on this TPM2 compliant
    cryptoprocessor. A bank is a set of PCRs which are
    extended using a particular hash algorithm.";
  reference
    "TPM2.0-Structures:
    https://www.trustedcomputinggroup.org/wp-content/uploads/
    TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.9.7";
  leaf tpm20-hash-algo {
    type identityref {
      base taa:hash;
    }
    must '/tpm:rats-support-structures'
      + '/tpm:attester-supported-algos'
      + '/tpm:tpm20-hash' {
      error-message "This platform does not support tpm20-hash-algo";
    }
    description
      "The hash scheme actively being used to hash a
```

```
        one or more TPM2.0 PCRs.";
    }
    leaf-list pcr-index {
        type tpm:pcr;
        description
            "Defines what TPM2 PCRs are available to be extracted.";
    }
}
leaf status {
    type enumeration {
        enum operational {
            value 0;
            description
                "The TPM currently is running normally and
                 is ready to accept and process TPM quotes.";
            reference
                "TPM2.0-Arch:
                 https://trustedcomputinggroup.org/wp-content/uploads/TCG\_TPM2\_r1p59\_Part1\_Architecture\_pub.pdf
                 Section 12";
        }
        enum non-operational {
            value 1;
            description
                "TPM is in a state such as startup or shutdown which
                 precludes the processing of TPM quotes.";
        }
    }
}
config false;
mandatory true;
description
    "TPM chip self-test status.";
}
container certificates {
    description
        "The TPM's certificates, including EK certificates
         and Attestation Key certificates.";
    list certificate {
        key "name";
        description
            "Three types of certificates can be accessed via
             this statement, including Initial Attestation
             Key Certificate, Local Attestation Key Certificate or
             Endorsement Key Certificate.";
        leaf name {
            type string;
            description
                "An arbitrary name uniquely identifying a certificate
```

```
        associated within key within a TPM.";
    }
    leaf keystore-ref {
        if-feature "ks:asymmetric-keys";
        type leafref {
            path "/ks:keystore/ks:asymmetric-keys/ks:asymmetric-key"
                + "/ks:name";
        }
        description
            "A reference to a specific certificate of an
            asymmetric key in the Keystore.";
    }
    leaf type {
        type enumeration {
            enum endorsement-certificate {
                value 0;
                description
                    "Endorsement Key (EK) Certificate type.";
                reference
                    "TPM2.0-Key:
                    https://trustedcomputinggroup.org/wp-content/
                    uploads/TPM-2p0-Keys-for-Device-Identity-
                    and-Attestation_v1_r12_publ0082021.pdf
                    Section 3.11";
            }
            enum initial-attestation-certificate {
                value 1;
                description
                    "Initial Attestation key (IAK) Certificate type.";
                reference
                    "TPM2.0-Key:
                    https://trustedcomputinggroup.org/wp-content/
                    uploads/TPM-2p0-Keys-for-Device-Identity-
                    and-Attestation_v1_r12_publ0082021.pdf
                    Section 3.2";
            }
            enum local-attestation-certificate {
                value 2;
                description
                    "Local Attestation Key (LAK) Certificate type.";
                reference
                    "TPM2.0-Key:
                    https://trustedcomputinggroup.org/wp-content/
                    uploads/TPM-2p0-Keys-for-Device-Identity-
                    and-Attestation_v1_r12_publ0082021.pdf
                    Section 3.2";
            }
        }
    }
}
```

```

        description
            "Function supported by this certificate from within the
            TPM.";
    }
}
}
}
container attester-supported-algos {
    description
        "Identifies which TPM algorithms are available for use on an
        attesting platform.";
    leaf-list tpm12-asymmetric-signing {
        when "../../tpm:tpms"
            + "/tpm:tpm[tpm:firmware-version='taa:tpm12']";
        type identityref {
            base taa:asymmetric;
        }
        description
            "Platform Supported TPM12 asymmetric algorithms.";
    }
    leaf-list tpm12-hash {
        when "../../tpm:tpms"
            + "/tpm:tpm[tpm:firmware-version='taa:tpm12']";
        type identityref {
            base taa:hash;
        }
        description
            "Platform supported TPM12 hash algorithms.";
    }
    leaf-list tpm20-asymmetric-signing {
        when "../../tpm:tpms"
            + "/tpm:tpm[tpm:firmware-version='taa:tpm20']";
        type identityref {
            base taa:asymmetric;
        }
        description
            "Platform Supported TPM20 asymmetric algorithms.";
    }
    leaf-list tpm20-hash {
        when "../../tpm:tpms"
            + "/tpm:tpm[tpm:firmware-version='taa:tpm20']";
        type identityref {
            base taa:hash;
        }
        description
            "Platform supported TPM20 hash algorithms.";
    }
}

```



```

    }
  }
}
<CODE ENDS>

```

Figure 1

2.1.2. 'ietf-tcg-algs'

This document has encoded the TCG Algorithm definitions of [TCG-Algos], revision 1.32. By including this full table as a separate YANG file within this document, it is possible for other YANG models to leverage the contents of this model. Specific references to [RFC2104], [RFC8017], [ISO-IEC-9797-1], [ISO-IEC-9797-2], [ISO-IEC-10116], [ISO-IEC-10118-3], [ISO-IEC-14888-3], [ISO-IEC-15946-1], [ISO-IEC-18033-3], [IEEE-Std-1363-2000], [IEEE-Std-1363a-2004], [NIST-PUB-FIPS-202], [NIST-SP800-38C], [NIST-SP800-38D], [NIST-SP800-38F], [NIST-SP800-56A], [NIST-SP800-108], [bios-log], as well as Appendix A and Appendix B exist within the YANG Model.

2.1.2.1. Features

There are two types of features supported: 'TPM12' and 'TPM20'. Support for either of these features indicates that a cryptoprocessor supporting the corresponding type of TCG TPM API is present on an Attester. Most commonly, only one type of cryptoprocessor will be available on an Attester.

2.1.2.2. Identities

There are three types of identities in this model:

1. Cryptographic functions supported by a TPM algorithm; these include: 'asymmetric', 'symmetric', 'hash', 'signing', 'anonymous_signing', 'encryption_mode', 'method', and 'object_type'. The definitions of each of these are in Table 2 of [TCG-Algos].
2. API specifications for TPM types: 'tpm12' and 'tpm20'
3. Specific algorithm types: Each algorithm type defines what cryptographic functions may be supported, and on which type of API specification. It is not required that an implementation of a specific TPM will support all algorithm types. The contents of each specific algorithm mirrors what is in Table 3 of [TCG-Algos].

2.1.2.3. YANG Module

```
<CODE BEGINS> file "ietf-tcg-algs@2022-03-23.yang"
module ietf-tcg-algs {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-tcg-algs";
  prefix taa;

  organization
    "IETF RATS (Remote ATtestation procedureS) Working Group";
  contact
    "WG Web:  <https://datatracker.ietf.org/wg/rats/>
    WG List:  <mailto:rats@ietf.org>
    Author:   Eric Voit <mailto:evoit@cisco.com>";
  description
    "This module defines identities for asymmetric algorithms.

    Copyright (c) 2022 IETF Trust and the persons identified as
    authors of the code. All rights reserved.
    Redistribution and use in source and binary forms, with
    or without modification, is permitted pursuant to, and
    subject to the license terms contained in, the Revised
    BSD License set forth in Section 4.c of the IETF Trust's
    Legal Provisions Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); see the RFC itself
    for full legal notices.

    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 (RFC 2119)
    (RFC 8174) when, and only when, they appear in all
    capitals, as shown here."

  revision 2022-03-23 {
    description
      "Initial version";
    reference
      "RFC XXXX: A YANG Data Model for Challenge-Response-based Remote
      Attestation Procedures using TPMs";
  }

  /*****
  /*  Features  */
  *****/
```

```
feature tpm12 {
  description
    "This feature indicates algorithm support for the TPM 1.2 API
    as per Section 4.8 of TPM1.2-Structures:
    TPM Main Part 2 TPM Structures
    https://trustedcomputinggroup.org/wp-content/uploads/TPM-
    Main-Part-2-TPM-Structures_v1.2_rev116_01032011.pdf";
}

feature tpm20 {
  description
    "This feature indicates algorithm support for the TPM 2.0 API
    as per Section 11.4 of Trusted Platform Module Library
    Part 1: Architecture. See TPM2.0-Arch:
    https://trustedcomputinggroup.org/wp-content/uploads/
    TCG_TPM2_r1p59_Part1_Architecture_pub.pdf";
}

/*****/
/* Identities */
/*****/

identity asymmetric {
  description
    "A TCG recognized asymmetric algorithm with a public and
    private key.";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2,
    https://trustedcomputinggroup.org/resource/
    tcg-algorithm-registry/TCG-Algorithm_Registry_r1p32_pub";
}

identity symmetric {
  description
    "A TCG recognized symmetric algorithm with only a private key.";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
}

identity hash {
  description
    "A TCG recognized hash algorithm that compresses input data to
    a digest value or indicates a method that uses a hash.";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
}

identity signing {
```

```
    description
      "A TCG recognized signing algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
  }

  identity anonymous_signing {
    description
      "A TCG recognized anonymous signing algorithm.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
  }

  identity encryption_mode {
    description
      "A TCG recognized encryption mode.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
  }

  identity method {
    description
      "A TCG recognized method such as a mask generation function.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
  }

  identity object_type {
    description
      "A TCG recognized object type.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 2";
  }

  identity cryptoprocessor {
    description
      "Base identity identifying a cryptoprocessor.";
  }

  identity tpml2 {
    if-feature "tpml2";
    base cryptoprocessor;
    description
      "Supportable by a TPM1.2.";
    reference
      "TPM1.2-Structures:
      https://trustedcomputinggroup.org/wp-content/uploads/
      TPM-Main-Part-2-TPM-Structures\_v1.2\_rev116\_01032011.pdf

```

```
        TPM_ALGORITHM_ID values, Section 4.8";
    }

    identity tpm20 {
        if-feature "tpm20";
        base cryptoprocessor;
        description
            "Supportable by a TPM2.";
        reference
            "TPM2.0-Structures:
            https://trustedcomputinggroup.org/wp-content/uploads/
            TPM-Rev-2.0-Part-2-Structures-01.38.pdf";
    }

    identity TPM_ALG_RSA {
        if-feature "tpm12 or tpm20";
        base tpm12;
        base tpm20;
        base asymmetric;
        base object_type;
        description
            "RSA algorithm";
        reference
            "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
            RFC 8017. ALG_ID: 0x0001";
    }

    identity TPM_ALG_TDES {
        if-feature "tpm12";
        base tpm12;
        base symmetric;
        description
            "Block cipher with various key sizes (Triple Data Encryption
            Algorithm, commonly called Triple Data Encryption Standard)
            Note: was banned in TPM1.2 v94";
        reference
            "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
            ISO/IEC 18033-3. ALG_ID: 0x0003";
    }

    identity TPM_ALG_SHA1 {
        if-feature "tpm12 or tpm20";
        base hash;
        base tpm12;
        base tpm20;
        description
            "SHA1 algorithm - Deprecated due to insufficient cryptographic
            protection. However, it is still useful for hash algorithms
```

```
        where protection is not required.";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
        ISO/IEC 10118-3. ALG_ID: 0x0004";
}

identity TPM_ALG_HMAC {
    if-feature "tpm12 or tpm20";
    base tpm12;
    base tpm20;
    base hash;
    base signing;
    description
        "Hash Message Authentication Code (HMAC) algorithm";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3,
        ISO/IEC 9797-2 and RFC2104. ALG_ID: 0x0005";
}

identity TPM_ALG_AES {
    if-feature "tpm12";
    base tpm12;
    base symmetric;
    description
        "The AES algorithm with various key sizes";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3,
        ISO/IEC 18033-3. ALG_ID: 0x0006";
}

identity TPM_ALG_MGF1 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    base method;
    description
        "hash-based mask-generation function";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3,
        IEEE Std 1363-2000 and IEEE Std 1363a-2004.
        ALG_ID: 0x0007";
}

identity TPM_ALG_KEYEDHASH {
    if-feature "tpm20";
    base tpm20;
    base hash;
    base object_type;
```

```
    description
      "An encryption or signing algorithm using a keyed hash.  These
      may use XOR for encryption or an HMAC for signing and may
      also refer to a data object that is neither signing nor
      encrypting.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3,
      ALG_ID: 0x0008";
  }

  identity TPM_ALG_XOR {
    if-feature "tpm12 or tpm20";
    base tpm12;
    base tpm20;
    base hash;
    base symmetric;
    description
      "The XOR encryption algorithm.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3.
      ALG_ID: 0x000A";
  }

  identity TPM_ALG_SHA256 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 256 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      ISO/IEC 10118-3. ALG_ID: 0x000B";
  }

  identity TPM_ALG_SHA384 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 384 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      ISO/IEC 10118-3. ALG_ID: 0x000C";
  }

  identity TPM_ALG_SHA512 {
    if-feature "tpm20";
    base tpm20;
```

```
    base hash;
    description
      "The SHA 512 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      ISO/IEC 10118-3. ALG_ID: 0x000D";
  }

  identity TPM_ALG_NULL {
    if-feature "tpm20";
    base tpm20;
    description
      "NULL algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3.
      ALG_ID: 0x0010";
  }

  identity TPM_ALG_SM3_256 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SM3 hash algorithm.";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      ISO/IEC 10118-3:2018. ALG_ID: 0x0012";
  }

  identity TPM_ALG_SM4 {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    description
      "SM4 symmetric block cipher";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3.
      ALG_ID: 0x0013";
  }

  identity TPM_ALG_RSASSA {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
      "RFC 8017 Signature algorithm defined in section 8.2
      (RSASSAPKCS1-v1_5)";
```



```
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      RFC 8017. ALG_ID: 0x0014";
  }

  identity TPM_ALG_RSAES {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base encryption_mode;
    description
      "RFC 8017 Signature algorithm defined in section 7.2
      (RSAES-PKCS1-v1_5)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      RFC 8017. ALG_ID: 0x0015";
  }

  identity TPM_ALG_RSAPSS {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
      "Padding algorithm defined in section 8.1 (RSASSA PSS)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      RFC 8017. ALG_ID: 0x0016";
  }

  identity TPM_ALG_OAEP {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base encryption_mode;
    description
      "Padding algorithm defined in section 7.1 (RSASSA OAEP)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      RFC 8017. ALG_ID: 0x0017";
  }

  identity TPM_ALG_ECDSA {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
```

```
    "Signature algorithm using elliptic curve cryptography (ECC)";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    ISO/IEC 14888-3. ALG_ID: 0x0018";
}

identity TPM_ALG_ECDH {
  if-feature "tpm20";
  base tpm20;
  base asymmetric;
  base method;
  description
    "Secret sharing using ECC";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    NIST SP800-56A. ALG_ID: 0x0019";
}

identity TPM_ALG_ECDSA {
  if-feature "tpm20";
  base tpm20;
  base asymmetric;
  base signing;
  base anonymous_signing;
  description
    "Elliptic-curve based anonymous signing scheme";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    TCG TPM 2.0 library specification. ALG_ID: 0x001A";
}

identity TPM_ALG_SM2 {
  if-feature "tpm20";
  base tpm20;
  base asymmetric;
  base signing;
  base encryption_mode;
  base method;
  description
    "SM2 - depending on context, either an elliptic-curve based,
    signature algorithm, an encryption scheme, or a key exchange
    protocol";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3.
    ALG_ID: 0x001B";
}

identity TPM_ALG_ECSCHNORR {
```

```
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
        "Elliptic-curve based Schnorr signature";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3.
        ALG_ID: 0x001C";
}

identity TPM_ALG_ECMQV {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base method;
    description
        "Two-phase elliptic-curve key";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
        NIST SP800-56A. ALG_ID: 0x001D";
}

identity TPM_ALG_KDF1_SP800_56A {
    if-feature "tpm20";
    base tpm20;
    base hash;
    base method;
    description
        "Concatenation key derivation function";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
        NIST SP800-56A (approved alternative1) section 5.8.1.
        ALG_ID: 0x0020";
}

identity TPM_ALG_KDF2 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    base method;
    description
        "Key derivation function";
    reference
        "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
        IEEE 1363a-2004 KDF2 section 13.2. ALG_ID: 0x0021";
}
```

```
identity TPM_ALG_KDF1_SP800_108 {
  base TPM_ALG_KDF2;
  description
    "A key derivation method";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    NIST SP800-108 - Section 5.1 KDF. ALG_ID: 0x0022";
}

identity TPM_ALG_ECC {
  if-feature "tpm20";
  base tpm20;
  base asymmetric;
  base object_type;
  description
    "Prime field ECC";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    ISO/IEC 15946-1. ALG_ID: 0x0023";
}

identity TPM_ALG_SYMCIPHER {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  base object_type;
  description
    "Object type for a symmetric block cipher";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    TCG TPM 2.0 library specification. ALG_ID: 0x0025";
}

identity TPM_ALG_CAMELLIA {
  if-feature "tpm20";
  base tpm20;
  base symmetric;
  description
    "The Camellia algorithm";
  reference
    "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
    ISO/IEC 18033-3. ALG_ID: 0x0026";
}

identity TPM_ALG_SHA3_256 {
  if-feature "tpm20";
  base tpm20;
  base hash;
```

```
    description
      "ISO/IEC 10118-3 - the SHA 256 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      NIST PUB FIPS 202. ALG_ID: 0x0027";
  }

  identity TPM_ALG_SHA3_384 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 384 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      NIST PUB FIPS 202. ALG_ID: 0x0028";
  }

  identity TPM_ALG_SHA3_512 {
    if-feature "tpm20";
    base tpm20;
    base hash;
    description
      "The SHA 512 algorithm";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      NIST PUB FIPS 202. ALG_ID: 0x0029";
  }

  identity TPM_ALG_CMAC {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    description
      "block Cipher-based Message Authentication Code (CMAC)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      ISO/IEC 9797-1:2011 Algorithm 5. ALG_ID: 0x003F";
  }

  identity TPM_ALG_CTR {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Counter mode";
```

```
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0040";
  }

  identity TPM_ALG_OFB {
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Output Feedback mode";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0041";
  }

  identity TPM_ALG_CBC {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Cipher Block Chaining mode";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0042";
  }

  identity TPM_ALG_CFB {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Cipher Feedback mode";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
      ISO/IEC 10116. ALG_ID: 0x0043";
  }

  identity TPM_ALG_ECB {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base encryption_mode;
    description
      "Electronic Codebook mode";
    reference
```

```
        "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
        ISO/IEC 10116. ALG_ID: 0x0044";
    }

    identity TPM_ALG_CCM {
        if-feature "tpm20";
        base tpm20;
        base symmetric;
        base signing;
        base encryption_mode;
        description
            "Counter with Cipher Block Chaining-Message Authentication
            Code (CCM)";
        reference
            "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
            NIST SP800-38C. ALG_ID: 0x0050";
    }

    identity TPM_ALG_GCM {
        if-feature "tpm20";
        base tpm20;
        base symmetric;
        base signing;
        base encryption_mode;
        description
            "Galois/Counter Mode (GCM)";
        reference
            "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
            NIST SP800-38D. ALG_ID: 0x0051";
    }

    identity TPM_ALG_KW {
        if-feature "tpm20";
        base tpm20;
        base symmetric;
        base signing;
        base encryption_mode;
        description
            "AES Key Wrap (KW)";
        reference
            "TCG-Algos:TCG Algorithm Registry Rev1.32  Table 3 and
            NIST SP800-38F. ALG_ID: 0x0052";
    }

    identity TPM_ALG_KWP {
        if-feature "tpm20";
        base tpm20;
        base symmetric;
```

```
    base signing;
    base encryption_mode;
    description
      "AES Key Wrap with Padding (KWP)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      NIST SP800-38F. ALG_ID: 0x0053";
  }

  identity TPM_ALG_EAX {
    if-feature "tpm20";
    base tpm20;
    base symmetric;
    base signing;
    base encryption_mode;
    description
      "Authenticated-Encryption Mode";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      NIST SP800-38F. ALG_ID: 0x0054";
  }

  identity TPM_ALG_EDDSA {
    if-feature "tpm20";
    base tpm20;
    base asymmetric;
    base signing;
    description
      "Edwards-curve Digital Signature Algorithm (PureEdDSA)";
    reference
      "TCG-Algos:TCG Algorithm Registry Rev1.32 Table 3 and
      RFC 8032. ALG_ID: 0x0060";
  }
}
<CODE ENDS>
```

Note that not all cryptographic functions are required for use by ietf-tpm-remote-attestation.yang. However the full definition of Table 3 of [TCG-Algos] will allow use by additional YANG specifications.

3. IANA Considerations

This document registers the following namespace URIs in the [xml-registry] as per [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation

Registrant Contact: The IESG.

XML: N/A; the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-tcg-algs

Registrant Contact: The IESG.

XML: N/A; the requested URI is an XML namespace.

This document registers the following YANG modules in the registry [yang-parameters] as per Section 14 of [RFC6020]:

Name: ietf-tpm-remote-attestation

Namespace: urn:ietf:params:xml:ns:yang:ietf-tpm-remote-attestation

Prefix: tpm

Reference: draft-ietf-rats-yang-tpm-charra (RFC form)

Name: ietf-tcg-algs

Namespace: urn:ietf:params:xml:ns:yang:ietf-tcg-algs

Prefix: taa

Reference: draft-ietf-rats-yang-tpm-charra (RFC form)

4. Security Considerations

The YANG module ietf-tpm-remote-attestation.yang specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., `_config true`, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., `_edit-config`) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees and data nodes as well as their sensitivity/vulnerability:

Container `'/rats-support-structures/attester-supported-algos'`: `'tpm12-asymmetric-signing'`, `'tpm12-hash'`, `'tpm20-asymmetric-signing'`, and `'tpm20-hash'`. All could be populated with algorithms that are not supported by the underlying physical TPM installed by the equipment vendor. A vendor should restrict the ability to configure unsupported algorithms.

Container: `'/rats-support-structures/tpms'`: `'name'`: Although shown as `'rw'`, it is system generated. Therefore, it should not be possible for an operator to add or remove a TPM from the configuration.

`'tpm20-pcr-bank'`: It is possible to configure PCRs for extraction which are not being extended by system software. This could unnecessarily use TPM resources.

`'certificates'`: It is possible to provision a certificate which does not correspond to an Attestation Identity Key (AIK) within the TPM 1.2, or an Attestation Key (AK) within the TPM 2.0 respectively. In such a case, calls to an RPC requesting this specific certificate could result in either no response or a response for an unexpected TPM.

RPC `'tpm12-challenge-response-attestation'`: The receiver of the RPC response must verify that the certificate is for an active AIK, i.e., the certificate has been confirmed by a third party as being able to support Attestation on the targeted TPM 1.2.

RPC `'tpm20-challenge-response-attestation'`: The receiver of the RPC response must verify that the certificate is for an active AK, i.e., the private key confirmation of the quote signature within the RPC response has been confirmed by a third party to belong to an entity legitimately able to perform Attestation on the targeted TPM 2.0.

RPC `'log-retrieval'`: Requesting a large volume of logs from the attester could require significant system resources and create a denial of service.

Information collected through the RPCs above could reveal that specific versions of software and configurations of endpoints that could identify vulnerabilities on those systems. Therefore, RPCs should be protected by NACM [RFC8341] with a default setting of deny-all to limit the extraction of attestation data by only authorized Verifiers.

For the YANG module `ietf-tcg-algs.yang`, please use care when selecting specific algorithms. The introductory section of [TCG-Algos] highlights that some algorithms should be considered legacy, and recommends implementers and adopters diligently evaluate available information such as governmental, industrial, and academic research before selecting an algorithm for use.

5. References

5.1. Normative References

- [bios-log] "TCG PC Client Platform Firmware Profile Specification, Section 9.4.5.2", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/PC-ClientSpecific_Platform_Profile_for_TPM_2p0_Systems_v51.pdf>.
- [BIOS-Log-Event-Type] "TCG PC Client Platform Firmware Profile Specification", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/TCG_PCClient_PFP_r1p05_v23_pub.pdf>.
- [cel] "Canonical Event Log Format, Section 4.3", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/TCG_IWG_CEL_v1_r0p41_pub.pdf>.
- [I-D.ietf-netconf-keystore] Watsen, K., "A YANG Data Model for a Keystore", Work in Progress, Internet-Draft, draft-ietf-netconf-keystore-24, 7 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-netconf-keystore-24.txt>>.
- [I-D.ietf-rats-architecture] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-15, 8 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-architecture-15.txt>>.
- [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://www.ietf.org/archive/id/draft-ietf-rats-tpm-based-network-device-attest-14.txt>>.

- [IEEE-Std-1363-2000]
"IEEE 1363-2000 - IEEE Standard Specifications for Public-Key Cryptography", n.d.,
<<https://standards.ieee.org/standard/1363-2000.html>>.
- [IEEE-Std-1363a-2004]
"1363a-2004 - IEEE Standard Specifications for Public-Key Cryptography - Amendment 1: Additional Techniques", n.d.,
<<https://ieeexplore.ieee.org/document/1335427>>.
- [ISO-IEC-10116]
"ISO/IEC 10116:2017 - Information technology", n.d.,
<<https://www.iso.org/standard/64575.html>>.
- [ISO-IEC-10118-3]
"Dedicated hash-functions - ISO/IEC 10118-3:2018", n.d.,
<<https://www.iso.org/standard/67116.html>>.
- [ISO-IEC-14888-3]
"ISO/IEC 14888-3:2018 - Digital signatures with appendix", n.d., <<https://www.iso.org/standard/76382.html>>.
- [ISO-IEC-15946-1]
"ISO/IEC 15946-1:2016 - Information technology", n.d.,
<<https://www.iso.org/standard/65480.html>>.
- [ISO-IEC-18033-3]
"ISO/IEC 18033-3:2010 - Encryption algorithms", n.d.,
<<https://www.iso.org/standard/54531.html>>.
- [ISO-IEC-9797-1]
"Message Authentication Codes (MACs) - ISO/IEC 9797-1:2011", n.d.,
<<https://www.iso.org/standard/50375.html>>.
- [ISO-IEC-9797-2]
"Message Authentication Codes (MACs) - ISO/IEC 9797-2:2011", n.d.,
<<https://www.iso.org/standard/51618.html>>.
- [NIST-PUB-FIPS-202]
"SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions", n.d.,
<<https://csrc.nist.gov/publications/detail/fips/202/final>>.

- [NIST-SP800-108]
"Recommendation for Key Derivation Using Pseudorandom Functions", n.d.,
<<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-108.pdf>>.
- [NIST-SP800-38C]
"Recommendation for Block Cipher Modes of Operation: the CCM Mode for Authentication and Confidentiality", n.d.,
<<https://csrc.nist.gov/publications/detail/sp/800-38c/final>>.
- [NIST-SP800-38D]
"Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC", n.d.,
<<https://csrc.nist.gov/publications/detail/sp/800-38d/final>>.
- [NIST-SP800-38F]
"Recommendation for Block Cipher Modes of Operation: Methods for Key Wrapping", n.d.,
<<https://csrc.nist.gov/publications/detail/sp/800-38f/final>>.
- [NIST-SP800-56A]
"Recommendation for Pair-Wise Key-Establishment Schemes Using Discrete Logarithm Cryptography", n.d.,
<<https://csrc.nist.gov/publications/detail/sp/800-56a/rev-3/final>>.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997,
<<https://www.rfc-editor.org/info/rfc2104>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004,
<<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010,
<<https://www.rfc-editor.org/info/rfc6020>>.

- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.
- [RFC6933] Bierman, A., Romascanu, D., Quittek, J., and M. Chandramouli, "Entity MIB (Version 4)", RFC 6933, DOI 10.17487/RFC6933, May 2013, <<https://www.rfc-editor.org/info/rfc6933>>.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.
- [RFC8017] Moriarty, K., Ed., Kaliski, B., Jonsson, J., and A. Rusch, "PKCS #1: RSA Cryptography Specifications Version 2.2", RFC 8017, DOI 10.17487/RFC8017, November 2016, <<https://www.rfc-editor.org/info/rfc8017>>.
- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/info/rfc8032>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC8348] Bierman, A., Bjorklund, M., Dong, J., and D. Romascanu, "A YANG Data Model for Hardware Management", RFC 8348, DOI 10.17487/RFC8348, March 2018, <<https://www.rfc-editor.org/info/rfc8348>>.

- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [TCG-Algos]
"TCG Algorithm Registry", n.d.,
<https://trustedcomputinggroup.org/wp-content/uploads/TCG-Algorithm_Registry_r1p32_pub.pdf>.
- [TPM1.2] TCG, ., "TPM 1.2 Main Specification", 2 October 2003,
<<https://trustedcomputinggroup.org/resource/tpm-main-specification/>>.
- [TPM1.2-Commands]
"TPM Main Part 3 Commands", n.d.,
<https://trustedcomputinggroup.org/wp-content/uploads/TPM-Main-Part-3-Commands_v1.2_rev116_01032011.pdf>.
- [TPM1.2-Structures]
"TPM Main Part 2 TPM Structures", n.d.,
<https://trustedcomputinggroup.org/wp-content/uploads/TPM-Main-Part-2-TPM-Structures_v1.2_rev116_01032011.pdf>.
- [TPM2.0] TCG, ., "TPM 2.0 Library Specification", 15 March 2013,
<<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.
- [TPM2.0-Arch]
"Trusted Platform Module Library - Part 1: Architecture", n.d., <https://trustedcomputinggroup.org/wp-content/uploads/TCG_TPM2_r1p59_Part1_Architecture_pub.pdf>.
- [TPM2.0-Key]
TCG, ., "TPM 2.0 Keys for Device Identity and Attestation, Rev12", 8 October 2021,
<https://trustedcomputinggroup.org/wp-content/uploads/TPM-2p0-Keys-for-Device-Identity-and-Attestation_v1_r12_pub10082021.pdf>.
- [TPM2.0-Structures]
"Trusted Platform Module Library - Part 2: Structures", n.d., <<https://trustedcomputinggroup.org/wp-content/uploads/TPM-Rev-2.0-Part-2-Structures-01.38.pdf>>.
- [UEFI-Secure-Boot]
"Unified Extensible Firmware Interface (UEFI) Specification Version 2.9 (March 2021), Section 32.1

(Secure Boot)", n.d.,
<https://uefi.org/sites/default/files/resources/UEFI_Spec_2_9_2021_03_18.pdf>.

5.2. Informative References

- [I-D.ietf-rats-reference-interaction-models]
Birkholz, H., Eckel, M., Pan, W., and E. Voit, "Reference Interaction Models for Remote Attestation Procedures", Work in Progress, Internet-Draft, draft-ietf-rats-reference-interaction-models-05, 26 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-reference-interaction-models-05.txt>>.
- [IMA-Kernel-Source]
"Linux Integrity Measurement Architecture (IMA): Kernel Sourcecode", n.d., <<https://github.com/torvalds/linux/blob/df0cc57e057f18e44dac8e6c18aba47ab53202f9/security/integrity/ima/>>.
- [NIST-915121]
"True Randomness Can't be Left to Chance: Why entropy is important for information security", n.d., <https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=915121>.
- [xml-registry]
"IETF XML Registry", n.d., <<https://www.iana.org/assignments/xml-registry/xml-registry.xhtml>>.
- [yang-parameters]
"YANG Parameters", n.d., <<https://www.iana.org/assignments/yang-parameters/yang-parameters.xhtml>>.

Appendix A. Integrity Measurement Architecture (IMA)

IMA extends the principles of Measured Boot [TPM2.0-Arch] and Secure Boot [UEFI-Secure-Boot] to the Linux operating system, applying it to operating system applications and files. IMA has been part of the Linux integrity subsystem of the Linux kernel since 2009 (kernel version 2.6.30). The IMA mechanism represented by the YANG module in this specification is rooted in the kernel version 5.16 [IMA-Kernel-Source]. IMA enables the protection of system integrity by collecting (commonly referred to as measuring) and storing measurements (called Claims in the context of IETF RATS) of files before execution so that these measurements can be used later, at

system runtime, in remote attestation procedures. IMA acts in support of the appraisal of Evidence (which includes measurement Claims) by leveraging reference integrity measurements stored in extended file attributes.

In support of the appraisal of Evidence, IMA maintains an ordered list of measurements in kernel-space, the Stored Measurement Log (SML), for all files that have been measured before execution since the operating system was started. Although IMA can be used without a TPM, it is typically used in conjunction with a TPM to anchor the integrity of the SML in a hardware-protected secure storage location, i.e., Platform Configuration Registers (PCRs) provided by TPMs. IMA provides the SML in both binary and ASCII representations in the Linux security file system `_securityfs_ (/sys/kernel/security/ima/)`.

IMA templates define the format of the SML, i.e., which fields are included in a log record. Examples are file path, file hash, user ID, group ID, file signature, and extended file attributes. IMA comes with a set of predefined template formats and also allows a custom format, i.e., a format consisting of template fields supported by IMA. Template usage is typically determined by boot arguments passed to the kernel. Alternatively, the format can also be hard-coded into custom kernels. IMA templates and fields are extensible in the kernel source code. As a result, more template fields can be added in the future.

IMA policies define which files are measured using the IMA policy language. Built-in policies can be passed as boot arguments to the kernel. Custom IMA policies can be defined once during runtime or be hard-coded into a custom kernel. If no policy is defined, no measurements are taken and IMA is effectively disabled.

A comprehensive description of the content fields in native Linux IMA TLV format can be found in Table 16 of the Canonical Event Log (CEL) specification [cel]. The CEL specification also illustrates the use of templates to enable extended or customized IMA TLV formats in Section 5.1.6.

Appendix B. IMA for Network Equipment Boot Logs

Network equipment can generally implement similar IMA-protected functions to generate measurements (Claims) about the boot process of a device and enable corresponding remote attestation. Network Equipment Boot Logs combine the measurement and logging of boot components and operating system components (executables and files) into a single log file in a format identical to the IMA format. Note that the format used for logging measurement of boot components in this scheme differs from the boot logging strategy described

elsewhere in this document.

During the boot process of the network device, i.e., from BIOS to the end of the operating system and user-space, all files executed can be measured and logged in the order of their execution. When the Verifier initiates a remote attestation process (e.g., challenge-response remote attestation as defined in this document), the network equipment takes on the role of an Attester and can convey to the Verifier Claims that comprise the measurement log as well as the corresponding PCR values (Evidence) of a TPM.

The verifier can appraise the integrity (compliance with the Reference Values) of each executed file by comparing its measured value with the Reference Value. Based on the execution order, the Verifier can compute a PCR reference value (by replaying the log) and compare it to the Measurement Log Claims obtained in conjunction with the PCR Evidence to assess their trustworthiness with respect to an intended operational state.

Network equipment usually executes multiple components in parallel. This holds not only during the operating system loading phase, but also even during the BIOS boot phase. With this measurement log mechanism, network equipment can take on the role of an Attester, proving to the Verifier the trustworthiness of its boot process. Using the measurement log, Verifiers can precisely identify mismatching log entries to infer potentially tampered components.

This mechanism also supports scenarios that modify files on the Attester that are subsequently executed during the boot phase (e.g., updating/patching) by simply updating the appropriate Reference Values in Reference Integrity Manifests that inform Verifiers about how an Attester is composed.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany
Email: henk.birkholz@sit.fraunhofer.de

Michael Eckel
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: michael.eckel@sit.fraunhofer.de

Shwetha Bhandari
ThoughtSpot
Email: shwetha.bhandari@thoughtspot.com

Eric Voit
Cisco Systems
Email: evoit@cisco.com

Bill Sulzen
Cisco Systems
Email: bsulzen@cisco.com

Liang Xia (Frank)
Huawei Technologies
101 Software Avenue, Yuhuatai District
Nanjing
Jiangsu, 210012
China
Email: Frank.Xialiang@huawei.com

Tom Laffey
Hewlett Packard Enterprise
Email: tom.laffey@hpe.com

Guy C. Fedorkow
Juniper Networks
10 Technology Park Drive
Westford
Email: gfedorkow@juniper.net

IETF
Internet-Draft
Intended status: Standards Track
Expires: 13 June 2022

K. Moriarty
Center for Internet Security (CIS)
A. Fontes
Dell Technologies
10 December 2021

Scalable Remote Attestation for Systems, Containers, and Applications
draft-moriarty-attestationsets-04

Abstract

This document establishes an architectural pattern whereby a remote attestation could be issued for a complete set of benchmarks or controls that are defined and grouped by an external entity, preventing the need to send over individual attestations for each item within a benchmark or control framework. This document establishes a pattern to list sets of benchmarks and controls within CWT and JWT formats for use as an Entity Attestation Token (EAT).

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 13 June 2022.

Copyright Notice

Copyright (c) 2021 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
2. Policy and Measurement Set Definitions	4
3. Supportability and Re-Attestation	4
4. Configuration Sets	5
5. Remediation	5
6. Security Considerations	6
7. IANA Considerations	6
8. Contributors	6
9. References	6
9.1. Normative References	6
9.2. Informative References	6
Appendix A. Change Log	7
Appendix B. Open Issues	7
Authors' Addresses	7

1. Introduction

Posture assessment has long been desired, but has been difficult to achieve due to complexities of customization requirements at each organization. By using policy and measurement sets that may be offered at various assurance levels, automating posture assessment through attestation becomes achievable for organizations of all sizes. The measurement and policy groupings may be provided by the vendor or by a neutral third party to enable ease of use and consistent implementations. This provides simpler options to enable posture assessment at selected levels by organizations without the need to have in-house expertise. The measurement and policy sets may also be customized, but not necessary to achieve posture assessment to predefined options. This document describes a method to use existing attestation formats and protocols while allowing for profiles of policies, benchmarks, and measurements at defined assurance levels that scale to provide transparency to posture assessment results with remote attestation.

By way of example, the Center for Internet Security (CIS) hosts recommended configuration settings to secure operating systems, applications, and devices in CIS Benchmarks developed with industry

experts. Attestations aligned to the CIS Benchmarks or other configuration guide such as a DISA STIG could be used to assert the configuration meets expectations. This has already been done for multiple platforms to demonstrate assurance for firmware according to NIST SP 800-193, Firmware Resiliency Guidelines. In order to scale remote attestation, a single attestation for a set of Benchmarks or policies being met may be sent to the remote atteststation management system.

On traditional servers, assurance to NIST SP 800-193 is provable through attestation from a root of trust (RoT), using the Trusted Computing Group (TCG) Trusted Platform Module (TPM) chip and attestation formats. At boot, policy and measurement expectations are verified against a set of "golden policies" from collected and attested evidence. Device identity and measurements can also be attested at runtime. The attestations on evidence (e.g. hash of boot element) and verification of attestations are typically contained within a system and are limited to the control plane for management. The policy and measurement sets for comparison are protected to assure the result in the attestation verification process for boot element. Event logs and PCR values may be exposed to provide transparency into the verified attestations. Remote attestation on systems is intended to provide an assessment of posture for all managed systems and across various layers in each of these systems in an environment.

There is a balance of exposure and evidence needed to assess posture when providing assurance of controls and system state. Currently, logs and TPM PCR values may be passed to provide assurance of verification of attestation evidence meeting set requirements. Providing the assurance can be accomplished with a remote attestation format such as the Entity Attestation Token (EAT) [I-D.ietf-rats-eat] and a RESTful interface such as ROLIE or RedFish. Policy definition blocks may be scoped to control measurement sets, where the EAT asserts compliance to the policy or measurement block specified and may include claims with the log and PCR value evidence. Measurement and Policy sets may be published and maintained by separate entities (e.g. CIS Benchmarks, DISA STIGs). The policy and measurement sets should be maintained separately even if associated with the same benchmark or control set. This avoids the need to transition the verifying entity to a remote system for individual policy and measurements which are performed locally for more immediate remediation as well as other functions.

Examples of measurement and policy sets include, but are not limited to:

- * Hardware attribute certificates
- * Hardware Attribute Certificate Comparison Results
- * Reference Integrity Measurements for firmware
- * Operating system benchmarks at Specified Assurance Levels
- * Application hardening Benchmarks at Specified Assurance Levels
- * Container security benchmarks at Specified Assurance Levels

Scale, ease of use, full automation, and consistency for customer consumption of a remote attestation function or service are essential toward the goal of consistently securing systems against known threats and vulnerabilities. Mitigations may be baked into policy. Measurement verification sets and the attestation that the sets meet expected policies and measurements are conveyed in an Entity Attestation Token made available to a RESTful interface in aggregate for the systems managed.

2. Policy and Measurement Set Definitions

This document defines EAT claims in the JWT [RFC7519] and CWT [RFC8392] registries to provide attestation to a set of verified claims within a defined grouping. The trustworthiness will be conveyed on original verified evidence as well as the attestation on the grouping.

```

{
+-----+-----+-----+
+-----+
| Claim | Long Name | Description | Fo
rmat
+-----+-----+-----+
+-----+
| MPS | Measurement or Policy Set | Name for the MPS |
| LEM | Log Evidence of MPS | Log File or URI |
| PCR | TPM PCR Values | |
| FMA | Format of MPS Attestations | Format of included attestations |
| HSH | Hash Value/Message Digest | Hash value of configuration set |
+-----+-----+-----+
+-----+
}

```

3. Supportability and Re-Attestation

The remote attestation framework shall include provisions within the system and attestation authority to allow for Product modification.

Over its lifecycle, the Product may experience modification due to: maintenance, failures, upgrades, expansion, moves, etc..

The customer can chose to:

- * Run remote attestation after product modification, or
- * Not take action and remain un-protected

In the case of Re-Attestation:

- * framework needs to invalidate previous TPM PCR values and tokens,
- * framework needs to collect new measurements,
- * framework needs to maintain history or allow for history to be logged to enable change traceability attestation, and
- * framework needs to notify that the previous attestation has been invalidated

4. Configuration Sets

In some cases, it may be difficult to attest to configuration settings for the initial or subsequent attestation and verification processes. The use of an expected hash value for configuration settings can be used to compare the attested configuration set. In this case, the creator of the attestation verification measurements would define a set of values for which a message digest would be created and then signed by the attestor. The expected measurements would include the expected hash value for comparison. The configuration set could be the full attestation set to a Benchmark or a defined subset.

5. Remediation

If policy and configration settings or measurements attested do not meet expected values, remediation is desireable. Automated remediation performed with alignment to zero trust architecture principles would require that the remeidation be performed prior to any relying component executing. The relying component would verifiy before continuing in a zero trust architecture.

Ideally, remediation would occur on system as part of the process to attest to a set of attestations, similar to how attestation is performed for firmware in the boot process. If automated remediation is not possible, an alert should be generated to allow for notification of the variance from expected values.

6. Security Considerations

This document establishes a pattern to list sets of benchmarks and controls within CWT and JWT formats. The contents of the benchmarks and controls are out of scope for this document. This establishes an architectural pattern whereby a remote attestation could be issued for a complete set of benchmarks or controls as defined and grouped by external entities, preventing the need to send over individual attestations for each item within a benchmark or control framework. This document does not add security consideration over what has been described in the EAT, JWT, or CWT specifications.

7. IANA Considerations

This memo includes no request to IANA, yet. This will list the initial registration sets to the JWT and CWT registries if adopted.

8. Contributors

Thank you to reviewers and contributors who helped to improve this document. Thank you to Nick Grobelney, Dell Technologies, for your review and contribution to separate out the policy and measurement sets. Thank you, Samant Kakarla and Huijun Xie from Dell Technologies, for your detailed review and corrections on boot process details. Section 3 has been contributed by Rudy Bauer from Dell as well and an author will be added on the next revision.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392, DOI 10.17487/RFC8392, May 2018, <<https://www.rfc-editor.org/info/rfc8392>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.

9.2. Informative References

[I-D.ietf-rats-eat]

Lundblade, L., Mandyam, G., and J. O'Donoghue, "The Entity Attestation Token (EAT)", Work in Progress, Internet-Draft, draft-ietf-rats-eat-11, 24 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-rats-eat-11.txt>>.

Appendix A. Change Log

Note to RFC Editor: if this document does not obsolete an existing RFC, please remove this appendix before publication as an RFC.

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Authors' Addresses

Kathleen M. Moriarty
Center for Internet Security (CIS)
31 Tech Valley Drive
East Greenbush, NY,
United States of America

Email: Kathleen.Moriarty.ietf@gmail.com

Antonio Fontes
Dell Technologies
176 South Street
Hopkinton, MA,
United States of America

Email: Antonio.Fontes@dell.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 11, 2022

E. Voit
Cisco
H. Birkholz
Fraunhofer SIT
T. Hardjono
MIT
T. Fossati
Arm Limited
V. Scarlata
Intel
September 07, 2021

Attestation Results for Secure Interactions
draft-voit-rats-attestation-results-02

Abstract

This document defines reusable Attestation Result information elements. When these elements are offered to Relying Parties as Evidence, different aspects of Attester trustworthiness can be evaluated. Additionally, where the Relying Party is interfacing with a heterogeneous mix of Attesting Environment and Verifier types, consistent policies can be applied to subsequent information exchange between each Attester and the Relying Party.

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 March 11, 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Notation	4
1.2. Terminology	4
2. Attestation Results for Secure Interactions	5
2.1. Information driving a Relying Party Action	5
2.2. Non-repudiable Identity	6
2.2.1. Attester and Attesting Environment	7
2.2.2. Verifier	9
2.2.3. Communicating Identity	10
2.3. Trustworthiness Claims	10
2.3.1. Design Principles	10
2.3.2. Enumeration Encoding	12
2.3.3. Assigning a Trustworthiness Claim value	13
2.3.4. Specific Claims	13
2.3.5. Trustworthiness Vector	17
2.3.6. Trustworthiness Vector for a type of Attesting Environment	17
2.4. Freshness	18
3. Secure Interactions Models	18
3.1. Pure Background-Check retrieval	19
3.2. Attestation Result Augmented Evidence	19
3.3. Mutual Attestation	24
3.4. Transport Protocol Integration	24
4. Privacy Considerations	24
5. Security Considerations	24
6. IANA Considerations	24
7. References	24
7.1. Normative References	24
7.2. Informative References	25
Appendix A. Supportable Trustworthiness Claims	26
A.1. Supportable Trustworthiness Claims for HSM-based CC . . .	26
A.2. Supportable Trustworthiness Claims for process-based CC .	28
A.3. Supportable Trustworthiness Claims for VM-based CC . . .	29
Appendix B. Some issues being worked	30
Appendix C. Contributors	31
Authors' Addresses	31

1. Introduction

The first paragraph of the May 2021 US Presidential Executive Order on Improving the Nation's Cybersecurity [US-Executive-Order] ends with the statement "the trust we place in our digital infrastructure should be proportional to how trustworthy and transparent that infrastructure is." Later this order explores aspects of trustworthiness such as an auditable trust relationship, which it defines as an "agreed-upon relationship between two or more system elements that is governed by criteria for secure interaction, behavior, and outcomes."

The Remote Attestation procedureS (RATS) architecture [I-D.ietf-rats-architecture] provides a useful context for programmatically establishing and maintaining such auditable trust relationships. Specifically, the architecture defines conceptual messages conveyed between architectural subsystems to support trustworthiness appraisal. The RATS conceptual message used to convey evidence of trustworthiness is the Attestation Results. The Attestation Results includes Verifier generated appraisals of an Attester including such information as the identity of the Attester, the security mechanisms employed on this Attester, and the Attester's current state of trustworthiness.

Generated Attestation Results are ultimately conveyed to one or more Relying Parties. Reception of an Attestation Result enables a Relying Party to determine what action to take with regards to an Attester. Frequently, this action will be to choose whether to allow the Attester to securely interact with the Relying Party over some connection between the two.

When determining whether to allow secure interactions with an Attester, a Relying Party is challenged with a number of difficult problems which it must be able to handle successfully. These problems include:

- o What Attestation Results (AR) might a Relying Party be willing to trust from a specific Verifier?
- o What information does a Relying Party need before allowing interactions or choosing policies to apply to a connection?
- o What are the operating/environmental realities of the Attesting Environment where a Relying Party should only be able to associate a certain confidence regarding Attestation Results out of the Verifier? (In other words, different types of Trusted Execution Environments (TEE) need not be treated as equivalent.)

- o How to make direct comparisons where there is a heterogeneous mix of Attesting Environments and Verifier types.

To address these problems, it is important that specific Attestation Result information elements are framed independently of Attesting Environment specific constraints. If they are not, a Relying Party would be forced to adapt to the syntax and semantics of many vendor specific environments. This is not a reasonable ask as there can be many types of Attesters interacting with or connecting to a Relying Party.

The business need therefore is for common Attestation Result information element definitions. With these definitions, consistent interaction or connectivity decisions can be made by a Relying Party where there is a heterogeneous mix of Attesting Environment types and Verifier types.

This document defines information elements for Attestation Results in a way which normalizes the trustworthiness assertions that can be made from a diverse set of Attesters.

1.1. Requirements Notation

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.

1.2. Terminology

The following terms are imported from [I-D.ietf-rats-architecture]: Appraisal Policy for Attestation Results, Attester, Attesting Environment, Claims, Evidence, Relying Party, Target Environment and Verifier.

[I-D.ietf-rats-architecture] also describes topological patterns that illustrate the need for interoperable conceptual messages. The two patterns called "background-check model" and "passport model" are imported from the RATS architecture and used in this document as a reference to the architectural concepts: Background-Check Model and Passport Model.

Newly defined terms for this document:

AR-augmented Evidence: a bundle of Evidence which includes at least the following:

1. Verifier signed Attestation Results. These Attestation Results must include Identity Evidence for the Attester, a Trustworthiness Vector describing a Verifier's most recent appraisal of an Attester, and some Verifier Proof-of-Freshness (PoF).
2. A Relying Party PoF which is bound to the Attestation Results of (1) by the Attester's Attesting Environment signature.
3. Sufficient information to determine the elapsed interval between the Verifier PoF and Relying Party PoF.

Identity Evidence: Evidence which unambiguously identifies an identity. Identity Evidence could take different forms, such as a certificate, or a signature which can be appraised to have only been generated by a specific private/public key pair.

Trustworthiness Claim: a specific quanta of trustworthiness which can be assigned by a Verifier based on its appraisal policy.

Trustworthiness Tier: a categorization of the levels of trustworthiness which may be assigned by a Verifier to a specific Trustworthiness Claim. These enumerated categories are: Affirmed, Warning, Contraindicated, and None.

Trustworthiness Vector: a set of zero to many Trustworthiness Claims assigned during a single appraisal procedure by a Verifier using Evidence generated by an Attester. The vector is included within Attestation Results.

2. Attestation Results for Secure Interactions

A Verifier generates the Attestation Results used by a Relying Party. When a Relying Party needs to determine whether to permit communications with an Attester, these Attestation Results must contain a specific set of information elements. This section defines those information elements, and in some cases encodings for information elements.

2.1. Information driving a Relying Party Action

When the action is a communication establishment attempt with an Attester, there is only a limited set of actions which a Relying Party might take. These actions include:

- o Allow or deny information exchange with the Attester. When there is a deny, reasons should be returned to the Attester.

- o Establish a transport connection between an Attester and a specific context within a Relying Party (e.g., a TEE, or Virtual Routing Function (VRF).)
- o Apply policies on this connection (e.g., rate limits).

There are three categories of information which must be conveyed to the Relying Party (which also is integrated with a Verifier) before it determines which of these actions to take.

1. Non-repudiable Identity Evidence - Evidence which undoubtably identifies one or more entities involved with a connection.
2. Trustworthiness Claims - Specifics a Verifier asserts with regards to its trustworthiness findings about an Attester.
3. Claim Freshness - Establishes the time of last update (or refresh) of Trustworthiness Claims.

The following sections detail requirements for these three categories.

2.2. Non-repudiable Identity

Identity Evidence must be conveyed during the establishment of any trust-based relationship. Specific use cases will define the minimum types of identities required by a particular Relying Party as it evaluates Attestation Results, and perhaps additional associated Evidence. At a bare minimum, a Relying Party MUST start with the ability to verify the identity of a Verifier it chooses to trust. Attester identities may then be acquired through signed communications with the Verifier identity and/or the pre-provisioning Attester public keys in the Attester.

During the Remote Attestation process, the Verifier's identity will be established with a Relying Party via a Verifier signature across recent Attestation Results. This Verifier identity could only have come from a key pair maintained by a trusted developer or operator of the Verifier.

Additionally, each set of Attestation Results must be provably and non-reputably bound to the identity of the original Attesting Environment which was evaluated by the Verifier. This will be accomplished via two items.

First the Verifier signed Attestation Results MUST include sufficient Identity Evidence to ensure that this Attesting Environment signature refers to the same Attesting Environment appraised by the Verifier. Second, where the passport model is used as a subsystem, an Attesting

Environment signature which spans the Verifier signature MUST also be included. As the Verifier signature already spans the Attester Identity as well as the Attestation Results, this restricts the viability of spoofing attacks.

In a subset of use cases, these two pieces of Identity Evidence may be sufficient for a Relying Party to successfully meet the criteria for its Appraisal Policy for Attestation Results. If the use case is a connection request, a Relying Party may simply then establish a transport session with an Attester after a successful appraisal. However an Appraisal Policy for Attestation Results will often be more nuanced, and the Relying Party may need additional information. Some Identity Evidence related policy questions which the Relying Party may consider include:

- o Does the Relying Party only trust this Verifier to make Trustworthiness Claims on behalf a specific type of Attesting Environment? Might a mix of Verifiers be necessary to cover all mandatory Trustworthiness Claims?
- o Does the Relying Party only accept connections from a verified-authentic software build from a specific software developer?
- o Does the Relying Party only accept connections from specific preconfigured list of Attesters?

For any of these more nuanced appraisals, additional Identity Evidence or other policy related information must be conveyed or pre-provisioned during the formation of a trust context between the Relying Party, the Attester, the Attester's Attesting Environment, and the Verifier.

2.2.1. Attester and Attesting Environment

Per [I-D.ietf-rats-architecture] Figure 2, an Attester and a corresponding Attesting Environment might not share common code or even hardware boundaries. Consequently, an Attester implementation needs to ensure that any Evidence which originates from outside the Attesting Environment MUST have been collected and delivered securely before any Attesting Environment signing may occur. After the Verifier performs its appraisal, it will include sufficient information in the Attestation Results to enable a Relying Party to have confidence that the Attester's trustworthiness is represented via Trustworthiness Claims signed by the appropriate Attesting Environment.

This document recognizes three general categories of Attesters.

1. HSM-based: A Hardware Security Module (HSM) based cryptoprocessor which continually hashes security measurements in a way which prevents an Attester from lying about measurements which have been extended into the Attesting Environment (e.g., TPM2.0.)
2. Process-based: An individual process which has its runtime memory encrypted by an Attesting Environment in a way that no other processes can read and decrypt that memory (e.g., [SGX] or [I-D.tschofenig-rats-psa-token].)
3. VM-based: An entire Guest VM (or a set of containers within a host) have been encrypted as a walled-garden unit by an Attesting Environment. The result is that the host operating system cannot read and decrypt what is executing within that VM (e.g., [SEV-SNP] or [TDX].)

Each of these categories of Attesters above will be capable of generating Evidence which is protected using private keys / certificates which are not accessible outside of the corresponding Attesting Environment. The owner of these secrets is the owner of the identity which is bound within the Attesting Environment. Effectively this means that for any Attester identity, there will exist a chain of trust ultimately bound to a hardware-based root of trust in the Attesting Environment. It is upon this root of trust that unique, non-repudiable Attester identities may be founded.

There are several types of Attester identities defined in this document. This list is extensible:

- o chip-vendor: the vendor of the hardware chip used for the Attesting Environment (e.g., a primary Endorsement Key from a TPM)
- o chip-hardware: specific hardware with specific firmware from an 'ae-vendor'
- o target-environment: a unique instance of a software build running in an Attester (e.g., MRENCLAVE [SGX], an Instance ID [I-D.tschofenig-rats-psa-token], an Identity Block [SEV-SNP], or a hash which represents a set of software loaded since boot (e.g., TPM based integrity verification.))
- o target-developer: the organizational unit responsible for a particular 'target-environment' (e.g., MRSIGNER [SGX])
- o instance: a unique instantiated instance of an Attesting Environment running on 'chip-hardware' (e.g., an LDevID [IEEE802.1AR])

Based on the category of the Attesting Environment, different types of identities might be exposed by an Attester.

Attester Identity type	Process-based	VM-based	HSM-based
chip-vendor	Mandatory	Mandatory	Mandatory
chip-hardware	Mandatory	Mandatory	Mandatory
target-environment	Mandatory	Mandatory	Optional
target-developer	Mandatory	Optional	Optional
instance	Optional	Optional	Optional

It is expected that drafts subsequent to this specification will provide the definitions and value domains for specific identities, each of which falling within the Attester identity types listed above. In some cases the actual unique identities might be encoded as complex structures. An example complex structure might be a 'target-environment' encoded as a Software Bill of Materials (SBOM).

With the identity definitions and value domains, a Relying Party will have sufficient information to ensure that the Attester identities and Trustworthiness Claims asserted are actually capable of being supported by the underlying type of Attesting Environment. Consequently, the Relying Party SHOULD require Identity Evidence which indicates of the type of Attesting Environment when it considers its Appraisal Policy for Attestation Results.

2.2.2. Verifier

For the Verifier identity, it is critical for a Relying Party to review the certificate and chain of trust for that Verifier. Additionally, the Relying Party must have confidence that the Trustworthiness Claims being relied upon from the Verifier considered the chain of trust for the Attesting Environment.

There are two categorizations of Verifier identities defined in this document.

- o verifier build: a unique instance of a software build running as a Verifier.
- o verifier developer: the organizational unit responsible for a particular 'verifier build'.

Within each category, communicating the identity can be accomplished via a variety of objects and encodings.

2.2.3. Communicating Identity

Any of the above identities used by the Appraisal Policy for Attestation Results needed to be pre-established by the Relying Party before, or provided during, the exchange of Attestation Results. When provided during this exchange, the identity may be communicated either implicitly or explicitly.

An example of explicit communication would be to include the following Identity Evidence directly within the Attestation Results: a unique identifier for an Attesting Environment, the name of a key which can be provably associated with that unique identifier, and the set of Attestation Results which are signed using that key. As these Attestation Results are signed by the Verifier, it is the Verifier which is explicitly asserting the credentials it believes are trustworthy.

An example of implicit communication would be to include Identity Evidence in the form of a signature which has been placed over the Attestation Results asserted by a Verifier. It would be then up to the Relying Party's Appraisal Policy for Attestation Results to extract this signature and confirm that it only could have been generated by an Attesting Environment having access to a specific private key. This implicit identity communication is only viable if the Attesting Environment's public key is already known by the Relying Party.

One final step in communicating identity is proving the freshness of the Attestation Results to the degree needed by the Relying Party. A typical way to accomplish this is to include an element of freshness be embedded within a signed portion of the Attestation Results. This element of freshness reduces the identity spoofing risks from a replay attack. For more on this, see Section 2.4.

2.3. Trustworthiness Claims

2.3.1. Design Principles

Trust is not absolute. Trust is a belief in some aspect about an entity (in this case an Attester), and that this aspect is something which can be depended upon (in this case by a Relying Party.) Within the context of Remote Attestation, believability of this aspect is facilitated by a Verifier. This facilitation depends on the Verifier's ability to parse detailed Evidence from an Attester and

then to assert conclusions about this aspect in a way interpretable by a Relying Party.

Specific aspects for which a Verifier will assert trustworthiness are defined in this section. These are known as Trustworthiness Claims. These claims have been designed to enable a common understanding between a broad array of Attesters, Verifiers, and Relying Parties. The following set of design principles have been applied in the Trustworthiness Claim definitions:

1. Expose a small number of Trustworthiness Claims.

Reason: 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.

Reason: 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 in order to accomplish the goals of (1) and (2).

Reason: 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 for (1) and (2).

Reason: 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.

These design principles are important to keep the number of Verifier generated claims low, and to retain the complexity in the Verifier rather than the Relying Party.

2.3.2. Enumeration Encoding

Per design principle (2), each Trustworthiness Claim will only expose specific encoded values. To simplify the processing of these enumerations by the Relying Party, the enumeration will be encoded as a single signed 8 bit integer. These value assignments for this integer will be in four Trustworthiness Tiers which follow these guidelines:

Affirming: The Verifier affirms the Attester support for this aspect of trustworthiness

- o Values 1 to 31: A standards enumerated reason for affirming.
- o Values -2 to -32: A non-standard reason for affirming.

Warning: The Verifier warns about this aspect of trustworthiness.

- o Values 32 to 95: A standards enumerated reason for the warning.
- o Values -33 to -96: A non-standard reason for the warning.

Contraindicated: The Verifier asserts the Attester is explicitly untrustworthy in regard to this aspect.

- o Values 96 to 127: A standards enumerated reason for the contraindication.
- o Values -97 to -128: A non-standard reason for the contraindication.

None: The Verifier makes no assertions about this Trustworthiness Claim.

- o Value 0: Note: this should always be always treated equivalently by the Relying Party as no claim being made. I.e., the RP's Appraisal Policy for Attestation Results SHOULD NOT make any distinction between a Trustworthiness Claim with enumeration '0', and no Trustworthiness Claim being provided.
- o Value -1: An unexpected error occurred during the Verifier's appraisal processing. Note: while no claim is being made, the Relying Party MAY make a distinction between a Trustworthiness

Claim with enumeration '-1', and no Trustworthiness Claim being provided.

This enumerated encoding listed above will simplify the Appraisal Policy for Attestation Results. Such a policies may be as simple as saying that a specific Verifier has recently asserted Trustworthiness Claims, all of which are Affirming.

2.3.3. Assigning a Trustworthiness Claim value

In order to simplify design, only a single encoded value is asserted by a Verifier for any Trustworthiness Claim within a using the following process.

1. If applicable, a Verifier MUST assign a standardized value from the Contraindicated tier.
2. Else if applicable, a Verifier MUST assign a non-standardized value from the Contraindicated tier.
3. Else if applicable, a Verifier MUST assign a standardized value from the Warning tier.
4. Else if applicable, a Verifier MUST assign a non-standardized value from the Warning tier.
5. Else if applicable, a Verifier MUST assign a standardized value from the Affirming tier.
6. Else if applicable, a Verifier MUST assign a non-standardized value from the Affirming tier.
7. Else a Verifier MAY assign a 0 or -1.

2.3.4. Specific Claims

Following are the Trustworthiness Claims and their supported enumerations which may be asserted by a Verifier:

configuration: A Verifier has appraised an Attester's configuration, and is able to make conclusions regarding the exposure of known vulnerabilities

0: No assertion

1: The configuration is a known and approved config

2: The configuration includes or exposes no known vulnerabilities

32: The configuration includes or exposes known vulnerabilities

96: The configuration is unsupportable as it exposes unacceptable security vulnerabilities

-1: Unexpected error

executables: A Verifier has appraised and evaluated relevant runtime files, scripts, and/or other objects which have been loaded into the Target environment's memory.

0: No assertion

1: Only a recognized genuine set of approved executables, scripts, files, and/or objects have been loaded during and after the boot process.

2: Only a recognized genuine set of approved executables have been loaded during the boot process.

32: Only a recognized genuine set of executables, scripts, files, and/or objects have been loaded. However the Verifier cannot vouch for a subset of these due to known bugs or other known vulnerabilities.

33: Runtime memory includes executables, scripts, files, and/or objects which are not recognized.

96: Runtime memory includes executables, scripts, files, and/or object which are contraindicated.

-1: Unexpected error

file-system: A Verifier has evaluated the Attester's file system.

0: No assertion

1: Only a recognized set of approved files are found.

32: The file system includes unrecognized executables, scripts, or files.

96: The file system includes contraindicated executables, scripts, or files

-1: Unexpected error

hardware: A Verifier has appraised any Attester hardware and firmware which are able to expose fingerprints of their identity and running code.

0: No assertion

1: An Attester has passed its hardware and/or firmware verifications needed to demonstrate that these are genuine/supported.

32: An Attester contains only genuine/supported hardware and/or firmware, but there are known security vulnerabilities.

96: Attester hardware and/or firmware is recognized, but its trustworthiness is contraindicated.

97: A Verifier does not recognize an Attester's hardware or firmware, but it should be recognized.

-1: Unexpected error

instance-identity: A Verifier has appraised an Attesting Environment's unique identity based upon private key signed Evidence which can be correlated to a unique instantiated instance of the Attester. (Note: this Trustworthiness Claim should only be generated if the Verifier actually expects to recognize the unique identity of the Attester.)

0: No assertion

1: The Attesting Environment is recognized, and the associated instance of the Attester is not known to be compromised.

96: The Attesting Environment is recognized, and but its unique private key indicates a device which is not trustworthy.

97: The Attesting Environment is not recognized; however the Verifier believes it should be.

-1: Unexpected error

runtime-opaque: A Verifier has appraised the visibility of Attester objects in memory from perspectives outside the Attester.

0: No assertion

1: the Attester's executing Target Environment and Attesting Environments are encrypted and within Trusted Execution

Environment(s) opaque to the operating system, virtual machine manager, and peer applications. (Note: This value corresponds to the protections asserted by O.RUNTIME_CONFIDENTIALITY from [GP-TEE-PP])

32: the Attester's executing Target Environment and Attesting Environments inaccessible from any other parallel application or Guest VM running on the Attester's physical device. (Note that unlike "1" these environments are not encrypted in a way which restricts the Attester's root operator visibility. See O.TA_ISOLATION from [GP-TEE-PP].)

96: The Verifier has concluded that in memory objects are unacceptably visible within the physical host that supports the Attester.

-1: Unexpected error

sourced-data: A Verifier has evaluated of the integrity of data objects from external systems used by the Attester.

0: No assertion

1: All essential Attester source data objects have been provided by other Attester(s) whose most recent appraisal(s) had both no Trustworthiness Claims of "0" where the current Trustworthiness Claim is "Affirming", as well as no "Warning" or "Contraindicated" Trustworthiness Claims.

32: Attester source data objects come from unattested sources, or attested sources with "Warning" type Trustworthiness Claims.

96: Attester source data objects come from contraindicated sources.

-1: Unexpected error

storage-opaque: A Verifier has appraised that an Attester is capable of encrypting persistent storage. (Note: Protections must meet the capabilities of [OMTP-ATE] Section 5, but need not be hardware tamper resistant.)

0: No assertion

1: the Attester encrypts all secrets in persistent storage via using keys which are never visible outside an HSM or the Trusted Execution Environment hardware.

32: the Attester encrypts all persistently stored secrets, but without using hardware backed keys

96: There are persistent secrets which are stored unencrypted in an Attester.

-1: Unexpected error

It is possible for additional Trustworthiness Claims and enumerated values to be defined in subsequent documents. At the same time, the standardized Trustworthiness Claim values listed above have been designed so there is no overlap within a Trustworthiness Tier. As a result, it is possible to imagine a future where overlapping Trustworthiness Claims within a single Trustworthiness Tier may be defined. Wherever possible, the Verifier SHOULD assign the best fitting standardized value.

Where a Relying Party doesn't know how to handle a particular Trustworthiness Claim, it MAY choose an appropriate action based on the Trustworthiness Tier under which the enumerated value fits.

It is up to the Verifier to publish the types of evaluations it performs when determining how Trustworthiness Claims are derived for a type of any particular type of Attester. It is out of the scope of this document for the Verifier to provide proof or specific logic on how a particular Trustworthiness Claim which it is asserting was derived.

2.3.5. Trustworthiness Vector

Multiple Trustworthiness Claims may be asserted about an Attesting Environment at single point in time. The set of Trustworthiness Claims inserted into an instance of Attestation Results by a Verifier is known as a Trustworthiness Vector. The order of Claims in the vector is NOT meaningful. A Trustworthiness Vector with no Trustworthiness Claims (i.e., a null Trustworthiness Vector) is a valid construct. In this case, the Verifier is making no Trustworthiness Claims but is confirming that an appraisal has been made.

2.3.6. Trustworthiness Vector for a type of Attesting Environment

Some Trustworthiness Claims are implicit based on the underlying type of Attesting Environment. For example, a validated MRSIGNER identity can be present where the underlying [SGX] hardware is 'hw-authentic'. Where such implicit Trustworthiness Claims exist, they do not have to be explicitly included in the Trustworthiness Vector. However, these implicit Trustworthiness Claims SHOULD be considered as being present

by the Relying Party. Another way of saying this is if a Trustworthiness Claim is automatically supported as a result of coming from a specific type of TEE, that claim need not be redundantly articulated. Such implicit Trustworthiness Claims can be seen in the tables within Appendix A.2 and Appendix A.3.

Additionally, there are some Trustworthiness Claims which cannot be adequately supported by an Attesting Environment. For example, it would be difficult for an Attester that includes only a TPM (and no other TEE) from ever having a Verifier appraise support for 'runtime-opaque'. As such, a Relying Party would be acting properly if it rejects any non-supportable Trustworthiness Claims asserted from a Verifier.

As a result, the need for the ability to carry a specific Trustworthiness Claim will vary by the type of Attesting Environment. Example mappings can be seen in Appendix A.

2.4. Freshness

A Relying Party will care about the recentness of the Attestation Results, and the specific Trustworthiness Claims which are embedded. All freshness mechanisms of [I-D.ietf-rats-architecture], Section 10 are supportable by this specification.

Additionally, a Relying Party may track when a Verifier expires its confidence for the Trustworthiness Claims or the Trustworthiness Vector as a whole. Mechanisms for such expiry are not defined within this document.

There is a subset of secure interactions where the freshness of Trustworthiness Claims may need to be revisited asynchronously. This subset is when trustworthiness depends on the continuous availability of a transport session between the Attester and Relying Party. With such connectivity dependent Attestation Results, if there is a reboot which resets transport connectivity, all established Trustworthiness Claims should be cleared. Subsequent connection re-establishment will allow fresh new Trustworthiness Claims to be delivered.

3. Secure Interactions Models

There are multiple ways of providing a Trustworthiness Vector to a Relying Party. This section describes two alternatives.

3.1. Pure Background-Check retrieval

It is possible to for a Relying Party to follow the Background-Check Model defined in Section 5.2 of [I-D.ietf-rats-architecture]. In this case, a Relying Party will receive Attestation Results containing the Trustworthiness Vector directly from a Verifier. These Attestation Results can then be used by the Relying Party in determining the appropriate treatment for interactions with the Attester.

While applicable in some cases, the utilization of the Background-Check Model without modification has potential drawbacks in other cases. These include:

- o Verifier scale: if the Attester has many Relying Parties, a Verifier appraising that Attester could be frequently be queried based on the same Evidence.
- o Information leak: Evidence which the Attester might consider private can be visible to the Relying Party. Hiding that Evidence could devalue any resulting appraisal.
- o Latency: a Relying Party will need to wait for the Verifier to return Attestation Results before proceeding with secure interactions with the Attester.

An implementer should examine these potential drawbacks before selecting this alternative.

3.2. Attestation Result Augmented Evidence

There is a hybrid alternative for the establishment and maintenance of trustworthiness between an Attester and a Relying Party which is not adversely impacted by the potential drawbacks with pure background-check. In this alternative, a Verifier evaluates an Attester and returns signed Attestation Results back to this original Attester no less frequently than a well-known interval. This interval may also be asynchronous, based on the changing of certain Evidence as described in [I-D.birkholz-rats-network-device-subscription].

When a Relying Party is to receive information about the Attester's trustworthiness, the Attesting Environment assembles the minimal set of Evidence which can be used to confirm or refute whether the Attester remains in the state of trustworthiness represented by the AR. To this Evidence, the Attesting Environment appends the signature from the most recent AR as well as a Relying Party Proof-of-Freshness. The Attesting Environment then signs the combination.

The Attester then assembles AR Augmented Evidence by taking the signed combination and appending the full AR. The assembly now consists of two independent but semantically bound sets of signed Evidence.

The AR Augmented Evidence is then sent to the Relying Party. The Relying Party then can appraise these semantically bound sets of signed Evidence by applying an Appraisal Policy for Attestation Results as described below. This policy will consider both the AR as well as additional information about the Attester within the AR Augmented Evidence the when determining what action to take.

This alternative combines the [I-D.ietf-rats-architecture] Sections 5.1 Passport Model and Section 5.2 Background-Check Model. Figure 1 describes this flow of information. The flows within this combined model are mapped to [I-D.ietf-rats-architecture] in the following way. "Verifier A" below corresponds to the "Verifier" Figure 5 within [I-D.ietf-rats-architecture]. And "Relying Party/Verifier B" below corresponds to the union of the "Relying Party" and "Verifier" boxes within Figure 6 of [I-D.ietf-rats-architecture]. This union is possible because Verifier B can be implemented as a simple, self-contained process. The resulting combined process can appraise the AR-augmented Evidence to determine whether an Attester qualifies for secure interactions with the Relying Party. The specific steps of this process are defined later in this section.

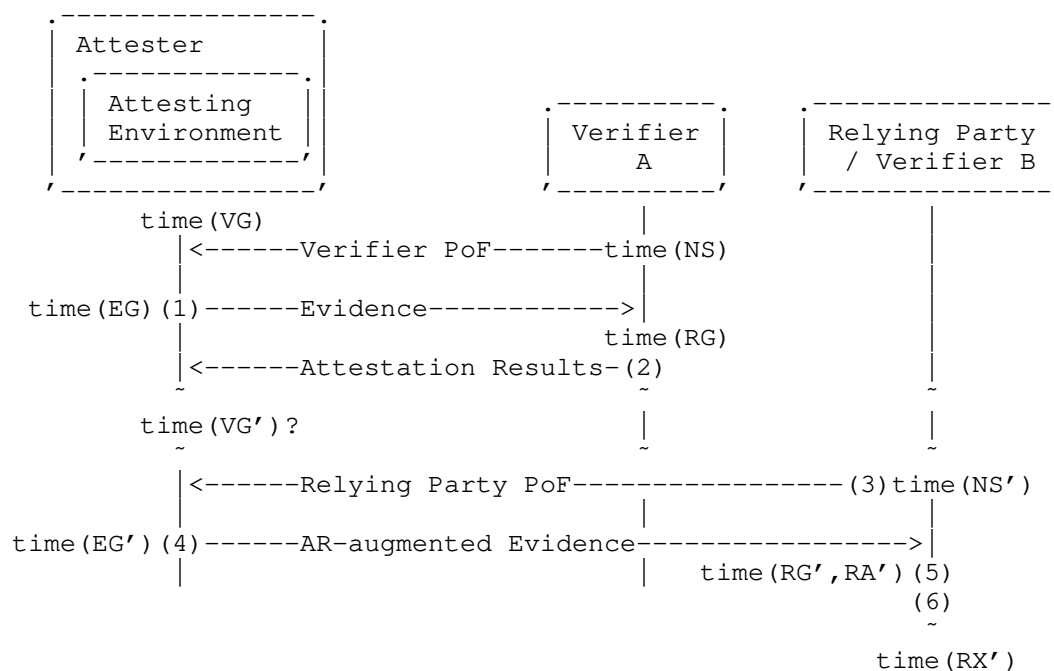


Figure 1: Secure Interactions Model

The interaction model depicted above includes specific time related events from Appendix A of [I-D.ietf-rats-architecture]. With the identification of these time related events, time duration/interval tracking becomes possible. Such duration/interval tracking can become important if the Relying Party cares if too much time has elapsed between the Verifier PoF and Relying Party PoF. If too much time has elapsed, perhaps the Attestation Results themselves are no longer trustworthy.

Note that while time intervals will often be relevant, there is a simplified case that does not require a Relying Party's PoF in step (3). In this simplified case, the Relying Party trusts that the Attester cannot be meaningfully changed from the outside during any reportable interval. Based on that assumption, and when this is the case then the step of the Relying Party PoF can be safely omitted.

In all cases, appraisal policies define the conditions and prerequisites for when an Attester does qualify for secure interactions. To qualify, an Attester has to be able to provide all of the mandatory affirming Trustworthiness Claims and identities needed by a Relying Party's Appraisal Policy for Attestation Results, and none of the disqualifying detracting Trustworthiness Claims.

More details on each interaction step are as follows. The numbers used in this sequence match to the numbered steps in Figure 1:

1. An Attester sends Evidence which is provably fresh to Verifier A at time(EG). Freshness from the perspective of Verifier A MAY be established with Verifier PoF such as a nonce.
2. Verifier A appraises (1), then sends the following items back to that Attester within Attestation Results:
 1. the verified identity of the Attesting Environment,
 2. the Verifier A appraised Trustworthiness Vector of an Attester,
 3. a freshness proof associated with the Attestation Results,
 4. a Verifier signature across (2.1) through (2.3).
3. At time(EG') a Relying Party PoF (such as a nonce) known to the Relying Party is sent to the Attester.
4. The Attester generates and sends AR-augmented Evidence to the Relying Party/Verifier B. This AR-augmented Evidence includes:
 1. The Attestation Results from (2)
 2. Any (optionally) new incremental Evidence from the Attesting Environment
 3. Attestation Environment signature which spans a hash of the Attestation Results (such as the signature of (2.4)), the proof-of-freshness from (3), and (4.2). Note: this construct allows the delta of time between (2.3) and (3) to be definitively calculated by the Relying Party.
5. On receipt of (4), the Relying Party applies its Appraisal Policy for Attestation Results. At minimum, this appraisal policy process must include the following:
 1. Verify that (4.3) includes the nonce from (3).
 2. Use a local certificate to validate the signature (4.1).
 3. Verify that the hash from (4.3) matches (4.1)
 4. Use the identity of (2.1) to validate the signature of (4.3).

5. Failure of any steps (5.1) through (5.4) means the link does not meet minimum validation criteria, therefore appraise the link as having a null Verifier B Trustworthiness Vector. Jump to step (6.1).
6. When there is large or uncertain time gap between time(EG) and time(EG'), the link should be assigned a null Verifier B Trustworthiness Vector. Jump to step (6.1).
7. Assemble the Verifier B Trustworthiness Vector
 1. Copy Verifier A Trustworthiness Vector to Verifier B Trustworthiness Vector
 2. Add implicit Trustworthiness Claims inherent to the type of TEE.
 3. Prune any Trustworthiness Claims unsupportable by the Attesting Environment.
 4. Prune any Trustworthiness Claims the Relying Party doesn't accept from this Verifier.
6. The Relying Party takes action based on Verifier B's appraised Trustworthiness Vector:
 1. Prune any Trustworthiness Claims not used in the Appraisal Policy for Attestation Results.
 2. Allow the information exchange from the Attester into a Relying Party context where the Verifier B appraised Trustworthiness Vector includes all the mandatory Trustworthiness Claims of value "1", and none of the disqualifying Trustworthiness Claims containing values that are not "0" or "1".
 3. Disallow any information exchange into a Relying Party context for which that Verifier B appraised Trustworthiness Vector is not qualified.

As link layer protocols re-authenticate, steps (1) to (2) and steps (3) to (6) will independently refresh. This allows the Trustworthiness of Attester to be continuously re-appraised. There are only specific event triggers which will drive the refresh of Evidence generation (1), Attestation Result generation (2), or AR-augmented Evidence generation (4):

- o life-cycle events, e.g. a change to an Authentication Secret of the Attester or an update of a software component.
- o uptime-cycle events, e.g. a hard reset or a re-initialization of an Attester.
- o authentication-cycle events, e.g. a link-layer interface reset could result in a new (4).

3.3. Mutual Attestation

In the interaction models described above, each device on either side of a secure interaction may require remote attestation of its peer. This process is known as mutual-attestation. To support mutual-attestation, the interaction models listed above may be run independently on either side of the connection.

3.4. Transport Protocol Integration

Either unidirectional attestation or mutual attestation may be supported within the protocol interactions needed for the establishment of a single transport session. While this document does not mandate specific transport protocols, messages containing the Attestation Results and AR Augmented Evidence can be passed within an authentication framework such the EAP protocol [RFC5247] over TLS [RFC8446].

4. Privacy Considerations

Privacy Considerations Text

5. Security Considerations

Security Considerations Text

6. IANA Considerations

See Body.

7. References

7.1. Normative References

[GP-TEE-PP]

"Global Platform TEE Protection Profile v1.3", September 2020, <<https://globalplatform.org/specs-library/tee-protection-profile-v1-3/>>.

[I-D.ietf-rats-architecture]

Fraunhofer SIT, Microsoft, Sandelman Software Works, Intel Corporation, and Huawei Technologies, "Remote Attestation Procedures Architecture", draft-ietf-rats-architecture-12 (work in progress), April 2021.

[OMTP-ATE]

"Open Mobile Terminal Platform - Advanced Trusted Environment", May 2009, <<https://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptrlv11.pdf>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

7.2. Informative References

[I-D.birkholz-rats-network-device-subscription]

Fraunhofer SIT, Cisco Systems, Inc., and Huawei Technologies, "Attestation Event Stream Subscription", draft-birkholz-rats-network-device-subscription-03 (work in progress), August 2021.

[I-D.tschofenig-rats-psa-token]

Arm Limited, Arm Limited, Arm Limited, Arm Limited, and Arm Limited, "Arm's Platform Security Architecture (PSA) Attestation Token", draft-tschofenig-rats-psa-token-08 (work in progress), March 2021.

[IEEE802.1AR]

"802.1AR: Secure Device Identity", August 2018, <<https://ieeexplore.ieee.org/document/8423794>>.

[RFC5247] Aboba, B., Simon, D., and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", RFC 5247, DOI 10.17487/RFC5247, August 2008, <<https://www.rfc-editor.org/info/rfc5247>>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

- [SEV-SNP] "AMD SEV-SNP: Strengthening VM Isolation with Integrity Protection and More", 2020, <<https://www.amd.com/system/files/TechDocs/SEV-SNP-strengthening-vm-isolation-with-integrity-protection-and-more.pdf>>.
- [SGX] "Supporting Third Party Attestation for Intel SGX with Intel Data Center Attestation Primitives", 2017, <<https://software.intel.com/content/dam/develop/external/us/en/documents/intel-sgx-support-for-third-party-attestation-801017.pdf>>.
- [TDX] "Intel Trust Domain Extensions", 2020, <<https://software.intel.com/content/dam/develop/external/us/en/documents/tdx-whitepaper-final9-17.pdf>>.
- [TPM-ID] "TPM Keys for Platform Identity for TPM 1.2", August 2015, <https://www.trustedcomputinggroup.org/wp-content/uploads/TPM_Keys_for_Platform_Identity_v1_0_r3_Final.pdf>.
- [US-Executive-Order] "Executive Order on Improving the Nation's Cybersecurity", May 2021, <<https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>>.

Appendix A. Supportable Trustworthiness Claims

The following is a table which shows what Claims are supportable by different Attesting Environment types. Note that claims MAY BE implicit to an Attesting Environment type, and therefore do not have to be included in the Trustworthiness Vector to be considered as set by the Relying Party.

A.1. Supportable Trustworthiness Claims for HSM-based CC

Following are Trustworthiness Claims which MAY be set for a HSM-based Confidential Computing Attester. (Such as a TPM [TPM-ID].)

Trustworthiness Claim	Required?	Appraisal Method
configuration	Optional	Verifier evaluation of Attester reveals no configuration lines which expose the Attester to known security vulnerabilities. This may be done with or without the involvement of a TPM PCR.
executables	Yes	Checks the TPM PCRs for the static operating system, and for any tracked files subsequently loaded
file-system	No	Can be supported, but TPM tracking is unlikely
hardware	Yes	If TPM PCR check ok from BIOS checks, through Master Boot Record configuration
instance-identity	Optional	Check IDevID
runtime-opaque	n/a	TPMs are not recommended to provide a sufficient technology base for this Trustworthiness Claim.
sourced-data	n/a	TPMs are not recommended to provide a sufficient technology base for this Trustworthiness Claim.
storage-opaque	Minimal	With a TPM, secure storage space exists and is writeable by external applications. But the space is so limited that it often is used just be used to store keys.

Setting the Trustworthiness Claims may follow the following logic at the Verifier A within (2) of Figure 1:

Start: Evidence received starts the generation of a new Trustworthiness Vector. (e.g., TPM Quote Received, log received, or appraisal timer expired)

Step 0: set Trustworthiness Vector = Null

Step 1: Is there sufficient fresh signed evidence to appraise?

(yes) - No Action

(no) - Goto Step 6

Step 2: Appraise Hardware Integrity PCRs

if (hardware NOT "0") - push onto vector

if (hardware NOT affirming or warning), go to Step 6

Step 3: Appraise Attesting Environment identity

if (instance-identity <> "0") - push onto vector

Step 4: Appraise executable loaded and filesystem integrity

if (executables NOT "0") - push onto vector

if (executables NOT affirming or warning), go to Step 6

Step 5: Appraise all remaining Trustworthiness Claims

Independently and set as appropriate.

Step 6: Assemble Attestation Results, and push to Attester

End

A.2. Supportable Trustworthiness Claims for process-based CC

Following are Trustworthiness Claims which MAY be set for a process-based Confidential Computing based Attester. (Such as a SGX Enclaves and TrustZone.)

Trustworthiness Claim	Required?	Appraisal Method
instance-identity	Optional	Internally available in TEE. But keys might not be known/exposed to the Relying Party by the Attesting Environment.
configuration	Optional	If done, this is at the Application Layer. Plus each process needs its own protection mechanism as the protection is limited to the process itself.
executables	Optional	Internally available in TEE. But keys might not be known/exposed to the Relying Party by the Attesting Environment.
file-system	Optional	Can be supported by application, but process-based CC is not a sufficient technology base for this Trustworthiness Claim.
hardware	Implicit in signature	At least the TEE is protected here. Other elements of the system outside of the TEE might need additional protections is used by the application process.
runtime-opaque	Implicit in signature	From the TEE
storage-opaque	Implicit in signature	Although the application must assert that this function is used by the code itself.
sourced-data	Optional	Will need to be supported by application code

A.3. Supportable Trustworthiness Claims for VM-based CC

Following are Trustworthiness Claims which MAY be set for a VM-based Confidential Computing based Attester. (Such as SEV, TDX, ACCA, SEV-SNP.)

Trustworthiness Claim	Required?	Appraisal Method
instance-identity	Optional	Internally available in TEE. But keys might not be known/exposed to the Relying Party by the Attesting Environment.
configuration	Optional	Requires application integration. Easier than with process-based solution, as the whole protected machine can be evaluated.
executables	Optional	Internally available in TEE. But keys might not be known/exposed to the Relying Party by the Attesting Environment.
file-system	Optional	Can be supported by application
hardware	Chip dependent	At least the TEE is protected here. Other elements of the system outside of the TEE might need additional protections is used by the application process.
runtime-opaque	Implicit in signature	From the TEE
storage-opaque	Chip dependent	Although the application must assert that this function is used by the code itself.
sourced-data	Optional	Will need to be supported by application code

Appendix B. Some issues being worked

It is possible for a cluster/hierarchy of Verifiers to have aggregate AR which are perhaps signed/endorsed by a lead Verifier. What should be the Proof-of-Freshness or Verifier associated with any of the aggregate set of Trustworthiness Claims?

There will need to be a subsequent document which documents how these objects which will be translated into a protocol on a wire (e.g. EAP

on TLS). Some breakpoint between what is in this draft, and what is in specific drafts for wire encoding will need to be determined. Questions like architecting the cluster/hierarchy of Verifiers fall into this breakdown.

For some Trustworthiness Claims, there could be value in identifying a specific Appraisal Policy for Attestation Results applied within the Attester. One way this could be done would be a URI which identifies the policy used at Verifier A, and this URI would reference a specific Trustworthiness Claim. As the URI also could encode the version of the software, it might also act as a mechanism to signal the Relying Party to refresh/re-evaluate its view of Verifier A. Do we need this type of structure to be included here? Should it be in subsequent documents?

Expand the variant of Figure 1 which requires no Relying Party PoF into its own picture.

In what document (if any) do we attempt normalization of the identity claims between different types of TEE. E.g., does MRSIGNER plus extra loaded software = the sum of TrustZone Signer IDs for loaded components?

Appendix C. Contributors

Guy Fedorkow

Email: gfedorkow@juniper.net

Dave Thaler

Email: dthaler@microsoft.com

Ned Smith

Email: ned.smith@intel.com

Authors' Addresses

Eric Voit

Cisco Systems

Email: evoit@cisco.com

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
Darmstadt 64295
Germany

Email: henk.birkholz@sit.fraunhofer.de

Thomas Hardjono
MIT

Email: hardjono@mit.edu

Thomas Fossati
Arm Limited

Email: Thomas.Fossati@arm.com

Vincent Scarlata
Intel

Email: vincent.r.scarlata@intel.com

RATS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2022

E. Voit
C. Gaddam
Cisco
G. Fedorkow
Juniper
H. Birkholz
Fraunhofer SIT
M. Chen
China Mobile
2 March 2022

Trusted Path Routing
draft-voit-rats-trustworthy-path-routing-05

Abstract

There are end-users who believe encryption technologies like IPSec alone are insufficient to protect the confidentiality of their highly sensitive traffic flows. These end-users want their flows to traverse devices which have been freshly appraised and verified for trustworthiness. This specification describes Trusted Path Routing. Trusted Path Routing protects sensitive flows as they transit a network by forwarding traffic to/from sensitive subnets across network devices recently appraised as trustworthy.

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 3 September 2022.

Copyright Notice

Copyright (c) 2022 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
2. Terminology	3
2.1. Terms	3
2.2. Requirements Notation	4
3. Implementation Prerequisites	4
4. End-to-end Solution	5
4.1. Network Topology Assembly	5
4.2. Attestation Information Flows	6
4.2.1. Step 1	9
4.2.2. Step 2	9
4.2.3. Step 3	13
4.2.4. Step 4	13
4.2.5. Step 5	15
4.2.6. Step 6	17
5. YANG Module	17
6. Security Considerations	28
7. References	28
7.1. Normative References	28
7.2. Informative References	29
Appendix A. Acknowledgements	30
Appendix B. Change Log	30
Appendix C. Open Questions	32
Authors' Addresses	32

1. Introduction

There are end-users who believe encryption technologies like IPSec alone are insufficient to protect the confidentiality of their highly sensitive traffic flows. These customers want their highly sensitive flows to be transported over only network devices recently verified as trustworthy.

By using a router's embedded TPM based cryptoprocessors in conjunction with the Remote Attestation context established by [attestation-results], a network provider can identify potentially compromised devices as well as potentially exploitable (or even exploited) vulnerabilities. Using this knowledge, it is then possible to redirect sensitive flows around these devices while other remediations are potentially considered by Network Operations.

Trusted Path Routing allows the establishing Trusted Topologies which only include trust-verified network devices. Membership in a Trusted Topology is established and maintained via an exchange of Stamped Passports at the link layer between peering network devices. As links to Attesting Devices are appraised as meeting at least a minimum set of formally defined Trustworthiness Claims, the links are then included as members of this Trusted Topology. Routing protocols are then used to propagate topology state throughout a network.

IP Packets to and from end-user designated Sensitive Subnets are then forwarded into this Trusted Topology at each network boundary. This is done by an end user identifying sensitive IP subnets where flows with applications using these IP subnets need enhanced privacy guarantees. Trusted Path Routing passes flows to/from these Sensitive Subnets over a Trusted Topology able to meet these guarantees. The Trusted Topology itself consists of the interconnection of network devices where each potentially transited device has been verified as achieving a specific set of Trustworthiness Claims during its most recent trustworthiness appraisal. Interesting sets of Trustworthiness Claims might be marketed to end-users in the following ways:

- * all transited devices have booted with known hardware and firmware
- * all transited devices are from a specific set of vendors and are running known software containing the latest patches
- * no guarantees provided

2. Terminology

2.1. Terms

The following terms are imported from [RATS-Arch]: Attester, Evidence, Passport, Relying Party, and Verifier.

The following terms are imported from [attestation-results]: Trustworthiness Claim, Trustworthiness Vector, AR-augmented Evidence

Newly defined terms for this document:

Attested Device -- a network connected Attester where a Verifier's most recent appraisal of Evidence has returned a Trustworthiness Vector.

Stamped Passport -- AR-augmented Evidence which can take two forms. First if the Attester uses a TPM2, the the Verifier Proof-of-Freshness includes the <clock>, <reset-counter>, <restart-counter> and <safe> objects from a recent TPM2 quote made by that Attester, and the Relying Party Proof-of-Freshness is returned along with the timeticks as objects embedded within the most recent TPM quote signed by the same TPM2. Second, if the Attester uses a TPM1.2: the Verifier Proof-of-Freshness includes a global timestamp from that Verifier, and the Relying Party Proof-of-Freshness is embedded within a more recent TPM quote signed by the same TPM Attesting Environment.

Sensitive Subnet -- an IP address range where IP packets to or from that range desire confidentially guarantees beyond those of non-identified subnets. In practice, flows to or from a Sensitive Subnet must only have their IP headers and encapsulated payloads accessible/visible only by Attested Devices supporting one or more Trustworthiness Vectors.

Transparently-Transited Device -- a network device within an network domain where any packets originally passed into that network domain are completely opaque on that network device at Layer 3 and above.

Trusted Topology -- a topology which includes only Attested Devices and Transparently-Transited Devices.

2.2. Requirements Notation

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.

3. Implementation Prerequisites

The specification is a valid instance of [attestation-results]. This specification works under the following protocol and preconfiguration prerequisite assumptions:

- * All Attested Devices support the TPM remote attestation profile as laid out in [RATS-Device], and include either [TPM2.0] or [TPM1.2].

- * One or more Verifier A's as defined in [attestation-results] 'Interaction Model' continuously appraise each of the Attested Devices in a network domain, and these Verifiers return the Attestation Results back to each originating Attested Device.
- * The Attested Devices are connected via link layer protocols such as [MACSEC] or [IEEE-802.1X].
- * Each Attester can pass a Stamped Passport to a Relying Party / Verifier B as defined in [attestation-results] 'Interaction Model' within [RFC3748] over that link layer protocol.
- * A Trusted Topology such as [I-D.ietf-lsr-flex-algo] exists in an IGP domain for the forwarding of Sensitive Subnet traffic. This Topology will carry traffic across a set of Attested Devices which currently meet at a defined set of Trustworthiness Vectors.
- * A Relying Party is able to use mechanisms such as [I-D.ietf-lsr-flex-algo]'s affinity to include/exclude links as part of the Trusted Topology based on the appraisal of a Stamped Passport.
- * Customer designated Sensitive Subnets and their requested Trustworthiness Vectors have been identified and associated with external interfaces to/from Attested Devices at the edge of a network. Traffic to a Sensitive Subnet can be passed into the Trusted Topology by the Attested Device.
- * Relying Party/Verifier B trusts information signed by Verifier A. Verifier B has also been pre-provisioned with certificates or public keys necessary to confirm that Stamped Passports came from Verifier A.

4. End-to-end Solution

4.1. Network Topology Assembly

To be included in a Trusted Topology, Stamped Passports are shared between Attested Devices (such as routers) as part of link layer authentication. Upon receiving and appraising the Stamped Passport during the link layer authentication phase, the Relying Party Attested Device decides if this link should be added as an active adjacency for a particular Trusted Topology. In Figure 1 below, this might be done by applying an Appraisal Policy for Attestation Results. The policy within each device might specify the evaluation of a 'hardware' claim as defined in [attestation-results], Section 2.3.4. With the appraisal, an Attesting Device be most recently appraised with the 'hardware' Trustworthiness Claim in the

'affirming' range. If Attested Device has been appraised outside that range, it would not become part of the Trustworthy Topology.

When enough links have been successfully added, the Trusted Topology will support edge-to-edge forwarding as routing protocols flood the adjacency information across the network domain.

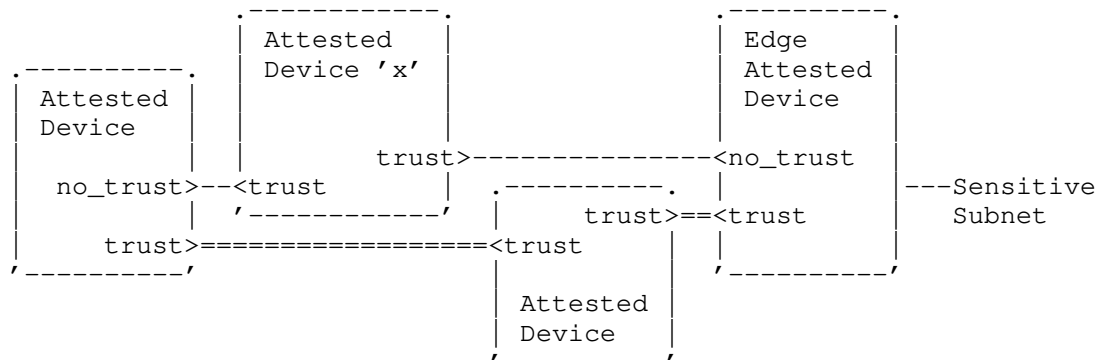


Figure 1: Trusted Path Topology Assembly

As the process described above repeats over time across the set of links within a network domain, Trusted Topologies can be extended and maintained. Traffic to and from Sensitive Subnets is then identified at the edges of the network domain and passed into this Trusted Topology. Traffic exchanged with Sensitive Subnets can then be forwarded across that Trusted Topology from all edges of the network domain. After the initial Trusted Topology establishment, new and existing devices will continue to provide incremental Stamped Passports. As each link is added/removed from the Trusted Topology, the topology will adjust itself accordingly.

Ultimately from an operator and users point of view, the delivered network will be more secure and therefore the service provided more valuable. As network operators attach great importance to the innate security of links, also delivering security for transited network and networking devices will also prove valuable.

4.2. Attestation Information Flows

Critical to the establishment and maintenance of a Trusted Topology is the Stamped Passport. A Stamped Passport is comprised of Evidence from both an Attester and a Verifier. A Stamped Passport is a valid type of AR-augmented evidence as described in [attestation-results].

Stamped Passports are exchanged between adjacent network devices over a link layer protocols like 802.1x or MACSEC. As both sides of a link may need might need to appraise the other, independent Stamped Passports will often be transmitted from either side of the link. Additionally, as link layer protocols will continuously re-authenticate the link, this allows for fresh Stamped Passports to be constantly appraised by either side of the connection.

Each Stamped Passport will include the most recent Verifier provided Attestation Results, as well as the most recent TPM Quote for that Attester. Upon receiving this information as part of link layer authentication, the Relying Party Router appraises the results and decides if this link should be added to a Trusted Topology.

Figure 2 describes this flow of information using the time definitions described in [RATS-Arch], and the information flows defined in Section 7 of [RATS-Interactions]. This figure is also a valid embodiment of the "Interaction Model" described within [attestation-results]. (Note that the Relying Party must also be an Attested Device in order to attract Sensitive Subnet traffic which may flow from the Attester.)

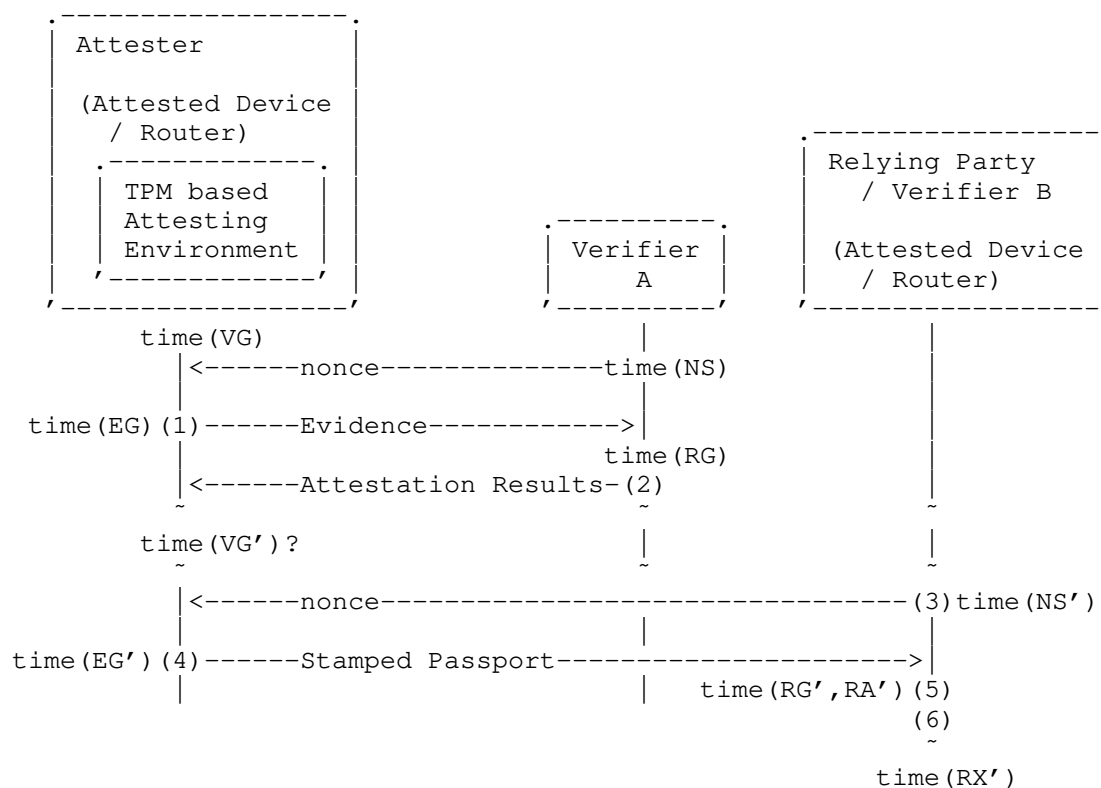


Figure 2: Trusted Path Timing

To summarize Figure 2 above, Evidence about a specific Attester is generated. Some subset of this evidence will be in the form of PCR quotes which are signed by a TPM that exists as the Attester's Attesting Environment. This Evidence will be deliberated to and appraised by Verifier A. Verifier A will then appraise the Attester and give it a Trustworthiness Vector. This Trustworthiness Vector is then signed by Verifier A and be returned as Attestation Results to the Attester. Later, when a request comes in from a Relying Party, the Attester assembles and returns a Stamped Passport. The Stamped Passport contains all the information necessary for Verifier B to appraise the most recent Trustworthiness Vector of the Attester. Based on the Verifier B appraisal, the link will be included or not in a Trusted Topology maintained on the Relying Party.

More details on the mechanisms used in the construction, verification, and transmitting of the Stamped Passport are listed below. These numbers match to both the numbered steps of Figure 2 and numbered steps described in Section 3 of [attestation-results]:

4.2.1. Step 1

Evidence about and Attester is generated. A portion of this Evidence will include a PCR quote signed by a TPM private LDevID key that exists within the Attester's TPM based Attesting Environment. The Attester sends a signed TPM Quote which includes PCR measurements to Verifier A at time(EG).

There are two alternatives for Verifier A to acquire this signed Evidence:

- * Subscription to the <attestation> stream defined in [stream-subscription]. Note: this method is recommended as it will minimize the interval between when a PCR change is made in a TPM, and when the PCR change appraisal is incorporated within a subsequent Stamped Passport.
- * Periodic polling of RPC <tpm20-challenge-response-attestation> or the RPC <tpm12-challenge-response-attestation> which are defined in [RATS-YANG].

4.2.2. Step 2

Verifier A appraises the Evidence from Step 1. A portion of this appraisal process will follow the appraisal process flow described below. This appraisal process MUST be able to set at least the following set of Trustworthiness Claims from [attestation-results]: 'hardware', 'instance-identity', and 'executables'. The establishment of a Trustworthiness Vector uses the following Figure 3 logic on the Verifier:

Start: TPM Quote Received, log received, or appraisal timer expired
for the the Attesting network device.

Appraisal 0: set Trustworthiness Vector = Null

Appraisal 1: Is there sufficient fresh signed evidence to appraise?
(yes) - No Action
(no) - Goto End

Appraisal 2: Appraise Hardware Integrity PCRs
if (hardware NOT "0") - push onto vector
if (hardware NOT affirming or warning), go to End

Appraisal 3: Appraise Attesting Environment identity
if (instance-identity <> "0") - push onto vector

Appraisal 4: Appraise executable loaded and filesystem integrity
if (executables NOT "0") - push onto vector
if (executables NOT affirming or warning), go to End

Appraisal 5: Appraise all remaining Trustworthiness Claims
Independently and set as appropriate.

End

Figure 3: Verifier A Appraisal Flow

After the appraisal and generation of the Trustworthiness Vector, the following are assembled as the set of Attestation Results from this particular appraisal cycle:

(2.1) the Public Attestation Key which was used to validate the TPM Quote of Step 1. This is encoded by <public-key>, <public-key-format>, and <public-key-algorithm-type>.

(2.2) the appraised Trustworthiness Vector of the Attester as calculated in Figure 3

(2.3) the PCR state information from the TPM Quote of (1) plus the time information associated with the TPM Quote of (1). Specifically if the Attester has a TPM2, then the values of the TPM PCRs are included (i.e., <TPM2B_DIGEST>, <tpm20-hash-algo>, and <pcr-index>), as are the timing counters from the TPM (i.e., <clock>, <reset-counter>, <restart-counter>, and <safe>). Likewise if the Attester has a TPM1.2, the TPM PCR values of the <pcr-index> and <pcr-value> are included. Timing information comes from the Verifier itself via the <timestamp> object.

(2.4) a Verifier A signature across (2.1) through (2.3). This signature is encoded by <verifier-signature>, <verifier-key-algorithm-type>, and <verifier-signature-key-name>.

Immediately subsequent to each Verifier appraisal cycle of an Attester, these Attestation Results MUST be pushed to the Attesting Router. This is done via a datastore write to the following YANG model on the Attester. A YANG tree showing the relevant YANG objects is below. The YANG model describing each of these objects is described later in the document. Note however that although the YANG model shows the specific objects which are needed, the specific set of objects needs to be encoded in CDDL. This makes the payload going over TLS more efficient. Look for this encoding in a new version of the draft which is pending.

```

module: ietf-trustworthiness-claims
+--rw attestation-results!
+--rw (tpm-specification-version)?
+--:(tpm20-attestation-results-cddl) {taa:tpm20}?
|   +--rw tpm20-attestation-results-cddl
|   |   +--rw trustworthiness-vector
|   |   |   +--rw hardware?          hardware
|   |   |   +--rw instance-identity?  instance-identity
|   |   |   +--rw executables?        executables
|   |   |   +--rw configuration?      configuration
|   |   +--rw tpm20-pcr-selection* [tpm20-hash-algo]
|   |   |   +--rw tpm20-hash-algo      identityref
|   |   |   +--rw pcr-index*          tpm:pcr
|   |   +--rw TPM2B_DIGEST              binary
|   |   +--rw clock                     uint64
|   |   +--rw reset-counter             uint32
|   |   +--rw restart-counter           uint32
|   |   +--rw safe                     boolean
|   |   +--rw attester-certificate-name
|   |   |   tpm:certificate-name-ref
|   |   +--rw appraisal-timestamp
|   |   |   yang:date-and-time
|   |   +--rw verifier-algorithm-type   identityref
|   |   +--rw verifier-signature        binary
|   |   +--rw verifier-certificate-keystore-ref
|   |   |   tpm:certificate-name-ref
+--:(tpm12-attestation-results-cddl) {taa:tpm12}?
|   +--rw tpm12-attestation-results-cddl
|   |   +--rw trustworthiness-vector
|   |   |   +--rw hardware?          hardware
|   |   |   +--rw instance-identity?  instance-identity
|   |   |   +--rw executables?        executables
|   |   |   +--rw configuration?      configuration
|   |   +--rw pcr-index*              pcr
|   |   +--rw tpm12-pcr-value*         binary
|   |   +--rw tpm12-hash-algo          identityref
|   |   +--rw TPM12-quote-timestamp
|   |   |   yang:date-and-time
|   |   +--rw attester-certificate-name
|   |   |   tpm:certificate-name-ref
|   |   +--rw appraisal-timestamp
|   |   |   yang:date-and-time
|   |   +--rw verifier-algorithm-type   identityref
|   |   +--rw verifier-signature        binary
|   |   +--rw verifier-certificate-keystore-ref
|   |   |   tpm:certificate-name-ref

```

Figure 4: Attestation Results Tree

4.2.3. Step 3

At time(NS') some form of time-based freshness (such as a nonce or Epoch Handle [RATS-Interactions]) will be generated in a way which makes it available to the Relying Party. Soon after time(NS'), a Relying Party will make a Link Layer authentication request to an Attester via either [MACSEC] or [IEEE-802.1X]. This connection request MUST expect the return of [RFC3748] credentials from the Attester.

4.2.4. Step 4

Upon receipt of the Link Layer request from Step 3, a Stamped Passport is generated and sent to the Relying Party. The Stamped Passport MUST include the following:

(4.1) The Attestation Results from Step 2

(4.2) New signed, verifiably fresh PCR measurements based on a TPM quote at time(EG') which incorporates the freshness information known by the Relying Party from Step 3. If it is a nonce, the freshness information will have been delivered as part of the link layer connection request in Steps 3.

Stamped Passports contain following objects, defined in this document via YANG. A subsequent draft will convert the objects below into CDDL format so that the objects can efficiently be passed over EAP.

If an Attester includes a TPM2, these YANG objects are:


```

+---n tpm20-stamped-passport
+--ro attestation-results
|   +--ro trustworthiness-vector
|   |   +--ro hardware?          hardware
|   |   +--ro instance-identity? instance-identity
|   |   +--ro executables?       executables
|   |   +--ro configuration?     configuration
|   +--ro tpm20-pcr-selection* [tpm20-hash-algo]
|   |   +--ro tpm20-hash-algo    identityref
|   |   +--ro pcr-index*         tpm:pcr
|   +--ro TPM2B_DIGEST           binary
|   +--ro clock                  uint64
|   +--ro reset-counter          uint32
|   +--ro restart-counter        uint32
|   +--ro safe                   boolean
|   +--ro attester-certificate-name
|   |   tpm:certificate-name-ref
|   +--ro appraisal-timestamp
|   |   yang:date-and-time
|   +--ro verifier-algorithm-type identityref
|   +--ro verifier-signature     binary
|   +--ro verifier-certificate-keystore-ref
|   |   tpm:certificate-name-ref
+--ro tpm20-quote
|   +--ro TPMS_QUOTE_INFO        binary
|   +--ro quote-signature        binary

```

Figure 5: YANG Tree for a TPM2 Stamped Passport

Note that where a TPM2.0 is used, the PCR numbers and hash algorithms quoted in Step 1 MUST match the PCR numbers and hash algorithms quoted in this step.

And if the Attester is a TPM1.2, the YANG object are:

```

+---n tpm12-stamped-passport
+--ro attestation-results
|   +--ro trustworthiness-vector
|   |   +--ro hardware?          hardware
|   |   +--ro instance-identity? instance-identity
|   |   +--ro executables?       executables
|   |   +--ro configuration?     configuration
|   +--ro pcr-index*             pcr
|   +--ro tpm12-pcr-value*       binary
|   +--ro tpm12-hash-algo        identityref
|   +--ro TPM12-quote-timestamp
|   |   yang:date-and-time
|   +--ro attester-certificate-name
|   |   tpm:certificate-name-ref
|   +--ro appraisal-timestamp
|   |   yang:date-and-time
|   +--ro verifier-algorithm-type identityref
|   +--ro verifier-signature      binary
|   +--ro verifier-certificate-keystore-ref
|   |   tpm:certificate-name-ref
+--ro tpm12-quote
+--ro TPM_QUOTE2?    binary

```

Figure 6: YANG Tree for a TPM1.2 Stamped Passport

With either of these passport formats, if the TPM quote is verifiably fresh, then the state of the Attester can be appraised by a network peer.

Note that with [MACSEC] or [IEEE-802.1X], Step 3 plus Step 4 will repeat periodically independently of any subsequent iteration Steps 1 and Step 2. This allows for periodic reauthentication of the link layer in a way not bound to the updating of Verifier A's Attestation Results.

4.2.5. Step 5

Upon receipt of the Stamped Passport generated in Step 4, the Relying Party appraises this Stamped Passport as per its Appraisal Policy for Attestation Results. The result of this application will determine how the Stamped Passport will impact adjacencies within a Trusted Topology. The decision process is as follows:

- (5.1) Verify that (4.2) includes the freshness context from Step 3.
- (5.2) Use a local certificate to validate the signature (4.1).
- (5.3) Verify that the hash from (4.2) matches (4.1)

(5.4) Use the identity of (2.1) to validate the signature of (4.2).

(5.5) Failure of any steps (5.1) through (5.4) means the link does not meet minimum validation criteria, therefore appraise the link as having a null Verifier B Trustworthiness Vector. Jump to Step 6.

(5.6) Compare the time(EG) TPM state to the time(EG') TPM state

* If TPM2.0

1. If the <TPM2B_DIGEST>, <reset-counter>, <restart-counter> and <safe> are equal between the Attestation Results and the TPM Quote at time(EG') then Relying Party can accept (2.1) as the link's Trustworthiness Vector. Jump to Step 6.
2. If the <reset-counter>, <restart-counter> and <safe> are equal between the Attestation Results and the TPM Quote at time(EG'), and the <clock> object from time(EG') has not incremented by an unacceptable number of seconds since the Attestation Result, then Relying Party can accept (2.1) as the link's Trustworthiness Vector. Jump to Step 6.)
3. Assign the link a null Trustworthiness Vector.

* If TPM1.2

1. If the <pcr-index>'s and <tpm12-pcr-value>'s are equal between the Attestation Results and the TPM Quote at time(EG'), then Relying Party can accept (2.1) as the link's Trustworthiness Vector. Jump to step (6).
2. If the time hasn't incremented an unacceptable number of seconds from the Attestation Results <timestamp> and the system clock of the Relying Party, then Relying Party can accept (2.1) as the link's Trustworthiness Vector. Jump to step 6.)
3. Assign the link a null Trustworthiness Vector.

(5.7) Assemble the Verifier B Trustworthiness Vector

1. Copy Verifier A Trustworthiness Vector to Verifier B Trustworthiness Vector
2. Prune any Trustworthiness Claims the Relying Party doesn't accept from this Verifier.

4.2.6. Step 6

After the Trustworthiness Vector has been validated or reset, based on the link's Trustworthiness Vector, the Relying Party adjusts the link affinity of the corresponding ISIS [I-D.ietf-lsr-flex-algo] topology. ISIS will then replicate the link state across the IGP domain. Traffic will then avoid links which do not have a qualifying Trustworthiness Vector.

5. YANG Module

This YANG module imports modules from [RATS-YANG], [crypto-types] and [RFC6021].

```
<CODE BEGINS> ietf-trustworthiness-claims@2021-11-03.yang
module ietf-trustworthiness-claims {
  yang-version 1.1;
  namespace
    "urn:ietf:params:xml:ns:yang:ietf-trustworthiness-claims";
  prefix tc;

  import ietf-yang-types {
    prefix yang;
  }
  import ietf-tcg-algs {
    prefix taa;
    reference
      "draft-ietf-rats-yang-tpm-charra";
  }
  import ietf-tpm-remote-attestation {
    prefix tpm;
    reference
      "draft-ietf-rats-yang-tpm-charra";
  }

  organization "IETF";
  contact
    "WG Web:  <http://tools.ietf.org/wg/rats/>
    WG List:  <mailto:rats@ietf.org>

    Editor:    Eric Voit
               <mailto:evoit@cisco.com>";

  description
    "This module contains conceptual YANG specifications for
    subscribing to attestation streams being generated from TPM chips.

    Copyright (c) 2020 IETF Trust and the persons identified as
```

authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
revision 2021-11-03 {
  description
    "Initial version.";
  reference
    "draft-voit-rats-trustworthy-path-routing";
}

/*
 * TYPEDEF
 */

typedef trustworthiness-claim {
  type int8;
  description
    "A Verifier asserted value designed to enable a common
    understanding of a Verifier trustworthiness appraisal. The
    value assignments for this 8 bit signed integer will follow
    these guidelines:

    Affirming: The Verifier affirms the Attester support for this
    aspect of trustworthiness
      - Values 2 to 31: A standards enumerated reason for affirming.
      - Values -2 to -32: A non-standard reason for affirming.

    Warning: The Verifier warns about this aspect of trustworthiness.
      - Values 32 to 63: A standards enumerated reason for the
        warning.
      - Values -33 to -64: A non-standard reason for the warning.

    Contraindicated: The Verifier asserts the Attester is explicitly
    untrustworthy in regard to this aspect.
      - Values 64 to 127: A standards enumerated reason for the
        contraindication.
      - Values -65 to -128: A non-standard reason for the
        contraindication.
```

None: The Verifier makes no assertions about this Trustworthiness Claim. The following values are reserved with the following meanings across all instances of trustworthiness-claim.

- Value 0: Note: this should always be always treated equivalently by the Relying Party as no claim being made. I.e., the RP's Appraisal Policy for Attestation Results SHOULD NOT make any distinction between a Trustworthiness Claim with enumeration '0', and no Trustworthiness Claim being provided.
- Value 1: The Evidence received contains unexpected elements which the Verifier is unable to parse. An example might be that the wrong type of Evidence has been delivered.
- Value -1: An unexpected error occurred during the Verifier's appraisal processing. Note: while no claim is being made, the Relying Party MAY make a distinction between a Trustworthiness Claim with enumeration '-1', and no Trustworthiness Claim being provided."

}

typedef hardware {
 type trustworthiness-claim;
 description
 "A Verifier has appraised any Attester hardware and firmware which are able to expose fingerprints of their identity and running code.

The following are specific reserved values of hardware and the meanings of these reserved values:

0: No assertion

1: The Verifier cannot parse unexpected Evidence

-1: Verifier malfunction

2: An Attester has passed its hardware and/or firmware verifications needed to demonstrate that these are genuine/supported.

32: An Attester contains only genuine/supported hardware and/or firmware, but there are known security vulnerabilities.

96: Attester hardware and/or firmware is recognized, but its trustworthiness is contraindicated.

97: A Verifier does not recognize an Attester's hardware or firmware, but it should be recognized."

}

```
typedef instance-identity {
    type trustworthiness-claim;
    description
        "A Verifier has appraised an Attesting Environment's unique
        identity based upon private key signed Evidence which can be
        correlated to a unique instantiated instance of the Attester.
        (Note: this Trustworthiness Claim should only be generated if
        the Verifier actually expects to recognize the unique identity
        of the Attester.)

        The following are specific reserved values of instance-identity
        and the meanings of these reserved values:

        0: No assertion

        1: The Verifier cannot parse unexpected Evidence

        -1: Verifier malfunction

        2: The Attesting Environment is recognized, and the associated
        instance of the Attester is not known to be compromised.

        96: The Attesting Environment is recognized, and but its unique
        private key indicates a device which is not trustworthy.

        97: The Attesting Environment is not recognized; however the
        Verifier believes it should be."
}

typedef executables {
    type trustworthiness-claim;
    description
        "A Verifier has appraised and evaluated relevant runtime files,
        scripts, and/or other objects which have been loaded into the
        Target environment's memory.

        The following are specific reserved values of executables and
        the meanings of these reserved values:

        0: No assertion

        1: The Verifier cannot parse unexpected Evidence

        -1: Verifier malfunction

        2: Only a recognized genuine set of approved executables,
        scripts, files, and/or objects have been loaded during
        and after the boot process.
```

```
3: Only a recognized genuine set of approved executables have
  been loaded during the boot process.

32:Only a recognized genuine set of executables, scripts, files,
  and/or objects have been loaded. However the Verifier cannot
  vouch for a subset of these due to known bugs or other known
  vulnerabilities.

33:Runtime memory includes executables, scripts, files, and/or
  objects which are not recognized.

96:Runtime memory includes executables, scripts, files, and/or
  object which are contraindicated.

99:Cryptographic validation of the Evidence has failed.";
}

typedef configuration {
  type trustworthiness-claim;
  description
    "A Verifier has appraised an Attester's configuration, and is
    able to make conclusions regarding the exposure of known
    vulnerabilities.

    The following are specific reserved values of configuration and
    the meanings of these reserved values:

    0: No assertion

    1: The Verifier cannot parse unexpected Evidence

    -1:Verifier malfunction

    2: The configuration is a known and approved config

    3: The configuration includes or exposes no known vulnerabilities

    32:The configuration includes or exposes known vulnerabilities

    64:The configuration is unsupportable as it exposes unacceptable
    security vulnerabilities.";
}

/*
 * GROUPINGS
 */
```



```
grouping trustworthiness-vector {
  description
    "Allows the inclusion of a Trustworthiness Vector into
    other constructs.";
  container trustworthiness-vector {
    description
      "One or more Trustworthiness Claims assigned which expose the
      Verifiers evaluation of the Evidence associated with the
      AIK which signed as associated TPM Quote.";
    leaf hardware {
      type hardware;
      description
        "An 8 bit signed integter encoded per the typedef.";
    }
    leaf instance-identity {
      type instance-identity;
      description
        "An 8 bit signed integter encoded per the typedef.";
    }
    leaf executables {
      type executables;
      description
        "An 8 bit signed integter encoded per the typedef.";
    }
    leaf configuration {
      type configuration;
      description
        "An 8 bit signed integter encoded per the typedef.";
    }
  }
}

grouping verifier-evidence {
  description
    "Evidence generated by the Verifier.";
  leaf appraisal-timestamp {
    type yang:date-and-time;
    mandatory true;
    description
      "The timestamp of the Verifier's appraisal. This can be used
      by a Relying Party to determine the freshness of the
      attestation results.";
  }
  leaf verifier-algorithm-type {
    type identityref {
      base taa:asymmetric;
    }
    mandatory true;
  }
}
```

```
    description
      "Platform asymmetric algorithm used in the Verifier signature
       process.";
  }
  leaf verifier-signature {
    type binary;
    mandatory true;
    description
      "Signature of the Verifier across all the current objects in
       the attestation-results container except for 'verifier-
       signature' and 'verifier-certificate-keystore-ref'.
       This assumes CDDL encoding of the objects in the current
       order of this YANG model.";
  }
  leaf verifier-certificate-keystore-ref {
    type tpm:certificate-name-ref;
    mandatory true;
    description
      "A reference to a specific certificate to an asymmetric key
       in the Keystore for the Verifier which can be used to validate
       the 'verifier-signature'. Note that the
       'name' reference must be globally unique so that it can be
       read by the Relying Party in a way which identifies a
       specific Verifier.";
  }
}

grouping tpm20-cddl-attestation-results {
  description
    "Elements combined into a CDDL representation for TPM2.0.";
  uses trustworthiness-vector;
  list tpm20-pcr-selection {
    key "tpm20-hash-algo";
    description
      "Specifies the list of PCRs and Hash Algorithms used by the
       Verifier.";
    reference
      "https://www.trustedcomputinggroup.org/wp-content/uploads/
       TPM-Rev-2.0-Part-2-Structures-01.38.pdf  Section 10.9.7";
    uses tpm:tpm20-hash-algo;
    leaf-list pcr-index {
      type tpm:pcr;
      description
        "The numbers of the PCRs associated with the TPM2B_DIGEST.";
    }
  }
  leaf TPM2B_DIGEST {
    mandatory true;
  }
}
```

```
type binary;
description
  "A hash of the latest PCR values (and the hash algorithm used)
  which have been returned from a Verifier for the selected PCRs
  identified within TPML_PCR_SELECTION.";
reference
  "https://www.trustedcomputinggroup.org/wp-content/uploads/
  TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.12.1";
}
leaf clock {
  mandatory true;
  type uint64;
  description
    "Clock is a monotonically increasing counter that advances
    whenever power is applied to a TPM2. The value of Clock is
    incremented each millisecond.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.11.2";
}
leaf reset-counter {
  mandatory true;
  type uint32;
  description
    "This counter increments on each TPM Reset. The most common
    TPM Reset would be due to a hardware power cycle.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.11.3";
}
leaf restart-counter {
  mandatory true;
  type uint32;
  description
    "This counter shall increment by one for each TPM Restart or
    TPM Resume. The restartCount shall be reset to zero on a TPM
    Reset.";
  reference
    "https://www.trustedcomputinggroup.org/wp-content/uploads/
    TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.11.4";
}
leaf safe {
  mandatory true;
  type boolean;
  description
    "This parameter is set to YES when the value reported in Clock
    is guaranteed to be unique for the current Owner. It is set to
    NO when the value of Clock may have been reported in a previous
```

```
        attestation or access.";
    reference
        "https://www.trustedcomputinggroup.org/wp-content/uploads/
        TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.11.5";
}
leaf attester-certificate-name {
    mandatory true;
    description
        "The Attester is associated with these results.";
    type tpm:certificate-name-ref;
}
uses verifier-evidence;
}

grouping tpml2-cddl-attestation-results {
    description
        "Elements combined into a CDDL representation for TPM1.2.";
    uses trustworthiness-vector;
    uses tpm:tpml2-pcr-selection;
    leaf-list tpml2-pcr-value {
        type binary;
        description
            "The list of TPM_PCRVALUES from each PCR selected in sequence
            of tpml2-pcr-selection.";
        reference
            "https://www.trustedcomputinggroup.org/wp-content/uploads/
            TPM-Main-Part-2-TPM-Structures_v1.2_rev116_01032011.pdf
            Section 10.9.7";
    }
    uses tpm:tpml2-hash-algo {
        refine "tpml2-hash-algo" {
            mandatory true;
        }
    }
}
leaf TPM12-quote-timestamp {
    type yang:date-and-time;
    mandatory true;
    description
        "The timestamp for when the indicator of freshness (such as a
        nonce) was generated. This is the indicator of freshness
        which was used in the generation of the TPM1.2 quote. This
        timestamp can be used by a Relying Party to determine the
        freshness of the attestation results.";
}
leaf attester-certificate-name {
    mandatory true;
    description
        "The Attester is associated with these results.";
```

```
        type tpm:certificate-name-ref;
    }
    uses verifier-evidence;
}

/*
 * NOTIFICATIONS
 */

notification tpm20-stamped-passport {
    description
        "The augmentation of the most recent Attestation Results
        delivered from a Verifier with a TPM2.0 Quote.";
    container attestation-results {
        description
            "The latest Verifier delivered Attestation Results.";
        uses tpm20-cddl-attestation-results;
    }
    container tpm20-quote {
        description
            "The TPM2.0 quote delivered in response to a connectivity
            request.";
        leaf TPMS_QUOTE_INFO {
            type binary;
            mandatory true;
            description
                "A hash of the latest PCR values (and the hash algorithm
                used) which have been returned from a Verifier for the
                selected PCRs and Hash Algorithms.";
            reference
                "https://www.trustedcomputinggroup.org/wp-content/uploads/
                TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 10.12.1";
        }
        leaf quote-signature {
            type binary;
            mandatory true;
            description
                "Quote signature returned by TPM Quote. The signature was
                generated using the key associated with the
                certificate 'name'.";
            reference
                "https://www.trustedcomputinggroup.org/wp-content/uploads/
                TPM-Rev-2.0-Part-2-Structures-01.38.pdf Section 11.2.1";
        }
    }
}

notification tpm12-stamped-passport {
```

```
description
  "The augmentation of the most recent Attestation Results
  delivered from a Verifier with a TPM1.2 Quote.";
container attestation-results {
  description
    "The latest Verifier delivered Attestation Results.";
  uses tpm12-cddl-attestation-results;
}
container tpm12-quote {
  description
    "The TPM1.2 quote delivered in response to a connectivity
    request.";
  leaf TPM_QUOTE2 {
    type binary;
    description
      "Result of a TPM1.2 Quote2 operation. This includes PCRs,
      signatures, locality, the provided nonce and other data which
      can be further parsed to appraise the Attester.";
    reference
      "TPM1.2 commands rev116 July 2007, Section 16.5";
  }
}
}

/*
 * DATA NODES
 */

container attestation-results {
  presence
    "Indicates that Verifier has appraised the security posture of
    the Attester, and returned the results within this container.";
  description
    "Retains the most recent Attestation Results for this Attester.
    It must only be written by a Verifier which is to be trusted by a
    Relying Party.";

  choice tpm-specification-version {
    description
      "Identifies the cryptoprocessor API set which drove the
      Attestation Results.";
    case tpm20-attestation-results-cddl {
      if-feature "taa:tpm20";
      description
        "Attestation Results which are returned from the
        evaluation of Evidence which includes a TPM2 quote.";
      container tpm20-attestation-results-cddl {
        description
```

```
        "Attestation Results which are returned from the
          evaluation of Evidence which includes a TPM2 quote.";
        uses tpm20-cddl-attestation-results;
      }
    }
  case tpm12-attestation-results-cddl {
    if-feature "taa:tpm12";
    description
      "Attestation Results which are returned from the
        evaluation of Evidence which includes a TPM1.2 quote.";
    container tpm12-attestation-results-cddl {
      description
        "Attestation Results which are returned from the
          evaluation of Evidence which includes a TPM1.2 quote.";
      uses tpm12-cddl-attestation-results;
    }
  }
}
}
}
}
<CODE ENDS>
```

6. Security Considerations

Verifiers are limited to the Evidence available for appraisal from a Router. Although the state of the art is improving, some exploits may not be visible via Evidence.

Only security measurements which are placed into PCRs are capable of being exposed via TPM Quote at time(EG').

Successful attacks on an Verifier have the potential of affecting traffic on the Trusted Topology.

For Trusted Path Routing, links which are part of the FlexAlgo are visible across the entire IGP domain. Therefore a compromised device will know when it is being bypassed.

Access control for the objects in Figure 4 should be tightly controlled so that it becomes difficult for the Stamped Passport to become a denial of service vector.

7. References

7.1. Normative References

- [attestation-results]
"Attestation Results for Connectivity", 2 December 2021,
<<https://datatracker.ietf.org/doc/draft-ietf-rats-ar4si/>>.
- [crypto-types]
"Common YANG Data Types for Cryptography", 17 December
2021, <<https://datatracker.ietf.org/doc/draft-ietf-netconf-crypto-types/>>.
- [RATS-Arch]
"Remote Attestation Procedures Architecture", 8 February
2022, <<https://datatracker.ietf.org/doc/draft-ietf-rats-architecture/>>.
- [RATS-YANG]
"A YANG Data Model for Challenge-Response-based Remote
Attestation Procedures using TPMs", 28 February 2022,
<<https://datatracker.ietf.org/doc/draft-ietf-rats-yang-tpm-charra/>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119,
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6021] Schoenwaelder, J., Ed., "Common YANG Data Types",
RFC 6021, DOI 10.17487/RFC6021, October 2010,
<<https://www.rfc-editor.org/info/rfc6021>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [TPM1.2] TCG, "TPM 1.2 Main Specification", 2 October 2003,
<<https://trustedcomputinggroup.org/resource/tpm-main-specification/>>.
- [TPM2.0] TCG, "TPM 2.0 Library Specification", 15 March 2013,
<<https://trustedcomputinggroup.org/resource/tpm-library-specification/>>.

7.2. Informative References

- [I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-18, 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-18.txt>>.
- [IEEE-802.1X]
Parsons, G., "802.1AE: MAC Security (MACsec)", 1 January 2020, <https://standards.ieee.org/standard/802_1X-2010.html>.
- [MACSEC] Seaman, M., "802.1AE: MAC Security (MACsec)", 1 January 2006, <<https://1.ieee802.org/security/802-1ae/>>.
- [RATS-Device]
"Network Device Remote Integrity Verification", n.d., <<https://datatracker.ietf.org/doc/draft-ietf-rats-tpm-based-network-device-attest>>.
- [RATS-Interactions]
"Reference Interaction Models for Remote Attestation Procedures", 26 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-rats-reference-interaction-models>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<https://www.rfc-editor.org/info/rfc3748>>.
- [stream-subscription]
"Attestation Event Stream Subscription", 16 October 2021, <<https://datatracker.ietf.org/doc/draft-ietf-rats-network-device-subscription/>>.

Appendix A. Acknowledgements

Peter Psenak, Shwetha Bhandari, Adwaith Gautham, Annu Singh, Sujal Sheth, Nancy Cam Winget, and Ned Smith.

Appendix B. Change Log

[THIS SECTION TO BE REMOVED BY THE RFC EDITOR.]

v04-v05

* Tweaks to text

- * Added text which was morphed from that provided by Meiling

v03-v04

- * YANG model updated in concert with draft-voit-rats-attestation-results as Trustworthiness Claim values are added.

v03-v04

- * Moved in concert with draft-voit-rats-attestation-results as Trustworthiness Claims became 8 bit signed integers.

v02-v03

- * Integrated [attestation-results] as prerequisite context.
- * Totally rearranged content. But there were not meaningful process changes.
- * Redid YANG model, and highlighted CDDL needs.

v01-v02

- * Minor tweaks such as renaming and removal of a few trustworthiness-claims

v00-v01

- * Minor tweaks

v02-v00 of draft-voit-rats-trustworthy-path-routing-00

- * file rename was due to an IETF tool submission glitch
- * The Attester's AIK is included within the Stamped Passport. This eliminates the need to provision to AIK certificate on the Relying Party.
- * Removed Centralized variant
- * Added timing diagram, and moved content around to match

v01-v02 of draft-voit-rats-trusted-path-routing

- * Extracted the attestation stream, and placed into draft-birkholz-rats-network-device-subscription
- * Introduced the Trustworthiness Vector

v00-v01 of draft-voit-rats-trusted-path-routing

- * Move all FlexAlgo terminology to allow passport definition to be more generic.
- * Edited Figure 1 so that (4) points to the egress router.
- * Added text freshness mechanisms, and articulated configured subscription support.
- * Minor YANG model clarifications.
- * Added a few open questions which Frank thinks interesting to work.

Appendix C. Open Questions

(1) When there is no available Trusted Topology?

Do we need functional requirements on how to handle traffic to/from Sensitive Subnets when no Trusted Topology exists between IGP edges? The network typically can make this unnecessary. For example it is possible to construct a local IPSec tunnel to make untrusted devices appear as Transparently-Transited Devices. This way Secure Subnets could be tunneled between FlexAlgo nodes where an end-to-end path doesn't currently exist. However there still is a corner case where all IGP egress points are not considered sufficiently trustworthy.

(2) Extension of the Stamped Passport?

Format of the reference to the 'verifier-certificate-name' based on WG desire to include more information in the Stamped Passport. Also we need to make sure that the keystore referenced names are globally unique, else we will need to include a node name in the object set.

(3) Encoding of objects in CDDL. A Verifier will want to sign encoded objects rather than YANG structures. It is most efficient to encode the Attestation Results once on the Verifier, and push these down via a YANG model to the Attester.

Authors' Addresses

Eric Voit
Cisco Systems, Inc.
8135 Maple Lawn Blvd
Fulton, Maryland 20759
United States of America
Email: evoit@cisco.com

Chennakesava Reddy Gaddam
Cisco Systems, Inc.
Cessna Business Park, Kadubeesanahalli
Bangalore 560103
Karnataka
India
Email: chgaddam@cisco.com

Guy C. Fedorkow
Juniper Networks
10 Technology Park Drive
Westford, Massachusetts 01886
United States of America
Email: gfedorkow@juniper.net

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany
Email: henk.birkholz@sit.fraunhofer.de

Meiling Chen
China Mobile
Email: chenmeiling@chinamobile.com