

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 10 October 2026

K. M. Moriarty
Transforming Information Security LLC
M. Wiseman
Beyond Identity
A.J. Stein

C. Nelogal
Dell
8 April 2026

Remote Posture Assessment for Systems, Containers, and Applications at
Scale
draft-ietf-rats-posture-assessment-04

Abstract

This document establishes an architectural pattern whereby Attestation Results could be produced for a complete set of benchmarks or controls that are defined and grouped by an external entity, eliminating the need to convey individual Evidence items 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). While the discussion below pertains mostly to TPM, other Roots of Trust such as TCG DICE and non-TCG defined components will also be included.

About This Document

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

The latest revision of this draft can be found at <https://ietf-rats-wg.github.io/draft-moriarty-attestationsets/draft-moriarty-rats-posture-assessment.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-rats-posture-assessment/>.

Discussion of this document takes place on the Remote ATtestation ProcedureS (rats) Working Group mailing list (<mailto:rats@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/rats/>. Subscribe at <https://www.ietf.org/mailman/listinfo/rats/>.

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

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 10 October 2026.

Copyright Notice

Copyright (c) 2026 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. Conventions and Definitions	4
3. Posture Assessment Scenarios	4
4. Policy and Measurement Set Definitions	6
5. Supportability and Re-Attestation	7
6. Configuration Sets	8
7. Remediation	8
8. Security Considerations	9
9. IANA Considerations	9
9.1. Reuse of CBOR and JSON Web Token (CWT and JWT) Claims Registries	9
9.2. CWT and JWT Claims Registered by This Document	9
9.3. Additions to the JWT and CWT registries requested	10
9.4. MPS (Measurement or Policy Set) Claim	10
10. Appendix A Extended Claims Table with RoT Variants	11

10.1.	A.1 Chained Attestation and Measurement Exposure Across Hardware Roots of Trust	11
10.2.	A.2 Extended Claims Table with DICE, Apple Secure Enclave, OpenTitan, and Amazon Nitro	12
11.	References	15
11.1.	Normative References	15
11.2.	Informative References	16
	Contributors	16
	Authors' Addresses	17

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, local appraisal of Evidence can be performed to continuously assess compliance.

For example, the Trusted Computing Group's Trusted Platform Module (TPM) format and assessment method can provide this kind of compliance. This and other methods employ a secured log for transparency on the results of the appraised Evidence against Reference Values.

In order to support continuous monitoring of posture assessment and integrity in an enterprise or large data center, local appraisal and remediation are useful to reduce load on the network and remote resources. This is currently done today in measured boot mechanisms.

It is useful to be able to share these results in order to gain a big picture view of the governance, risk, and compliance posture for a network.

As such, communicating a summary result as Evidence including a link to supporting logs within an Entity Attestation Token (EAT) profile [I-D.ietf-rats-eat] provides a way to accomplish that goal.

This level of integration, which includes the ability to remediate, makes posture assessment through remote attestation achievable for organizations of all sizes. This is enabled through integration with existing toolsets and systems, built as an intrinsic capability.

The measurement and policy grouping results summarized in an EAT profile may be provided by the vendor or by a neutral third party, acting as a Reference Value Provider, to enable ease of use and consistent implementations.

The local system or server host, acting as a local Verifier, performs the appraisal of posture Evidence and remediation. 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 remote attestation formats and protocols. The method described allows for defined profiles of policies, benchmarks, and measurements for specific assurance levels. This provides transparency on posture assessment results summarized as Attestation Results.

2. Conventions and Definitions

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. Posture Assessment Scenarios

By way of example, the Center for Internet Security (CIS) hosts recommended configuration settings to secure operating systems, applications, and devices in CIS Benchmarks [BENCHMARKS] developed with industry experts. Attestation Results aligned to the CIS Benchmarks or other configuration guide such as one of the Defense Information Systems Agency's Security Technical Implement Guides [STIG] could be used to assert that 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 [FIRMWARE]. In order to scale remote attestation, a single set of Attestation Results for a set of benchmarks or policies being met, with a link to the verification logs from the local appraisals, is the Evidence that may be conveyed to the Verifier and then the Relying Party. On traditional servers, assurance to NIST SP 800-193 is provable through Evidence generated from a Root of Trust (RoT), using the Trusted Computing Group (TCG) Trusted Platform Module (TPM) chip and attestation formats. However, this remains local and one knows the policies and measurements have been met when other functions that rely on the assurance are running.

At boot, policy and measurement expectations are verified against a set of Reference Values from collected Evidence and are verified to meet expected values. Device identity and measurements can also be attested at runtime, producing Evidence for appraisal. The

generation of Evidence (e.g. hash of boot element) and appraisal of Evidence 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 Evidence appraisal process for boot elements. Event logs and PCR values may be exposed to provide transparency into the appraised Evidence. The Attestation Results defined in this document provide a summary of a local appraisal of posture for managed systems and across various layers (operating system, application, containers) in each of these systems in a managed environment as Evidence. The Relying Party uses the Attestation Results to understand posture of interconnected operating systems, applications, and systems that are communicated in summary Attestation Results.

There is a balance of exposure and evidence needed to assess posture when providing assurance of controls and system state. Currently, if using the TPM, logs and TPM PCR values may be passed to provide assurance that appraised Evidence meets set requirements. Providing the set of Evidence as assurance to a policy set can be accomplished with a remote attestation format such as the Entity Attestation Token (EAT) [I-D.ietf-rats-eat] aligned to the RATS Architecture [RFC9334] and Attestation Results for Secure Interactions [I-D.ietf-rats-ar4si] and a RESTful interface such as ROLIE [RFC8322] or RedFish [REDFISH]. Policy definition blocks may be scoped to control measurement sets, where the EAT profile asserts compliance to the policy or measurement block specified and may include claims with the log and PCR value Evidence. Measurement and Policy sets, referenced in an EAT profile 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 that could be defined in EAT profiles include, but are not limited to:

- * Hardware attribute certificates, TCG
- * Hardware Attribute Certificate Comparison Results, TCG
- * Reference Integrity Measurements for firmware, TCG
- * Operating system benchmarks at Specified Assurance Levels, CIS
- * Application hardening Benchmarks at Specified Assurance Levels, CIS, DISA STIG

* Container security benchmarks at Specified Assurance Levels, CIS

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. Claim sets of measurements and policy verified to meet or not meet Endorsed Values [I-D.ietf-rats-eat] are conveyed in an Entity Attestation Token made available to a RESTful interface in aggregate for the systems managed as Evidence for the Attestation Results. The Measurement or Policy Set may be registered in the IANA registry created in this document (Section 9), detailing the specific configuration policies and measurements required to adhere or prove compliance to the associated document to enable interoperability. Levels (e.g. high, medium, low, 1, 2, 3) or vendor specific instances of the policy defined in code required to verify the policy and measurements would be registered using a name for the policy set, that would also be used in the reporting EAT that includes the MPS along with other artifacts to prove compliance.

4. Policy and Measurement Set Definitions

This document defines EAT claims in the JWT [RFC7519] and CWT [RFC8392] registries to provide Attestation Results for a set of verified claims within a defined grouping. The trustworthiness will be conveyed on original appraised Evidence as well as the Attestation Results on the grouping. The claims provide the additional information needed for an EAT to convey compliance to a defined policy or measurement set to a system or application collecting evidence on policy and measurement assurance, for instance a Governance, Risk, and Compliance (GRC) system.

Claim	Long Name	Claim Description	Format
mps	Measurement or Policy Set	Name for the MPS	
lem	Log Evidence of MPS	Log File or URI	
pcr	TPM PCR Values	URI	
fma	Format of MPS Attestation Results	Format of included Attestation Results	
hsh	Hash Value/Message Digest	Hash value of claim-set	

Table 1

5. Supportability and Re-Attestation

The remote attestation framework shall include provisions for a Verifier Owner and Relying Party Owner to declare an Appraisal Policy for Attestation Results and Evidence that allows for modification of the Target Environment (e.g. a product, system, or service).

Over its lifecycle, the Target Environment may experience modification due to: maintenance, failures, upgrades, expansion, moves, etc..

The Relying Party Owner managing the Target Environment (e.g. customer using the product) can choose to:

- * Update the Appraisal Policy for Attestation Results and re-assess posture with this updated policy, summarizing with a remote attestation to the new policy or level, or
- * Run remote attestation after modification of the Target Environment as external validation, or
- * Continue operation of the Target Environment as-is, without verification, potentially increasing risk

In the case of Re-Attestation:

- * framework needs to invalidate previous Reference Values (e.g. TPM PCR values and tokens),

- * framework needs to specify an Appraisal Policy for Evidence that requires fresh Evidence,
- * 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 Results has been invalidated

6. 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 against the attested configuration set. In this case, the creator of the Evidence appraisal measurements would define a set of values for which a message digest would be created and then signed by the Attester. 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. These configuration sets can be registered for general use to reduce the need to replicate the policy and measurement assessments by others aiming to assure at the same level for a benchmark or hardening guide.

This document creates an IANA registry for this purpose, creating consistency between automated policy and measurement set levels and the systems used to collect and report aggregate views for an organization across systems and applications, such as a GRC platform.

7. Remediation

If policy and configuration settings or measurements attested do not meet expected values, remediation is desirable. Automated remediation performed with alignment to zero trust architecture principles would require that the remediation be performed prior to any relying component executing. The relying component would verify before continuing in a zero trust architecture.

Ideally, remediation would occur on-system as part of the process to generate Evidence for a set of claims, similar to how Evidence is generated 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.

8. 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 Attestation Results could be produced for a complete set of benchmarks or controls as defined and grouped by external entities, preventing the need to convey individual Evidence items 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.

9. IANA Considerations

9.1. Reuse of CBOR and JSON Web Token (CWT and JWT) Claims Registries

Claims defined in this document are a profile of EAT. Like the base claims of EAT, the claims below 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.

9.2. CWT and JWT Claims Registered by This Document

This document requests the creation of a Measurement and Policy Set (MPS) registry. The MPS registry will contain the names of the Benchmarks, Policy sets, DISA STIGS, controls, or other groupings as a policy and measurement set that MAY correlate to standards documents containing assurance guidelines, compliance requirements, or other defined claim sets for verification of posture assessment to that MPS. The MPS registry will include the policy definition for specific levels of MPS assurance to enable interoperability between Attestation Results asserting compliance (or lack thereof) and reporting systems.

MPS Name	MPS Description	File with MPS definition
Ubuntu-CIS-L1	Ubuntu CIS Benchmark, level 1 assurance	http:// /Ubuntu-CIS-L1.txt

Table 2

The MPS name includes versions or level information, allowing for distinct policy or measurement sets and definitions of those sets (including the supporting formats used to write the definitions).

9.3. Additions to the JWT and CWT registries requested

This document requests the following JWT claims per the specification requirement required for the JSON Web Token (JWT) registry defined in RFC7519.

Claim	Long Name	Claim Description
MPS	Measurement or Policy Set	Name for the MPS
LEM	Log Evidence of MPS	Log File or URI
PCR	TPM PCR Values	URI
FMA	Format of MPS Attestation Results	Format of included Attestation Results
HSH	Hash Value/Message Digest	Hash value of claim-set

Table 3

9.4. MPS (Measurement or Policy Set) Claim

The MPS (Measurement or Policy Set) claim identifies the policy and measurement set being reported. The MPS MAY be registered to the MPS IANA registry. The MPS may be specified to specific levels of assurance to hardening, loosening guides or benchmarks to provide interoperability in reporting. The processing of this claim is generally application specific. The MPS value is a case-sensitive string containing a StringOrURI value. Use of this claim is OPTIONAL.

This document requests the following CWT claims per the specification requirement required for the CBOR Web Token (CWT) registry defined in RFC8392.

Claim	Long Name	Claim Description	JWT Claim Name
MPS	Measurement or Policy Set	Name for the MPS	MPS
LEM	Log Evidence of MPS	Log File or URI	LEM
PCR	TPM PCR Values	URI	PCR
FMA	Format of MPS Attestation Results	Format of included Attestation Results	FMA
HSH	Hash Value/Message Digest	Hash value of claim-set	HSH

Table 4

10. Appendix A Extended Claims Table with RoT Variants

10.1. A.1 Chained Attestation and Measurement Exposure Across Hardware Roots of Trust

Platform	Hardware Root of Trust	Measurement Chaining	Attestation Evidence	PCR/ Measurement Exposure
DICE	Minimal hardware + UDS	Yes	Stage certificates, attestation chain	CDIs (not standard PCRs)
Apple	Boot ROM & Secure Enclave	Yes	Internal logs, biometric/auth evidence	Not standard PCRs (internal use)
OpenTitan	Open-source silicon chip	Yes	Logs, certificates	TPM-like PCRs
Amazon Nitro	Nitro Security Chip	Yes	NitroTPM logs, AWS-signed attestation	NitroTPM PCRs

Nitro Enclaves	Nitro Security Chip + Hypervisor	Yes	AWS-signed attestation tokens	Enclave PCRs
----------------	----------------------------------	-----	-------------------------------	--------------

Table 5

10.2. A.2 Extended Claims Table with DICE, Apple Secure Enclave, OpenTitan, and Amazon Nitro

Platform	Claim	Long Name	Claim Description (with RoT Mechanism)	Format
DICE	mps	Measurement or Policy Set	DICE CDI lineage: Minimalist hardware root of trust (UDS in ROM/ SoC), performs measured boot, derives Compound Device Identifiers (CDIs) at each boot stage, binding device and firmware state.	CDI derivation chain; certificate-based attestation
DICE	lem	Log Evidence of MPS	Implementation-specific logs and attestation certificates; evidence is the chain of CDIs and alias certificates, often used in TLS client authentication.	Log file, attestation certificate, or URI
DICE	pcr	TPM PCR Values	Not applicable; DICE does not use TPM PCRs. Instead, it uses derived cryptographic	—

			identities (CDIs) per boot stage.	
Apple	mps	Measurement or Policy Set	Secure Enclave root identity and enclave measurements: Hardware root of trust (immutable Boot ROM in SoC), secure boot of SEP OS, device-unique UID, monotonic counter for anti-replay/rollback, local attestation.	Local attestation (UID + monotonic counter); enclave measurement logs
Apple	lem	Log Evidence of MPS	SEP logs, anti-replay/rollback counters, internal secure boot logs (not externally exposed); used for local security and policy enforcement.	Internal log file or counter
Apple	pcr	TPM PCR Values	Not applicable; Apple does not expose TPM PCRs. Integrity is enforced using UID, monotonic counters, and secure boot chain.	—
OpenTitan	mps	Measurement or Policy Set	Boot-time firmware validation and Creator Identity checks: Open-source silicon root of trust (boot ROM in silicon), measured boot, stable Creator Identity, hardware-backed key management,	Hardware-backed attestation logs and certificates

			remote attestation.	
OpenTitan	lem	Log Evidence of MPS	Firmware integrity logs, certificate-based attestation, logs tied to Creator Identity and measured boot state.	Certificate-based attestation, log file
OpenTitan	pcr	TPM PCR Values	Custom attestation logs; supports TPM-like PCRs for remote attestation, but uses a stable Creator Identity rather than evolving device identity at each firmware layer.	Hardware-backed boot state, TPM-like PCRs
Amazon Nitro	mps	Measurement or Policy Set	Secure boot and firmware validation: Nitro Security Chip as hardware root of trust, NitroTPM provides virtual TPM instance with PCRs, measured boot, and remote attestation for cloud servers and Nitro Enclaves.	NitroTPM PCR values, AWS-signed attestation
Amazon Nitro	lem	Log Evidence of MPS	NitroTPM logs, attestation evidence, logs signed by AWS; used for remote attestation and verification of cloud instance or enclave state.	AWS-signed attestation, NitroTPM log file
Amazon	pcr	TPM PCR	NitroTPM PCR	Cryptographically

Nitro		Values	values (cryptographically signed, used for attestation and sealing secrets to boot state); standard TPM 2.0 interface for EC2 instances and Nitro Enclaves.	signed PCR values, standard TPM 2.0 format
-------	--	--------	--	--

Table 6

11. References

11.1. Normative References

[I-D.ietf-rats-ar4si]
Fossati, T., Voit, E., and S. Trofimov, "Attestation Results for Secure Interactions", Work in Progress, Internet-Draft, draft-ietf-rats-ar4si-07, March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-rats-ar4si-07>>.

[I-D.ietf-rats-eat]
Lundblade, L., Mandyam, G., O'Donoghue, J., and C. Wallace, "The Entity Attestation Token (EAT)", Work in Progress, Internet-Draft, draft-ietf-rats-eat-31, 6 September 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-rats-eat-31>>.

[IANA.CWT.Claims]
IANA, "CBOR Web Token (CWT) Claims", <<https://www.iana.org/assignments/cwt>>.

[IANA.JWT.Claims]
IANA, "JSON Web Token (JWT)", <<https://www.iana.org/assignments/jwt>>.

[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>>.

[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>>.

- [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>>.
- [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>>.
- [RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <<https://www.rfc-editor.org/rfc/rfc9334>>.

11.2. Informative References

- [BENCHMARKS] "Center for Internet Security Benchmarks List", n.d., <<https://www.cisecurity.org/cis-benchmarks>>.
- [FIRMWARE] Regenscheid, A., "Platform firmware resiliency guidelines", National Institute of Standards and Technology, DOI 10.6028/nist.sp.800-193, May 2018, <<https://doi.org/10.6028/nist.sp.800-193>>.
- [REDFISH] "Redfish Specification Version 1.20.0", n.d., <https://www.dmtf.org/sites/default/files/standards/documents/DSP0266_1.20.0.pdf>.
- [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>>.
- [STIG] "Defense Information Systems Agency Security Technical Implementation Guides", n.d., <<https://public.cyber.mil/stigs/>>.

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. IANA section added in version 7 by Kathleen Moriarty, expanding the claims registered and adding a proposed registry to define policy and

measurement sets. Thank you to Henk Birkholz for his review and edits. Thanks to Thomas Fossati, Michael Richardson, and Eric Voit for their detailed reviews on the mailing list. Thank you to A.J. Stein for converting the XMLMind workflow to Markdown and GitHub, editorial contributions, and restructuring of the document.

Authors' Addresses

Kathleen M. Moriarty
Transforming Information Security LLC
MA
United States of America
Email: kathleen.Moriarty.ietf@gmail.com

Monty Wiseman
Beyond Identity
3 Park Avenue
NY, NY 10016
United States of America
Email: monty.wiseman@beyondidentity.com

A.J. Stein
United States of America
Email: ajstein.standards@gmail.com
URI: <https://orcid.org/0000-0003-1092-2642>

Chandra Nelogal
Dell Technologies
176 South Street
MA 01748
United States of America
Email: chandra.nelogal@dell.com