JPEG Status and Progress Report 2026

By Antonio Pinheiro, Frederik Temmermans, Joao Ascenso, Sabrina Caldwell, Saeed Mahmoudpour, Siegfried Foessel, Stuart Perry, Tim Bruylants, and Touradj Ebrahimi

Abstract

This article presents an overview of the most recent specifications produced by the JPEG International Standardization Committee. It also provides an overview of the JPEG Trust and JPEG XE standards, addressing requirements for emerging applications, particularly those arising from recent developments in artificial intelligence. JPEG Trust defines new mechanisms for authenticity, provenance, attribution, intellectual property rights, and integrity of media assets throughout their life cycle. JPEG XE is the first international standard that specifies a lossless coding method for data on visual events, typically generated by event-based vision sensors, enabling interoperability between event sensing, storage, and processing applications. In addition, this article summarizes a recent exploration of radiance-field coding known as JPEG RF.

The JPEG Committee, under the auspices of ISO, IEC and ITU, continues to produce international standards for visual information in response to evolving technologies and applications. Recent activities address emerging challenges in digital media trust through JPEG Trust, event-based vision through JPEG XE, and immersive three-dimensional scene representation as Radiance Fields through JPEG RF.

JPEG Trust addresses the growing need for reliable and interoperable mechanisms to establish trust in digital media. Its framework supports the representation and evaluation of authenticity, provenance, attribution, intellectual property rights, and integrity throughout the life cycle of a media asset, while allowing trust information to be derived from diverse sources and evaluated according to application-specific requirements.

JPEG XE addresses the interoperability challenges created by the increasing adoption of event-based vision sensors. By defining a common representation for visual event-based data, JPEG XE enables interoperable sensing, transmission, storage, and processing while supporting the stringent latency, complexity, and power constraints of sensor and edge implementations. Its initial standardization work has established a core lossless coding system, with later phases targeting higher compression efficiency, including lossy coding.

JPEG RF addresses the emerging requirements associated with radiance-field representations, including technologies such as Neural Radiance Fields and 3D Gaussian Splatting. The activity investigates coding and quality assessment methodologies for these representations, with particular attention to compact and interoperable representations and visual consistency during viewpoint navigation.

These activities complement ongoing work across the broader JPEG standards portfolio, including **JPEG DNA** for data and image storage on DNA support, **JPEG AIC** for image quality assessment, and **JPEG Pleno** for light-field image coding and quality assessment, alongside continued maintenance and extensions of established JPEG standards such as **JPEG 2000**, **JPEG XS** and **JPEG XL**.

JPEG Trust

The rapid development of generative AI and increasingly sophisticated media manipulation technologies has intensified the need for interoperable mechanisms to establish trust in digital media. JPEG Trust (ISO/IEC 21617) addresses this challenge through a comprehensive framework covering authenticity, provenance, attribution, intellectual property rights and integrity throughout the media asset lifecycle. Rather than defining a single model of trust, the JPEG Trust framework, illustrated in **Fig. 1**, accommodates differing real-world trust needs.[1]

A range of Trust Indicators can be derived from a media asset, originating from different sources. These sources can include traditional metadata (for example, EXIF or XMP metadata, including unsigned information such as creation date, camera settings, or author), provenance information (for example, from Content Credentials, including signed assertions about content creation or subsequent modifications), as well as indicators derived from the media content itself (for example, the results of a GenAI detection algorithm). In JPEG Trust, these indicators are collected and serialized together in a Trust Indicator Set.

To make use of this information in a way that reflects the requirements of a particular use case, different actors, including companies, public authorities, organizations, and individuals, can define a Trust Profile containing statements that use and combine Trust Indicators to derive the information relevant to that specific context. The Trust Profile is then evaluated against the Trust Indicator Set derived from a given media asset, resulting in a Trust Report tailored to that

1545-0279/26©2026SMPTE

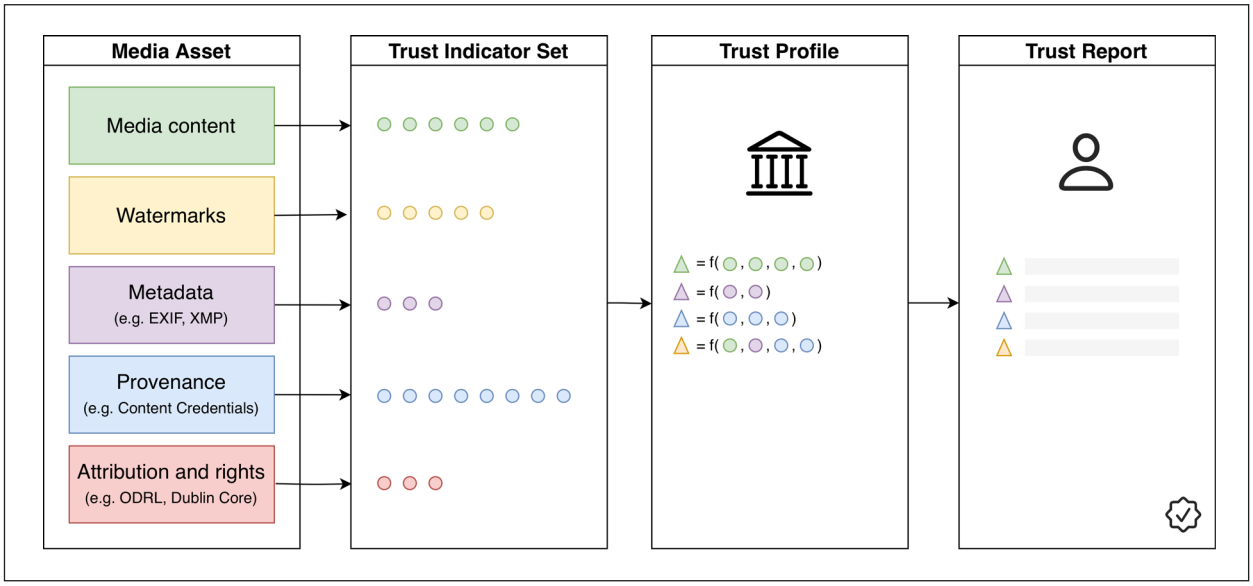

FIGURE 1: The JPEG Trust framework.

context. The Trust Report can be expressed in human- and/or machine-readable form. This architecture enables different communities and applications to establish their own trust requirements while relying on a common interoperable framework **(Fig. 1).**

Significant progress has been made in the ISO/IEC 21617 series of standards. Part 1: Core foundation, second edition, was published in August 2026. The new edition further develops the trust framework and introduces support for attribution and rights declarations. Rights holders can express rights and permissions using W3C ODRL, while established metadata vocabularies including Dublin Core and IPTC metadata can provide complementary descriptive and attribution information. The framework also addresses declarations concerning the use of media for AI/ML training and data mining.

Part 2: Trust Profiles and Reports is, at the time of writing, undergoing FDIS ballot. It extends the Trust Profile and Trust Report mechanisms defined in Part 1 with additional features, including variables, reusable expressions, templates, and metadata blocks, enabling more expressive and reusable profiles. The specification also provides a rich collection of profile snippets illustrating how these mechanisms can be applied across different trust scenarios. Part 3: Media Asset Watermarking is undergoing DIS ballot and specifies a framework for invisible watermarking without standardizing individual watermarking algorithms. It defines signaling mechanisms and their integration with Trust Indicators, Trust Profiles, and Trust Reports, as well as standardized evaluation procedures and metrics for assessing properties such as watermark robustness. This provides a common, reproducible and technology-neutral basis for evaluating and comparing watermarking technologies. Part 4: Reference software is also undergoing DIS ballot, with further implementations planned for future editions.

Interoperability with existing technologies and standards is a central objective of JPEG Trust. Both the first and second editions of Part 1 are aligned with C2PA Content Credentials, using compatible structures for manifests, assertions, claims, and signatures, while extending their use within the broader JPEG Trust framework. At the same time, the framework can be used without provenance information, for example by extracting Trust Indicators from sources such as media content and metadata, and is flexible enough to support various representations of trust information. Ongoing JPEG activities continue the alignment with C2PA, while the committee also investigates collaboration with other relevant standards organizations.

Recent exploration work investigates further support for integrating outputs from media detection mechanisms as Trust Indicators. Similar to the approach taken for watermarking, the objective is not to standardize individual detection algorithms, but to provide interoperable mechanisms through which their results can be represented, evaluated by Trust Profiles and communicated through Trust Reports. This allows provenance, metadata, rights declarations, watermarking, detection results, and other sources of information to contribute to contextual assessments of media trustworthiness.

JPEG XE

JPEG XE has progressed from an exploration activity to a defined core coding system for interoperable event-based vision. The current focus has been on constrained lossless coding for implementation within sensors and at the acquisition edge. Evaluation of 27 test sequences covering diverse applications demonstrates effective compression, while higher-efficiency lossless and lossy coding remain on the roadmap.[2]

Scope and Applications

Event-based sensors asynchronously capture changes in the logarithm of incident light intensity. Each event contains a timestamp, pixel coordinates, and binary polarity. This approach provides microsecond-scale temporal resolution, high dynamic range, low latency, and low power consumption, supporting applications including robotics, autonomous driving, industrial inspection, eye tracking, augmented and virtual reality, and mobile imaging.

The use of proprietary event-data formats limits interoperability between sensors, processing platforms, and applications. JPEG XE provides a common representation for sensing, transmission, storage, and processing, while allowing the encoder to trade compression efficiency against computational and resource requirements.

Development and Standardization Progress

JPEG XE development is organized into three phases. Phase 1 addresses constrained lossless coding with minimal complexity, latency, and power consumption for sensor and edge implementations. Phase 2 targets unconstrained lossless coding with higher compression efficiency for storage and subsequent processing. Phase 3 will address lossy coding for bandwidth-constrained applications, including the development of evaluation methods related to downstream machine-vision performance.

The JPEG Committee issued a draft Call for Proposals for the first two phases in April 2024, followed by the final call in July 2024, supported by use cases, requirements, and Common Test Conditions. The work has involved industry and academia within ISO/IEC JTC 1/SC 29/WG 1, in collaboration with ITU-T SG21, and includes liaison with the MIPI Alliance. The JPEG XE standards are organized as ISO/IEC 26112 and Recommendation ITU-T T.841. Part 1 specifies the core syntax and decoding process and is awaiting publication. Further parts address profiles and levels, reference software, conformance testing, and file formats and metadata.

Phase 1 Codec

The Phase 1 codec combines lossless event representation with ultra-low latency and low implementation complexity. It preserves event types, field values, and chronological ordering while allowing spatial reordering of events with identical timestamps. The design supports sensor-to-host transmission, including over interfaces such as MIPI CSI-2, and implementation on resource-constrained platforms such as microcontrollers and FPGAs.

The codec supports contrast detection events, external triggers, frame synchronization, and metadata such as inertial measurement unit readings. Sensor dimensions up to 4096 × 2048 pixels are supported, with a configurable timescale defaulting to one microsecond. Profiles define syntax subsets and levels impose numerical constraints for different application requirements. The normative decoding process allows encoders to select coding modes according to compression, buffering, and latency requirements.[2]

The codestream uses markers, regular 32-bit messages, and variable-length Huffman-Based Prefix (HBP) messages. Different message types efficiently represent isolated events, local clusters, and dense spatial activity. HBP coding further reduces redundancy by reusing timestamp and coordinate information available at the decoder. An extension mechanism supports vendor information and metadata, while reserved message types and markers provide flexibility for future extensions.

Performance Evaluation

The Common Test Conditions define the sequences, operating scenarios, methodology, and objective metrics used for JPEG XE evaluation. A canonical raw format with 48 bits (6 bytes) per event provides the uncompressed reference. Compression is characterized by average bits per event and compression ratio, defined as: CR = 48 / (bits per event).

The evaluation dataset contains 27 sequences from industrial, automotive, mobile, scientific, and other applications, covering different event densities, sensor resolutions, and temporal activity patterns. The Phase 1 codec achieves compression ratios ranging from 1.48 to 20.68, with an average of 5.18, corresponding to compressed event sizes from 2.33 to 32.53 bits per event, with an average of 15.08 bits per event.[2] Compression efficiency depends strongly on the spatial and temporal structure of the event data, with dense and structured sequences generally providing the largest gains.

Future Work and Outlook

Future work will focus on unconstrained lossless coding using more advanced entropy coding and context modeling, followed by lossy coding for bandwidth-constrained applications. Further activities include conformance testing and support for JPEG XE codestreams in the ISO Base Media File Format. Part 5 is also expected to support synchronization of event timestamps with real-world time, enabling high-precision synchronization between event sensors, conventional cameras, and other data sources. These developments will extend JPEG XE from sensor-level interoperability toward broader storage, transmission, and application support.

Exploration on JPEG RF

JPEG Radiance Fields (JPEG RF) is an exploratory standardization activity within the JPEG Committee that addresses the representation, compression, and evaluation of radiance-field content. Radiance fields describe a scene as a continuous function of position and viewing direction, which allows entirely new perspectives to be rendered on demand. Unlike more classical methods, such as point clouds, radiance fields have a close coupling between the representation and the act of rendering the scene. This facilitates high-fidelity rendering, but adds additional considerations in terms of coding and quality evaluation. Recent radiance field representations, including Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS), have demonstrated significant potential for immersive media, free-viewpoint video, extended reality, cultural heritage, telepresence, and 3D scene communication. However, these representations often require large amounts of data

1545-0279/26©2026SMPTE

and currently lack common coding, evaluation, and interoperability frameworks.

The scope of JPEG RF is to investigate a coding standard for radiance fields that provides a compact representation of radiance field models while enabling immersive depiction of 3D scenes from a set of input images captured from different viewpoints. The intended standard is expected to accommodate modern radiance field generation and rendering techniques, provide efficient compression of radiance field models, and minimize the impact of compression on the visual quality of rendered views. The activity also aims to support the development of a royalty-free baseline for future radiance field coding solutions.

As part of its ongoing exploratory work, JPEG RF has focused on establishing the technical basis for possible future standardization of radiance field coding. The activity has investigated use cases and requirements, common training and test conditions, anchor technologies, subjective assessment procedures, and objective quality metrics for rendered views and navigation trajectories. These efforts aim to support a fair, reproducible, technology-neutral evaluation framework for future compression technologies targeting radiance field representations.

A key aspect of JPEG RF is designing quality assessment methodologies suited to radiance field content, beyond evaluating individual rendered images. Since radiance field content is typically consumed through viewpoint navigation, evaluation must also consider visual consistency across views and along camera trajectories. For this reason, JPEG RF has been developing evaluation methodologies that address the selection of rendered camera trajectories, visualization methods for test stimuli, subjective rating protocols, and objective metrics suitable for immersive and free-viewpoint visual media.

Summary

The activities reported in this article demonstrate the continued expansion of the JPEG family of standards to address significant changes in how visual information is captured, represented, processed, stored, and consumed. While each initiative targets different technical challenges, they share a common objective: establishing interoperable, efficient frameworks that enable emerging technologies to deploy across diverse applications and ecosystems.

JPEG XE, JPEG Trust, and JPEG RF illustrate three complementary dimensions of the JPEG Committee's current work: enabling interoperability for new sensing technologies, establishing trust and transparency in digital media, and preparing standards for emerging representations of visual content. Continued work on coding efficiency, interoperability, conformance, metadata, synchronization, and quality evaluation will be essential to transform these developments into mature and broadly deployable standards.

The JPEG Committee continues to pursue these activities while advancing the broader portfolio of JPEG standards. Through this combination of established standards, new standardization efforts, and exploratory activities, JPEG aims to provide an internationally endorsed framework for interoperable visual information technologies as the requirements of imaging and multimedia systems evolve.

References

1. F. Temmermanns, S. Caldwell, Ph. Rixhon, April 2025, "JPEG Trust White Paper version 2," URL: https://ds.jpeg.org/whitepapers/jpeg-trust-whitepaper.pdf, downloaded 08.10.2026.
2. V. Schwambach, D. Saito, D. Joubert, W. Yang, X. Hu, T. Bruylants, C. Poscha, T. Ebrahimi, J. Ascenso, "JPEG XE: Enabling Interoperable Event-based Vision Systems", *Applications of Digital Image Processing XLIX*, SPIE Optics + Photonics, paper 14217-37, Aug. 2026.

About the Authors

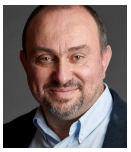

Antonio Pinheiro is a professor at Universidade da Beira Interior and a researcher at Instituto de Telecomunicações, Portugal. He is the chair of the JPEG Communication Subgroup.



Frederik Temmermanns is a guest professor at the Department of Electronics and Informatics (ETRO) associated with the Vrije Universiteit Brussel and IMEC. He chairs the JPEG Systems and Integration subgroup and is the project lead of JPEG Trust.



João Ascenso is a professor at University of Lisbon and researcher at Instituto de Telecomunicações. He currently chairs the Coding and Performance for Machines subgroup and the JPEG AI Ad-Hoc Group.



Sabrina Caldwell is a senior lecturer in the School of Computing at the Australian National University, where she uses biometric signal processing to investigate affective reasoning. She is a co-chair of JPEG Trust Ad hoc Group.



Saeed Mahmoudpour is a guest professor at VUB-imec, Belgium, chairs JPEG Radiance Fields, and leads JPEG Pleno Light Field Quality Assessment.



Siegfried Foessel is Head of Moving Picture Technologies at Fraunhofer IIS in Erlangen and is a professor at HFF Film School in Munich.



Stuart Perry is a professor of engineering and IT at the University of Technology Sydney. He is the Co-Director of the Perceptual Imaging Laboratory (PILab) and the Co-Chair of the JPEG Plenoptic Coding and Quality Subgroup.



Touradj Ebrahimi is a professor of image processing at Ecole Polytechnique Fédérale de Lausanne (EPFL) and the current Convenor (Chairman) of the JPEG Committee.

A contribution received from the authors for the Progress Report.

Copyright © 2026 by SMPTE.

DOI: 10.5594/JMI.2026/QNAJ9845
Date of publication: 15 November 2026